TECHNISCHE UNIVERSITÄT MÜNCHEN

TUM School of Computation, Information and Technology

Improving the Applicability of Differential Privacy in Data Sharing and Analytics Applications

Gonzalo Munilla Garrido

Vollständiger Abdruck der von der TUM School of Computation, Information and Technology der Technischen Universität München zur Erlangung eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz:		Prof. Dr. Stefanie Rinderle-Ma
Prüfer der Dissertation:	1.	Prof. Dr. Florian Matthes
	2.	Prof. Dr. Gilbert Fridgen

Die Dissertation wurde am 17.04.2023 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information and Technology am 21.02.2024 angenommen.

Abstract

Motivation: As the internet has evolved, so too have the methods for collecting and processing personal information. From static websites to social media platforms and mobile devices, the scope of accessible data attributes and their potential misuse has expanded. In response, researchers and practitioners have focused on developing privacy-enhancing technologies (PETs) to balance privacy and utility, allowing for the extraction of valuable information from user data while preserving privacy to a controllable degree. Among PETs, differential privacy (DP) has become the golden standard, providing a provable privacy guarantee that bounds the amount of new information an attacker can gain from observing a function's output. Despite their promise, many PETs, including DP, are not yet ready for widespread adoption in practice. In this dissertation, we aim to help tackle the lack of widespread adoption of PETs, particularly DP, by improving their applicability in data sharing and analytics applications (DSAA).

Research Design: Throughout this dissertation, we employed various research methods to address the identified gaps in the literature. These methods included *systematic literature reviews*, *expert interviews*, and *design science*.

Contribution: This dissertation presents three contributions aimed at improving the applicability of PETs, particularly DP, in practice. (I) The first contribution reveals opportunities and challenges in the applicability of PETs in DSAAs. (II) The second contribution helps in improving the applicability of DP algorithms, and (III) the third contributions delves into improving the applicability of DP systems. Together, these contributions help advance the state of the art in PETs and DP in particular, and facilitate their wider adoption in practice.

Results: In this dissertation, we draw on the results of four publications to contribute to the field of PETs and DP. Through systematic literature reviews, we identified encryption, secure and outsourced computation, and anonymization and plausible deniability as the main techniques used in PETs. We also identified the challenges associated with applying PETs and found no established solution for building privacy-enhancing DSAAs, highlighting the need for improved applicability of PETs. Additionally, through expert interviews, we studied the use cases where PETs can be most valuable and provided guidance on how researchers can select the appropriate technology based on the characteristics of the use case. As a result of these initial findings, we proposed a verifiable DP algorithm that uses zero-knowledge proofs to attest to the correctness of a DP query output while maintaining practical performance. We also explored the limitations of DP tooling and highlighted the gaps that practitioners need to address to bring DP into practice. Our work also contributes a range of artifacts, including classifications, mapping, non-functional requirements, and two software components.

Limitations: The findings presented in this dissertation may be subject to certain limitations, including those related to internal validity, external validity, construct validity, and reliability. It is important to carefully consider these limitations and interpret the results accordingly. To minimize potential threats to the validity of the findings, we implemented appropriate countermeasures when possible.

Future Research: We outline the potential for further research in each of the three contribution streams and discuss ongoing work moving Contribution III forward and establishing a new research stream focused on improving the applicability of differential privacy in virtual reality applications. This section presents opportunities for researchers and academics to investigate applied DP and suggests avenues for further research, such as studying systematic methods for selecting suitable PETs, exploring new combinations of PETs, and designing models to quantify data leakage in data exchanges and analytics.

Zusammenfassung

Motivation: Mit der fortschreitenden Verbreitung des Internets verändern sich auch die Methoden zum Sammeln und Verarbeiten von persönlichen Informationen. Von statischen Websites über soziale Medienplattformen bis hin zu mobilen Geräten hat sich die Datenmenge und auch ihr möglicher Missbrauch stetig ausgeweitet. Um dem entgegenzuwirken, entwickelten Forscher und Praktiker Privacy-Enhancing Technologies (PETs). Diese balancieren den Datenschutz und die Nützlichkeit von verarbeiteten Daten aus. Dadurch können wertvolle Informationen aus Nutzerdaten extrahiert werden, doch der Datenschutz bleibt in einem kontrollierbaren Maße erhalten. Aus unterschiedlichen PETs entwickelte sich *Differential Privacy (DP)* zum Goldstandard. Die Technologie liefert eine nachweisbare Privatsphärengarantie, die die Menge an neuen Informationen begrenzt, die ein Angreifer aus der Beobachtung des Ergebnisses einer Funktion gewinnen kann. Trotz ihres Versprechens sind viele PETs, einschließlich DP, noch nicht bereit für eine breitere Anwendung in der Praxis. In dieser Dissertation möchten wir das Problem des mangelnden Einsatzes von PETs, insbesondere DP, in *Data Sharing and Analytics Applications (DSAA)* beleuchten.

Forschungsdesign: In dieser Dissertation verwenden wir die Forschungsmethoden *Systematische Literaturüberprüfungen*, *Experteninterviews* und *Design Science*, um die von uns identifizierten Lücken in der Literatur zu schließen.

Ergebnisse: Diese Dissertation stellt drei Beiträge vor, die darauf abzielen, die Anwendbarkeit von PETs, insbesondere DP, in der Praxis zu verbessern. (I) Der erste Beitrag zeigt Chancen und Herausforderungen bei der Anwendung von PETs in DSAAs auf. (II) Der zweite Beitrag hilft dabei, die Anwendbarkeit von DP-Algorithmen zu verbessern, und (III) der dritte Beitrag befasst sich mit der Verbesserung der Anwendbarkeit von DP-Systemen. Diese Beiträge tragen insgesamt zur Weiterentwicklung des Standes der Technik bei PETs und DP und erleichtern ihre breitere Anwendung in der Praxis.

Beitrag: In dieser Dissertation nutzen wir die Ergebnisse von vier unserer Veröffentlichungen, um zum Forschungsfeld der PETs und DP beizutragen. Durch systematische Literaturübersichten identifizierten wir Verschlüsselung, sichere und ausgelagerte Berechnung sowie Anonymisierung und plausible Leugnung als die Haupttechniken, die in PETs verwendet werden. Wir identifizierten auch die Herausforderungen im Zusammenhang mit der Anwendung von PETs und fanden dabei keine etablierte Lösung für DSAAs mit verbessertem Datenschutz. Dies verdeutlicht die Notwendigkeit einer verbesserten Anwendbarkeit von PETs. Zusätzlich haben wir durch Experteninterviews untersucht, in welchen Anwendungsfällen PETs am wertvollsten sind und Anleitung dazu gegeben, wie Forscher die geeignete Technologie basierend auf den Eigenschaften ihres Anwendungsfalls auswählen können. Basierend auf diesen ersten Ergebnissen haben wir einen überprüfbaren DP-Algorithmus vorgeschlagen, der *Zero-Knowledge Proofs* verwendet, um effizient die Korrektheit eines DP-Abfrageergebnisses zu belegen. Weiterhin haben wir die Grenzen von DP-Tools untersucht und Lücken identifiziert, die in der Praxis angegangen werden müssen, um DP einzusetzen. Unsere Arbeit umfasst Klassifikationen, Mapping, nicht-funktionale Anforderungen und zwei Softwarekomponenten.

Limitationen: Die Ergebnisse dieser Dissertation können gewisse Limitierungen aufweisen, beispielsweise in der Übertragbarkeit auf andere betrachtete Anwendungsfälle. Es ist wichtig, diese Einschränkungen sorgfältig zu berücksichtigen und die Ergebnisse entsprechend zu interpretieren. Um potenzielle Limitierungen bei der Validität unserer Ergebnisse zu minimieren, haben wir geeignete Gegenmaßnahmen ergriffen.

Ausblick: Wir stellen Möglichkeiten für künftige Forschung in jedem der drei Beitragsthemen vor und diskutieren laufende Arbeiten zur Weiterentwicklung von Beitrag III. Darüber hinaus stellen wir einen neuen Forschungsbereich vor, der sich auf die verbesserte Anwendbarkeit von DP in Virtual-Reality-Anwendungen konzentriert. Dieser Abschnitt präsentiert Möglichkeiten für Wissenschaftler, angewandte DP zu untersuchen und schlägt Wege für weitere Forschung vor, wie z.B. die systematische Untersuchung geeigneter PETs, die Erkundung neuer Kombinationen von PETs und die Gestaltung von Modellen zur Quantifizierung von Datenlecks bei Datenaustausch und -analyse.

Acknowledgment

This dissertation emerged from my work at the sebis chair at TUM, with my industry partner: the BMW Group, and my research period at UC Berkeley's Computer Science department. I am grateful for all the fantastic people I have met throughout this journey and all I learned from them.

I am deeply grateful to Prof. Dr. Florian Matthes for his invaluable support and guidance as my doctoral father and supervisor. I appreciate the research freedom he has provided, as well as the time and care he dedicated to shaping my dissertation. Additionally, I would like to thank Prof. Dr. Gilbert Fridgen and Prof. Dr. Stefanie Rinderle-Ma for reviewing and supervising my dissertation and offering guidance during the concluding phase of my academic journey. Furthermore, I am grateful for the friendship and support of Sascha Nägele and Oleksandra Klymenko, who have made my Ph.D. journey an enjoyable and memorable experience.

I am grateful to all my friends and colleagues at the sebis chair for their support, including Gloria Bondel, Ulrich Gallersdörfer, Ingo Glaser, Nektarios Machner, Pascal Maurice Philipp, and Fatih Yılmaz. I am likewise thankful to my colleagues at BMW for showing me the ropes of the industry and for their support, including Geraldine Bous, Ann Christin, Djerekarov Emil, Eduard Hez, Johannes Klepsch, Andre Luckow, Lukas Müller, and Philipp Ross. I would also like to thank my friends and colleagues at UC Berkeley for making my stay there a wonderful experience and teaching me so much. Special thanks to Ishaq Aden-Ali, Seri Khoury, Norman Mu, Julien Piet, Vamsi Policharla, Mayank Rathee, and Deevashwer Rathee.

Among my co-authors, I would like to send my special thanks to Johannes Sedlmeir, who has taught me so much about what it means to be a Ph.D. student and has been an incredible source of support and guidance throughout the key publications in this dissertation. I am also grateful to Joseph P. Near, with whom I submitted my first paper and who has demonstrated incredible knowledge, kindness, and patience in my research endeavors. I have greatly enjoyed working with Xiaoyuan Liu, from whom I have learned a lot about systems and with whom I have had countless enlightening and enjoyable discussions. I am honored he contributed to one of this dissertation's key publications. I am also grateful to Vivek Nair for his determination, time management, and passion for technology, which have been incredibly inspiring and have made me grow as a computer scientist. Christopher Harth-kitzerow's adeptness in developing software and our discussions about technology, privacy, and life have been a great source of energy and inspiration. Lastly, I am thankful to Prof. Dawn Song for the confidence she placed in me by bringing me to UC Berkeley and for her valuable feedback on my publications. I could not have asked for better co-authors and I hope our paths cross again in the future. I wish them all the very best in their future endeavors.

I am grateful to the students who wrote their theses under my guidance or assisted me in my research as student assistants. I want to express my appreciation for their dedication.

I would like to extend my sincere thanks to my family, particularly my parents Maria Pilar Garrido and Santiago Pablo Munilla, my sister Elena Isabel Munilla, my grandmother María Angeles López, and my aunts and uncles for their unwavering support throughout my life. Their

love and encouragement have been a constant source of strength for me, and I am forever thankful for their presence in my life. Thank you, family.

Disclaimer: The information and views set out in this dissertation are those of the author and do not necessarily reflect the official opinion of the institutions.

Table of Contents

List of Figures

List of Tables

Part A

CHAPTER 1

Introduction

This dissertation investigates the opportunities and challenges of privacy-enahncing technology (PET) in data sharing and analytics applications (DSAAs) and delves deeper into improving the applicability of one of these technologies, differential privacy (DP), in such applications. The embedded papers comprising this cumulative dissertation identify and fill the research gaps by systematically reviewing the state-of-the-art of PETs and DP, and conducting expert interviews to identify their use cases and adoption roadblocks. Additionally, we outline the outstanding gaps in DP tooling, elicit a series of non-functional requirements to fill such gaps, and develop two software components to improve the applicability of DP in the broader industry. We motivate the dissertation in section 1.1, introduce the research questions that guide this dissertation and discuss related work in section 1.2, summarize the embedded papers and outline the contributions in section 1.3, and present the structure of the dissertation in section 1.4.

1.1. Motivation

Through the mists of rapid changes in technology paradigms and consumer preferences throughout the life of the internet, a clear pattern has remained true: with every new technology and medium for accessing the internet come new and enhanced methods for collecting and processing personal information. From static websites in the early stages of the internet, to social media platforms in the early 2000s and mobiles phones, smart wearables, and virtual assistants more recently, users have experienced an increase in the scope of accessible data attributes and their unwarranted use for purposes other than the strictly required [NGS23]. Thus, as new internet technology expands the attack surface on user privacy, a suite of defensive technologies must correspond to these new threats, the so-called PETs. In short, PETs are tools designed to balance privacy and utility, i.e., they enable the extraction of valuable information from user

data while preserving their privacy to a (potentially adjustable) degree. Thus, privacy becomes a continuum rather than a binary state of a user's information disclosure. At one extreme, data remains siloed and never shared, which would throttle innovation and the economy; at the other end, data is open to the scrutiny and use of any entity, which would incur profound social consequences and political unrest. Striking a balance between the two extremes is a delicate endeavor, and consequently, PETs require special attention in their deployment in practice. However, despite the efforts from academics and researchers, many of the most promising PETs are not yet ready for broad adoption.

Among PETs, one particular technique has become the golden privacy standard in academia: differential privacy (DP), which Dwork et al. [DMNS06] introduced in 2006. DP is a context-agnostic provable privacy guarantee that bounds the amount of new information an attacker can gain from observing a function's output. A function (e.g., an aggregation) fulfills DP by adding calibrated random noise, typically to its deterministic outputs. In effect, with DP, the likelihood of obtaining the same output is essentially the same with or without the presence of an input data element (e.g., an individual's *age*). In other words, as the outputs are unlikely to diverge significantly as the input changes by one element, the presence or absence of a particular individual in the input is protected, providing plausible deniability. Moreover, one feature that makes DP so appealing is the possibility to tune such indistinguishability with a parameter ex-ante, which, depending on its value, privacy is more or less protected. Given this unprecedented mathematically quantifiable formulation of privacy, DP holds much promise. However, its full potential has yet to be realized.

A few organizations managed to deploy DP in productive environments [App17, AG21, DKY17, Joh21], demonstrating its value. Nevertheless, the vast majority of organizations have not adopted DP, despite such examples and the plethora of open-source tools available [Goo21a, Goo21b, Goo20, WZL+19, Goo19, Har21, GHV20, IBM20, Met21, Ope22, Tum21, NH12, RSK+10, BBG+21, HMM+16, JDL+21, TBG+20a, GHK+16, NBH+22]. Thus, a gap between academia and practice needs to be closed for DP to become widely adopted in the industry, which is the goal this dissertation helps tackle.

In this dissertation, after building an understanding of the state-of-the-art of PETs and DP, as well as their challenges, and opportunities, we tackle a set of research gaps to improve the applicability of DP in DSAAs. In particular, we think practitioners unfamiliar with privacy must be able to seamlessly use DP tooling, which still requires the improvement of the available open-source DP libraries and frameworks. Additionally, we find it essential to verify the correct execution of DP, which requires a detailed examination and careful adaption of existing differentially-private noise sampling algorithms. With this effort, organizations can be more willing to adopt DP as their privacy guarantee, and users can increase their trust in such systems. In summary, in this dissertation, we identify the need for applicability improvements in DP tooling and algorithms and realize a set of these improvements to help push the adoption of DP in the industry.

1.2. Research Questions & Related Work

This section highlights the research gaps in related work, which this dissertation aims to fill by answering the hereby presented research questions.

Research Gap 1: *There is a lack of comprehensive yet detailed systematic analysis of the body of knowledge, use cases, and challenges of privacy-enhancing technology.*

A few studies have provided a comprehensive overview or classification of PETs [SCS18, PHS+19a, TBG+20b], but they laack the depth and systematic rigor needed to fully understand the challenges and opportunities of these technologies. Among other less-comprehensive studies that organized existing knowledge [LYA+18, SLZ20, ZMWC19, DJG+18, PRW15, PGUXS16, ZMW14], some focused only on challenges of DSAA [ZMWC19, PRW15, ZMW14], while others reviewed a subset of PETs from a technical perspective and outlined their challenges and opportunities [DJG+18, PGUXS16]. Other studies delved into user privacy preferences [SLZ20] and technical design requirements [LYA+18] for DSAAs. Furthermore, some studies [DJG+18, PGUXS16, SCS18] briefly discussed secure and outsourced computation, syntactic privacy definitions like k-anonymity, and DP, and outlined their implications without regard for other existing PETs. For example, J. Pennekamp et al.'s [PHS+19a] overview of PETs and their associated challenges lacks depth, as they defined concepts briefly, presented a subset of the challenges we found during our research, and drew their results from use cases, and therefore, cannot provide the rigor of a systematic search and review of the extant literature. The rest of the aforementioned studies provided privacy policies, commented on digital rights, scratched the surface of PETs challenges, or provided a user perspective on DSAAs. Something to note is that only some of the technologies included in these extant reviews enhance privacy. Particularly, these other technologies focus on authenticating information, e.g., distributed ledger technologies and version controls, which some researchers have confounded with PETs [LF20, DKJG17, DJG+18].

Regarding PETs and classifications of DSAAs, some of the above papers contained frameworks. For example, S. Sharma et al. [SCS18] classified PETs into *information sharing* and *outsourced computation*, and Pennekamp at al. [PHS+19a] proposed a structure for the layers of DSAAs: *data security*, *data processing*, *proving support*, *platform capabilities*, and *external measures*. Moreover, A. Trask et al. [TBG+20b], heavily inspired by Nissenbaum's contextual integrity [Nis09], looked at PETs from the lenses of an information flow segmented into *input*, *computation*, and *output* steps, and assessed *privacy*, *verifiability*, and *governance* at each step. We enhanced and expanded their frameworks for our classifications and mappings.

Looking at other extant studies, we found works focused on a single PET. Most of such studies either applied the technology to a specific use case, optimized the solution in terms of performance, or added new guarantees. For example, there is work on applying secure multiparty computation (MPC) to train deep learning models [BCD21], apply DP to health data [CGDS+20], or map trusted execution environments (TEE) to different use cases [AGT14]. However, these publications do not provide an overview of PETs or their associated privacy use cases. In contrast, other researchers stayed at a high level by surveying general privacy requirements and how PETs can fulfill them [HZNF15, PHS+19b], highlighted the business problems that PETs can tackle (e.g., building trust and competitive advantage [JTBN12]), and proposed industry use

cases without explicitly mapping them to PETs (e.g., for the supply chain [GKHD20], predictive maintenance in the automotive industry [TPVKE21], and smart homes [CAEK19]). Moreover, a few researchers have produced surveys of PETs and their applications; namely, a repository with use cases across different industries like health, finance, and transportation [CDE21], and case studies that leveraged PETs in the financial sector [FFI20]. However, these publications did not identify suitable technology capabilities to map use cases to PETs systematically.

Summary. Overall, the extant literature had not provided a systematic, holistic, and detailed overview of the scientific body of knowledge in PET, which signals a research gap and the lack of academic rigor of the limited-scope, extant reviews on PETs. We address this gap by conducting a systematic literature review of PETs and their challenges and draw a mapping to use cases. We start by identifying the most prominent PETs, study the literature to outline the outstanding challenges, and interview experts to reveal the use cases where PETs can contribute the most. Against this backdrop, we formulate the first research question:

Research question 1 (RQ1)

What are the most relevant privacy-enhancing technologies and their corresponding challenges and use cases in the scope of data sharing and analytics applications?

Answering RQ1, revealed the two remaining gaps we tackle during this dissertation.

Research Gap 2: *There is a lack of cryptographic primitives and applications that combine the unique properties of secure computation and techniques that offer formal privacy guarantees.*

Despite researchers acknowledging the benefits of combining *anonymization* with *secure and outsourced computation* techniques, it had been rare to find such publications. At the time of answering RQ1 partly with a systematic literature review, only two of the 37 collected studies—focused on implementations—designed a DSAA using both technology types. For example, some techniques enable data and computation verification, others can hide the inputs and the computation itself, while some can protect the outputs; in combination, the application would be close to achieving an end-to-end privacy-enhancing data sharing and analysis workflow. In particular, we found it critical to verify the correct use of formal privacy guarantees that require randomness (e.g., with a cryptographically verifiable process), as proving the correct execution of a stochastic process to third parties is not trivial and has been understudied in the literature.

While there exists a significant body of research in formal privacy guarantees (DP) [HLM17a, KSK14, KOV14, HKR12, HLM17b, BV19] and in verification technologies (zero-knowledge proofs (ZKP)) [GOS06, BSCG+13, BSBHR19], there are only a few that connect them. Rückel et al. [RSH22] designed an architecture to verify the sampling of differentially-private noise in a federated learning setting. However, the authors did not acknowledge that their solution only fulfills a particular type of DP, omitted the corresponding bound of a critical parameter (δ), and did not contemplate a high-precision approach as they used an approximate inverse cumulative distribution function. Moreover, Tsaloli et al. [TM19] limited their work to providing a high-level motivation for using ZKP for verifiable DP without an associated implementation or design. Furthermore, although Kato et al. [KCY21] elicited details on the creation of randomness with a similar technology (MPC), they did not have a verifiability step, i.e., there is

no cryptographic check of whether the value is truthful prior to noise addition. Lastly, while Narayan et al. [NFPH15] discussed the upsides of verifiable DP, they did not provide details about their ZKP-based implementation, e.g., the specifics of rounding or achieving accuracy guarantees. Additionally, their implementation has impractical performance, requiring 2 hours of proof generation for 32 servers.

Summary. Altogether, there is no proposal for a performant verification of DP in the literature, a gap we aim to fill by exploring, combining, and improving existing algorithms related to sampling differentially-private noise so that ZKP can verify their execution. Given this background, we aim to answer the following research question:

Research question 2 (RQ2)

What new algorithm can verify the use of formal privacy guarantees with practical performance and bounded guarantees?

Research Gap 3: *Researchers and industry practitioners have not widely adopted existing open-source formal privacy-enhancing tools for developing their work.*

As with research gap 2, answering RQ1 also revealed the third gap we aim to tackle in this dissertation: the studied publications rarely used open-source libraries or built upon previously peer-reviewed systems. This lack of adoption is the case despite the multiple open-source tools providing MPC, DP, ZKP, and homomorphic encryption (HE) functionality. Building upon our knowledge in DP obtained by answering RQ2 and because of the increased maturity of DP primitives relative to other PETs, i.e., practical DP is within striking distance, we targeted the open-source libraries, frameworks, and systems dedicated to lowering the entry barrier of DP to help with the last yet critical development efforts. Particularly, we aim to guide library designers and privacy practitioners by outlining the remaining challenges to bridge the gap between theory and practice in DP.

Adopting DP in productive environments is possible, as Apple [App17, AG21], Google [AG21], Microsoft [DKY17], and the US Census Bureau [Joh21] have demonstrated. However, no other organizations have showcased the use of DP. The options are ample and have been designed by a diverse set of institutions: Google [Goo21a, Goo21b, Goo20, WZL+19, Goo19], Harvard [Har21, GHV20], IBM [IBM20], Meta [Met21], OpenMined [Ope22] (experimental product), Tumult Labs [Tum21], and the University of Pennsylvania [NH12], and Texas [RSK+10]. Additionally, some open-source tools are focused on visualizations: Bittner et. al [BBG+21], DPcomp [HMM+16], DPP [JDL+21], Overlook [TBG+20a], PSI (Ψ) [GHK+16], and ViP [NBH+22]. In this dissertation, we distill a set of key requirements from this collection of tools that a holistic and formal privacy-enhancing analytics tool should fulfill and outline the remaining gaps.

Concerning existing work related to the applicability issues of current DP tooling, we find surveys of DP applications in social networks [JPY+21], a user survey regarding privacy in data sharing applications [CZ19], cyber physical systems like the IoT [HRC20], location-based services [KEK+21], statistical learning [SWW+20], and lessons learned from employing DP in the US Census [GAP18]. Prominently, Kifer et al. [KMR+20] compiled best practices and lessons learned from their experience implementing DP systems at Meta, an organization familiar with

DP. In contrast, to fill research gap 3, we focused on companies unfamiliar with DP and discussed whether DP has potential in the broader industry. Lastly, Dwork et al. [DKM19] produced the closest work to ours. We differentiate from their work in that we interview practitioners without an in-depth knowledge of DP. Furthermore, we note that if DP is widely adopted, it will depend on these practitioners and their organizations unfamiliar with the technicalities of DP, which are the vast majority.

Summary. Overall, we aim to fill this research gap by conducting interviews at organizations without knowledge in DP to qualitatively understand the practicality and adaptability of DP to legacy data sharing and analytics workflows. With the learnings and our understanding of the available DP tooling, we elicit a set of key system desiderata for holistic DP tools and draw the attention of library designers to the remaining gaps in their tools.

Research question 3 (RQ3)

Which are the existing gaps and key requirements a formal privacy-enhancing analytics tool should fulfil to become closer to a broad industry adoption?

1.3. Contributions & Publications Summary

This dissertation resulted in **four first-author publications providing three main contributions to answer the three research questions**. Outside this dissertation, we expand these contributions with additional related publications comprising one equal-contribution publication, two first-author and one equal-contribution pre-prints (PP), and four second-author publications (2AP). Overall, we contributed a total of 12 manuscripts to the research community. The rankings of the publication outlets were measured on 07.12.2022 using the conference ranking tools CORE [Edu] for P1, P3-4, 2AP1, and 2AP4 and the journal ranking tool Scopus Preview [Sco] in the subject of *computer science applications* for P2. For the publications not ranked in CORE (2AP2-3), we used ERA-sourced conference ranks [Sch]. We describe this dissertation's contributions and summarize the papers in the following.

1.3.1. Dissertation Contributions & Publications

We provide in Table 1.1 an overview of the publications associated with the three dissertation's contributions.

The contributions are three-fold:

Contribution I: Revealing Opportunities and Challenges in the Applicability of Privacy-Enhancing Technologies in Data Sharing and Analytics Applications.

Given the lack of comprehensive, detailed, and systematic analyses of the literature regarding the use of PETs (research gap 1), we conducted a systematic literature review [fGS+21] and unstructured expert interviews [GSU+22] to shed light on the current opportunities and challenges of applying PETs to DSAAs—answering research question 1.

RQ	No.	Title	Outlet	Type	Ranking	Pages
Contribution I: *Opportunities and Challenges in the Applicability of PETs*						
RQ1	P1	Exploring privacy-enhancing technologies in the automotive value chain†	Big Data	CON	**B**	**8**
RQ1	P2	Revealing the Landscape of Privacy-Enhancing Technologies in the Context of Data Markets for the IoT: A Systematic Literature Review	JNCA	JNL	**Top 3%**	**43**
Contribution II: *Improving the Applicability of Differential Privacy Algorithms*						
RQ2	P3	Towards Verifiable Differentially-Private Polling	ARES	CON	**B**	**10**
Contribution III: *Improving the Applicability of Differential Privacy Systems*						
RQ3	P4	Lessons Learned: Surveying the Practicality of Differential Privacy in the Industry	PETS	CON	**A**	**20**

†Short paper. Abbreviations: P = Publication, CON = Conference, JNL = Journal. Outlet: Big Data = 2021 IEEE International Conference on Big Data (Big Data), JNCA = Journal of Network and Computer Applications, ARES = 17th International Conference on Availability, Reliability and Security (2022), PETS = 23rd Privacy Enhancing Technologies Symposium (2023). Ranking: Conferences use CORE and journals use Scopus Preview.

Table 1.1.: Overview of (embedded) dissertation publications.

In our first work (P1) [fGS+21] primarily focuses on opportunities in the application of PETs in DSAA, namely, in the most promising industry use cases where PETs can make a difference. We conducted unstructured expert interviews to identify such opportunities and searched for open-source tools that practitioners can readily use to start deploying PETs in the identified use cases. Moreover, we selected a few reference use cases to elaborate on the role and importance of PETs and elicited a series of technology capabilities required by use cases so that mapping PETs to use cases becomes easier for practitioners. The concrete output artifacts of our publication are (i) a list of use cases classified per application domain, (ii) a list of the most prominent PETs and associated open-source tools, (iii) a mapping between PETs and use case characteristics, and (iv) a mapping between PETs and selected use cases per application domain.

Complementary to P1, P2 [GSU+22] sheds light on the broad and narrow challenges in the application of PETs in DSAA. We conducted a systematic literature review to classify PETs and outline the remaining challenges for the adoption of PETs in DSAAs, providing a starting point ofr practitioners looking to improve the applicability of PETs. The specific output artifacts of this work are (i) a taxonomy of PETs, (ii) privacy challenges in IoT data markets, (iii) an examination of the IoT's negative impacts on privacy, (iv) a mapping between PETs and the associated privacy challenges, (v) a reference model for a privacy-enhancing IoT data market, and (vi) a metadata aggregation of the selected publications.

Contribution II: Improving the Applicability of Differential Privacy Algorithms.

Motivated by the lack of studies bridging formal privacy guarantees and cryptographic verification tools, we designed, implemented, and evaluated an algorithm capable of verifying the use of DP—answering research question 2.

Given the challenges and opportunities outlined by P1 and P2, P3 [MGSB22] tackled the challenge of the verifiability of DP algorithms. In particular, we explored the literature to find

algorithms related to the problem and proposed and adapted a series of algorithms until we found the logic we could successfully verify cryptographically. The output artifacts of P3 consist of (i) an algorithm to cryptographically verify the use of DP and (ii) its open-source implementation.

Contribution III: Improving the Applicability of Differential Privacy Systems.

The gap between the theory and practice of DP is closing; however, there are remaining critical challenges that require the attention of library designers and practitioners. Thus, we conducted expert interviews and outlined the future work items to help close the gap—answering research question 3.

Motivated by the challenge revealed in Contribution I, P4 [GLMS23] contains the results of structured expert interviews regarding the usefulness of DP in practice and examined the state-of-the-art of open-source DP tooling to highlight the aspects library designers should work on to bridge the gap between theory and the practical use of DP in the industry. The artifacts of this publication are (i) a set of ten functional requirements for holistic and formal privacy-enhancing analytics systems, (ii) a list of outstanding research and engineering gaps blocking the wider adoption of DP, (iii) an early-stage blueprint for the design of a holistic privacy-enhancing analytics system, (iv) an early-stage open-source implementation of a benchmarking tool to test the suitability of DP libraries and frameworks for some of the blueprint's components.

1.3.2. Additional Related Publications

Adding to the embedded publications of this dissertation, we wrote eight additional related publications—many led by co-authors—expanding the answers to the three research questions and adding a new research stream (see Table 1.2). Even though the insights of these additional publications help expand the answers to the research question, we selected P1-4 as the primary research contributions for this dissertation.

Expanding Contribution I

A series of second-author publications expand the search for opportunities and challenges in the use of PETs in the DSAA. Specifically, we improved further the mappings between PETs and challenges of P2 (2AP1) [SMGMM22], proposed future work avenues in developing systems for the de-identification of datasets (2AP2) and in evaluating such a tool [BGBM20], (2AP3) [BGBM], and propose a verification scheme for outsourced computation (2AP4) [HKMG].

Expanding Contribution III

We upgraded the benchmarking tool initially designed in P4 from an early-stage prototype to a more comprehensive system. With such a system, by selecting an input dataset and a query, the tool indicates the user which is the best-performing library in terms of accuracy, precision, and execution time [GNM+21]. Additionally, we evaluated the most prominent open-source tools based on these metrics and highlighted areas of future work for library designers.

New Research Stream: Improving the Applicability of Differential Privacy in Virtual Reality Applications.

RQ	No.	Title	Outlet	Type	Ranking	Pages
Expanding Contribution I: *Opportunities and Challenges in the Applicability of PETs*						
RQ1	2AP1	Mitigating Sovereign Data Exchange Challenges: A Mapping to Apply Privacy- and Authenticity-Enhancing Technologies	TB	CON	**B**	**16**
RQ1	2AP2	The Use of De-identification Methods for Secure and Privacy-enhancing Big Data Analytics in Cloud Environments	ICEIS	CON	**C**	**7**
RQ1	2AP3	Towards a Privacy-Enhancing Tool Based on De-Identification Methods†	PACIS	CON	**A**	**8**
RQ1	2AP4	Verifying Outsourced Computation in an Edge Computing Marketplac	CMLA	CON	**C**	**19**
Expanding Contribution III: *Improving the Applicability of DP Systems*						
RQ3	PP1	Do I get the privacy I need? Benchmarking utility in differential privacy libraries	UR	CON	**-**	**20**
New Research Stream: *Improving the Applicability of DP in Virtual Reality Applications*						
-	PP2	SoK: Data Privacy in Virtual Reality	PETS	CON	**A**	**18**
-	PP3	Exploring the Privacy Risks of Adversarial VR Game Design*	PETS	CON	**A**	**19**
-	PP4	Going Incognito in the Metaverse*	UIST	CON	**A***	**16**

*Equal Contribution. †Position paper. Abbreviations: DP = Differential Privacy, 2AP = Second-Author Publication, P = Publication, PP = Pre-Print, CON = Conference, JNL = Journal. Outlet: UR = Under review, TB = Trust, Privacy and Security in Digital Business (2022), ICEIS = International Conference on Enterprise Information Systems (2020); PACIS = Pacific Asia Conference on Information Systems (2020); CMLA = International Conference on Machine Learning & Applications (2022), NDSS = Symposium on Network and Distributed System Security (2023). Ranking: CORE ranking for all manuscripts except for 2AP3 and 2AP3, which used Conference Ranks.

Table 1.2.: Overview of additional related publications.

Having expanded our knowledge in DP, we hypothesized that emerging systems such as virtual reality (VR) could benefit from its formal privacy guarantees. Thus, we conducted a systematization of knowledge to draw the landscape of privacy attacks and defenses in VR [NGS24] and demonstrated in a case study the unprecedented scope and depth of new attack vectors in VR [NGS23]. With the acquired practical offensive knowledge in VR [NGS23], an overview of existing privacy defenses [NGS24], and our expertise in DP, we developed the first proposal for an "incognito mode" for VR: MetaGuard [NMGS23]. We provide more details in ongoing work in section 6.2.

1.4. Structure of the Dissertation

This dissertation comprises four publications that aim to answer the three research questions. We structure the dissertation into three parts as per Fig. 1.1:

Part A includes three chapters. Chapter 1 introduces this dissertation by motivating the

importance of PET in DSAA, particularly of DP (see section 1.1), presents the three research questions we tackle in this dissertation (see section 1.2), and introduces the three contributions and four publications (see section 1.3). Afterward, we compile in Chapter 2 all the necessary background to follow this dissertation; specifically, we discuss the concept of privacy (see section 2.1), which systems are included in DSAA (see section 2.2), the diverse PETs available to practitioners (see section 2.3), and present in more depth the concept of DP (see section 1.1). Lastly, in Chapter 3, we explain our research strategy (see section 3.1) and research methods (see section 3.2).

Part B compiles the fact sheets of the four embedded publications. The first two publications (P1-2) compose Contribution I and tackle the first research question, the third publication (P3) provides Contribution II and answers the second research question, and the fourth study (P4) sheds light on the third research questions by realizing Contribution III.

Part C comprises of three chapters. In Chapter 5, we discuss the key findings of this dissertation (see section 5.1, answer the research questions (see section 5.2), and reflect on the dissertation's limitations (see section 5.3). Consecutively, we provide an overview of future and ongoing work in Chapter 6 by introducing the possible expansions of Contributions I, II, and III (see section 6.1) and the new research stream (see section 6.2). We conclude this dissertation with our conclusions and answers to the research questions in Chapter 7.

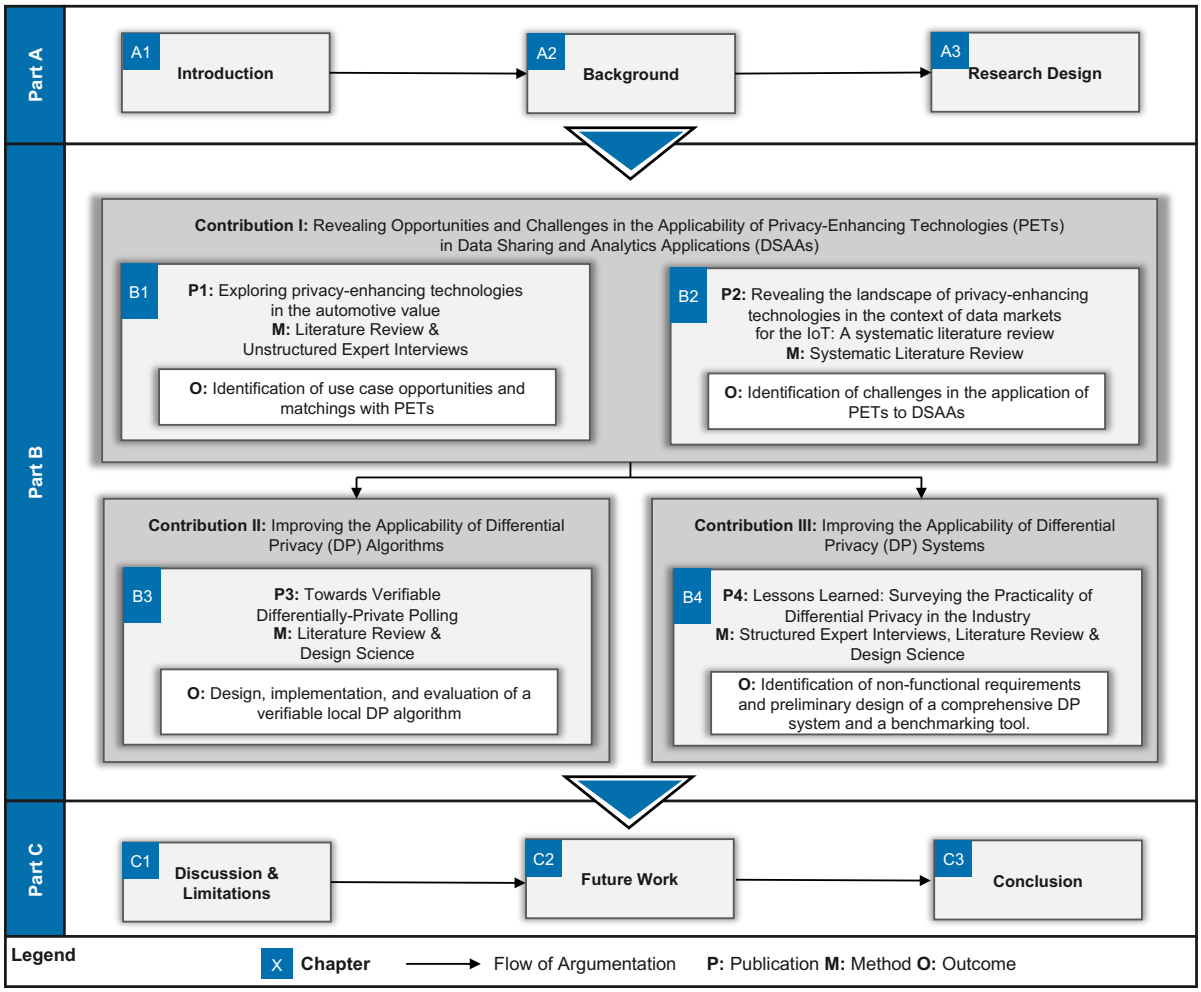

Figure 1.1.: Structure of the dissertation.

CHAPTER 2

Background

In this section, we provide the necessary descriptions of the fundamental concepts appearing throughout this dissertation and the embedded publications. We first describe *Privacy* (§2.1), a key concept that frames our work, followed by a description of *Data Sharing & Analytics Applications* (§2.2), the target of our differential privacy applicability improvements. Consecutively, we provide an overview of *Privacy-Enhancing Technologies* (§2.3) to contextualize *Differential Privacy* (§2.4), which is at the core of Contributions II and III.

2.1. Privacy

Throughout human history, the evolution of technology and the activities created around it affected how humans and institutions perceived and acted concerning privacy, shaping, in turn, human culture [DFF14]. However, it was not until the emergence of the first information and communications technologies (ICT) such as the radio, television, first computers, and early mobile phones that privacy became a major concern for society. Such concerns led governments to create the first data protection laws in the world: the Swedish Data Act enacted in 1973, and one year later, the Privacy Act in the USA. Since then, the increasing attention and importance of privacy have encouraged practitioners to provide many acknowledged privacy definitions. For example, Fink et al. [FSJ18]: "[...] *freedom from observation, disturbance, or unwanted public attention* [...]", from Westin [Wes67]: "[...] *the claim of individuals* [...] *to determine for themselves when, how and to what extent information about them is communicated* [...]", or from Renaud and Galvez-Cruz [RGC10]: *"Privacy is the faculty and right that a person has to define, preserve and control the boundaries that limit the extent to which the rest of society can interact with or intrude upon. At the same time, he or she retains complete control over information generated by, and related to, him or her."* Despite these efforts, there is no

concise and agreed-upon definition. Solove describes the privacy community's inability to form a consensus and argues that any attempt to distill a unique, timeless definition is infeasible due to the multifaceted concept of privacy [Sol15].

Nonetheless, in computer science, adopting a *threat model* perspective narrows the possible definitions, as privacy would likely not have emerged if adversaries would not exist: transgressors of one's sensitive information gave reason to define privacy. Hence, Wu [Wu12] may provide a suitable privacy definition in the context of computer science: *"[Privacy] is defined not by what it is, but by what it is not – it is the absence of a privacy breach that defines a state of privacy,"*. Wu defined privacy based on the threat model of Deng et al. [DWS+11] (Lindunn), which provides a framework to assess the privacy risks of an application and guides practitioners in determining what information to protect and from whom before defining an application. Following Wu's [Wu12] privacy philosophy, for the context of this dissertation, we consider privacy as the prevention of an individual's re-identification by an adversary. Further, note that in the context of this work, we refer to *security* as the measures for blocking *unauthorized* data access, while *privacy* focuses on limiting harm by *authorized* entities.

Once ICTs become privacy-forward, several advantages emerge for society. For example, from an economic standpoint, privacy features enable cross-organizational data exchange and fair products and services that prevent price discrimination [GO18]. Moreover, enhancing privacy may increase in number the data collection sources because such features may help overcoming regulatory barriers, in addition to mitigating the risk of fines, appreciating and differentiating a brand [McK20], and increasing customers' willingness to pay [SLZ20]. Moreover, researchers argue that privacy may be the only way towards unobtrusive forms of governments that enjoy political freedom and stability [CC02], privacy-first journalism, and less pervasive digital platforms (e.g., social media) that can enable pernicious social engineering [ZMW14]. Furthermore, academics indicate that disregarding privacy may result in long-term economic adversities [LLS20].

Despite the potential benefits of using privacy-first applications, and while consumers emphasize the importance of privacy, they typically do not take small additional efforts or pay for privacy [Kok17], the so-called *privacy-paradox*. Therefore, in the interest of business and end users, governments have enacted more data protection directives, rules, and laws in the past decades to protect users against privacy violations, specifically for data collection through advanced ICTs such as personal computers, smartphones, home assistants, or wearables. Examples of these regulations include the European Data Protection Directive in 1995, the HIPAA Privacy and Security Rule in 1996, the APEC cross-border privacy rules in 2011, the GDPR in 2016, and the Consumer Privacy Act in 2020 in the USA, which comprises of Acts such as the CCPA of 2018.

2.2. Data Sharing & Analytics Applications

Using the broad umbrella term "*Data Sharing & Analytics Applications*" is intentional, as this dissertation's contributions apply to many current applications that manage and analyze sensitive data and, therefore, require privacy measures. For example, devices such as smartphones,

wearables, home assistants, and personal computers provide a plethora of useful applications that interact with users and gather and analyze data daily. Examples of these applications are browsers, data markets, social media, e-commerce, or media entertainment, among others. Moreover, corporate-grade systems and sensors that collect, store, communicate, and process large scale data flows from users and critical business operations are also subject to privacy risks. Thus, in this dissertation, we consider any device with a CPU connected to the Internet that acts as the gateway to applications that share and analyze user or business-critical data.

Contribution I [fGS+21, GSU+22] explores the challenges and opportunities of privacy technologies in the broad landscape of IoT [ORRK18] and its applications, which encompass all the aforementioned devices. These applications rely on several privacy-critical layers, described in Table 2.1, which require distinct privacy data protections—briefly introduced in § 2.3. *Storage* and *processing* allow systems to utilize information when required, the *communication* layer comprises a set of technologies and protocols to exchange such information, *verification* allows the recipients to validate the authenticity of the data and identities, and the *governance* layer comprises a set of rules that govern which (sensitive) data may be exposed to which entities, usually in the form of pre-determined privacy policies, e.g., blocklists and allowlists.

Layer	**Description**
Storage	Stores data.
Processing	Accesses stored data and executes algorithms to extract valuable information.
Communication	Shares data between machines.
Verification	Checks the authenticity of data and the identities involved.
Sovereignty	Governs the privacy policies of the other layers.

Table 2.1.: Description of the privacy-relevant layers in data sharing and analytics applications.

Contribution II [MGSB22] targets user devices such as personal computers or smartphones that communicate with servers interested in performing aggregate statistics. Lastly, Contribution III [GLMS23] focuses on industry-grade analytics tools that act as the gateway for practitioners to access multiple large-scale user datasets.

2.3. Privacy-Enhancing Technologies

Since the advent of the internet in the 90s, the market has experienced a surge of digital platforms that require an increasing amount of data to offer products, services, and plan their corresponding future iterations. Accordingly, companies have deployed scalable infrastructure for the generation, collection, storage, processing, distribution, and analysis of *big data* to realize the economic potential of users' information. However, these lucrative practices bring a set of obligations and prohibitions towards the public's privacy, as their sensitive information can be compromised and used against their private interests. Thus, institutions face the challenge of enhancing user privacy while extracting value from their data. This contemporary challenge sparked the development of a new suit of tools called PET. The term PET was coined by the Dutch Data Protection Authority and the Ontario Information Commissioner [HBNC98] in 1995, which explored a novel approach to privacy protection [Opp05]. PETs may adopt various forms, some are architectures built with privacy-by-design policies and principles [PMB+16, SN15],

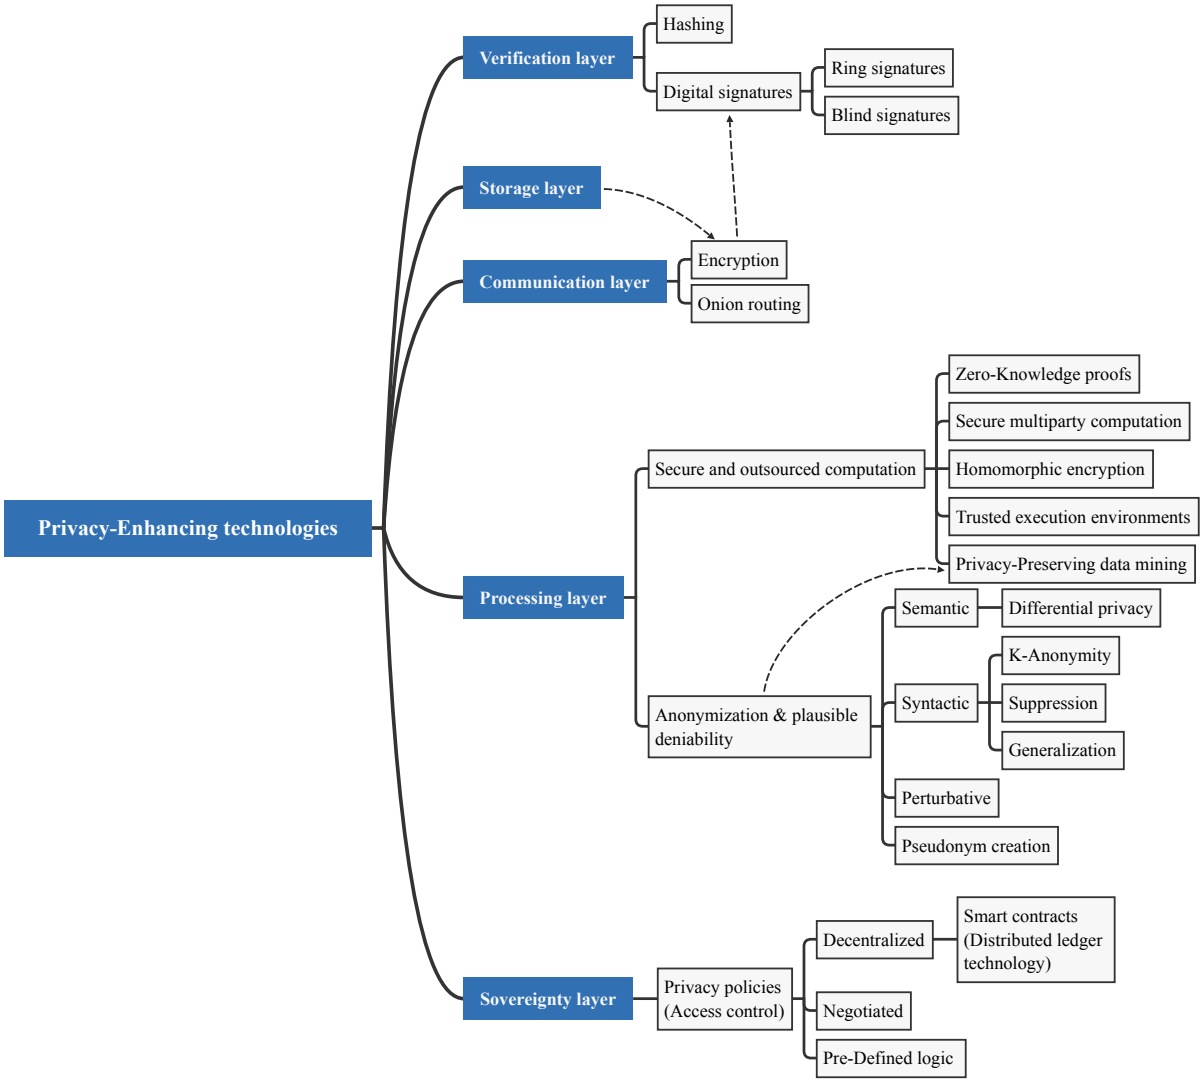

Figure 2.1.: Simplified privacy-enhancing technology classification by removing the leaf nodes of Fig.5.2 (cf. [GSU+22]).

others rely on cryptographic primitives or on data alterations stemming from heuristics or mathematical privacy guarantees. Fig. 2.1 (cf. adapted [GSU+22]) depicts a simplified classification of PETs based on their deployment layer: *verification*, *storage*, *communication*, *processing*, and *sovereignty* (see section 5.2 for details).

In the following, we briefly define the most representative PETs per layer in Fig. 2.1 with the aim to (i) provide context to industry practitioners not familiar with the topic (a recurrent request throughout the recent past years), and (ii) contextualize DP by describing what the other PETs can accomplish.

PETs Shielding Data

Encryption is a foundational building block of confidential storage and transmission of data [PHS+19a], digital signatures, and is at the core of most of the other PETs shielding data and altering computation. Encryption schemata are either symmetric (a single key for both encryption and decryption) or asymmetric (a public key for encryption and a private key for decryption, or viceversa).

Homomorphic Encryption (HE) enables the computation of functions over encrypted data (ciphertext), effectively hiding the inputs and allowing only the parties with the decryption key to decrypt the outputs [Che16]. Researchers classify HE based on the breadth of operations allowed by the corresponding scheme [WK15, CGSM19]: Fully HE supports addition and multiplication, partially HE allows for one of these two operations typically in exchange for improved performance, and somewhat HE schemata lay in between the two [SCS18]. Celebrated HE schemata are Paillier [Pai99] Boneh-Goh-Nissim [BGN05], and Hash-ElGamal [NIW+13]. While this PET holds promise, HE incurs significantly higher computational complexity and comparatively large cyphertext storage requirements [BGV14].

Trusted executions environments (TEE) were introduced in 2009 by the Open Mobile Terminal Platform as hardware and software components support for applications that require protection against adversaries aiming to steal cryptographic key material or other critical information [OMT09a]. TEEs are unique in that the considered adversaries include the legitimate hardware owners or remote access to the system running the TEE. In practice, a TEE allows a user to define a secure area of a CPU that impedes any code outside the secure environment to record the data or tamper with its computation, i.e., TEEs ensure confidentiality of inputs, outputs, and computation integrity. Precisely, TEEs associate hardware with unique encryption keys, and thus, hardware tampering is as hard as attacking the software layer. Nonetheless, TEEs still holds limitations, such as the limited memory offers and the vulnerabilities unveiled on multiple occasions by researchers [KM22, AVBS+22, SYG+19].

PETs altering computation

Zero-Knowledge Proofs (ZKP) facilitate a *verifier* to validate the data authenticity and computation integrity of a *prover* without requiring revealing such data or replicating the computation [GO94]. Moreover, if the *prover*'s claims are attested by a digital certificate signed by a trusted third party, ZKPs can also verify identity authenticity. ZKPs enable such verifications thanks to the following properties: (i) *zero-knowledgeness*, i.e., the *verifier* only learns the correctness of the statement, (ii) *completeness*, i.e., the *prover*'s attestation is successful with high probability, (iii) *soundness*, i.e., the *prover*'s attestation of a wrong statement has a low probability of succeeding [Sim02, KPC+20]. Furthermore, practitioners can choose between *interactive* and *non-interactive* ZKP protocols. The former requires engaging in a sequence of messages, while the latter allows a *prover* to convince the *verifiers* of a claim with a single message [Sim02]. Regarding the limitations of ZKPs, researchers agree the need to keep improving its computational complexity (primarily for the *prover*) and close the gap between cryptographers and software engineers [BBK+09].

Secure Multiparty Computation (MPC) allows many parties to jointly compute a func-

tion over obliviously exchanged data, i.e., without revealing any inputs [Yao82a, Sta96]. MPC implementations are based in either *secret sharing* or *garbled circuits*. In *secret sharing*, each party splits the secret input into shares broadcasted to the other computing parties. Upon receiving the sharded information, each party computes arithmetic operations independently and communicates the outputs to the rest of the parties to reconstruct the result [Yao82a]. On the other hand, *garbled circuits* achieve the same result by transforming functions into Boolean circuits, i.e., a combination of logic gates like AND, XOR, and OR that can construct any function [Yao82b, Yak17]. Oblivious transfer [GIR20] is the core of garbled circuits, i.e., a party shares one of many potential inputs to other parties while remaining oblivious of which data point has been sent. While MPC is more computationally efficient than fully HE schemata [YHL+19], it incurs high processing and communication costs and high sensitivity to latency, which decreases performance [KLB20].

Privacy-Preserving data mining (PPDM) is a means to extract useful information from data by the use of machine learning (ML) or conventional statistical analyses such as aggregations while enhancing the privacy of the process. In short, PPDM performs computations in the data owner's machine and can further protect the computation with cryptographic and data alteration techniques. The most popular PPDM tool is federated learning (FL) [KMR15, LSTS20, YLC+19], which collaboratively trains a base ML model on each client's local data; thereafter, the trained weights are shared with a server that aggregates them to form a unique model. Alternatives to FL are *split learning* approaches [VGSR18, GR18], which decompose the neural network's layers across different machines to separate data inputs and labels, and *gossip learning* [GG19, OHJ], whereby ML models perform a random walk over clients, where they are trained and merged. Additionally, practitioners can enhance the privacy of these PPDM approaches by broadcasting the weights with MPC [BIK+17] protocols, protect the client selection mechanisms with HE [ZLC+21], and enhance the privacy of the underlying data by altering the input data or the ML training process.

PETs Altering Data

DP, the targeted technology of this dissertation, provides a unique privacy guarantee among the rest of PETs, which is discussed in detailed in §2.4.

K-Anonymity is the primordial privacy definition among a suit of subsequent syntactic privacy models such as l-diversity, t-closeness, β-likeness, and δ-presence, among others. Unlike DP, k-anonymity and its variations define a property of the data itself [DFLS12]. For example, k-anonymization, i.e., processing data so that the output dataset fulfills k-anonymity, clusters attribute values into equivalence classes of size k, where each individual is indistinguishable from $k-1$ others. Typically, a k-anonymization process that minimizes information loss is an NP-hard problem. Therefore, researchers have studied heuristics to achieve near-optimal results [MW04].

PETs Managing Data

Privacy policies are *soft* privacy measures that compile the obligations and prohibitions of a data governance model for an application. Practitioners should define privacy policies ex-ante and consider the requirements of the user. For example, access control embodies a policy

that can dictate *when* and *who* can access *which* data. While policies do not require advanced technologies, they face challenges. There is a lack of a global standard for electronic privacy policies [SN15], applications must prioritize among conflicting policies based on the circumstances [GTD18], and policy enforcement in conventional systems typically requires a human in the loop (as observed in the interviewed organizations [GLMS23]) due to the organizations' lack of automated privacy tooling, potentially causing delays in data processing [GTD18].

We should note that there is no "one-size-fits-all" PET; it is only the combination of these technologies that can strive to guarantee an end-to-end privacy-enhancing solution. For example, while MPC might maintain the inputs of a computation secret, its outputs are typically known to the participants. These outputs could leak information about the inputs; thus, definitions such as DP are necessary to protect the parties' privacy completely.

2.4. Differential Privacy

The privacy community has demonstrated how unsafe traditional privacy techniques are (e.g., suppressing names, generalizing values, and syntactic privacy definitions like k-anonymity), primarily because they are vulnerable to background knowledge attacks [DSSU17, SAW13, GFS$^+$14, KHdMR20, NS08, AGKZ18]. In contrast, DP, originally introduced by Dwork et al. [DMNS06] in 2006, proposes a context-agnostic provable privacy guarantee, i.e., the privacy guarantee holds despite available past, present, or future auxiliary information [DR13]. A function $\mathcal{M}(\cdot)$ (e.g., a summary statistic) fulfils DP if it effectively bounds how much new information an adversary can gain from observing its outputs. Specifically, such (randomized) function $\mathcal{M}(\cdot)$ satisfies DP by adding calibrated random noise, typically to a deterministic function's output: $\mathcal{M}(\mathcal{D}) = \mathcal{W}(\mathcal{D}) + Noise$. The similarity of likelihoods is bounded by the parameter ε, which is inversely proportional to the privacy guarantee's strength. In practice, the outputs of $\mathcal{M}(\cdot)$ are similarly likely with or without an individual's input contribution (dataset D vs. D'). Symmetrically, individuals absent in a dataset have "essentially" the same privacy guarantee as if they were in the dataset. Formally, Dwork et al. [DR13] defined DP as:

> **Definition 1.** *(ε, δ)-Differential privacy). A randomized function $\mathcal{M}(\cdot)$ is (ε, δ)-differentially private iff for any two neighboring datasets D and D' differing on at most one element, and any set of possible outputs $\mathcal{S} \subseteq Range(\mathcal{M})$:*
>
> $$\Pr[\mathcal{M}\ (D) \in \mathcal{S}] \leq e^{\varepsilon} \times \Pr[\mathcal{M}\ (D') \in \mathcal{S}] + \delta.$$

Having briefly introduced the fundamental concept of DP, the following provides a series of critical aspects to consider when working with DP and necessary to follow Contribution III [GLMS23].

Pure & Approximate DP. Definition 1 is considered pure DP when $\delta = 0$. In contrast, if $\delta \neq 0$, the DP function provides more utility (e.g., output accuracy) in exchange for lowering the privacy guarantee of the individual [DR13]. Specifically, the parameter δ is the probability that the information gained by an adversary is not bounded by ε, i.e., distinguishing between D and D' is trivial.

Boundedness. Depending on how the neighboring datasets D and D' differ in one individual, the randomized function fulfills either *unbounded* or *bounded* DP. In *unbounded* DP, an individual record is removed ($|D'| = |D| - 1$) or added ($|D'| = |D| + 1$), whereas in *bounded* DP, an individual record is changed ($|D'| = |D|$). In both settings, D and D' are at a Hamming distance of $d_h(D, D') = 1$, i.e., differ in one individual. A higher Hamming distance is desired for DP guarantees in group privacy [Mir17].

Sensitivity. The noise added by a randomized function $\mathcal{M}(\cdot)$ is sampled from a random variable (r.v.). The scale of this r.v. (noise) is affected by the value ε and the *global* ℓ_1-sensitivity of the deterministic function, which corresponds to the maximum difference of the function's outputs over all possible neighboring datasets D and D'. Picking the maximum ensures that any other individual's contribution with a lower impact on the output is protected. Depending on the function type, the ℓ_1-sensitivity may vary, e.g., a mean query could have a lower ℓ_1-sensitivity than a count query and, therefore, the scale of the added noise is smaller. Formally, Dwork et al. [DR13] defines ℓ_1-sensitivity as:

> **Definition 2.** *(ℓ_1-sensitivity). The ℓ_1-sensitivity of an algorithm $\mathcal{W} : \mathbb{R}^m \to \mathbb{R}^n$, executed over datasets $\mathcal{D}$, $\mathcal{D}' \in \mathbb{R}^k$ at a Hamming distance of $d_h(\mathcal{D}, \mathcal{D}') = 1$, is:*
>
> $$\Delta f = \max_{\substack{\mathcal{D},\mathcal{D}' \in \mathbb{R}^k \\ d_h(\mathcal{D},\mathcal{D}')=1}} \|\mathcal{W}(\mathcal{D}) - \mathcal{W}(\mathcal{D}')\|_1.$$

Depending on whether DP was defined as *bounded* or *unbounded*, the value of ℓ_1-sensitivity varies, e.g., a bounded multi-dimensional count query has an ℓ_1-sensitivity$= 2$, while its unbounded counterpart has a ℓ_1-sensitivity$= 1$. Additionally, if we define ℓ_1-sensitivity for a fixed D, Definition 2 describes *local* sensitivity. The *local* sensitivity is smaller than the *global* sensitivity (the upper bound) because it does not consider all the possible values and combinations of D and D', resulting in better utility (lower noise scale). However, using a function's *local* sensitivity limits its applicability to one particular D, as it does not account for any other possible larger individual contribution.

DP Mechanisms. There exist a myriad of mechanisms fulfilling Definitions 1 and 2. The most popular are the Laplace mechanism [Dwo08] and the Gaussian mechanism [DR13] for numerical data, and the Exponential mechanism [MT07] for categorical and numerical data. The original Laplace and Exponential mechanisms fulfill *pure* DP, while the Gaussian mechanism follows *approximate* DP. Formally, the Laplace mechanism is defined as [Dwo08]:

> **Definition 3** *(Laplace mechanism). For an algorithm $\mathcal{W}$ executed over a dataset $\mathcal{D}$, its differentially private version $\mathcal{M}$ adds Laplace noise: $\mathcal{M}(\mathcal{D}) = \mathcal{W}(\mathcal{D}) + Lap(x|\mu, b)$, with*
>
> $$\text{Lap}\left(\mathcal{W}(\mathcal{D}) = x|\mu = 0, b = \frac{\Delta f}{\varepsilon}\right) = \frac{\varepsilon}{2\Delta f} \exp\left(-\frac{\varepsilon|x|}{\Delta f}\right).$$

The Gaussian mechanism is equivalent; however, the noise is sampled from a normal r.v. In contrast, the Exponential mechanism returns the category that approximately maximizes a utility function u, i.e., unlike the Laplace and Gaussian mechanisms, it does not add noise directly to the output of a deterministic function.

> **Definition 4** *(The Exponential mechanism). Given a utility function $u : (D \times R) \to R$, and a dataset D, the exponential mechanism $\mathcal{M}(D, u)$ returns outputs $r \in \mathcal{R}$ with probability $p_r \propto \exp(\frac{\varepsilon}{2\Delta u} u_r)$, where u_r is he utility of the element r and Δu is the ℓ_1-sensitivity of u.*

Definitions 3 and 4 show the impact of the value ε: the larger the ε, the larger the standard deviation of the sampled noise, and thus, the stronger the privacy guarantee and the weaker the utility of the mechanism. Nonetheless, note that the magnitude of the noise is independent of the dataset size (the number of records). Therefore, analyzing larger datasets provides better relative utility.

Central/Local Model. Practitioners enhance the privacy of applications by deploying DP in the *local* mode, i.e., imbuing noise to each client's data point individually (using techniques such as randomized response [War65], based on two coin flips), or in the *central* mode, i.e., adding noise after aggregating the data points across clients. The *central* model is typically less noisy than the *local* model; however, the *local* model requires less trust assumption with the server.

Privacy Levels. Practitioners can choose the scope of their DP implementation at three levels [NR20]: (i) *user level privacy*, when all the records linked with a user are either absent or present, (ii) *event level privacy*, when all the records associated with an event or a group of events are either present or absent, and *w-event level privacy*, when a set of w occurrences of records produced by an event or groups of events are either absent or present.

Sequential Composition. DP algorithms obey *sequential composition* [DR13], i.e., the resulting ε of the execution of a sequence of n (possibly different) DP mechanisms over D with ε_i is the consumed *privacy budget* $\varepsilon = \sum \varepsilon_i$.

Privacy Budget Tracking. Given that the sampling noise distribution is centered around 0, an adversary could reverse engineer the n query results by averaging out the noise. Therefore, DP tool designers should implement built-in budget trackers to prevent this type of attack.

Floating-Point Vulnerabilities. The original proofs for the Gaussian and Laplace mechanisms rely on distributions on the real numbers line; however, computers approximate real numbers with floating-point representations that depend on the input values. In contrast, DP outputs should be independent of the input data, and therefore, practical DP implementations make outputs distinguishable—breaking DP [Mir12]. Moreover, DP implementations are vulnerable to precision-based attacks, where outputs can be distinguished based on the fact that some will be multiples of a specific power of two while others may not [HDH+22].

Quality of Randomness. DP implementations using low-quality random number generators may lead to privacy leakage [GL20]. Thus, developers should be mindful of using cryptographic-secure random number generators.

Semantic Consistency. An important practical consideration of DP applications is preserving the *output consistency* of the analysis, e.g., a count query result should not be negative or have decimal values, the variance should not result in negative or zero values, and dividing a DP sum by the count query result should yield similar values to the DP mean. While there are solutions to constraint DP query outputs, they can also introduce errors or biases.

Outstanding Challenges. While DP has become the privacy golden standard for academics

and the industry and governmental institutions strive to deploy DP in practice [App17, AG21, DKY17, Joh21], there still exist remaining challenges [GLMS23]:

1. **Choosing Privacy Parameters**. The most limiting roadblock is choosing an appropriate ε value, as there is not yet a consensus and generalizable guidelines [DKM19].
2. **DP tooling maturity**. While at a striking distance, practitioners still have issues deploying current open-source DP tooling [GLMS23].
3. **High Adaption Effort**. While academics have adapted DP to numerous algorithms (e.g., ML and aggregations) and use cases using the *local* or *global* models, these adaptions require significant effort and extensive expertise [DFSBJ21].
4. **Fairness**. In use cases whose outcome significantly impacts users' well-being, e.g., distributing budget across social groups based on aggregated economic data, DP calculations could over- or under-estimate the allocation of resources or reduce the correctness of data-driven decisions [KKM$^+$21].
5. **User Data Streams**. Users produce and share data on a daily basis, which, because of *sequential composition*, it is a challenge to guarantee *user-level* privacy [LSV$^+$19].
6. **Privacy Budget Tracking Across Systems**. Users employ different siloed platforms and produce linkable and redundant data, giving the opportunity to adversaries to switch between platforms to exploit untethered privacy budgets.
7. **DP Verifiability**. Verifying the correctness of DP algorithms is a complex endeavour, as it requires proving and fitting a mathematical model to the system's semantics or the use of technologies such as ZKP to prove validity [KMR$^+$20].

Throughout this dissertation, we survey and work with tools focused on facilitating the use of DP (Contribution III), which offer choice across the aspects introduced in this section, e.g., available tools allow for both *pure* and *approximate* DP and the two *boundedness* definitions, some calculate sensitivity automatically, employ cryptographic-secure random number generation, have protections to circumvent floating-point vulnerabilities, and most offer a range of diverse DP mechanisms, privacy budget odometers, and semantic consistency features. Mainly, we focus our work on *user-level* privacy, and we provide contributions in the *local* model (Contribution II) and the *central* model (Contribution III) of DP.

Concerning the challenges of DP, in short, we help tackle Challenge (2) *DP Tooling Maturity* in Contribution III by outlining the remaining gap for the deployment of DP tooling and Challenge (7) *DP Verifiability* in Contribution II by proving the use of DP in the *local* model with a cryptographically verifiable proof based on ZKPs (see section 2.3).

CHAPTER 3

Research Design

This chapter describes the research design we used to conduct our research endeavors and answer our research questions and fill the gaps satisfactorily [Bha12]. The two aspects of a research design are defining a research strategy (see Section 3.1) and selecting research methods appropriately [Bha12] (see Section 3.2).

3.1. Research Strategy

Scientific inquiries are either *inductive*, i.e., the derivation of patterns from observations and theoretical knowledge for those patterns, or *deductive*, i.e., developing a hypothesis and conducting experiments to validate it [Bha12]. The type of inquiry influences the research design and data collection [Mye97]. We rely on an *inductive* process for contributions I and III and a *deductive* process for contribution II. Furthermore, research could be *primary*, i.e., studies that contain original data collected by the authors or their designs, artifacts, and experiments, or *secondary*, i.e., work that reviews or systematizes existing knowledge. Additionally, we could use several research approaches: *qualitative*, *quantitative*, and *mixed* strategies [CC18]:

- *Qualitative strategies* strive to understand and explain complex social phenomena that often cannot be generalized [SC90]. Researchers use this strategy when conducting case studies and expert interviews [Bha12].
- *Quantitative strategies* aim to measure variables and understand such data to derive patterns or test hypotheses [Yıl13, CC18]. Researchers rely on this strategy typically for conducting surveys and experiments.
- *Mixed strategies* are rooted in the philosophy that combining quantitative and qualitative

information provides synergistic insights, i.e., the researcher can learn more new information than with the qualitative and quantitative data alone [CC18]. Thus, the researcher collects, integrates, and interprets both types of data [JOT07, CC18].

Based on the research gaps of section 1.2, we decided to follow a *mixed* strategy by using a three-fold approach combining (i) semi-structured expert interviews and (ii) (systematic) literature reviews aimed to understand the state-of-the-art (*qualitative* research), and (iii) design-science research verified with experimental results (*quantitative* research). This resulted in three of our papers to be mainly *primary* (P1, P3, P4) and one *secondary* study (P2). Through our three-fold approach, we aim to reveal unique insights and solutions [VBB13] aimed to improve the applicability of Privacy-Enhancing Technology (PETs) in Data Sharing and Analytics Applications (DSAAs), specifically, in the use of DP—a technique providing privacy guarantees unique among PETs.

Fig. 3.1 illustrates the overarching strategy we followed in this dissertation. First, we used evidence-based research methods, namely, gray and systematic literature reviews, to identify the problems and research gaps in applying PETs to DSAAs (P1, P2). Secondly, we conducted other minor literature reviews (P3, P4) to elicit the necessary algorithm components (P3) and functional requirements to create tools that can tackle the identified problems (P4). In combination with expert interviews (P1, P4) and the industry experience in privacy of the authors of the research papers (P4), we designed and developed two software components: a verifiable differential privacy algorithm (P3) and an initial benchmarking tool (introduced in P4 and expanded upon in PP1). By working on PP1, we built knowledge and expertise that also helped writing P4. We note that the benchmarking tool is at an initial stage, and its purpose is to compare existing DP tools. With the comparison results, we can choose the tools that best suit the architecture of a holistic privacy-enhancing analytics tool, which we designed in P4.

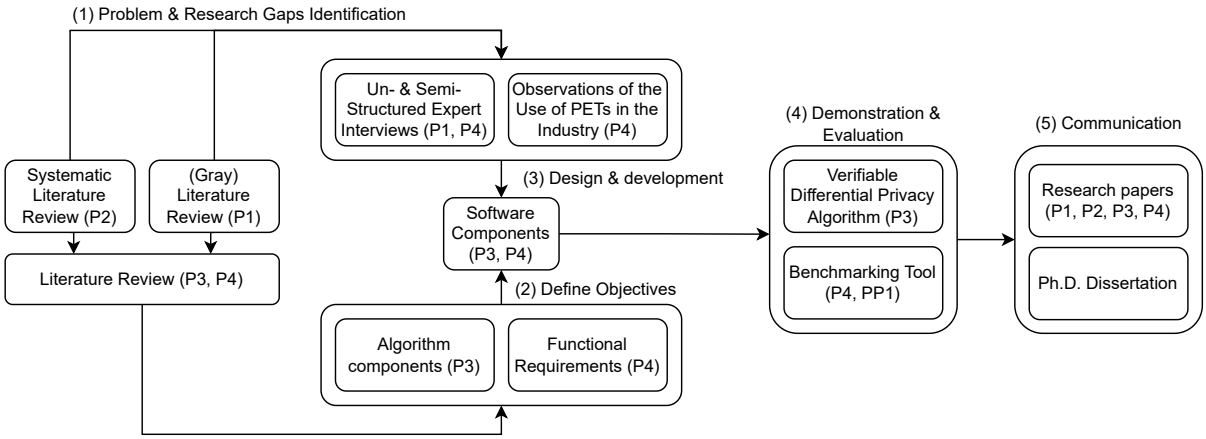

Figure 3.1.: Overview of our research approach. P = Embedded paper.

No.	Inquiry Type	Strategy type	Study type	Research Method
Contribution I: *Opportunities and Challenges in the Applicability of PETs*				
P1	Inductive	Qualitative	Primary Secondary	Unstructured Expert Interviews (Gray) Literature Review
P2	Inductive	Qualitative	Secondary	Systematic Literature Review
Contribution II: *Improving the Applicability of Differential Privacy Algorithms*				
P3	Deductive	Quantitative	Primary	Literature Review Design Science
Contribution III: *Improving the Applicability of Differential Privacy Systems*				
P4	Inductive	Qualitative	Primary Secondary	Structured Expert Interviews Literature Review Design Science

Table 3.1.: Overview of research types, strategies, and methods applied in the embedded publications.

3.2. Research Methods

In this section, we briefly describe the three research methods used throughout this dissertation: *systematic literature reviews*, *expert interviews*, and *research design*. Particularly, we describe their general characteristics, process, and how they contributed to the results of this dissertation. Each of the embedded publications P1-4 describe in detail the corresponding used method. In summary, we lay out the different research strategies and methods each of the embedded papers follows in Table 3.1. The descriptions of this section are inspired by the work of one of the co-authors of P2, Ömer Uludağ [Ulu22].

Systematic Literature Reviews. Literature reviews aim to provide an overview of the target research field by dissecting the content of a series of extant publications and deriving insights, taxonomies, or other artifacts in a systematic and reproducible manner [Coo88, Bak00, WW02, RS04]. Moreover, a literature review highlights, assesses, and structures the seminal studies in the field, and summarizes their content for researchers to quickly understand the state-of-the-art [LE06, TDS03]. Furthermore, a literature review extracts future work avenues for researchers looking to advance the field and facilitates theory development [WW02]. The transparency of the review process, i.e., search, filtering, and analysis, is critical for other researchers to verify that the findings are not biased or lack depth [VBSN$^+$09]. In response, researchers conduct literature reviews with a systematic process detailed in their publications [TDS03, VBSN$^+$09].

For the systematic literature review of P2, we used the goal-question-metric paradigm [BCR94]. We formulated the goal of as follows:
(P2) "*We systematically analyze peer-reviewed literature to provide an overview of the state-of-the-art concerning available research and trade-offs on privacy-enhancing data markets for the IoT as well as potential research gaps from the point of view of both scholars and practitioners.*"
From that standpoint, we created two research questions to guide our search for seminal studies. The literature reviews for P1, P3, and P4 had a similar approach and were conducted mainly to find relevant related work and were not the main means to achieve their associated contributions. The goals of the literature reviews of P1, P3, and P4 were:
(P1) "*We analyze gray literature to provide an overview of existing open-source tools to support*

use cases handling critically sensitive information,"
(P3) "...[provide an overview] *of a series of differentially private noise sampling algorithms that we could use for implementing their verifiable version*," and
(P4) "...[provide an overview] *of existing tools offering differential privacy functionality and profiling the gap our research aims to fill.*"

Expert Interviews. Researchers conduct empirical studies to extract new information or validate hypotheses, which the field of software engineering also requires [GPRN18]. Conducting a set of expert interviews is one such empirical method. These studies aim to gather data from a target population sample by personally or impersonally interacting with the interviewees [IM95, Kas05].

Two of the studies in this dissertation used expert interviews. P1 performs unstructured interviews, i.e., the goal and outcome of the interview are set, but the researchers do not specify questions ex-ante. The goal of the interviews in P1 was to find industry use cases where PETs could provide value. In contrast, the interviews in P4 were structured, i.e., the questions were defined before the interview. The questions in P4 aimed to understand the role that DP could play in companies unfamiliar with its formal privacy guarantee.

Design Science. Design science aims to create novel artifacts, i.e., solutions to critical problems or roadblocks in a target research field [HMPR04, PTRC07], with rigor, and relevance [HMPR04, RV08, Sta95, BW96]. If the artifacts proved relevant and have been evaluated and validated, early-stage design artifacts could positively impact organizations' localized problems. To create these artifacts, e.g., software components, researchers must rely on rigorous methodologies in close collaboration with the industry [BMSS13a, BMSS13b].

We followed the guidelines of design science research [HMPR04] and its method [PTRC07] to develop two software components: a verifiable differentially private noise sampler (P3), and a benchmarking tool (P4); however, note that the benchmarking tool is only introduced in P4, while it is in PP1 [GNM$^+$21] where the tool is the focus. Fig. 3.1 follows the method of Peffers et al. [PTRC07]: (1) identify a problem, (2) define objectives, (3) design and develop the artifact, (4) demonstrate and evaluate the artifact, and (5) communicate the results. Note that to identify the problem this artifacts tackle and fine-tune their requirements and characteristics, we used the aforementioned systematic literature reviews and expert interviews.

Part B

CHAPTER 4

Contributions

4.1. Contribution I: Revealing Opportunities and Challenges in the Applicability of Privacy-Enhancing Technologies in Data Sharing and Analytics Applications

(P1) Exploring Privacy-Enhancing Technologies in the Automotive Value Chain

Authors. Munilla Garrido, Gonzalo* and Schmidt, Kaja and Harth-Kitzerow, Christopher and Klepsch Johanne and Luckow, Andre and Matthes, Florian*.
*Technische Universität München, Chair of Software Engineering for Business Information Systems, Boltzmannstraße 3, D-85748 Garching, Germany.

Outlet. 2021 IEEE International Conference on Big Data.

Page Number. 8.

Status. Published.

Paper Version. Accepted Version.

Contribution of first author. Problem definition, research design, data collection, data analysis, interpretation, writing, and reporting.

Abstract. Privacy-Enhancing Technologies (PETs) are becoming increasingly crucial for addressing customer needs, security, privacy (e. g., enhancing anonymity and confidentiality), and regulatory requirements. However, applying PETs in organizations requires a precise understanding of use cases, technologies, and limitations. This paper investigates several industrial

use cases, their characteristics, and the potential applicability of PETs to these. We conduct expert interviews to identify and classify uses cases, a gray literature review of relevant open-source PET tools, and discuss how the use case characteristics can be addressed using PETs' capabilities. While we focus mainly on automotive use cases, the results also apply to other use case domains.

Research Gap. While there is a wealth of research applying or optimizing PET to particular use cases, there is a lack of studies outlining opportunities and guiding practitioners to identify actionable privacy-related use cases in the industry and select the appropriate PET for a given use case.

Method. We conducted a survey of 17 experts and a gray literature review that revealed 76 open-source PET tools and mapped the identified use cases with PET.

Artifacts.

1. List of use cases classified per application domain.
2. List of the most prominent PETs and associated open-source tools.
3. Mapping between PETs and use case characteristics.
4. Mapping between PETs and selected use cases per application domain.

Results. We define a guideline to select a PET based on use case characteristics and provide representative examples to showcase the reasoning for selecting a particular PETs.

Conclusion. We conclude that there is "no-size-fits-all" PET and the level of knowledge to deploy them is significantly high, and, thus, we call for caution at the time of integrating PETs in legacy or new systems and encourage researchers to produce more guidance and improve open-source tools to bridge the gap between theory and practice.

(P2) Revealing the landscape of privacy-enhancing technologies in the context of data markets for the IoT: A systematic literature review

Authors. Munilla Garrido, Gonzalo* and Sedlmeir, Johannes and Uludağ, Ömer* and Soto Alaoui, Ilias* and Luckow, Andre and Matthes, Florian*.
*Technische Universität München, Chair of Software Engineering for Business Information Systems, Boltzmannstraße 3, D-85748 Garching, Germany.

Outlet. Volume 207, Journal of Network and Computer Applications.

Page Number. 43.

Status. Published.

Paper Version. Published Version.

Contribution of first author. Problem definition, research design, data collection, data analysis, interpretation, writing, and reporting.

Abstract. IoT data markets in public and private institutions have become increasingly relevant

in recent years because of their potential to improve data availability and unlock new business models. However, exchanging data in markets bears considerable challenges related to disclosing sensitive information. Despite considerable research focused on different aspects of privacy-enhancing data markets for the IoT, none of the solutions proposed so far seems to find a practical adoption. Thus, this study aims to organize the state-of-the-art solutions, analyze and scope the technologies that have been suggested in this context, and structure the remaining challenges to determine areas where future research is required. To accomplish this goal, we conducted a systematic literature review on privacy enhancement in data markets for the IoT, covering 50 publications dated up to July 2020, and provided updates with 24 publications dated up to May 2022. Our results indicate that most research in this area has emerged only recently, and no IoT data market architecture has established itself as canonical. Existing solutions frequently lack the required combination of anonymization and secure computation technologies. Furthermore, there is no consensus on the appropriate use of blockchain technology for IoT data markets and a low degree of leveraging existing libraries or reusing generic data market architectures. We also identified significant challenges remaining, such as the copy problem (i.e., once an entity releases data, the data is no longer under the original owner's control) and the recursive enforcement problem (i.e., the recursive need to supervise a supervising entity). These challenges—while solutions have been suggested to some extent—are often not sufficiently addressed in proposed designs. We conclude that privacy-enhancing technologies need further improvements to positively impact data markets so that, ultimately, the value of data is preserved through data scarcity and users' privacy and businesses-critical information are protected.

Research Gap. Despite the increasing number of publications dedicated to studying PET, there is a lack of comprehensive yet detailed review, classification, and analysis of PET and challenges in the broad context of IoT.

Method. We conducted a structured literature review covering 74 publications, from which we extracted taxonomies, opportunities, and challenges of applying PET in the IoT.

Artifacts.

1. Taxonomy of PETs.
2. Taxonomy of privacy challenges in IoT data markets.
3. Examination of the IoT's negative impacts on privacy.
4. Mapping between PETs and the associated privacy challenges.
5. Reference model for a privacy-enhancing IoT data market.
6. Metadata aggregation of the selected publications.

Results. The main results indicate that many researchers tend to re-invent the wheel instead of improving and contributing to existing open-source tools, combinations of PETs offering secure computation and provable privacy guarantees are rare, and blockchains appear in many publications despite the lack of data exchange applications in production.

Conclusion. We conclude that there is not yet a canonical solution for building privacy-enhancing IoT data markets, for which the practicality of PET needs to improve.

4.2. Contribution II: Improving the Applicability of Differential Privacy Algorithms

(P3) Towards Verifiable Differentially-Private Polling

Authors. Munilla Garrido, Gonzalo* and Sedlmeir, Johannes and Babel, Matthias.
*Technische Universität München, Chair of Software Engineering for Business Information Systems, Boltzmannstraße 3, D-85748 Garching, Germany.

Outlet. ARES '22: Proceedings of the 17th International Conference on Availability, Reliability and Security, Article No.: 6.

Page Number. 10.

Status. Published.

Paper Version. Published Version.

Contribution of first author. Problem definition, research design, algorithm design, verification, writing, and reporting.

Abstract. Analyses that fulfill differential privacy provide plausible deniability to individuals while allowing analysts to extract insights from data. However, beyond an often acceptable accuracy tradeoff, these statistical disclosure techniques generally inhibit the verifiability of the provided information, as one cannot check the correctness of the participants' truthful information, the differentially private mechanism, or the unbiased random number generation. While related work has already discussed this opportunity, an efficient implementation with a precise bound on errors and corresponding proofs of the differential privacy property is so far missing. In this paper, we follow an approach based on zero-knowledge proofs (ZKPs), in specific succinct non-interactive arguments of knowledge, as a verifiable computation technique to prove the correctness of a differentially private query output. In particular, we ensure the guarantees of differential privacy hold despite the limitations of ZKPs that operate on finite fields and have limited branching capabilities. We demonstrate that our approach has practical performance and discuss how practitioners could employ our primitives to verifiably query individuals' age from their digitally signed ID card in a differentially private manner.

Research Gap. The few extant works bridging DP and ZKP fail to acknowledge their solution guarantees the *approximate* definition of DP and not in its *pure* form, do not provide a cryptographic validity guarantee, their performance is impractical, or have scant implementation details.

Method. We performed a literature survey to find DP algorithms that could be adapted so that ZKP can verify their use. Consecutively, we implemented and evaluated our adapted DP algorithm.

Artifacts.

1. Algorithm to cryptographically verify the use of DP on the local mode.

2. The algorithm's open-source implementation (software component I): `https://github.com/applied-crypto/DPfeatZKP`.

Results. We provide the privacy community with primitives for implementing cryptographically verifiable DP in the local model.

Conclusion. Enabling a server to cryptographically verify the correct execution of DP noise sampling in the local setting (i.e., a provable privacy guarantee) can prevent malicious clients from adding "garbage" noise to their local data, which can deteriorate the accuracy of a server's aggregate statistics.

4.3. Contribution III: Improving the Applicability of Differential Privacy Systems

(P4) Lessons Learned: Surveying the Practicality of Differential Privacy in the Industry

Authors. Munilla Garrido, Gonzalo* and Liu, Xiaoyua and Song, Dawn and Matthes, Florian*. *Technische Universität München, Chair of Software Engineering for Business Information Systems, Boltzmannstraße 3, D-85748 Garching, Germany.

Outlet. 23rd Privacy Enhancing Technologies Symposium (PETS 2023).

Page Number. 20.

Status. Published.

Paper Version. Published Version.

Contribution of first author. Problem definition, research design, data collection, data analysis, interpretation, requirements elicitation, writing, and reporting

Abstract. Since its introduction in 2006, differential privacy has emerged as a predominant statistical tool for quantifying data privacy in academic works. Yet despite the plethora of research and open-source utilities that have accompanied its rise, with limited exceptions, differential privacy has failed to achieve widespread adoption in the enterprise domain. Our study aims to shed light on the fundamental causes underlying this academic-industrial utilization gap through detailed interviews of 24 privacy practitioners across 9 major companies. We analyze the results of our survey to provide key findings and suggestions for companies striving to improve privacy protection in their data workflows and highlight the necessary and missing requirements of existing differential privacy tools, with the goal of guiding researchers working towards the broader adoption of differential privacy. Our findings indicate that analysts suffer from lengthy bureaucratic processes for requesting access to sensitive data, yet once granted, only scarcely-enforced privacy policies stand between rogue practitioners and misuse of private information. We thus argue that differential privacy can significantly improve the processes of requesting and conducting data exploration across silos, and conclude that with a few of the

improvements suggested herein, the practical use of differential privacy across the enterprise is within striking distance.

Research Gap. Given the lack of widespread adoption of DP tooling, we research the suitability of DP in organizations' workflows by interviewing practitioners unfamiliar with the technology, in contrast to other works focusing on interviewing experts.

Method. We conducted a survey of 24 experts. We distilled key findings, proposed functional requirements for a holistic privacy-enhancing analytics tool, and outlined the gaps in existing open-source DP tooling.

Artifacts.

1. Set of 10 functional requirements for holistic privacy-enhancing analytics systems.
2. List of outstanding research and engineering gaps blocking the wider adoption of DP.
3. Early-Stage blueprint for the design of a holistic privacy-enhancing analytics system.
4. Early-Stage open-source implementation of a benchmarking tool to test the suitability of DP libraries and frameworks for some of the blueprint's components (software component II): `https://github.com/camelop/dp_lab`.

Results. We provide a comprehensive summary of and key findings from the practitioners' answers to the 24 questions revolving around analytics use cases, data access requests, and dealing with noisy results, which may particularly interest *privacy officers* and *legal practitioners*. Additionally, the artifacts provided throughout P4 may be helpful to *software engineers* and *developers*.

Conclusion. We argue that DP could reduce data access request times by allowing the exploration of critically sensitive data across silos, reducing the access restrictions thanks to its stronger privacy guarantees, and, depending on the use case requirements, a DP aggregation analysis may directly fulfill the business request.

Part C

CHAPTER 5

Discussion

We move on to discuss the key findings of the embedded publications (section 5.1), answer the research questions (section 5.2), and examine the limitations of our studies (section 5.3).

5.1. Discussion of Key Findings

Our three contributions comprise four research studies with a series of key findings (see Table 5.1) discussed in the following:

Contribution I: Revealing Opportunities and Challenges in the Applicability of Privacy-Enhancing Technologies in Data Sharing and Analytics Applications.

Laying out the opportunities and challenges creates avenues for future work. In particular, P1 [fGS+21] conducted expert interviews to find such opportunities, and P2 [GSU+22] focused on revealing challenges by systematically reviewing 50 studies filtered from hundreds of search hits. These studies focused on DSAA architectures or data trading schemata (KF3.4), and characterized DSAAs based on the degree of decentralization, the types and number of data domains, and seller and consumer types (KF3.6). Moreover, the timeline of these 50 publications showed accelerated growth of publications dedicated to studying PETs in DSAAs. The updated search we conducted in May 2022 revealed another 24 publication: 3 more from 2020, 14 from 2021, and—as of May—7 from 2022, indicating the trend depicted in Fig. 5.1 has not reversed. Therefore, we can confirm that the attention of researchers in PET for DSAAs has increased notably in recent years (KF2.1). However, despite the surge in research, privacy-oriented DSAAs are still maturing, which is evident given the lack of production-grade implementations (KF2.2).

Based on our research, the reasons behind the lack of adoption could be the low maturity in available PET tools (KF1.2) and the increased complexity resulting from adding PETs to a

No.	Key Findings
Contribution I: Opportunities and Challenges in the Applicability of PETs	
P1	**(KF1.1)** There is "no-size-fits-all" PET. **(KF1.2)** Given that most open-source tools are not yet ready for deployment, using PETs requires an in-depth understanding of the technology, its limitations, and application domain. **(KF1.3)** PETs increase architectural and computational complexity.
P2	**(KF2.1)** The attention of scientists toward privacy-enhancing technologies in data markets for IoT devices has increased notably in recent years. **(KF2.2)** Privacy-oriented IoT data markets are still maturing and have not faced many production-grade implementations yet. **(KF2.3)** Studies rarely leveraged existing PETs libraries, and therefore, it may be beneficial for researchers to use or improve existing work instead of reinventing the wheel. **(KF2.4)** IoT data markets research is divided into architectures and data trading schemata. **(KF2.5)** Despite the acknowledged need for combining anonymization and secure and outsourced computation techniques, only two of the collected studies used a combination of PETs. **(KF2.6)** Data markets are characterized based on the degree of decentralization, the types and number of data domains, and seller and consumer types. **(KF2.7)** Blockchain is frequently used as the backbone of IoT data markets despite the lack of consensus on its use and blockchain-based applications in production.
Contribution II: Improving the Applicability of Differential Privacy Algorithms	
P3	**(KF3.1)** There is a lack of work on the cryptographic verifiability of DP. **(KF3.2)** Using Bernstein polynomials to approximate a probability density function in a closed interval and rejection sampling is not suitable for building a verifiable DP algorithm. **(KF3.3)** The truncated geometric mechanism relies on integers that could become so large that a ZKP circuit compiler cannot handle them.
Contribution III: Improving the Applicability of Differential Privacy Systems	
P4	**(KF4.1)** Data stewards seemed to be more concerned about security than privacy. **(KF4.2)** Running analysts' scripts without "seeing" the data is still a distant possibility. **(KF4.3)** DP could have a significant impact on dataset exploration. **(KF4.4)** Analysts could employ DP mechanisms to fulfill certain analytics use cases. **(KF4.5)** Companies do not have an ex-post human-supported privacy auditing step. **(KF4.6)** Given the six reasons analysts shared for fully accessing datasets, DP mechanisms could help in "obtaining a holistic understanding of data" by providing summary statistics. **(KF4.7)** Analysts are blocked for significant periods when requesting data access. **(KF4.8)** Differential privacy could arguably reduce the time to access data. **(KF4.9)** Most analysts employed aggregations and visualizations to successfully perform their assigned tasks, while machine learning was not as predominant. **(KF4.10)** SQL was more important than machine learning and was frequently used. **(KF4.11)** Analysts confirm that DP would be helpful for dataset exploration, fulfilling certain use cases, and for enabling privacy-enhancing dashboards for dataset visualization.

Table 5.1.: Overview of key findings in the embedded publications [GSU+22, MGSB22, fGS+21, GLMS23].

system (KF1.3). Consequently, a lack of confidence in available tools may lead researchers to re-invent the wheel in their studies (KF2.3). Therefore, we suggest researchers focus on improving existing open-source tools or explore the possibilities with existing tools before creating ad-hoc PET solutions. Thus, we decided to dedicate Contribution III to guide researchers in improving existing PET tools, specifically the ones providing DP functionality. Furthermore, we realized that there is "no-size-fits-all" PET (KF1.1), and yet, there is a lack of work focused on combining different PETs (KF3.5), which might be due to requiring an interdisciplinary team of researchers. Thus, we combined the expertise on DP with the knowledge ZKP of our co-authors, resulting in Contribution II, where we cryptographically verified the sampling of DP noise. Lastly, an interesting key finding was the use of blockchain, i.e., a tamper-proof distributed database whose state is stored, synchronized, and replicated by nodes in a P2P network following a consensus algorithm [BTH20], without clear reasoning and despite the lack of blockchain-based applications in production (KF3.7). This lack of rigor surfaces as questionable arguments in some of the reviewed studies: "[...] *researchers and technologists have found that blockchain can be a potential solution to the privacy problem by decentralizing information* [...] *Blockchain can be used to securely share private information* [...]" [LF20], "*Blockchain-based approaches provide decentralized security and privacy* [...]" [DKJG17], or "*Blockchain has been proven to possess security, immutability, and privacy properties, which has caused a lot of researchers to introduce it into the privacy and security concerned IoT*" [DJG+18]. Reviewers should be mindful and iron out these indiscretions related to blockchain technology in future work.

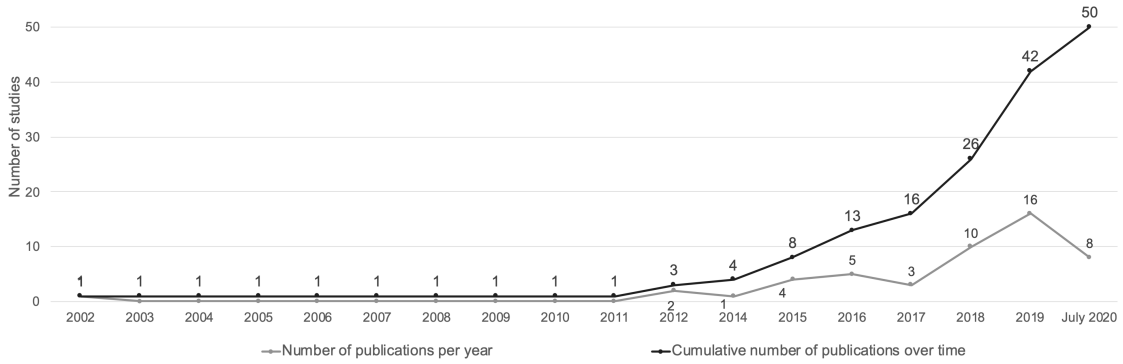

Figure 5.1.: Publications in the field of privacy-enhancing data markets for the IoT from January 2002 to July 2020 [GSU+22].

Contribution II: Improving the Applicability of Differential Privacy Algorithms.

Given the lack of work on combining PETs to enhance the privacy and security properties of systems beyond their isolated use (KF2.5), we worked on cryptographically proving the sampling of differentially-private noise with ZKP in P3 [MGSB22], which is particularly understudied (KF3.1). By providing a zero-trust algorithm construction, we improve the applicability of DP, as system designers will be more inclined to adopt DP if the machines involved can prove the use of DP. Additionally, users' trust in DSAAs can also increase. Through successfully designing such a verifiable algorithm, we extracted a series of lessons learned.

Designing an algorithm to sample (DP) exponentially distributed noise with ZKP has two significant challenges: (i) efficiently processing floats and (ii) the finite computers' inability to fulfill

the definition of DP on the real line. These challenges led to failure in several of our attempts: (1) The first (naive) attempt discretized the support of the Laplace probability distribution function (PDF) by sampling from its inverse cumulative distribution function (CDF) with a finite input range. However, we were unable to provide a provable privacy guarantee on δ, as the resulting algorithm would fulfill *approximate* DP. (2) Consecutively, we resorted to the Stone-Weierstrass theorem for approximating a polynomial to the Laplace PDF to the extent where the approximation error would be negligible. Specifically, we used Bernstein polynomials [Far12] for approximating the PDF in a closed interval and used rejection sampling thereafter. However, the flaws of our second approach were using a closed interval without condensing more probability density on the bounds and that the Bernstein polynomial coefficients are, in general, real numbers, which ZKP cannot process efficiently (KF3.2). (3) Subsequently, we tried the truncated geometric mechanism (TGM) [GRS12], as it fulfilled *pure* DP even when accounting for a truncated PDF and integers as support. However, the probabilities assigned to the integers were real numbers, which, even though we ensured they became rational by carefully choosing ε, ZKP could not handle them either because of its finite precision limitation. Likewise, we were unable to find the δ if we chose to approximate the real numbers with finite precision (KF3.3). We were successful in our fourth attempt, which we covered in more detail in section 5.2.

Contribution III: Improving the Applicability of Differential Privacy Systems.

Motivated by the potential low maturity of PET tooling (KF1.2), we further investigated to confirm such finding for DP tools specifically. Indeed, despite the broad offer of open-source tools, the expert interviews and literature review of P4 [GLMS23] revealed that DP is underused in the broad industry, and existing tools still have remaining gaps that reduce their applicability in systems. As the first steps towards precisely outlining the remaining gaps during this study, we made several observations that helped draw a picture of the current situation in the industry regarding analytics and privacy.

Firstly, we noted some shortcomings: data stewards seemed to be more careful with security measures than with privacy (KF4.1), the interviewed companies do not have a human-in-the-loop process for auditing ex-post the privacy precautions adopted by an analyst during their work, and analysts are blocked for significant periods when requesting data access (KF4.7). Additionally, we think that analyzing data without "seeing" will be standard practice in the future as it provides an additional layer of protection. Nevertheless, based on the interviewed experts, this vision is a distant reality. Regarding operations, most analysts did not use ML; instead, they relied on aggregations and visualizations to fulfill their use cases (KF4.9), and in particular, they deemed SQL as a meaningful and frequently used tool (KF4.10). These findings favor DP, as it is more easily adapted to aggregations and has shown more accuracy in general than when adapted to ML.

Most importantly, we made observations regarding the applicability of DP in practice, which were predominantly positive. We noticed that upon receiving a use case or a business-related question, analysts first had to explore many datasets, which many were inaccessible due to their sensitivity, and they had to request permission. In the worst-case scenario, after the lengthy request, the analyst might discover that the dataset is unsuitable. Thus, we deem DP an appropriate tool for sensitive dataset exploration so that analysts can explore them with reduced waiting time (KF4.8) and adversity (KF4.3) of the request process thanks to DP's

stronger privacy guarantees. Additionally, analysts shared some use cases that DP mechanisms could directly fulfill (e.g., aggregations for demographics or product analytics)(KF4.4), and we think DP can help practitioners with their need to fully access datasets "obtaining a holistic understanding of data" by providing summary statistics. In conclusion, analysts confirmed that DDPP would be helpful for dataset exploration, fulfilling certain use cases, and enabling privacy-enhancing dashboards for dataset visualization.

5.2. Discussion of Results

This section presents our results by answering the researcher questions (RQ) posed in section 1.2. Note that we created sub-RQs to help guide our discussion.

5.2.1. Contribution I: Revealing Opportunities and Challenges in the Applicability of Privacy-Enhancing Technologies in Data Sharing and Analytics Applications

With our two first publications: P1 [fGS+21] and P2 [GSU+22], we tackle the first research gap by answering RQ1.

Research question 1 (RQ1)

What are the most relevant privacy-enhancing technologies and their corresponding challenges and use cases in the scope of data sharing and analytics applications?

(RQ1.1) *Which relevant PETs can enhance DSAA?*
Garrido et al. [GSU+22] (P2) conducted a systematic literature review of publications focused on using PETs for improving the privacy measures in DSAAs. Expanding on the PETs classification of section 2.3, Fig. 5.2 is the result of the detailed study of 74 publications. Firstly, we identified that DSAA proposals had a common set of layers in their technology stack: *verification*, *storage*, *communication*, *processing*, and *sovereignty*. In all these layers, practitioners had enhanced privacy using different PETs. In summary, as expected, encryption was the backbone of the privacy enhancements at the *verification*, *storage*, and *communication* layer. In the *sovereignty* layer, where rules are established and potentially enforced, practitioners defined privacy policies and access controls. Notably, while blockchains are a privacy anti-pattern, they could be helpful to set privacy rules so that no central authority may change them ex-post.

Furthermore, we observed that most PETs focus on *processing* data. On the one hand, we identified cryptographic-based techniques dedicated to *secure and outsourced computation*, namely: ZKP, MPC, HE, and TEE (introduced in section 2.3). These techniques protect the input information (encryption) and the computation itself; however, they do not protect the outputs. Practitioners can protect the outputs with *anonymization and plausible deniability* techniques. The strongest guarantee is a *semantic* definition of privacy, i.e., the computation itself, but not the output, fulfills a privacy property that provides plausible deniability to individual inputs. DP is the quintessential semantic privacy definition and the focus of Contributions II and III.

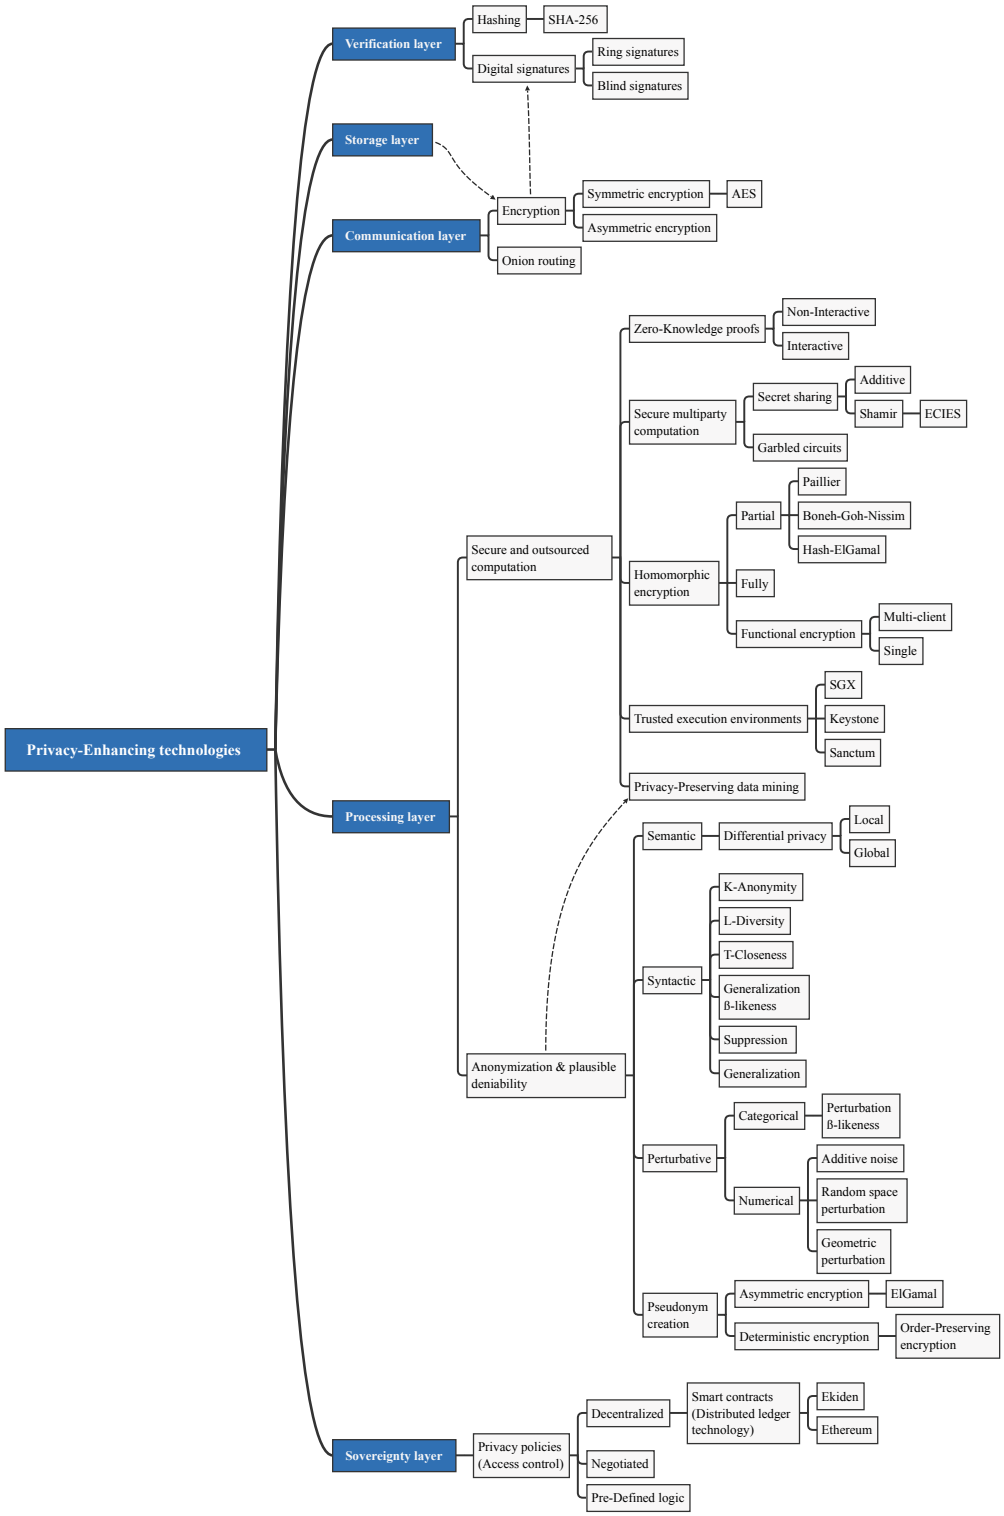

Figure 5.2.: Classification of PETs (cf. adapted [GSU+22])

Technologies	Tools
Differential privacy	Google-DP (Python wrapper: PyDP), SmartNoise, diffprivlib, DiffPriv, OpenDP, DPComp Core and Chorus (behind Uber's DP SQL). Focused on DP and deep learning: TensorFlow privacy and PyTorch Opacus.
K-anonymity	ARX, Amnesia, and Anonimatron.
Zero-Knowledge proof	emmy, dizk, zkMega, libsnark, libiop, ZKRollups, ZKRP, ckb-zkp, ginger-lib, OpenZKP, and gnark.
Secure multiparty computation	Multi-Protocol SPDZ, LIBSCAPI, MPyC, CrypTen, EMP-Toolkit, Multiparty, ZoKrates and MPC-SoK.
Homomorphic encryption	TFHE, fhe-toolkit-linux, Google FHE SEAL, Concrete, eclib, HElib, and PALISADE.
Trusted execution environments	mTower, Open Enclave SDK, Trusty, TrustZone, Mystikos, Open-TEE and Intel's Trusted Execution Technology.
Federated Learning	Fate, sherpa.ai, PaddleFL, PySft, Xaynet, fedn, FedML-AI, Flower, PyVertical, TensorFlow Federated, and federated-learning-lib.

Table 5.2.: Most prominent open-source tools for each PETs (cf. adapted [fGS+21]).

Other weaker techniques provide *syntactic* privacy definitions, i.e., the output (not the computation) fulfills a property, like k-anonymity or its variants, such as l-diversity or t-closeness. One of the main reasons why a *semantic* definition, namely DP, is stronger than *syntactic* ones is DP's context-agnostic mathematical formulation of privacy, which *syntactic* definitions lack. Other forms to anonymize data rely on perturbation without a formal guarantee or with pseudonyms. Lastly, one PET that lays in between *secure and outsourced computation* and *plausible deniability* is privacy-preserving data mining (PPDM), whose primary examples are the use of technologies such as HE or MPC to train ML models, using DP in the computation so that adversaries cannot reverse engineer the outputs, and outsourcing the training to the clients so that no central server collects potentially sensitive information from users. We refer to P2 [GSU+22] for more details, e.g., the inclusion, distinction, and discussion of authenticity-enhancing technologies like blockchains, and of other not-privacy related DSAA layers: *data auction*, *contractual*, *incentives*, and *consensus*.

(RQ1.2) *Which relevant privacy-enhancing tools are available?*
Amongst the most prominent PETs identified in P2 [GSU+22], Garrido et al. [fGS+21] (P1) conducted a gray literature review to find the most relevant open-source libraries (i.e., software that provides specific functions) and frameworks (i.e., software that provides abstractions used to build specific applications) that enable PET functionalities. With the list compiled in Table 5.2, we hope practitioners can quickly find the right tool for their privacy-forward systems.

(RQ1.3) *What are relevant use cases (opportunities) for PETs*?
In addition to identifying open-source tooling for implementing PETs, P1 [fGS+21] interviewed experts to find potential use cases where PETs can add the most value. Among the ones outlined in P1 [fGS+21], we selected a set of reference use cases in Table 5.3 as examples. While these use cases are centered around the automotive industry, they can be generalized to other sectors, e.g., data analytics in the financial sector, geoservices in mobile applications, sensitive data management in health clinics, computer vision in identification services, asset search across pharmaceutical companies, edge-computing in IoT, and sharing data across any organization. With this overview, we hope to guide practitioners in finding the right match between their use cases and a suitable PET.

#	Domain: Use Cases	Descriptions	PETs
1	Geoservices: charging	Discovering most frequent locations on an aggregated dataset where electric vehicles have low batteries.	Differential privacy, k-anonymity
2	Computer vision: attentiveness detection	Training ML models across multiple vehicles and devices.	Federated learning
3	Sensitive data management: automating anonymization	A practitioner automates the anonymization of ingested customer vehicle data.	K-anonymity
4	Data analytics: group statistics	Computing aggregate business KPIs for dashboards by querying various datasets without downloading the underlying data.	Differential privacy
5	Asset search: tracking components	Tracking components and parts across the value chain to optimize supply chain management (e. g., management of stock levels).	Secure multiparty computation
6	IoT: Connected car	Management of vast amounts of sensor data from vehicles and traffic infrastructure across the edge and cloud.	Homomorphic encryption
7	Cross-organizational data sharing: Logistics & supply chain	Track and share data across organizations to optimize business processes, e. g., for improved supply chain visibility.	Trusted execution environments

Table 5.3.: Reference use cases mapped to PETs (cf. adapted [fGS+21]).

Notably, the guide we procured in P1 [GSU+22] helped us identify a use case for an industry project (proof-of-concept) [Lab] between our industry partner and Oasis Labs that we ran in parallel and is closely related to this dissertation: #4 "Data analytics: group statistics." In P1 [GSU+22], we described such solution type as "*using a DP-aware SQL engine and a privacy budget that controls the number of queries allowed*," which corresponds to the middle-layer component of step (2) in Fig. 5.3. In practice, analysts would (1) write a query using the familiar SQL syntax. Then, (2) the middle layer component would rewrite the query so that it fulfills DP, (3) run the rewritten query in the database engine, and (4) return and (5) forward the output to the analyst. The proof-of-concept was successful, providing no more than 15% of computation overhead and maintaining the query output accuracy within 10%. With the correct adaptions, this DP middle layer could be deployed in any analytics pipeline (independently of the industry) to serve analysts explore sensitive datasets or potentially fulfill aggregation use cases with the highest standard for privacy.

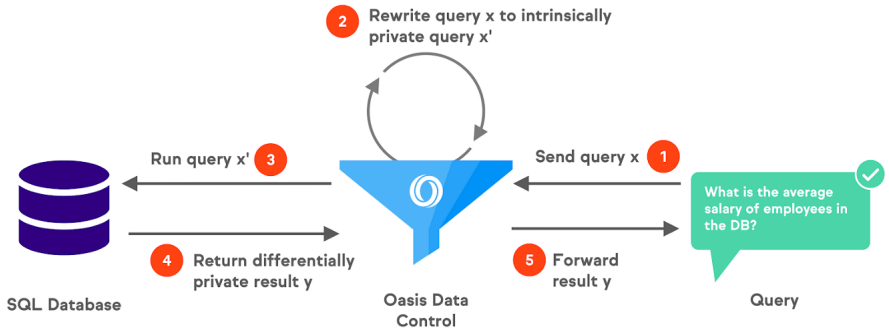

Figure 5.3.: Differential privacy component in the industry project (cf. [Lab])

(RQ1.4) *What challenges hinder privacy-enhancing DSAAs?*

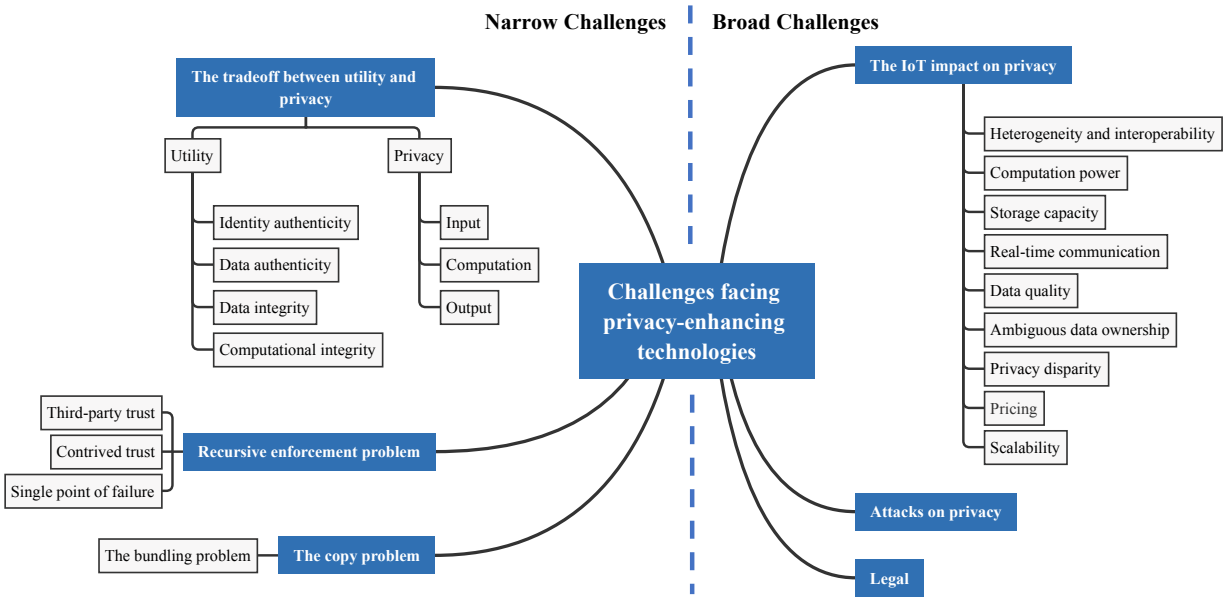

Figure 5.4.: Challenges facing PETs (cf. adapted [GSU+22]).

In P2 [GSU+22], we established the foundations for understanding and categorizing the most prominent PET. Building on this work, we then explored the remaining barriers preventing widespread PET adoption. As depicted in Fig. 5.4, we classified the challenges into two categories: *narrow* and *broad*. *Broad* challenges relate to the large-scale implications of the proliferation of more smart devices, stringent laws, and the continuous and varied attacks on privacy. For example, more devices increase the mosaic effect [Ker12], i.e., there will be more data sources available to extract significant new personal information, miscalibrated legal fines could increase the opportunity cost of deploying PETs, and the ever-present need to improve PETs as adversaries device new ways to attack personal data. On the other hand, *narrow* challenges are more concrete and actionable. The most critical one is the *copy problem*, i.e., an entity loses control of its data the moment it is shared, the *bundling problem*, i.e., subset of the *copy problem* where more data than strictly needed is shared, the *recursive enforcement problem*, i.e., the recursive need to include a new entity to oversee the incumbent parties, and striking the right balance between data utility and privacy. In P2 [GSU+22], we encourage researchers to tackle these narrow challenges, namely the *copy problem*. With the mapping between PETs and these narrow challenges in Table. 5.4, we hope to guide practitioners in navigating the challenges their applications might face. Some PETs can directly tackle the *copy problem*, e.g., TEEs allow computations without revealing the inputs. Other PETs tackle sub-problems, e.g., ZKP can verify that someone is of legal age using a digital certificate without revealing other identity information. Regarding balancing privacy and utility, some PETs can protect the inputs (identity or data), the computation, or the outputs.

Summary. Altogether, with the above, we covered the most relevant components of RQ1: (i) the most prominent PETs, (ii) their challenges, and (iii) use cases, and thus, we help **revealing the opportunities and challenges of PETs in DSAAs**. In summary, we guided practitioners looking into applying PETs to DSAA with a series of contributions: Practitioners can consult

Layer	Technology category	Technology	Narrow challenges: Privacy-utility tradeoff — Privacy [Confidentiality (Conf)] Input [Data (D)] [Identity (I)]	Privacy: Computation	Privacy: Output	Utility: Authenticity [Data (D)] [Identity (I)]	Utility: Integrity [Data (D)] [Computation (C)]	Recursive enforcement problem	Copy problem [Bundling problem (BP)]
		PET							
Processing	Secure and outsourced computation (SOC)	Zero knowledge proofs for computational integrity	++ D	++	++	++ D	++ D, C	Circumvents	Tackles (BP)
		ZKPs of anonymous credentials	++ I	++	++	++ I	++ D, C	Circumvents	Tackles (BP)
		Trusted execution environments	++ D	++	+- Conf	na	++ D, C	Circumvents	Tackles
		Partially homomorphic encryption	++ D	++	+- Conf	na	++ D, C	Circumvents	Tackles
		Fully homomorphic encryption	++ D	++	+- Conf	na	++ D, C	Circumvents	Tackles
		Secure multiparty computation	++ D	++	+- Conf	na	++ D, C	Circumvents	Tackles
		Privacy-preserving data mining — W/ SOC (w/ AN)	++ D (+ D)	++ (na)	+- , Conf (+)	na (+- D)	++ D, C (na)	Circumvents	Tackles
		Privacy-preserving data mining — Federated (w/ AN)				na (+- D)	na (na)		
	Anonymization (AN)	Differential privacy (DP)	+ D	na	+	+- D	na	na	Tackles (BP)
		K-anonymity	+- D	na	+-	- D	na	na	Tackles (BP)
		Perturbative	+- D	na	+-	- D	na	na	Tackles (BP)
		Pseudonym creation	+- I	na	+-	- I	na	na	Tackles (BP)
Storage / Communication		Encryption	Storage layer: ++ D, Conf; Communication layer: ++ D, Conf			na	++ *D	Circumvents	na
		Onion routing	++ I	+- D, Conf	na	na	na	++ D	Tackles	na
Verification		Hashing	+- D, Conf	na	na	na	++ D	Circumvents	na
		Ring digital signatures	+ I	na	na	- I	++ D	na	Tackles (BP)
		Blind digital signatures	-- *I	+ D	na	na	++ *I	na D	++ D	na	Tackles (BP)
Sovereignty		Smart contracts (for privacy policies)	Characteristics and challenges are pegged to distributed ledger technology.						
		Privacy policies (Access control)	Characteristics and challenges are pegged to the selected PETs and AETs employed to fulfil the privacy requirements.						

Legend: The extent of enhancement of privacy, utility, and characteristics of the different PETs and AETs varies from significantly increasing ++, over +, +-, - to significantly decreasing - -. na denotes *not applicable*. w/ denotes *with*.
*Considering a digital certificate when using digital signatures, if applicable. The privacy column assumes data and identity are authentic.

Table 5.4.: Mapping of PETs to the narrow challenges (cf. adapted [GSU+22]).

P1 [fGS+21] for an overview of possible use cases where PETs can add value; specifically, they can consult Table 5.3 for an *outline of use cases*, and Table 5.2 to *quickly select the right tool* (among other artifacts and insights in P1 [fGS+21]). Additionally, P2 [GSU+22] can serve as a *starting point for new practitioners in privacy* thanks to its detailed introductory explanations and classifications of PET, e.g., studying Fig. 5.2 and the detailed explanations included in [GSU+22]. Furthermore, practitioners can quickly identify a *challenge* (e.g., the copy problem) in Fig. 5.4 and the associated PET in Table 5.4 to start working on improving the applicability of PET.

5.2.2. Contribution II: Improving the Applicability of Differential Privacy Algorithms

Our third publication P3 [MGSB22] investigated the answer to RQ2. In the following, we discuss the selection of the technologies, the context of the application, and the solution's implementation and practicality.

Research question 2 (RQ2)

What new algorithm can verify the use of formal privacy guarantees with practical performance and bounded guarantees?

(RQ2.1) *Which verifiable computation technique and formal privacy guarantee are best suited for the task?*
Verifiable computation is a technique that allows proof of the execution of a particular algorithm using truthful inputs without revealing any private information [BSBHR19]. There are two main approaches to verifiable computation: trusted execution environments (TEEs) [OMT09b] and zero-knowledge proofs (ZKPs) [GMR89, Sim02, BSBHR19]. Given the known vulnerabilities and attacks on TEEs [KM22, AVBS+22, SYG+19, Int22], we decided to focus on ZKP-based approaches. *Non-interactive* ZKPs, in particular, do not require sequential messaging so that the prover can convince multiple parties of a claim with a single message [Sim02]. Therefore, we chose to use non-interactive ZKPs to enable the verification of the computational integrity of the selected privacy definition. For our implementation, we use Circom [ide18a] and SnarkJS [ide18b], a well-known open-source technology stack for implementing zero-knowledge proofs. Regarding a *formal privacy guarantee*, we unequivocally chose the strongest definition available today: DP. Specifically, we chose randomized response [War65] for binary attributes and exponentially distributed noise for continuous values [DKM+06].

(RQ2.2) *What is an appropriate context to showcase the solution?*
We propose a solution for a server running an anonymous poll on a series of client devices about a given attribute, such as calculating the average *age* of the application's user base. The clients and server are not necessarily trusted, so the server requires attestations of the shared attributes and computations. Moreover, the clients execute privacy-enhancing computations themselves, as they do not trust the server with their sensitive information. Fig. 5.5 depicts the interacting entities and Fig. 5.6 shows their interaction flow, where the *issuer* is a trusted entity (e.g., a government or institution trusted by the server and clients) that provides a digital certificate about a set of attributes (e.g., *identity*, *age*, *gender*) of the requester (client device), known as the *holder* of the certificate. In our case, this certificate is an anonymous credential, which allows an entity to attest to their attributes without revealing their identity through the use of digital signatures and ZKPs. Once the *holder* stores their credential in a digital wallet (a mobile application), the server, known as the *verifier*, sets the survey and cryptographic parameters for the poll and sends the survey to the clients. The clients, now the *provers*, use their anonymous credential to verify the use of their truthful polled value (e.g., their *age*) and employ our verifiable differentially private (DP) algorithm to create a ZKP of the DP result (i.e., *age* + Noise). In this context, where the server is also not trusted, we use *local* DP so that the clients can run

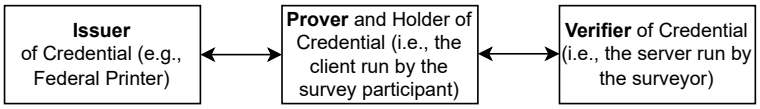

Figure 5.5.: Interacting entities of the privacy-enhanced survey.

the DP computation themselves. Finally, the *prover* sends the ZKP to the server, which verifies it and includes the result in the survey evaluation.

(RQ2.3) *How can zero-knowledge proof verify the use of differential privacy?*
In the following, we delve into the core of Contribution II; specifically, we cover applying verifiable *local* DP to (i) binary and (ii) numerical attributes.

(i) **Algorithm 1** contains a verifiable version of randomized response (RR) [War65, DR13]:

1. Flip a coin with bias p.
2. If $Bernoulli(p)$, answer truthfully.
3. Else, flip the coin again and respond "Yes" if $Bernoulli(p)$ and "No" otherwise.

RR provides plausible deniability in answers regarding one's, e.g., *gender*, *smoking habits*, *student status*, etc. In lines 1-4, we use two inputs (A, B) corresponding to the prover's private key and the verifier's challenge, and a hash function as a random oracle [CGH04]. We use two bits of the resulting hash as unbiased coin flips for the randomized response implementation of lines 5-12.

Algorithm 1: Verifiable randomized response (cf. adapted [MGSB22]).

Data: V: binary truthful value ("Yes" or "No"); A: prover contribution to randomness (secret key); B: verifier contribution to randomness (challenge).
Result: Differentially private answer.

```
1  Function VerifiableUnifRandomness(A, B):
2      S = sign(A, B)                  // sign challenge with secret key
3      R = hash(S)                     // R is an array of bits
4      return R
5  Function VerifiableRandomizedResponse(V, A, B):
6      R = VerifiableUnifRandomness(A, B)
7      if R[0] = 0 then
8          return V
9      else if R[1] = 0 then
10         return No
11     else
12         return Yes
```

(ii) **Algorithm 2** shows a verifiable version of sampling DP exponentially distributed noise. Following the discussion of key findings of section 5.1, Algorithm 2 required a solution that addresses the precision limitations and the difficulty of bounding δ. Thus, we decided to turn to simple sampling methods that provide bounds on δ. A good candidate for adaptation was Dwork et al.'s method [DKM$^+$06] for sampling exponentially distributed noise based on repeatedly flipping unbiased coins. This method is easily implemented in Circom and provides a bound for δ based on the precision achieved with Circom. However, the method from Dwork et al. [DKM$^+$06] had three limitations we had to address in our adaption: *limited precision representation*, *the zero probability* assigned to values larger than the ones we could represent, and *the double probability*

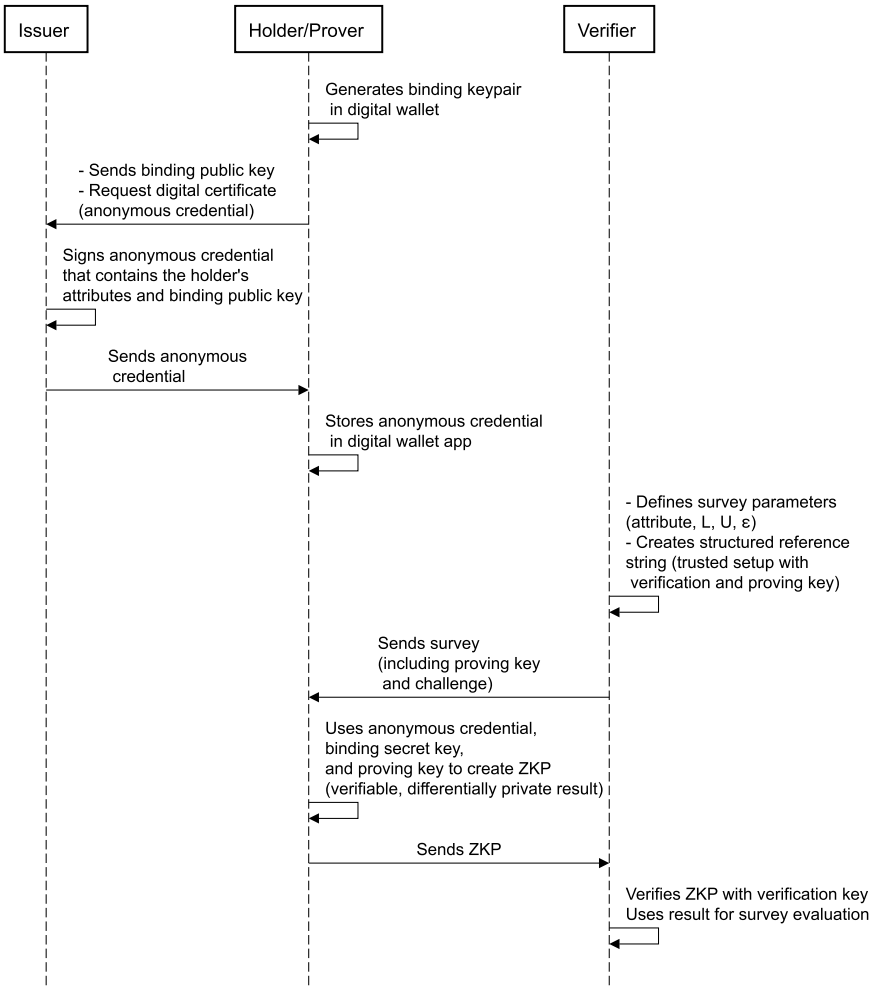

Figure 5.6.: Survey workflow with verifiable differential privacy (cf. adapted [MGSB22]).

on the distribution's center derived from flipping an unbiased coin to assign the sign of the noise. P3 [MGSB22] provides the details of our adaptations.

Algorithm 2: Verifiable exponentially distributed noise generation. By $X \bmod (L, U)$ we denote $L + (X \bmod (U - L))$ (cf. adapted [MGSB22]).

Data: V: integer-valued truthful value; U: upper bound; L, lower bound; $\Delta = |U - L| \geq 0$: sensitivity of query function; $\varepsilon \geq 0$: privacy parameter; $d \geq 0$: precision of binary expansion; A: prover contribution to randomness (secret key); B: verifier contribution to randomness (challenge).
Result: V + noise $\sim \text{Pois}(V \mid \frac{\varepsilon}{\Delta})$.

Function `VerifiableExponentialNoise`$(V, \Delta, \varepsilon, d)$:
- B_K = BinaryExpansion(Δ)
- B_v = BinaryExpansion(V)
- B_r = [] `// B_r stacks biased bits`
- **for** $k \leftarrow 0$ **to** NumBits(B_K) **do**
  - $p_k = \frac{1}{1+\exp(2^k \frac{\varepsilon}{\Delta})}$
  - B_{p_k} = BinaryExpansion(p_k)
  - `// R has at least d bits`
  - R = `VerifiableUnifRandomness`(A, B)
  - **for** $j \leftarrow 0$ **to** d **do**
    - `// Where d is the least significant bit`
    - R_j = r[j] `// R_j ∈ {0, 1}`
    - **if** $R_j = B_{j,p_k}$ **then**
      - **continue**
    - **else**
      - B_r.push(B_{j,p_k})
      - **break**
    - **if** $j = d$ **then**
      - **return** RaiseError
- noise = DecimalExpansion(B_r)
- sign = VerifiableUnifRandomness$(A, B)[0]$
- **if** (noise = 0 and sign = 0) **then**
  - **return** DecimalExpansion(VerifiableUnifRandomness(A, B)) mod (L, U)
- **else**
  - **return** $[(V + (2 \cdot \text{sign} - 1) \cdot \text{noise}]$ mod (L, U)

Fig. 5.7 shows the resulting discrete Laplace distribution used by Algorithm 2 to sample DP noise—the histogram was compiled by querying Algorithm 2 10000 times. Finally, Algorithms 1 and 2 (computational problems) are compiled by Circom [ide18a], i.e., Circom encodes them into a system of polynomial equations, with which SnarkJS [ide18b] can construct a ZKP.

(RQ2.4) *Are the proposed algorithms practical?*
By demonstrating the reasonable performance of the more complex Algorithm 2, we concluded that Algorithm 1 was also practical. In summary, the proof generation (in a commercial-grade laptop) for Algorithm 2 was of only 140 ms, its verification 0.8 s in JavaScript, and the size of the verification and proving keys are close to 3.5 kB and 3.4 MB. In addition, we tested deploying a smart contract verifier on Ethereum, which could be used for blockchain-based, incentivized, differentially private surveys and, therefore, GDPR-compliant applications on personal data. We measured the smart contract's deployment cost to be around $1,150,000$ gas and its invocation to be around $300,000$ gas. P3 [MGSB22] provides more details about the benchmark and machine specifications.

Summary. Overall, we answered RQ2 by proposing algorithms capable of verifying the use of DP so that practitioners are more willing to adopt DP for systems where mutually-untrusting entities share information, and thus, we help **improving the applicability of DP algorithms.** In summary, we proposed Algorithms 1 and 2 and their open-source implementation so that practitioners can replicate our evaluation and quickly deploy our verifiable DP solution (software component I): `https://github.com/applied-crypto/DPfeatZKP`.

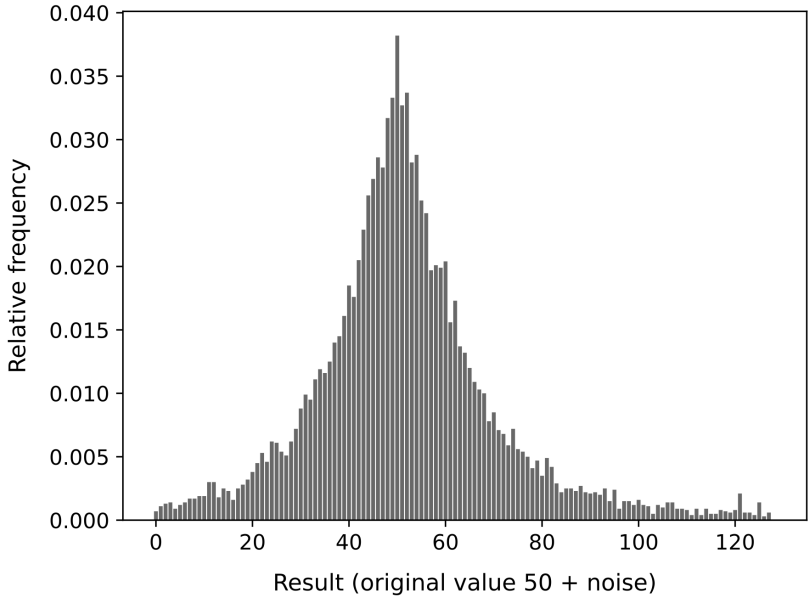

Figure 5.7.: Example histogram for $\mathbf{l} = \mathbf{0}$, $\mathbf{u} = \mathbf{128}$, $\mathbf{d} = \mathbf{20}$, $\boldsymbol{\varepsilon} = \mathbf{10}$, and true value $\mathbf{v} = \mathbf{50}$ with a sample size of **10 000** (cf. adapted [MGSB22]).

5.2.3. Contribution III: Improving the Applicability of Differential Privacy Systems

Our fourth publication, P4 [GLMS23], addresses RQ3 by identifying the remaining gaps in existing differentially private (DP) tools and proposing a set of key requirements that privacy-enhancing analytics tools should fulfill. In particular, P4 outlines the desiderata that such tools should strive to achieve to provide strong privacy guarantees while still enabling useful data analysis.

Research question 3 (RQ3)

Which are the existing gaps and key requirements a formal privacy-enhancing analytics tool should fulfill to become closer to a broad industry adoption?

In P4 [GLMS23], we conduct expert interviews to understand the applicability of differentially private (DP) methods to the analysis workflow in the broader industry. The first five sub-RQs provide context and summarize our conclusions from these interviews.

(RQ3.1) *What is the context of privacy protection in the targeted organization?*
The interviewed organizations invest more resources in security than in privacy-enhancing analysis. Additionally, stewards view privacy as an asset and strive to provide the best standards for their customers. However, these organizations still rely on traditional anonymization techniques, which have limitations in terms of protecting privacy. Furthermore, companies today are unable to effectively measure the privacy of their processes, and while they may have specific privacy and security measures in place, such as access controls, these are often difficult to quantify for-

mally. This lack of tangibility can make it challenging for organizations to accurately assess the privacy of their data and processes. Lastly, we observed that the companies we interviewed are far from having "*data at their fingertips*", largely due to onerous dataset request processes. This can hinder their ability to perform effective data analysis.

(RQ3.2) *Could differential privacy tackle the privacy-related pain points of an analysis workflow in an organization?*
Yes, to a large extent. In essence, the main problems that hinder effective data analysis are (i) lengthy and cumbersome dataset request processes, and (ii) the fact that once analysts are granted access to data, they can sometimes "see", download, and share the data, potentially even colluding with other co-workers to link sensitive datasets. In the latter case, only policies protect the data, which may not be sufficient. Based on our work, we argue that DP can reduce the time it takes to access data by enabling exploration of critically sensitive data or across third-party data sources. It can also relax current data access restrictions thanks to its formal privacy guarantee, and is applicable to some aggregation-based use cases. For certain use cases, engineers should consider building solutions that prevent analysts from "*seeing*" the data directly, in order to protect privacy.

(RQ3.3) *When does differential privacy impede an analysis?*
The answer to this RQ depends heavily on the specific use case and whether the analysts are willing to sacrifice some accuracy in order to protect privacy. On the one hand, noise addition-based DP is useful for some the types of aggregations performed by the interviewees (e.g., querying demographics or frequently used product features). Additionally, on average, the interviewees were comfortable with a 98% accuracy rate for their analyses. However, it is important to note that DP is not a perfect solution for all situations. Some of the interviewees' use cases should not rely on DP (e.g., error analyses or critical financial estimations). Therefore, we suggest building systems that enable DP while still providing the flexibility to allow non-differentially private queries when the use case requires it. This would allow organizations to balance the need for privacy with the need for accurate data analysis in critical use cases.

(RQ3.4) *How would differential privacy affect the workflow of an analyst?*
If DP were to enable previously unavailable data exploration and provide data for privacy-enhanced dashboards, analysts would likely have a better user experience in their workflow, with reduced time spent on data request processes and exploration. However, they would also need to become accustomed to working with noisy data, which can be challenging for some types of analysis.

(RQ3.5) *Can differential privacy be applied to the frequent SQL-like queries analysts execute?*
Yes, based on our research, around a third of the aggregations performed by the interviewees were amenable to DP. This suggests that DP could be a useful tool for improving the privacy of data analysis in some cases, but it may not be suitable for all types of aggregations. It is important to carefully consider the specific use case and the tradeoffs involved in order to determine whether DP is an appropriate solution.

Based on the research and conclusions presented in P4 [GLMS23], and the authors' industry experience in DP projects (see RQ1.3), we have extracted a set of key system desiderata for a holistic privacy-enhancing analytics tool.

(RQ3.6) *What is the set of critical requirements that a differential privacy (DP) analytics system should satisfy in order to be suitable for practical deployments?*
From P4 [GLMS23], converged in the following 10 key desiderata:

(I) *Differentially private analytics.* The system adds noise to the outputs of learning functions in order to preserve privacy, and it supports a variety of aggregation queries as well as machine learning features. It also allows for the storage of executed queries for future reference.

(II) *Usability.* The system ensures the (i) semantic consistency of queries, and offers options for (ii) estimating sensitivity and (iii) setting privacy parameters automatically.

(III) *Security.* The system (i) automatically verifies that the algorithm meets privacy requirements using a stochastic tester or other functions. It employs (ii) cryptographically secure pseudo-random number generation, and (iii) generates noise so that is resistant to floating-point vulnerabilities. Additionally, the system (iv) blocks users from accessing the data directly and (v) does not allow them to execute arbitrary code. The system executes (vi) heuristic optimizers only at post-processing and (vii) protects against timing attacks.

(IV) *Synthetic data generation.* The system produces synthetic data for testing and exploring machine learning models, and it allows analysts to proceed with real data without accessing it directly. Synthetic data can be generated using simple techniques, machine learning, or a combination of both. If the analyst is only interested in the data schema, the system produces dummy data that preserves only the schema and data types.

(V) *Visualization.* The system's dashboard provides interactive plots based on DP queries for easy dataset exploration and visualization of analysis' expected accuracy, disclosure risk, uncertainty, statistical inference, and budget splitting.

(VI) *Privacy budget.* The system (i) tracks the privacy budget spent, (ii) prevents further queries if the budget is exhausted, and (iii) accommodates for growing datasets. It also allows data stewards to (iv) specify budgets for teams, individual analysts, and use cases based on the sensitivity of the data.

(VII) *Accuracy adjustment.* The system allows the user to specify a desired accuracy level for their query, and it provides information about the noise scales or a confidence interval after the query is executed.

(VIII) *Query sensitivity.* The system allows practitioners to input the attributes' bounds as function parameters or in the dataset schema, enabling the system to calculate the sensitivity of the query.

(IX) *Privacy-sensitive data annotation.* The system allows data stewards to specify which attributes are accessible to different teams, roles, and use cases, and it automatically obfuscates attributes that are not on the allowlist.

(X) *Authentication and access controls.* As a security measure, most organizations require data stewards and analysts to authenticate before using tools that handle data. Thus, the system easily integrates with existing authentication and access control services and allows data stewards to define their own access policies.

Tool/Desiderata	(I) DP Analytics	(II) Usability	(III) Security	(IV) Synthetic Data	(V) Visuals	(VI) Privacy Budget	(VII) Accuracy Adjustment	(VIII) Query Sensitivity	(IX) Data Annotation	(X) Access Controls
Libraries[†]										
diffprivlib [53]	I.i, ii ✓	II.i ✓	III.ii, iii ✓	✗	N/A	VI.i ✓	✗	✓	N/A	N/A
Google DP [39]	I.i ✓	II.ii ✓	✓	✗	N/A	VI.i ✓	✗	✓	N/A	N/A
Opacus [74]	I.ii ✓	✗	III.ii ✓	✗	N/A	VI.i ✓	✗	✓	N/A	N/A
OpenDP [48]	I.i ✓	II.iii ✓	III.ii, iii ✓	✗	N/A	VI.i, ii ✓	✓	✓	N/A	N/A
TF Privacy [41]	I.ii ✓	✗	✗	✗	N/A	VI.i ✓	✗	✓	N/A	N/A
Frameworks[†]										
Chorus [61]	I.i, iii ✓	✗	✗	✗	N/A	VI.i ✓	✗	✓	✓	N/A
PipelineDP [83]	I.i ✓	✗	III.ii, iii ✓	✗	N/A	VI.i ✓	✗	✓	✗	N/A
P. on Beam [40]	I.i ✓	II.ii ✓	✓	✗	N/A	VI.i ✓	✗	✓	✗	N/A
Tumult Analy.[99]	I.i, iii ✓	✗	✓	✗	N/A	VI.i, ii ✓	✗	✓	N/A	N/A
ZetaSQL [43]	I.i, iii ✓	II.ii ✓	✓	✗	N/A	✗	✗	✓	✗	N/A
Systems										
Airavat [90]	I.i, ii ✓	✗	III.iv ✓	✗	✗	VI.i, ii ✓	✗	✓	✗	✓
DJoin [78]	I.i, iii ✓	✗	III.ii, iv, v ✓	✗	✗	VI.i, ii ✓	✓	✓	✗	✗

[†]Libraries' and frameworks' (III) Security scope is limited to three sub-desiderata (i), (ii), and (iii).

Table 5.5.: Mapping between open-source tools and the key system desiderata (cf. [GLMS23]).

(RQ3.6) *What are the current gaps in the state-of-the-art DP tools that prevent them from being fully practical?*
We examined the most relevant open-source libraries, frameworks, and systems to determine whether they fulfill the key system desiderata (see Table 5.5). Based on the missing building blocks we identified, we highlight seven gaps in usability, security, synthetic data generation, visualizations, privacy budget, and accuracy adjustment, as well as functionality for data stewards. We strongly encourage library designers to use secure random number generation, implement patches for the floating-point vulnerability, ensure output semantic consistency, and advance research and development of DP synthetic data generation tools. More details about the outstanding gaps can be found in P4 [GLMS23].

(RQ3.7) *What would a comprehensive DP system design that meets all the key system desiderata involve?*
To help practitioners in the adoption of the ten system requirements, in Fig. 5.8, we depict a blueprint of an initial design of a system that complies with all the proposed requirements—More details can be found in the Appendix of P4 [GLMS23]. Additionally, while not the focus of P4 [GLMS23], we developed an early-stage benchmarking tool (software component II, see Fig. 5.9) capable of comparing DP libraries and frameworks to help practitioners decide which one to use for some of the components in Fig. 5.8. For example, the benchmarking tool can output which library provided the best accuracy based on an input dataset and query.

Summary. Altogether, by answering RQ3, we have contextualized and assessed the applicability of DP in analytics workflows in the broader industry and have identified a series of critical system requirements and outstanding gaps in state-of-the-art DP tools. These findings will help to **improve the applicability of DP systems** in industry settings. In addition to delineating a road map for DP tool improvements, we provided in Fig. 5.8 a blueprint to spark the interest of practitioners looking into the DP systems to adopt the key desiderata and developed a tool (see Fig. 5.9) to help them compare different DP components for their system.

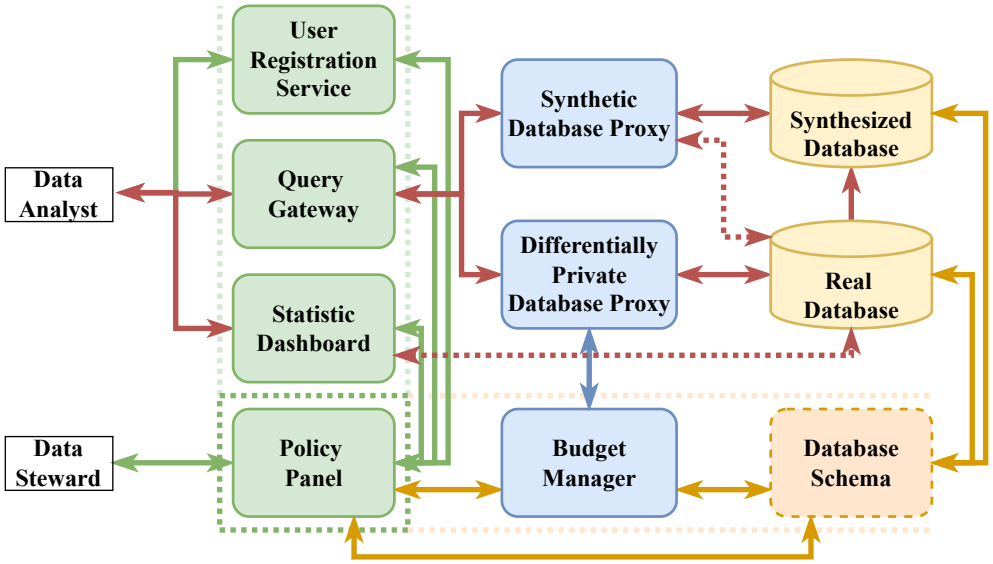

Legend: Solid lines represent communications between components triggered by all relevant query events, while dashed lines represent communications that happen periodically or only under certain circumstances. We specify those circumstances in the workflow description in the body.

Figure 5.8.: System design blueprint of a privacy-enhancing analytics tool (cf. [GLMS23]).

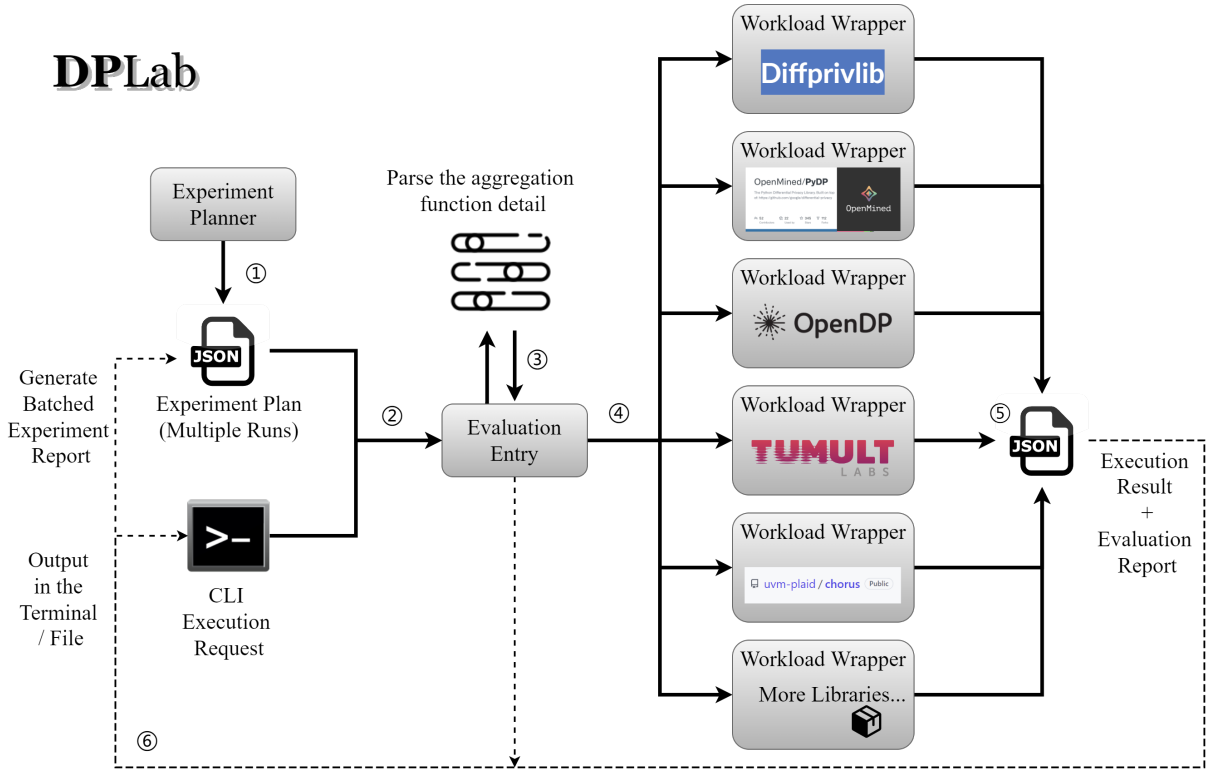

Figure 5.9.: DPLab: Benchmarking tool architecture.

5.3. Limitations

This section describes the limitations of our scientific work and the measures we took to maintain the validity of our studies. These limitations are more prevalent in the studies with a secondary component (e.g., reviews or interviews) (P1 [fGS+21], P2 [GSU+22], and P4 [GLMS23]) than in the proposed software components of P3 [MGSB22] and P4 [GLMS23], as these components are not as strongly dependent on reviews and interviews. The four major characteristics of a valid study, as described in previous research [Bha12, RH09, ESSD08], are:

(i) Internal validity examines how much the results of a study can be confidently attributed to the specific cause-and-effect relationship being studied [RH09]. The works with a secondary study type component (P1 [fGS+21], P2 [GSU+22], and P4 [GLMS23]) rely on the researchers' analyses to extract findings and conclusions from literature reviews and expert interviews, which can be affected by biases [ESSD08]. To counter this limitation, multiple researchers conducted literature reviews, resolved inclusion conflicts, and collaborated on creating artifacts. However, the software components in P3 [MGSB22] and P4 [GLMS23] are mainly free of this bias.

(ii) External validity relates to how generalizable are the study's results [RH09]. In the case of studies with a secondary study type component (P1, P2, P3), external validity concerns the representativeness of the body of works included in the study. In order to mitigate this risk, the authors of this study followed strict and inclusive search and filtering criteria. They also reached out to as many potential respondents from a variety of industries as possible in their expert interviews. While the number of interviewees was sufficient to extract insightful results (17 in P1 and 24 in P4), the authors acknowledge that more interviewees would have strengthened the findings. To address this limitation, they supported the findings and artifacts from the interviews with insights from literature reviews. Finally, the software components in P3 and P4 are predominantly free from this limitation.

(iii) Construct validity measures the degree to which a study's metrics accurately represent the theoretical concepts they are intended to assess [RH09]. The publications with a secondary study type component (P1, P2, and P4) may be subject to incompleteness due to the limited selection of search keywords and sources for their literature reviews. Multiple researchers curated the search strings for mitigation, and we used the major digital libraries available [KB13, PVK15], e.g., ACM Digital Library, Scopus, Springer Link, and IEEE Xplore, among others. Concerning the expert interviews, we ensured enough time to discuss as many of the details as the interviewees could share. The researchers also refined the questionnaires through multiple rounds of refactoring to ensure that the research questions were answered appropriately. A defective design could affect the software components in P3 and P4. To mitigate this risk, we designed the software components not only based on the authors' experience but also on the results of the interviews, literature reviews, and related software components.

(iv) Reliability in scientific work refers to the consistency and stability of a study's results. This means that a study is considered reliable if another researcher can replicate the results under the same conditions [RH09]. To mitigate this threat in all of our publications (P1-4), we included detailed *method* sections (inspired by renown studies [KC07, DKM+06, DKM19]) so that any researcher could conduct the same study. Additionally, we open-sourced the code for the software components I (P3) and II (P4) so that other researchers can curate the programs

and replicate the benchmarks guided by the detailed specifications of the machines we used to carry out the evaluations.

CHAPTER 6

Future Work

We outline the various avenues for future work in Table 6.1.

No.	Future Work
Expanding Contribution I: Opportunities and Challenges in the Applicability of PETs	
P1	• Extend the use cases classification to other applications and domains. • Experiment with concrete PET in productive environments. • Build a comprehensive decision tree to enhance the actionability of our guidance.
P2	• Tackle the *copy problem*, i.e., the problem that occurs when an entity releases data, the data is no longer under the original owner's control. • Improve IoT devices' computation and storage to tackle the limitations of deploying PETs. • Combine complementary PETs in meaningful use cases. • Develop data exchange and privacy standards and universal privacy-forward APIs. • Investigate whether monetizing privacy in a competitive market ultimately benefits society. • Disambiguate data ownership to define clear privacy and monetization policies. • Design an information-theoretic model of the data leakage in data exchanges and analytics.
Expanding Contribution II: Improving the Applicability of Differential Privacy Algorithms	
P3	• Obtain tighter δ bounds for using Poisson distribution in our solution. • Develop frameworks to bound δ in *approximate* DP. • Validate whether the verification key sizes are practical for web-based mobile applications. • Improve ZKP circuit compilers' precision limitations. • Study how users perceive built-in trust and how to convey these measures. • Design new primitives to verify other DP mechanisms. • Design a scheme that uses ZKP to verify the use of DP in the *central* model.
Expanding Contribution III: Improving the Applicability of Differential Privacy Systems	
P4	• Tackle the outlined engineering and research gaps in DP in practice. • Improve guidance to select ε. • Studying and communicating to non-experts how mechanism design decision affect utility. • Study the unpredictable artifacts introduced by existing complex DP algorithms (e.g., bias). • Increase the maturity of DP ML and synthetic data generation.

Table 6.1.: Overview of future work items (cf. [GSU+22, MGSB22, fGS+21, GLMS23]).

We propose future work within each of the three contribution streams of this dissertation and present ongoing work that pushes Contribution III forward and establishes a new research stream focused on improving the applicability of DP in VR applications. We hope this section inspires and acts as a starting point for researchers and academics looking into applied DP.

6.1. Expanding Contributions I, II, and III

Expanding Contribution I: *Opportunities and Challenges in the Applicability of PETs.* Current research in the practical use of PET does not explore systematic methods and tools for helping practitioners decide which PET is suitable for a given need and application domain [fGS+21]. Moreover, reports and performance studies on the use of PETs in productive applications are lacking [fGS+21]. Additionally, there are numerous outstanding challenges in the field of privacy that need urgent attention [GSU+22]. Primarily, solving the *copy problem*, i.e., the loss of control over information upon unprotected data sharing, by using solutions like TEEs (not-yet-matured [KM22, AVBS+22, SYG+19]) would empower users and enable digital platforms to use critically sensitive data for new applications, e.g., in the health industry [ZLC+21, CGDS+20, SCS18]. Tackling the *recursive enforcement problem*, i.e., the recursive need of additional third parties to supervise an incumbent party in distributed systems, by employing, e.g., data shielding PET, would eliminate some of the obstacles towards zero-trust data exchange applications [TBG+20b]. Moreover, researchers and practitioners can tackle broader challenges related to the current IoT devices' limitations to spare compute for running PETs [RRK+20, RPX+22], explore new combinations of PETs [BS23, DM23], and design models to quantify the data leakage of data exchanges and analytics [LRYY19]. In addition to future work items in the technical domain, there are numerous avenues for economics and legal experts to contribute [GSU+22]. Examples of critical questions revolve around the economic impact of too-protective privacy measures, whether monetizing privacy is beneficial for society, the ambiguity of data ownership, and developing standards for privacy-forward data exchange and APIs.

Expanding Contribution II: *Improving the Applicability of Differential Privacy Algorithms.* Throughout the writing of P3 [MGSB22], we iterated through several algorithms until we identified the most fitting to adapt for verifiability with ZKP. The main issue we encountered was the lack of guidance on how to bound δ for *approximate* DP and find tighter bounds. The clarity in these two aspects would allow researchers to contrive and adapt DP algorithms more quickly. Concerning practical aspects, validating whether the sizes of the verification keys are suitable for web-based mobile applications would shed light on the extended applicability of our solution. Moreover, tailoring the ZKP techniques to our DP algorithm would increase its performance, and overall, improving ZKP circuit compiler's precision would help in the design and performance of our algorithm. In terms of usability, it is worth studying what messages are most effective for users to understand the robust privacy measures of our solution and whether that could influence their willingness to participate. Additionally, researchers still have a plethora of DP algorithms to adapt for verifiability with ZKPs, e.g., the exponential or the Gaussian mechanisms. Lastly, adapting our scheme to the *central* model of DP would increase the accuracy of the outputs. With an MPC component, interviewees can share their deterministic value secretly (e.g., their age), and, together with the server's ZKP of the sampled DP noise verifiable by all

clients, the scheme can calculate a verifiable DP function (e.g., the average). If a party drops out, the process can fall back to the *local* mode of DP.

Expanding Contribution III: *Improving the Applicability of Differential Privacy Systems.* In P4 [GLMS23], we provide an explicit set of engineering and research gaps and the most critical DP general challenges that block the broad adoption of DP. We suggest practitioners, researchers, and library designers work in these gaps and challenges to close the gap between theory and practice. Among the gaps, we prompt practitioners first to ensure the security vulnerabilities of their tools, and subsequently, we suggest improving the tools' usability and visualizations and including synthetic data generation capabilities. However, beyond working on tool improvements, institutions should also become more flexible in their technology deployments. Regarding the significant challenges of DP as a technique, the community has yet to see guidelines for the appropriate selection of the privacy parameter ε, track privacy budgets across institutions, and develop robust methods to verify the fulfillment of DP. Additionally, the use of DP in ML and synthetic data generation is still in its initial stages. Lastly, we encourage researchers to study the artifacts introduced by new complex DP algorithms (e.g., biases) and understand and communicate how algorithm design decisions affect the output utility (e.g., bounding outputs for semantic consistency).

We are currently working on some of the above improvements in PP1 (see Table 6.2). DP is a complicated topic, and the number of current open-source tools is considerable enough to overwhelm practitioners unfamiliar with DP and its utilities. Thus, we propose a framework to benchmark DP libraries. A key feature of our framework is its extendability, i.e., other practitioners can add new libraries and functions to our open-source project, which is critical to prevent obsolescence. With this benchmarking tool, we aim to provide a comprehensive benchmark of the most relevant DP libraries and frameworks. These results will help practitioners better understand the available tools' capabilities and help library designers identify learning across libraries.

6.2. New Research Stream: Improving the Applicability of Differential Privacy in Virtual Reality Applications

VR and the so-called "metaverse" present a new hypothetical medium to access and interact with the internet. However, unlike the available tools to protect user privacy in web 2.0 like VPNs, Tor, or "incognito mode," there is yet not an equivalent suite of privacy tools to tackle the new attack vectors VR exposes. Particularly, given the wealth of new data VR devices can collect, VR can increase the ease at which VR users are profiled and tracked across internet sessions. To start tackling these issues, we drew the landscape of data privacy in VR by conducting a systematization of knowledge (PP2). We proposed threat and defense models so that researchers could easily frame their future work and provide a taxonomy of data attributes available in VR to better understand the scope and gravity of future privacy attacks. Moreover, based on existing attacks and their associated defenses, we highlighted the gaps in the research necessary to develop a holistic defensive VR framework. our ongoing work in privacy and VR is succinctly summarized in Table 6.2.

No.	Ongoing Work
Expanding Contribution III: Improving the Applicability of Differential Privacy Systems	
PP1	• Build an extendable framework to benchmark DP libraries and frameworks. • Benchmark existing open-source DP libraries and frameworks. • Draw insights for practitioners using and designers of DP libraries and frameworks.
New Research Stream: Improving the Applicability of Differential Privacy in VR Applications	
PP2	• Provide taxonomies of data attributes, and privacy attacks and defenses in VR. • Outline the gaps of privacy defenses in VR. • Outline the most vulnerable and easiest to protect data attributes in VR.
PP3	• Show how malicious VR application design could make users reveal information unknowingly. • Reveal concrete privacy attack vectors in VR applications. • Measure the effectiveness of such attacks.
PP4	• Design a suit of VR privacy protections with provable privacy guarantees. • Evaluate the effectiveness degradation of the attacks when the privacy defenses are enabled. • Measure the usability impact when the privacy defenses are enabled.

Table 6.2.: Overview of ongoing work items (cf. [GNM+21, NGS24, NMGS23, NGS23]).

VR Privacy Attacks. Having an overview of the field of data privacy in VR, we moved on to show how adversaries could carry out these privacy attacks effectively in practice. For this endeavor, we conducted a user study, *MetaData* (PP3) [NGS23]. In *MetaData*, thirty volunteers participated in an innocent-looking "escape room" VR game (Fig. 6.1 shows the virtual building where participants had to solve the escape room's puzzles). During the playthroughs, we collected in-game data to accurately infer over 25 personal data points, ranging from demographics such as gender and age to anthropometrics like wingspan and height, within minutes of starting the game. The attackers in a VR data flow may be other VR users or as strong as the server running multiplayer functionality or the VR device and application running locally. By calculating the accuracy of the demonstrated attacks, we showed the potential scale and breadth of data collection in VR could far exceed what is possible within the traditional web and mobile applications. Overall, we aimed to bring attention to the looming privacy risks of the metaverse by providing the first framework for understanding the potential pervasiveness of emerging VR ecosystems.

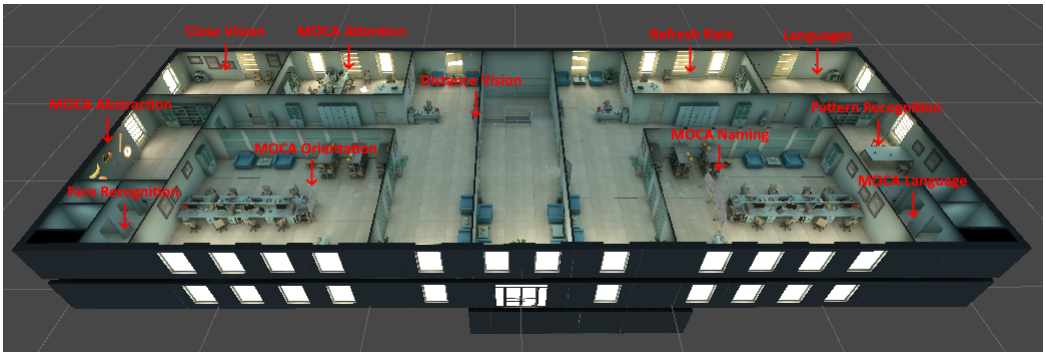

Figure 6.1.: Virtual office building hosting the escape room.

VR Privacy Defenses. Given that privacy defenses are still understudied in the field of VR and the increased attention VR has received by data-hungry companies, we find ourselves in a dangerous situation where the privacy tools of the current internet cannot cover the new attack

surface exposed by a potential surge in VR adoption. Consequently, we worked in a defense suite to tackle the privacy issues uncovered in *MetaData* (PP3) [NGS23], which could soon become a typical VR experience.

MetaGuard (PP4) [NMGS23] constitutes the implementation of a series of DP techniques to quantifiably obfuscate user sensitive data points. Specifically, we used local DP to add noise to geospatial telemetry data (i.e., to the headset's and controllers' X, Y, and Z coordinate positions) before streaming such information to servers and other users. As a result, the would-be adversaries could only read values that are DP. Thus, as we showed in our evaluation, adversaries can no longer accurately infer attributes such as height, wingspan, age, gender, and identity, effectively preventing profiling and identification (tracking across VR sessions). The underlying DP techniques we used are randomized response [War65] for binary attributes like handedness (i.e., we mirror the avatar based on the outcome of a series of coin flips), and the Bounded Laplace mechanism [HABMA19], which samples noise from a bounded Laplace distribution so that values can preserve their semantic consistency (e.g., we prevent sampling negative height values). DP has the additional unique property to optimally balance privacy and usability, which is critical in VR.

While protecting attributes like interpupillary distance (IPD) is straightforward because such value types do not change throughout a playthrough, other data attributes like height and wingspan demand more careful consideration. For IPD, we simply sampled a calibrated noise distribution centered around the deterministic value and use the sampled value as the new IPD for the VR session. If we employed the same technique for wingspan, when users would touch their hands in the real world, their virtual hands would not because of the DP offset. Thus, we added DP ly, i.e., as the user extends their arms apart, the noise becomes larger, becoming a DP wingspan value once their arms are fully extended. This technique effectively protects the DP offset and ameliorates the impact in user experience. In practice, the noise value would be maintained in the same VR session but would be re-sampled in new sessions so that they are statistically unlinkable. Assuming the use of complementary PET such as VPNs or proxies to protect users' IP, the risk of user tracking is effectively reduced.

We implemented our suite of DP techniques as a universal Unity (C#) plugin that virtually any user running an application that supports Melon Loader can enable. Because any user unfamiliar with privacy or DP should be able to use our "incognito mode," we developed a simple user interface that provides users with multiple options to protect their privacy. Fig. 6.2 shows *MetaGuard*'s user interface. The user can enable *MetaGuard* with the flick of a switch, and further select the different attributes to protect. Furthermore, we provide a privacy slider to choose the strength of the privacy protection—the slider controls the ε values. To select the default values set in *MetaGuard*, we further conducted a small empirical study to assess how the ε values affect the usability in different VR environments, ranging from competitive games to social VR.

Overall, we hope our taxonomies (PP2), attack framework (PP3) [NGS23], and provable privacy guarantees in *MetaGuard*'s defenses (PP4) [NMGS23] help and encourage researchers to continue developing defenses in emerging system like VR.

Figure 6.2.: Mixed reality photo of a user enabling “MetaGuard,” our implementation of the first proposal for a VR “incognito mode”.

CHAPTER 7

Conclusion

As the internet has evolved and the methods for collecting and processing personal information have become more complex, there has been a growing need for privacy-enhancing technologies (PETs) to balance privacy and utility. Among PETs, differential privacy (DP) has become the standard for providing a provable privacy guarantee. However, despite its promise, many PETs, including DP, have not yet been widely adopted in practice. In this dissertation, we addressed this issue by studying and developing the improvements necessary to increase the applicability of PETs, particularly DP, in industry.

To help tackle the lack of widespread adoption of PETs in practice, and of DP in particular, we make three contributions comprising four research papers (P1 [fGS+21], P2 [GSU+22], P3 [MGSB22], and P4 [GLMS23]). Our first contribution is to reveal the opportunities and challenges in the applicability of PETs in data sharing and analytics applications (DSAAs) (P1, P2). Our second contribution is to improve the applicability of DP algorithms (P3). Finally, our third contribution is to improve the applicability of DP systems (P4). Together, these contributions aim to tackle the barriers to adopting DP by studying and developing the improvements necessary to increase their applicability in the broader industry.

Conclusion of Contribution I. P1 [fGS+21] and P2 [GSU+22] investigate the challenges facing PETs and the opportunities of these technologies in DSAAs. P1 [fGS+21] provides guidelines for selecting PETs based on the capabilities required by the use cases collected from expert interviews. The paper highlights the need for understanding the capabilities and limitations of PETs as well as the characteristics of the use case when choosing a PET before deploying it in practice. We conclude in this paper that there is no one-size-fits-all PET and that caution should be exercised when integrating PETs into systems. Moreover, in P2's [GSU+22] systematic literature review of 74 publications, we found that the combination of secure computation and provable privacy guarantees are rare in PETs in DSAAs, and that blockchains are often

included without a clear set of reasons and despite the lack of practical examples in the industry. Moreover, the study found that many researchers tend to reinvent the wheel instead of improving and contributing to existing open-source tools, and concluded that there is not yet a canonical solution for building privacy-enhancing DSAAs, meaning the practicality of PETs needs to improve.

Conclusion of Contribution II. P3 [MGSB22] helps tackle one area of improving the applicability of DP algorithms—verifiability. We proposed an algorithm using zero-knowledge proofs to verify the correctness of a DP query output. Moreover, we showed that our approach had practical performance and can be used to verifiably query individuals' data in a DP manner. The study's findings provide the privacy community with primitives for implementing cryptographically verifiable DP in the local model, which can prevent malicious clients from adding noise to their local data, which could negatively impact the accuracy of a server's aggregate statistics.

Conclusion of Contribution III. Lastly, P4 [GLMS23] delves into DP systems and how to improve their applicability. Specifically, after interviewing 24 practitioners from 9 major companies, we found that analysts often face lengthy bureaucratic processes for requesting access to sensitive data and that once granted, only scarcely-enforced privacy policies stand between rogue practitioners and misuse of private information. In conclusion, we argue that DP could reduce data access request times by allowing the exploration of critically sensitive data across silos and reducing access restrictions thanks to its stronger privacy guarantees. Based on our research, with a series of feasible yet critical improvements, the practical use of DP across the enterprise is within reach.

As a result of our contributions, we have provided the privacy community with new insights and tools to improve the applicability of PETs and DP in particular. By revealing the opportunities and challenges of PETs in data sharing and analytics applications (Contribution I), improving the applicability of DP algorithms (Contribution II), and improving the applicability of DP systems (Contribution III), we aim to bring academia and industry closer together and advance the state of the art in PETs. These contributions not only support the ethical pursuit of protecting user privacy while enabling the extraction of valuable insights from data, but also provide practical tools for practitioners to adopt and integrate PETs into their systems.

Bibliography

[AG21] Apple and Google. Exposure notification privacy-preserving analytics (enpa). `https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ENPA_White_Paper.pdf`, 2021. Online; accessed 15 August 2022.

[AGKZ18] Maryam Archie, Sophie Gershon, Abigail Katcoff, and Aileen Zeng. Who′s watching? de-anonymization of netflix reviews using amazon reviews. 2018.

[AGT14] Ghada Arfaoui, Saïd Gharout, and Jacques Traoré. Trusted execution environments: A look under the hood. In *2014 2nd IEEE International Conference on Mobile Cloud Computing, Services, and Engineering*, pages 259–266, 2014.

[App17] Apple Differential Privacy Team. Learning with privacy at scale. `http://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html`, 2017. Online; accessed 18 February 2022.

[AVBS+22] Fritz Alder, Jo Van Bulck, Jesse Spielman, David Oswald, and Frank Piessens. Faulty point unit: Abi poisoning attacks on trusted execution environments. *Digital Threats: Research and Practice*, 3(2), 2022.

[Bak00] Michael Baker. Writing a literature review. *The Marketing Review (TMR)*, 1(2):219–247, 2000.

[BBG+21] Daniel M Bittner, Alejandro E Brito, Mohsen Ghassemi, Shantanu Rane, Anand D Sarwate, and Rebecca N Wright. Understanding privacy-utility tradeoffs in differentially private online active learning. *Journal of Privacy and Confidentiality*, 10(2), 2021.

[BBK+09] Endre Bangerter, Stefania Barzan, Stephan Krenn, Ahmad-Reza Sadeghi, Thomas Schneider, and Joe-Kai Tsay. Bringing Zero-Knowledge Proofs of Knowledge to Practice. page 12, 2009.

[BCD21] Kyle Bittner, Martine De Cock, and Rafael Dowsley. Private speech classification with secure multiparty computation, 2021.

[BCR94] Victor Basili, Gianluigi Caldiera, and Dieter Rombach. The goal question metric approach. *Encyclopedia of Software Engineering*, pages 528–532, 1994.

[BGN05] Dan Boneh, Eu-Jin Goh, and Kobbi Nissim. Evaluating 2-dnf formulas on ciphertexts. volume 3378, pages 325–341, 02 2005.

[BGV14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. *ACM Trans. Comput. Theory*, 6(3), July 2014.

[Bha12] Anol Bhattacherjee. *Social Science Research: Principles, Methods, and Practices*. CreateSpace, Scotts Valley, CA, USA, 2012.

[BIK+17] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, CCS '17, page 1175–1191, New York, NY, USA, 2017. Association for Computing Machinery.

[BMSS13a] Sabine Buckl, Florian Matthes, Alexander Schneider, and Christian Schweda. Pattern-based design research – an iterative research method balancing rigor and relevance. In *Proceedings of the 8th International Conference on Design Science Research in Information Systems and Technology (DESRIST)*, pages 73–87, Berlin, Heidelberg, Germany, 2013. Springer Berlin Heidelberg.

[BMSS13b] Sabine Buckl, Florian Matthes, Alexander Schneider, and Christian Schweda. Pattern-based design research in enterprise architecture management. In *Proceedings of the 25th International Conference on Advanced Information Systems Engineering Workshops (CAiSEW)*, pages 30–42, Berlin, Heidelberg, Germany, 2013. Springer Berlin Heidelberg.

[BS23] Matthias Babel and Johannes Sedlmeir. Bringing data minimization to digital wallets at scale with general-purpose zero-knowledge proofs, 2023.

[BSBHR19] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable zero knowledge with no trusted setup. In *Annual International Cryptology Conference*, pages 701–732. Springer, 2019.

[BSCG+13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *Annual Cryptology Conference*, pages 90–108. Springer, 2013.

[BTH20] Bert-Jan Butijn, Damian A Tamburri, and Willem-Jan van den Heuvel. Blockchains: a systematic multivocal literature review. *ACM Computing Surveys (CSUR)*, 53(3):1–37, 2020.

[BV19] Victor Balcer and Salil Vadhan. Differential Privacy on Finite Computers. *Journal of Privacy and Confidentiality*, 9(2), 2019.

[BW96] Izak Benbasat and Ron Weber. Research commentary: Rethinking "diversity" in information systems research. *Information Systems Research (ISR)*, 7(4):389–399, 1996.

[CAEK19] James Curzon, Abdulaziz Almehmadi, and Khalil El-Khatib. A survey of privacy enhancing technologies for smart cities. *Pervasive and Mobile Computing*, 55:76–95, 2019.

[CC02] John Edward Campbell and Matt Carlson. Panopticon.com: Online surveillance and the commodification of privacy. *Journal of Broadcasting & Electronic Media*, 46(4):586–606, 2002.

[CC18] John Creswell and David Creswell. *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*. SAGE, Thousand Oaks, CA, USA, 2018.

[CDE21] CDEI. Privacy enhancing technologies adoption guide. 2021.

[CGDS+20] Olivia Choudhury, Aris Gkoulalas-Divanis, Theodoros Salonidis, Issa Sylla, Yoonyoung Park, Grace Hsu, and Amar Das. Differential privacy-enabled federated learning for sensitive health data, 2020.

[CGH04] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. *Journal of the ACM*, 51(4):557–594, 2004.

[CGSM19] Pratibha Chaudhary, Ritu Gupta, Abhilasha Singh, and Pramathesh Majumder. Analysis and Comparison of Various Fully Homomorphic Encryption Techniques. *2019 International Conference on Computing, Power and Communication Technologies, GUCON 2019*, pages 58–62, 2019.

[Che16] Xiaofeng Chen. *Introduction to Secure Outsourcing Computation*. Morgan & Claypool publishers, 2016.

[Coo88] Harris Cooper. Organizing knowledge syntheses: A taxonomy of literature reviews. *Knowledge in Society*, 1(1):104–126, 1988.

[CZ19] Andre Calero Valdez and Martina Ziefle. The users' perspective on the privacy-utility trade-offs in health recommender systems. *International Journal of Human-Computer Studies*, 121:108–121, 2019. Advances in Computer-Human Interaction for Recommender Systems.

[DFF14] Michelle Finneran Dennedy, Jonathan Fox, and Thomas R. Finneran. Technology Evolution, People, and Privacy. In *The Privacy Engineer's Manifesto: Getting from Policy to Code to QA to Value*, pages 3–24. Apress, Berkeley, CA, 2014.

[DFLS12] Sabrina De Capitani Di Vimercati, Sara Foresti, Giovanni Livraga, and Pierangela Samarati. Data privacy: Definitions and techniques. *International Journal of Uncertainty, Fuzziness and Knowlege-Based Systems*, 20(6):793–817, 2012.

[DFSBJ21] Josep Domingo-Ferrer, David Sánchez, and Alberto Blanco-Justicia. The limits of differential privacy (and its misuse in data release and machine learning). *Communications of the ACM*, 64(7):33–35, 2021.

[DJG+18] Jun Du, Chunxiao Jiang, Erol Gelenbe, Lei Xu, Jianhua Li, and Yong Ren. Distributed Data Privacy Preservation in IoT Applications. *IEEE Wireless Communications*, 25(December):68–76, 2018.

[DKJG17] Ali Dorri, Salil S. Kanhere, Raja Jurdak, and Praveen Gauravaram. Blockchain for IoT security and privacy: The case study of a smart home. *2017 IEEE International Conference on Pervasive Computing and Communications Workshops, PerCom Workshops 2017*, pages 618–623, 2017.

[DKM+06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In Serge Vaudenay, editor, *Advances in Cryptology*, volume 4004, pages 486–503. Springer, 2006.

[DKM19] Cynthia Dwork, Nitin Kohli, and Deirdre Mulligan. Differential Privacy in Practice: Expose your Epsilons! *Journal of Privacy and Confidentiality*, 9(2), 2019.

[DKY17] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pages 3574-3583, Red Hook, NY, USA, 2017. Curran Associates Inc.

[DM23] Danielle Movsowitz Davidow and Yacov Manevich. Privacy-preserving payment system with verifiable local differential privacy. Cryptology ePrint Archive, Paper 2023/126, 2023. `https://eprint.iacr.org/2023/126`.

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography*, pages 265–284, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. Online; accessed 30 December 2021.

[DR13] Cynthia Dwork and Aaron Roth. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2013.

[DSSU17] Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. Exposed! a survey of attacks on private data. *Annual Review of Statistics and Its Application*, 4(1):61–84, 2017. Publisher: Annual Reviews.

[Dwo08] Cynthia Dwork. Differential Privacy: A Survey of Results. In Manindra Agrawal, Dingzhu Du, Zhenhua Duan, and Angsheng Li, editors, *Theory and Applications of Models of Computation*, volume 4978, pages 1–19. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. Series Title: Lecture Notes in Computer Science.

[DWS+11] Mina Deng, Kim Wuyts, Riccardo Scandariato, Bart Preneel, and Wouter Joosen. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering*, 16(1):3–32, March 2011.

[Edu] Computing Research & Education. Core. `http://portal.core.edu.au/conf-ranks/?search=pets&by=all&source=CORE2021&sort=atitle&page=1`. Online; accessed 7 December 2022.

[ESSD08] Steve Easterbrook, Janice Singer, Margaret-Anne Storey, and Daniela Damian. Selecting empirical methods for software engineering research. In *Guide to Advanced Empirical Software Engineering*, pages 285–311. Springer London, London, England, UK, 2008.

[Far12] Rida T. Farouki. The bernstein polynomial basis: A centennial retrospective. 29(6):379–419, 2012.

[FFI20] FFIS. Case studies of the use of privacy preserving analysis to tackle financial crime. 2020.

[FSJ18] Glenn A. Fink, Houbing Song, and Sabina Jeschke, editors. *Security and privacy in cyber-physical systems: Foundations, principles, and applications*. Wiley IEEE Press, Hoboken, NJ, first edition edition, 2018.

[GAP18] Simson L. Garfinkel, John M. Abowd, and Sarah Powazek. Issues encountered deploying differential privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*, WPES'18, pages 133-137, New York, NY, USA, 2018. Association for Computing Machinery.

[GFS+14] Xianyi Gao, Bernhard Firner, Shridatt Sugrim, Victor Kaiser-Pendergrast, Yulong Yang, and Janne Lindqvist. Elastic pathing: your speed is enough to track you. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '14 Adjunct*, pages 975–986, Seattle, Washington, 2014. ACM Press.

[GG19] Lodovico Giaretta and Sarunas Girdzijauskas. Gossip learning: Off the beaten path. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 1117–1124, 2019.

[GHK+16] Marco Gaboardi, James Honaker, Gary King, Kobbi Nissim, Jonathan Ullman, Salil Vadhan, and Jack Murtagh. Psi: a private data sharing interface. In *Theory and Practice of Differential Privacy*, New York, NY, 6 2016.

[GHV20] Marco Gaboardi, Michael Hay, and Salil Vadhan. A programming framework for OpenDP. page 31, 2020.

[GIR20] Ziya Alper Genç, Vincenzo Iovino, and Alfredo Rial. The simplest protocol for oblivious transfer. *Information Processing Letters*, 161:1–12, 2020.

[GKHD20] Peter Gonczol, Panagiota Katsikouli, Lasse Herskind, and Nicola Dragoni. Blockchain implementations and use cases for supply chains – a survey. *IEEE Access*, 8:11856–11871, 2020.

[GL20] Simson L Garfinkel and Philip Leclerc. Randomness concerns when deploying differential privacy. In *Proceedings of the 19th Workshop on Privacy in the Electronic Society*, pages 73–86, 2020.

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.

[GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, December 1994.

[GO18] Rod Garratt and Maarten R.C. van Oordt. Privacy as a public good: A case for electronic cash. *Journal of Political Economy*, 2018.

[Goo19] Google. TensorFlow Privacy repository. `https://github.com/tensorflow/privacy`, 2019. Online; accessed 8 May 2022.

[Goo20] Google. ZetaSQL repository. `https://github.com/google/zetasql`, 2020. Online; accessed 22 August 2022.

[Goo21a] Google. Google DP repository. `https://github.com/google/differential-privacy`, 2021. Online; accessed 22 August 2022.

[Goo21b] Google. Privacy on Beam repository. `https://github.com/google/differential-privacy/tree/main/privacy-on-beam`, 2021. Online; accessed 13 June 2022.

[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for NP. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 339–358. Springer, 2006.

[GPRN18] Ahmad Nauman Ghazi, Kai Petersen, Sri Sai Vijay Raj Reddy, and Harini Nekkanti. Survey research in software engineering: Problems and mitigation strategies. *IEEE Access*, 7:24703–24718, 2018.

[GR18] Otkrist Gupta and Ramesh Raskar. Distributed learning of deep neural network over multiple agents. *Journal of Network and Computer Applications*, 116:1–8, 2018.

[GRS12] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally Utility-maximizing Privacy Mechanisms. *SIAM Journal on Computing*, 41(6):1673–1693, 2012.

[GTD18] Michele Guerriero, Damian Andrew Tamburri, and Elisabetta Di Nitto. Defining, enforcing and checking privacy policies in data-intensive applications. *Proceedings - International Conference on Software Engineering*, pages 172–182, 2018.

[HABMA19] Naoise Holohan, Spiros Antonatos, Stefano Braghin, and Pól Mac Aonghusa. The Bounded Laplace Mechanism in Differential Privacy. *Journal of Privacy and Confidentiality*, 10(1), December 2019.

[Har21] Harvard. OpenDP repository. `https://github.com/opendp/opendp`, 2021. Online; accessed 22 August 2022.

[HBNC98] Ronald Hes, John J. Borking, Netherlands, and Information and Privacy Commissioner/Ontario, editors. *Privacy-enhancing technologies: the path to anonymity.* Number 11 in Achtergrondstudies en verkenningen. Registratiekamer, The Hague, rev. ed edition, 1998.

[HDH+22] Samuel Haney, Daniel Desfontaines, Luke Hartman, Ruchit Shrestha, and Michael Hay. Precision-based attacks and interval refining: how to break, then fix, differential privacy on finite computers. *Theory and Practice of Differential Privacy, ICML 2022*, 2022. Online; accessed 22 August 2022.

[HKR12] Justin Hsu, Sanjeev Khanna, and Aaron Roth. Distributed private heavy hitters. In *Proceedings of the 39th International Colloquium Conference on Automata, Languages, and Programming – Volume Part I*, page 461–472. Springer, 2012.

[HLM17a] Naoise Holohan, Douglas J. Leith, and Oliver Mason. Extreme points of the local differential privacy polytope. *Linear Algebra and its Applications*, 534:78–96, 2017.

[HLM17b] Naoise Holohan, Douglas J. Leith, and Oliver Mason. Optimal Differentially Private Mechanisms for Randomised Response. *IEEE Transactions on Information Forensics and Security*, 12(11):2726–2735, 2017.

[HMM+16] Michael Hay, Ashwin Machanavajjhala, Gerome Miklau, Yan Chen, Dan Zhang, and George Bissias. Exploring privacy-accuracy tradeoffs using dpcomp. In *Proceedings of the 2016 International Conference on Management of Data*, SIGMOD '16, page 2101–2104, New York, NY, USA, 2016. Association for Computing Machinery.

[HMPR04] Alan Hevner, Salvatore March, Jinsoo Park, and Sudha Ram. Design science in information systems research. *MIS Quarterly (MISQ)*, 28(1):75–105, 2004.

[HRC20] Muneeb Ul Hassan, Mubashir Husain Rehmani, and Jinjun Chen. Differential privacy techniques for cyber physical systems: A survey. *IEEE Communications Surveys Tutorials*, 22(1):746–789, 2020.

[HZNF15] Johannes Heurix, Peter Zimmermann, Thomas Neubauer, and Stefan Fenz. A taxonomy for privacy enhancing technologies. *Computers & Security*, 53:1–17, 2015.

[IBM20] IBM. diffprivlib repository. `https://github.com/IBM/differential-privacy-library`, 2020. Online; accessed 22 August 2022.

[ide18a] iden3. Circom, 2018.

[ide18b] iden3. Snarkjs, 2018.

[IM95] Stephen Isaac and William Michael. *Handbook in Research and Evaluation: A Collection of Principles, Methods, and Strategies useful in the Planning, Design, and Evaluation of Studies in Education and the Behavioral Sciences.* Edits Publishers, 3 edition, 1995.

[Int22] Intel. 12th generation intel core processors, 2022.

[JDL+21] Mark F. St. John, Grit Denker, Peeter Laud, Karsten Martiny, Alisa Pankova, and Dusko Pavlovic. Decision support for sharing data using differential privacy. In *2021 IEEE Symposium on Visualization for Cyber Security (VizSec)*, pages 26–35, 2021.

[Joh21] John M. Abowd, Gary L. Benedetto, Simson L. Garfinkel, Scot A. Dahl, Aref N. Dajani, Matthew Graham, Michael B. Hawes, et al. The modernization of statistical disclosure limitation at the u.s. census bureau. `https://www.census.gov/library/working-papers/2020/adrm/modernization-statistical-disclosure-limitation.html`, 2021. Online; accessed 18 February 2022.

[JOT07] Burke Johnson, Anthony Onwuegbuzie, and Lisa Turner. Toward a definition of mixed methods research. *Journal of Mixed Methods Research (JMMR)*, 1(2):112–133, 2007.

[JPY+21] Honglu Jiang, Jian Pei, Dongxiao Yu, Jiguo Yu, Bei Gong, and Xiuzhen Cheng. Applications of differential privacy in social network analysis: A survey. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–1, 2021.

[JTBN12] Martin Jaatun, Inger Anne Tondel, Karin Bernsmed, and Asmund Nyre. *Privacy enhancing technologies for information control*, pages 1–31. 2012.

[Kas05] Mark Kasunic. Designing an effective survey. Technical report, Carnegie-Mellon University, Pittsburgh, PA, USA, 2005.

[KB13] Barbara Kitchenham and Pearl Brereton. A systematic review of systematic review process research in software engineering. *Information and Software Technology (IST)*, 55(12):2049–2075, 2013.

[KC07] Barbara Kitchenham and Stuart Charters. Guidelines for performing systematic literature reviews in software engineering. Technical report, Keele University, Keele, England, UK, 2007.

[KCY21] Fumiyuki Kato, Yang Cao, and Masatoshi Yoshikawa. Preventing manipulation attack in local differential privacy using verifiable randomization mechanism, 2021.

[KEK+21] Jong Wook Kim, Kennedy Edemacu, Jong Seon Kim, Yon Dohn Chung, and Beakcheol Jang. A survey of differential privacy-based techniques and their applicability to location-based services. 111:102464, 2021.

[Ker12] Orin S. Kerr. The mosaic theory of the fourth amendment, 2012.

[KHdMR20] Daniel Kondor, Behrooz Hashemian, Yves-Alexandre de Montjoye, and Carlo Ratti. Towards Matching User Mobility Traces in Large-Scale Datasets. *IEEE Transactions on Big Data*, 6(4):714–726, December 2020.

[KKM+21] Christopher T. Kenny, Shiro Kuriwaki, Cory McCartan, Evan T. R. Rosenman, Tyler Simko, and Kosuke Imai. The use of differential privacy for census data and its impact on redistricting: The case of the 2020 u.s. census. 7(41):eabk3283, 2021.

[KLB20] Nesrine Kaaniche, Maryline Laurent, and Sana Belguith. Privacy enhancing technologies for solving the privacy-personalization paradox: Taxonomy and survey. *Journal of Network and Computer Applications*, 2020.

[KM22] Fatima Khalid and Ammar Masood. Vulnerability analysis of qualcomm secure execution environment. *Computers & Security*, 116:102628, 2022.

[KMR15] Jakub Konečný, Brendan McMahan, and Daniel Ramage. Federated optimization:distributed optimization beyond the datacenter, 2015.

[KMR+20] Daniel Kifer, Solomon Messing, Aaron Roth, Abhradeep Thakurta, and Danfeng Zhang. Guidelines for implementing and auditing differentially private systems, 2020.

[Kok17] Spyros Kokolakis. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64:122–134, 2017.

[KOV14] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal Mechanisms for Local Differential Privacy. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27, 2014.

[KPC+20] Vlasis Koutsos, Dimitrios Papadopoulos, Dimitris Chatzopoulos, Sasu Tarkoma, and Pan Hui. Agora: A Privacy-Aware Data Marketplace. page 13, 2020.

[KSK14] Vishesh Karwa, Aleksandra B Slavković, and Pavel Krivitsky. Differentially private exponential random graphs. In *International Conference on Privacy in Statistical Databases*, pages 143–155. Springer, 2014.

[Lab] Oasis Labs. Driving innovation with differential privacy. `https://medium.com/oasis-protocol-project/driving-innovation-with-differential-privacy-620a31e3d61f`. Online; accessed 10 December 2022.

[LE06] Yair Levy and Timothy Ellis. A systems approach to conduct an effective literature review in support of information systems research. *Informing Science Journal (InformingSciJ)*, 9:181–212, 2006.

[LF20] David López and Bilal Farooq. A multi-layered blockchain framework for smart mobility data-markets. *Transportation Research Part C: Emerging Technologies*, 111(June 2019):588–615, 2020.

[LLS20] Andreas Lichter, Max Löffler, and Sebastian Siegloch. The Long-Term Costs of Government Surveillance: Insights from Stasi Spying in East Germany. *Journal of the European Economic Association*, 19(2):741–789, 04 2020.

[LRYY19] Yanan Li, Xuebin Ren, Shusen Yang, and Xinyu Yang. Impact of prior knowledge and data correlation on privacy leakage: A unified analysis. *IEEE Transactions on Information Forensics and Security*, 14(9):2342–2357, 2019.

[LSTS20] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.

[LSV+19] Mathias Lécuyer, Riley Spahn, Kiran Vodrahalli, Roxana Geambasu, and Daniel Hsu. Privacy accounting and quality control in the sage differentially private ml platform. SOSP '19, page 181–195, New York, NY, USA, 2019. Association for Computing Machinery.

[LYA+18] Fan Liang, Wei Yu, Dou An, Qingyu Yang, Xinwen Fu, and Wei Zhao. A Survey on Big Data Market: Pricing, Trading and Protection. *IEEE Access*, 6(May):15132–15154, 2018.

[McK20] McKinsey & Company. Four ways to accelerate the creation of data ecosystems. `https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/four-ways-to-accelerate-the-creation-of-data-ecosystems`, November 2020.

[Met21] Meta. Opacus repository, 2021. Online; accessed 8 May 2022.

[Mir12] Ilya Mironov. On significance of the least significant bits for differential privacy. In *Proceedings of the 2012 ACM conference on Computer and communications security - CCS '12*, page 650, Raleigh, North Carolina, USA, 2012. ACM Press.

[Mir17] Ilya Mironov. Rényi Differential Privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275, Santa Barbara, CA, August 2017. IEEE.

[MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 94–103, 2007.

[MW04] Adam Meyerson and Ryan Williams. On the complexity of optimal k-anonymity. *Proceedings of the ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, 23:223–228, 2004.

[Mye97] Michael Myers. Qualitative research in information systems. *MIS Quarterly (MISQ)*, 21(2):241–242, 06 1997.

[NBH+22] Priyanka Nanayakkara, Johes Bater, Xi He, Jessica Hullman, and Jennie Rogers. Visualizing privacy-utility trade-offs in differentially private data releases. 2022(2):601–618, 2022.

[NFPH15] Arjun Narayan, Ariel Feldman, Antonis Papadimitriou, and Andreas Haeberlen. Verifiable differential privacy. In *Proceedings of the 10th European Conference on Computer Systems*, 2015.

[NH12] Arjun Narayan and Andreas Haeberlen. DJoin: Differentially Private Join Queries over Distributed Databases. page 14, 2012.

[Nis09] Helen Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.

[NIW+13] Valeria Nikolaenko, Stratis Ioannidis, Udi Weinsberg, Marc Joye, Nina Taft, and Dan Boneh. Privacy-preserving matrix factorization. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, page 801–812. Association for Computing Machinery, 2013.

[NR20] Boel Nelson and Jenni Reuben. Sok: Chasing accuracy and privacy, and catching both in differentially private histogram publication. *Trans. Data Priv.*, 13:201–245, 2020.

[NS08] Arvind Narayanan and Vitaly Shmatikov. Robust De-anonymization of Large Sparse Datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 111–125, Oakland, CA, USA, May 2008. IEEE. ISSN: 1081-6011.

[OHJ] Róbert Ormándi, István Hegedűs, and Márk Jelasity. Gossip learning with linear models on fully distributed data: EFFICIENT p2p ENSEMBLE LEARNING WITH LINEAR MODELS ON FULLY DISTRIBUTED DATA. 25(4):556–571.

[OMT09a] OMTP. Advanced trusted environment: Omtp tr1, May 2009.

[OMT09b] OMTP. Advanced trusted environment: Omtp tr1, May 2009.

[Ope22] OpenMined. PipelineDP repository. `https://github.com/OpenMined/PipelineDP`, 2022. Online; accessed 27 July 2022.

[Opp05] Rolf Oppliger. Privacy-enhancing technologies for the world wide web. *Computer Communications*, 28(16):1791–1797, October 2005.

[ORRK18] Anna Maria Oberländer, Maximilian Röglinger, Michael Rosemann, and Alexandra Kees. Conceptualizing business-to-thing interactions – a sociomaterial perspective on the internet of things. *European Journal of Information Systems*, 27(4):486–502, 2018.

[Pai99] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. *Eurocrypt*, 1999.

[PGUXS16] Layla Pournajaf, Daniel A. Garcia-Ulloa, Li Xiong, and Vaidy Sunderam. Participant Privacy in Mobile Crowd Sensing Task Management. *ACM SIGMOD Record*, 44(4):23–34, 2016.

[PHS+19a] Jan Pennekamp, Martin Henze, Simo Schmidt, Philipp Niemietz, Marcel Fey, Daniel Trauth, Thomas Bergs, Christian Brecher, and Klaus Wehrle. Dataflow Challenges in an Internet of Production. In *ACM Workshop on Cyber-Physical Systems Security & Privacy (CPS-SPC'19), November 11, 2019, London, United Kingdom. ACM*, pages 27–38, 2019.

[PHS+19b] Jan Pennekamp, Martin Henze, Simo Schmidt, Philipp Niemietz, Marcel Fey, Daniel Trauth, Thomas Bergs, Christian Brecher, and Klaus Wehrle. Dataflow challenges in an internet of production: A security and privacy perspective. In *Proceedings of the ACM Workshop on Cyber-Physical Systems Security and Privacy*, CPS-SPC'19, page 27–38. Association for Computing Machinery, 2019.

[PMB+16] Charith Perera, Ciaran McCormick, Arosha K. Bandara, Blaine A. Price, and Bashar Nuseibeh. Privacy-by-design framework for assessing internet of things applications and platforms. *ACM International Conference Proceeding Series*, 07-09-Nove:83–92, 2016.

[PRW15] Charith Perera, Rajiv Ranjan, and Lizhe Wang. End-to-end privacy for open big data markets. *IEEE Cloud Computing*, 2(4):44–53, 2015.

[PTRC07] Ken Peffers, Tuure Tuunanen, Marcus Rothenberger, and Samir Chatterjee. A design science research methodology for information systems research. *Journal of Management Information Systems (JMIS)*, 24(3):45–77, 2007.

[PVK15] Kai Petersen, Sairam Vakkalanka, and Ludwik Kuzniarz. Guidelines for conducting systematic mapping studies in software engineering: An update. *Information and Software Technology (IST)*, 64:1–18, 2015.

[RGC10] Karen Renaud and Dora Galvez-Cruz. Privacy: Aspects, definitions and a multi-faceted privacy preservation approach. pages 1–8, 09 2010.

[RH09] Per Runeson and Martin Höst. Guidelines for conducting and reporting case study research in software engineering. *Empirical Software Engineering (EMSE)*, 14(2):131–164, 2009.

[RPX+22] Deevashwer Rathee, Guru Vamsi Policharla, Tiancheng Xie, Ryan Cottone, and Dawn Song. Zebra: Anonymous credentials with practical on-chain verification and applications to kyc in defi. Cryptology ePrint Archive, Paper 2022/1286, 2022. `https://eprint.iacr.org/2022/1286`.

[RRK+20] Deevashwer Rathee, Mayank Rathee, Nishant Kumar, Nishanth Chandran, Divya Gupta, Aseem Rastogi, and Rahul Sharma. Cryptflow2: Practical 2-party secure inference. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, CCS '20, page 325–342, New York, NY, USA, 2020. Association for Computing Machinery.

[RS04] Jennifer Rowley and Frances Slack. Conducting a literature review. *Management Research News (MRN)*, 27(6):31–39, 2004.

[RSH22] Timon Rückel, Johannes Sedlmeir, and Peter Hofmann. Fairness, integrity, and privacy in a scalable blockchain-based federated learning system. *Computer Networks*, 202:108621, 2022.

[RSK+10] Indrajit Roy, Srinath T V Setty, Ann Kilzer, Vitaly Shmatikov, and Emmett Witchel. Airavat: Security and Privacy for MapReduce. page 16, 2010.

[RV08] Michael Rosemann and Iris Vessey. Toward improving the relevance of information systems research to practice: the role of applicability checks. *MIS Quarterly (MISQ)*, 32(1):1–22, 2008.

[SAW13] Latanya Sweeney, Akua Abu, and Julia Winn. Identifying Participants in the Personal Genome Project by Name. *SSRN Electronic Journal*, 2013.

[SC90] Anselm Strauss and Juliet Corbin. *Basics of Qualitative Research*. SAGE, Thousand Oaks, CA, USA, 1990.

[Sch] B. Schauerte. Conference ranks. `http://www.conferenceranks.com/`. Online; accessed 7 December 2022.

[Sco] Scopus. Scopus preview. `https://www.scopus.com/sources.uri`. Online; accessed 7 December 2022.

[SCS18] Sagar Sharma, Keke Chen, and Amit Sheth. Toward practical privacy-preserving analytics for IoT and cloud-based healthcare systems. *IEEE Internet Computing*, 22(2):42–51, 2018.

[Sim02] Gerardo I Simari. A Primer on Zero Knowledge Protocols. page 12, 2002.

[SLZ20] Eva Maria Schomakers, Chantal Lidynia, and Martina Ziefle. All of me? Users' preferences for privacy-preserving data markets and the importance of anonymity. *Electronic Markets*, 2020.

[SN15] Sarah Spiekermann and Alexander Novotny. A vision for global privacy bridges: Technical and legal measures for international data markets. *Computer Law and Security Review*, 31(2):181–200, 2015.

[Sol15] Daniel J. Solove. The meaning and value of privacy. In Beate Roessler and Dorota Mokrosinska, editors, *Social Dimensions of Privacy*, pages 71–82. Cambridge University Press, Cambridge, 2015.

[Sta95] Barry Staw. Repairs on the road to relevance and rigor. In *Publishing in the Organizational Sciences*, volume 2, pages 85–97. SAGE, Thousand Oaks, CA, USA, 1995.

[Sta96] Markus Stadler. Publicly verifiable secret sharing. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 1070:190–199, 1996.

[SWW+20] Lingling Shen, Xiaotong Wu, Datong Wu, Xiaolong Xu, and Lianyong Qi. A survey on randomized mechanisms for statistical learning under local differential privacy. In *2020 IEEE 22nd International Conference on High Performance Computing and Communications; IEEE 18th International Conference on Smart City; IEEE 6th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, pages 1195–1202, 2020.

[SYG+19] Dimitrios Skarlatos, Mengjia Yan, Bhargava Gopireddy, Read Sprabery, Josep Torrellas, and Christopher W. Fletcher. Microscope: Enabling microarchitectural replay attacks. In *Proceedings of the 46th International Symposium on Computer Architecture*, page 318–331. ACM, 2019.

[TBG+20a] Pratiksha Thaker, Mihai Budiu, Parikshit Gopalan, Udi Wieder, and Matei Zaharia. Overlook: Differentially private exploratory visualization for big data. In *Workshop on Theory and Practice of Differential Privacy*, 2020.

[TBG+20b] Andrew Trask, Emma Bluemke, Ben Garfinkel, Claudia Ghezzou Cuervas-Mons, and Allan Dafoe. Beyond privacy trade-offs with structured transparency, 2020.

[TDS03] David Tranfield, David Denyer, and Palminder Smart. Towards a methodology for developing evidence-informed management knowledge by means of systematic review. *British Journal of Management (BJM)*, 14(3):207–222, 2003.

[TM19] Georgia Tsaloli and Aikaterini Mitrokotsa. Differential privacy meets verifiable computation: Achieving strong privacy and integrity guarantees. In *6th International Joint Conference on e-Business and Telecommunications*, pages 425–430, 2019.

[TPVKE21] Andreas Theissler, Judith Pérez-Velázquez, Marcel Kettelgerdes, and Gordon Elger. Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry. *Reliability Engineering and System Safety*, page 107864, 2021.

[Tum21] Tumult Labs. Tumult Analytics repository. `https://gitlab.com/tumult-labs/analytics`, 2021. Online; accessed 15 August 2022.

[Ulu22] Ömer Uludag. Empirical analysis of the adoption of large-scale agile development methods, 2022.

[VBB13] Viswanath Venkatesh, Susan Brown, and Hillol Bala. Bridging the qualitative-quantitative divide: Guidelines for conducting mixed methods research in information systems. *MIS Quarterly (MISQ)*, 37(1):21–54, 2013.

[VBSN+09] Jan Vom Brocke, Alexander Simons, Bjoern Niehaves, Kai Riemer, Ralf Plattfaut, and Anne Cleven. Reconstructing the giant: On the importance of rigour in documenting the literature search process. In *Proceedings of the 17th European Conference on Information Systems (ECIS)*, pages 2206–2217, Atlanta, GA, USA, 2009. AIS.

[VGSR18] Praneeth Vepakomma, Otkrist Gupta, Tristan Swedish, and Ramesh Raskar. Split learning for health: Distributed deep learning without sharing raw patient data, 2018.

[War65] Stanley L Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

[Wes67] Alan Furman Westin. *Privacy and Freedom*. IG Publishing, New York, 1967.

[WK15] Mark A. Will and Ryan K.L. Ko. *A guide to homomorphic encryption*. Elsevier Inc., 2015.

[Wu12] Felix T Wu. Defining privacy and utility in data sets. *84 University of Colorado Law Review 1117 (2013); 2012 TRPC*, pages 1117–1177, 2012.

[WW02] Jane Webster and Richard Watson. Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly (MISQ)*, 26(2):13–23, 2002.

[WZL+19] Royce J Wilson, Celia Yuxin Zhang, William Lam, Damien Desfontaines, Daniel Simmons-Marengo, and Bryant Gipson. Differentially private sql with bounded user contribution, 2019.

[Yak17] Sophia Yakoubov. A Gentle Introduction to Yao' s Garbled Circuits. 2017.

[Yao82a] A. C. Yao. Protocols for secure computations. In *23rd Annual Symposium on Foundations of Computer Science (sfcs 1982)*, pages 160–164, 1982.

[Yao82b] Andrew C. Yao. Protocols for secure computations. In *23rd Annual Symposium on Foundations of Computer Science (sfcs 1982)*, pages 160–164, Chicago, IL, USA, November 1982. IEEE.

[YHL+19] Yang Yang, Xindi Huang, Ximeng Liu, Hongju Cheng, Jian Weng, Xiangyang Luo, and Victor Chang. A Comprehensive Survey on Secure Outsourced Computation and Its Applications. *IEEE Access*, 7:159426–159465, 2019.

[Yıl13] Kaya Yılmaz. Comparison of quantitative and qualitative research traditions: Epistemological, theoretical, and methodological differences. *European Journal of Education (EJE)*, 48(2):311–325, 2013.

[YLC+19] Qiang Yang, Yang Liu, Yong Cheng, Yan Kang, Tianjian Chen, and Han Yu. Federated Learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 13(3):1–207, December 2019.

[ZLC+21] Shulai Zhang, Zirui Li, Quan Chen, Wenli Zheng, Jingwen Leng, and Minyi Guo. Dubhe: Towards data unbiasedness with homomorphic encryption in federated learning client selection. In *50th International Conference on Parallel Processing*, ICPP 2021, New York, NY, USA, 2021. Association for Computing Machinery.

[ZMW14] Jan Henrik Ziegeldorf, Oscar Garcia Morchon, and Klaus Wehrle. Privacy in the internet of things: Threats and challenges. *Security and Communication Networks*, 7(12):2728–2742, 2014.

[ZMWC19] Zhenzhe Zheng, Weichao Mao, Fan Wu, and Guihai Chen. Challenges and opportunities in IoT data markets. *SocialSense 2019 - Proceedings of the 2019 4th International Workshop on Social Sensing*, pages 1–2, 2019.

Publications

[BGBM] Gloria Bondel, Gonzalo Munilla Garrido, Kevin Baumer, and Florian Matthes. Towards a Privacy-Enhancing Tool Based on De- Identification Methods. page 8.

[BGBM20] Gloria Bondel, Gonzalo Munilla Garrido, Kevin Baumer, and Florian Matthes. The use of de-identification methods for secure and privacy-enhancing big data analytics in cloud environments. In *ICEIS*, 2020.

[fGS+21] ©2021 IEEE. Reprinted, with permission, from, Gonzalo Munilla Garrido, Kaja Schmidt, Christopher Harth-Kitzerow, Johannes Klepsch, Andre Luckow, and Florian Matthes. Exploring privacy-enhancing technologies in the automotive value chain. In *2021 IEEE International Conference on Big Data (Big Data)*, pages 1265–1272, 2021.

[GLMS23] Gonzalo Munilla Garrido, Xiaoyuan Liu, Florian Matthes, and Dawn Song. Lessons learned: Surveying the practicality of differential privacy in the industry, 2023.

[GNM+21] Gonzalo Munilla Garrido, Joseph Near, Aitsam Muhammad, Warren He, Roman Matzutt, and Florian Matthes. Do i get the privacy i need? benchmarking utility in differential privacy libraries, 2021.

[GSU+22] Gonzalo Munilla Garrido, Johannes Sedlmeir, Ömer Uludağ, Ilias Soto Alaoui, Andre Luckow, and Florian Matthes. Revealing the landscape of privacy-enhancing technologies in the context of data markets for the iot: A systematic literature review. volume 207, page 103465, 2022.

[HKMG] Christopher Harth-Kitzerow and Gonzalo Munilla Garrido. Verifying outsourced computation in an edge computing marketplace. In *Machine Learning & Applications*, pages 139–157. Academy and Industry Research Collaboration Center (AIRCC).

[MGSB22] Gonzalo Munilla Garrido, Johannes Sedlmeir, and Matthias Babel. Towards verifiable differentially-private polling. In *Proceedings of the 17th International Con-*

ference on Availability, Reliability and Security, ARES '22, New York, NY, USA, 2022. Association for Computing Machinery.

[NGS23] Vivek Nair, Gonzalo Munilla Garrido, and Dawn Song. Exploring the privacy risks of adversarial vr game design, 2023.

[NGS24] Vivek Nair, Gonzalo Munilla Garrido, and Dawn Song. Sok: Data privacy in virtual reality, 2024.

[NMGS23] Vivek C Nair, Gonzalo Munilla-Garrido, and Dawn Song. Going incognito in the metaverse: Achieving theoretically optimal privacy-usability tradeoffs in vr. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST '23, New York, NY, USA, 2023. Association for Computing Machinery.

[SMGMM22] Kaja Schmidt, Gonzalo Munilla Garrido, Alexander Mühle, and Christoph Meinel. Mitigating sovereign data exchange challenges: A mapping to apply privacy- and authenticity-enhancing technologies. In Sokratis Katsikas and Steven Furnell, editors, *Trust, Privacy and Security in Digital Business*, pages 50–65, Cham, 2022. Springer International Publishing.

Abbreviations

DSAA	data sharing and analytics application
DP	differential privacy
FL	federated learning
HE	homomorphic encryption
ICT	information and communications technologies
ML	machine learning
MPC	secure multiparty computation
PET	privacy-enhancing technology
PPDM	privacy-preserving data mining
TEE	trusted execution environments
VR	virtual reality
ZKP	zero-knowledge proofs

APPENDIX A

Embedded Publications

Exploring privacy-enhancing technologies in the automotive value chain

Gonzalo Munilla Garrido*
Department of Informatics, TUM
Group IT, The BMW Group
Munich, Germany
gonzalo.munilla-garrido@tum.de

Kaja Schmidt
Department of Business Informatics
University of Potsdam
Potsdam, Germany
kaja.schmidt@uni-potsdam.de

Christopher Harth-Kitzerow
Department of Informatics, TUM
Group IT, The BMW Group
Munich, Germany
christopher.harth-kitzerow@tum.de

Johannes Klepsch
Group IT
The BMW Group
Munich, Germany
Andre.Luckow@bmwgroup.com

Andre Luckow
Group IT
The BMW Group
Munich, Germany
Andre.Luckow@bmwgroup.com

Florian Matthes
Department of Informatics
Technical University of Munich
Munich, Germany
Matthes@in.tum.de

***Abstract*—Privacy enhancing technologies (PETs) are becoming increasingly important for securing privacy, confidentiality and security, addressing both customer needs and regulatory requirements. However, applying privacy enhancing technologies in organizations requires a precise understanding of data, use cases, technologies and their limitations. This study investigates several industrial use cases and their characteristics, and the potential applicability of PETs to these. We conduct comprehensive expert interviews and a literature review to identify and classify uses cases and discuss how their requirements can be addressed using PETs as well as the trade-off associated with these solutions. While we focus particularly on automotive use cases, the results can be transferred to other industries.**

***Index Terms*—Privacy-enhancing technologies (PETs), anonymization, data exchange, platforms, gray literature review.**

I. Introduction

The volume of data generated by smart devices and harvested by institutions is increasing [5][31], and, coincidentally, the number of data breaches is simultaneously growing[1]. These trends are partially responsible for instigating the emergence of new privacy regulations in recent years, namely the European General Data Protection Regulation (GDPR) and the Consumer Privacy Act in the United States. The modern paradigm of privacy and the advent of more data breaches have driven institutions to increase their privacy-enhancing efforts when utilizing *big data* and improve the privacy-compliance of products and services [45].

This confluence of regulation, security risks, and public and private institutions' quest for privacy may steer society towards a new paradigm, where users would be more in control of their data [22]. Thus, privacy would become a foundational pillar of any modern digital platform that employs cloud services or machine learning (ML) [44]. For example, automotive service providers may find more restrictions than today to train recommender systems for personalized driving behavior to improve road safety or minimize emissions in millions of vehicles.

While users may have more control over their data and privacy is protected more strongly, privacy may also generate benefits for the public and private institutions at the forefront of privacy-enhancing innovation. Counter-intuitively, a careful deployment of privacy-enhancing technologies (PETs) may *increase* the amount of data collected, because PETs help to overcome customer concerns and address regulatory requirements [25]. As a consequence, an increase in data collection accelerates existing processes and unlocks new business models [15][30]. Furthermore, privacy-enhancing implementations mitigate the risk of fines [30], which may reach as high as 4% of a company's total annual turnover under the GDPR [1]. Lastly, after numerous data breaches across the globe [23], companies or institutions that wish to be perceived as trustworthy must include the latest privacy enhancements as their standard. Not taking such action could be interpreted as disregard for customers, thereby exposing the organization to reputational and financial risks [24]

Some prominent technology players are aware of these benefits, and therefore, push their brands to signal privacy as a core component of their products with stances, such as "*[...] customer trust is our top priority.*" or "*Your privacy is our priority.*". Despite these claims, there remains a certain level of dissonance, as the same privacy-aware industry players have committed the most significant privacy infringements to date, which can rise up to €887M or €120M. Therefore, a gap exists between specification and implementation. Given that this gap is present even in the largest and most advanced technology companies, one may infer that most organizations do not have the appropriate technology deployments to adequately enhance user privacy. Moreover, even if an organization adopts conventional de-identification techniques (e. g.,

*Corresponding author

[1]From 2014 to 2020, the average total cost of a data breach is around $3.8 million [23]

attribute deletion or pseudo-anonymization), it has been shown that such approaches do not provide sufficient protection for user data [32][14][26].

Goal. Because of the emerging need for institutions to increase their privacy efforts and the inherent complexity of such a feat [2], in this paper, we aim to provide practitioners with an overview of privacy-enhancing products that ease the implementation of privacy-related use cases, which we also uncover to motivate privacy-enhancement in organizations. While practitioners can apply our research to other industries, we conduct the study from an automotive domain perspective. Specifically, our results are shown in the framework of Fig. 1, which proposes an approach to map use cases with products in the domain of privacy. Furthermore, we have classified privacy products into solutions devised by startup vendors (see Table II) and open-source tooling (see Table III). Moreover, we have classified use cases between enablers (see Table IV) and business-oriented (see Table V). Overall, we underline the variety of use cases an institution can pursue and the importance of adopting PETs to remain competitive.

This paper is structured as follows. We provide terminology in Section II and describe our methodology and research questions in Section III. The results of our study appear in Sections IV and V. We discuss key findings in Section VI and related work in Section VII, concluding the paper with Section VIII.

II. Terminology

In this section, we define important concepts in the scope of this study. The term *product* refers to either a *tool* or a *solution*. A *tool* is a general-purpose implementation of an algorithm that abstracts the deployment of a specific technology, i.e., the user does not need to have expertise in the underlying technology for its use. For example, an open-source library that abstracts the use of differential privacy to aggregate personal data in a privacy-enhancing manner. A *solution* is a commercial implementation of a single or a set of technologies often focused on a specific purpose and domain. A solution could be a user-friendly platform for deploying federated learning models with secure computation to classify driving behavior.

A *use case* is a particular situation where a product could enable or directly leverage a business opportunity or improve an existing product or service. For example, the mood of a driver (situation) could be categorized on real time to accordingly adapt the environment within the vehicle for a better customer experience (business opportunity) with the use of privacy-enhancing machine learning (product).

III. Methodology

We outline the goal of this study as follows: interview experts to discover use cases and solutions in the domain of privacy and analyze the gray literature to identify open-source privacy-enhancing tools. To accomplish our goals, we devised two research questions (RQ).

RQ1. *What relevant privacy-enhancing products are available?* With RQ1, we aim to uncover, characterize, and classify the privacy-enhancing products that practitioners can leverage (see Section IV). We divide these products into two groups (i) solutions built by startups and (ii) open-source tooling.

(i) We include privacy-related startups because to survive as a new industry incumbent, startups must identify significant problems and solve them with novelty. To collect a curated list of startups in the domain of privacy, we conducted iterative unstructured expert interviews [36] at the end of 2020 and during the first half of 2021 with two venture capital professionals focused on privacy solutions (see Table I). Based on the discussions around their corpus of startups, which they built with dozens of expert interviews, screening calls and market searches, we created 18 categories and classified 47 startups accordingly after examining their solutions (see Table II).

(ii) Secondly, during June and July 2021, we searched for tools that implemented the most prevalent PETs included in seminal surveys or implementations in the domain of privacy [29][16][40]. We list the tools in Table III. Furthermore, each tool had to be open-source so that the scientific and engineering community could audit and freely access them. However, systematically collecting peer-reviewed publications would not capture all of the novel tools available [21]. Thus, for our purposes, S. Hopewell and M. Clarke and S. Mallett [21], and J. Vom Brocke et. al [41] indicated that a gray literature review (GLR) would be a more optimal strategy. Consequently, we included tools that appeared within the first 100 results of a Google search for the search string "*PET name* AND *open-source* AND *tool* AND *GitHub*". Two researchers independently conducted the search (one identified 48 tools while the other 47), and merged the results into 57 tools after removing duplicates (38).

RQ2.*What are the relevant use cases in the privacy domain of the automotive industry?* With the answer to RQ2, we provide use cases to motivate practitioners to enhance privacy in their institutions (see Section V).

To plan and conduct the interviews to answer this RQ, we followed guidelines from P. Runeson and M. Höst [36]. Specifically, throughout the end of 2020 and during the first half of 2021, we interviewed 17 interested practitioners who worked directly or indirectly in the automotive industry; all of the participants focused on cross-company data exchange or privacy protection (see Table I — the 17 interviewees do not include the venture capitalists in this RQ). 7 of the interviews were conducted verbally, while the remaining 10 through email correspondence. The interviews were semi-structured [36], i.e., while we initiated the conversation with a set of introductory questions about their background and followed up with RQ2 to collect a list of use cases, we promoted further exploration of their ideas to help us in the use case classification.

Afterward, we aggregated all of the use cases and analyzed the interviewees' discussions for outlining a classification. With the outline, we designed the framework of Fig. 1 to help mapping over 30 use cases to the privacy-enhancing products

TABLE I
STUDY PARTICIPANTS FROM THE AUTOMOTIVE INDUSTRY (ANONYMOUS IDENTITIES AND RESPONSES).

Participant's role	Years
Interviews	
Researcher	2 <
IT Project Manager	2 <
IT Product Owner	2 <
Privacy Officer	20 <
Privacy Officer	6 <
Data Management Officer	5 <
De-identification Specialist	2 <
Venture capital	6 <
Venture capital	6 <
Correspondence	
Cloud Security Architect	10 <
Head of Research and former Computer Science Professor	34 <
IT Program Lead	4 <
IT Product Owner	2 <
IT Product Owner	2 <
IT Product Owner	2 <
IT Project Manager	5 <
IT Project Manager	5 <
IT Product Manager	5 <
Product Manager	5 <

in Section IV (see Tables IV and V). Note that practitioners can also apply these use cases to other industries other than the automotive industry (e.g., building data markets for the financial, health, or telecommunications industry).

Limitations. Even though we have followed expert interview [36] and GLR [21][41] guidelines, this study has limitations that may have undermined its effectiveness. These threats include bias in classifications and human error in data collection. To mitigate these threats in the GLR, two researchers independently searched for tools, whose lists contained around 40% of duplicates. Regarding the expert interviews to discover use cases, we countered potential bias by interviewing 17 experts from different departments and institutions and summarized the findings before the conclusion of the interview to get feedback and thus avoid misinterpretation [36]. Likewise, to identify startups, we discussed with two venture capital experts who, in turn, had iterated interviews with dozens of other privacy domain experts to build the startup corpus. Furthermore, the authors of this paper were engaged in creating the startup classification to minimize bias.

IV. PRIVACY-ENHANCING PRODUCTS (RQ1)

Privacy-Enhancing Solutions. After aggregating and dissecting the startups found through our interviews with the venture capital experts, we classify in Table II the privacy solutions based on six capabilities: *Data management*. Enterprise applications that provide a holistic picture of an organization's data characteristics and life cycle. *Regulation compliance management*. Enterprise platforms that support and automate compliance with GDPR and other data privacy laws. *User governance*. Enterprise applications that orchestrate users' data management preferences and requests, complying with data privacy law. *Privacy integration*. Platforms and applications that bring to production privacy-focused design principles [6] and tasks such as anonymization. *Data analytics*. Platforms enabling sharing and analytics of sensitive data in a privacy-compliant manner. *Consumer-facing solutions*. Software that empowers users to profit from their data, provides users better understanding of their data on the Internet, or enables users to communicate or share information in a secure and privacy-enhancing environment.

The overlaps of the different solutions strengthened the inclusion of a category within each of the six capabilities. For example, one may observe clear overlaps between Onetrust and BigID. Onetrust offers in their product portfolio risk assessment automation, data mapping, or incident management. On the other hand, BigID provides incident reporting, data mapping or anonymization. Based on the overlaps and offerings of the collected startups, we decided for the 18 categories of Table II arranged across the six capabilities, which we use as the links between use cases and products. Lastly, note that some solution categories (e.g. data mapping or privacy managers) do not enhance privacy directly and often do not need a PET included in Table III. However, such solutions enable others such as anonymization managers to balance privacy and utility optimally by specifying the data type, provenance, and context, among others.

Privacy-Enhancing Tools While institutions can have paid access and ease integration of the startups' privacy-enhancing solutions, except for some open-source projects, institutions can also integrate freely available tools in their stack with engineering effort. The tools we included in Table III are open-source, or reference implementations based on a single PET or a cluster of PETs within the same category, for example, k-anonymity, l-diversity or generalization (ARX), or the Laplacian, Gaussian, or Snapping mechanism (Google DP). Furthermore, note that PETs offer two major functionalities to enhance privacy, anonymization (blurs the link between data and their provenance to some degree) and confidentiality (data are only shared with the intended parties) [16]. Notably, the solutions proposed in "privacy-enhancing analytics" in Table II are based on the PETs of Table III. Furthermore, the rest of the startups also rely on these PETs to some extent, such as the ones included in *Privacy integration*, while others do not rely on them like *data management* solutions.

V. MAPPING USE CASES IN THE AUTOMOTIVE VALUE CHAIN TO PRIVACY-ENHANCING PRODUCTS (RQ2)

After the discussions with the 17 experts regarding privacy-specific use cases and having collected the privacy-enhancing products, we designed a framework to perform the mapping depicted in Tables IV and V. Fig. 1 illustrates our framework and contains some use cases for exemplification. Furthermore, note the inclusion of benchmarks, which some interviewees

TABLE II
PRIVACY-ENHANCING SOLUTIONS AND THEIR DESCRIPTIONS CLASSIFIED BY CAPABILITIES.

Privacy-Enhancing Solution	**Description (Startups)**
Data management	
Activity monitoring*	Determines who has access to personal data and when such data are accessed or processed within an organization (e.g., Onetrust, and Pkware).
Data discovery*	Ascertains and classifies what kind of personal data an organization possesses to help manage privacy risk and compliance (e.g., Onetrust, and Pkware).
Data mapping*	Determines data flows throughout an enterprise (e.g., Onetrust, Datagrail, Trustarc, BigID, Wirewheel, Usercentrics, Didomi, Cookiebot, metomic, Pkware, and 1touch.io).
Regulation compliance management	
Incidence response*	Aids in the response to data breaches by providing information to relevant stakeholders of what was compromised and what regulatory obligations must be met (e.g., Onetrust, BigID, Wirewheel, Pkware, and 1touch.io).
Privacy managers*	Provide organizations with extensive and often automated information on the latest privacy laws around the world (e.g., Onetrust, Datagrail, Trustarc, Usercentrics, Didomi, Cookiebot, metomic, and 1touch.io).
Risk assessment*	Automate and scale privacy compliance processes, such as privacy impact assessment generation, locating risk gaps, and demonstrating compliance (e.g., Onetrust, Pkware, BigID, Datagrail, Privitar, Immuta, and Exate).
Website scanning*	Determines what cookies, web beacons, and other trackers are embedded in a website to identify potential disagreements with regulations and act accordingly (e.g., Onetrust, Datagrail, Trustarc, Usercentrics, Didomi, Cookiebot, metomic, 1touch.io, and Osano).
User governance	
Consent managers	Collects, tracks, demonstrates and manages user consent regarding their privacy preferences (e.g., Onetrust, Datagrail, Trustarc, BigID, Wirewheel, Usercentrics, Didomi, Cookiebot, metomic, Pkware, and Osano).
Data subject request managers	Facilitates inquires made by individuals who wish to exercise their data rights. These requests include the right to access, rectify, move, and erase data (e.g., Onetrust, Datagrail, Trustarc, BigID, Wirewheel, Usercentrics, Didomi, Cookiebot, metomic, Pkware, Osano, and 1touch.io).
Personal data account	Provides secure storage of personal data and gives ownership and control to the user over his or her data, based on which organizations can access such data through SDKs and APIs (e.g., Dataswift, and Infosum).
Privacy integration	
Anonymization managers	Automate and scale anonymization. (e.g., Onetrust, Pkware, BigID, Datagrail, Privitar, Immuta, and Exate).
Privacy as a service	APIs that enable the inclusion of privacy by design principles in an application or workflow, for example, performing encryption in a proxy before data reaches the backend of an application (e.g., StrongSalt, Evervault, and Skyflow).
Video / photo anonymization	Anti-facial recognition software to make organizations' photos and videos unrecognizable to facial recognition tools (e.g., D-ID, and BrighterAI).
Data analytics	
Privacy-enhancing analytics	Allows companies to strike a balance between privacy and data utility in analysis by leveraging secure computation (e.g., Inpher, Enveil, Arpa, Duality, Zama, Fortanix, Scontain, decentriq, Oasis Labs, Edgeless systems), anonymization like differential privacy, k-anonymity, masking or generalization (e.g., Aircloak, Privitar, Immuta, Pkware), or a combination of the previous technologies in addition to federated learning (e.g., OpenMined).
Synthetic data	Generation of a new dataset based on the properties and relationships of existing sensitive data, in contrast to anonymizing (e.g., Statice, Synthesized, Hazy, MostlyAI, and Tonic).
Consumer-facing applications	
Personal data economy*	Software that enables the user to profit from his or her data while preserving his sovereignty (e.g., digi.me, metame, and meeco).
Privacy assistants	Help consumers discover, understand, and effectively manage what the Internet knows about them, enabling users to reduce unnecessary online data exposure while using mobile- or web-based applications (e.g., Jumbo, and Mine).
Private messaging	Enables consumers to communicate and share information in a secure and privacy-enhancing environment (e.g., Signal, Yeo, and Misakey).

*Flags solutions that do not enhance privacy directly but provide a supporting role.

recommended conducting under realistic settings to compare products before deploying them.

We subdivided the use cases into two groups based on *use case types*: enablers (A and B) and business-oriented (C and D). While both groups are complementary, (A) and (B) facilitate (C) and (D) because the former two enable the accessibility to data for business-oriented analytics. Thus, for the most part, the use cases from (A) and (B) are not business-

TABLE III
PRIVACY-ENHANCING TECHNOLOGIES' DESCRIPTIONS AND THEIR MOST RELEVANT OPEN-SOURCE TOOLS.

Privacy-Enhancing Technology (PET)	Description	Tool
Anonymization		
Differential privacy (DP)	Mathematically guarantees that the output of a dataset analysis is "essentially" identical, despite the presence or absence of an individual in the dataset [11][10]. An analysis may be a statistical query like the mean or a ML model.	Google-DP (and its Python wrapper PyDP), SmartNoise, diffprivlib, DiffPriv, OpenDP, DPComp Core and Chorus (behind Uber's DP SQL). Focused on DP and deep learning: TensorFlow privacy and PyTorch Opacus
K-anonymity	Within the context of a dataset, k-anonymity guarantees the indistinguishability of a record with k-1 number of others [37]. K-anonymity is useful to anonymize datasets before realising them.	ARX, Amnesia, and Anonimatron.
Secure and outsourced computation (confidentiality)		
Zero-knowledge proof (ZKP)	Proves a claim by demonstrating the authenticity of data and integrity of a computation without revealing the data or computation [18][17].	emmy, dizk, zkMega, libsnark, libiop, ZKRollups, ZKRP, ckb-zkp, ginger-lib, OpenZKP, and gnark.
Secure multiparty computation (SMC)	Parties can jointly compute a function without disclosing their inputs by employing secret sharing or garbled circuits [43].	Multi-Protocol SPDZ, LIBSCAPI, MPyC, CrypTen, EMP-Toolkit, Multiparty, ZoKrates and MPC-SoK.
Homomorphic encryption (HE)	Allows computing functions on ciphertext without prior decryption [42][8].	TFHE, fhe-toolkit-linux, Google FHE SEAL, Concrete, eclib, HElib, and PALISADE.
Trusted execution environments (TEE)	Hardware and software that provide computation security against the unwarranted retrieval of sensitive information [33].	mTower, Open Enclave SDK, Trusty, TrustZone, Mystikos, Open-TEE and Intel's Trusted Execution Technology.
Federated Learning (FL)	Distributes ML models across data sources for training and ultimately averages the weights into one model [27][28].	Fate, sherpa.ai, PaddleFL, PySft, Xaynet, fedn, FedML-AI, Flower, PyVertical, TensorFlow Federated, and federated-learning-lib.

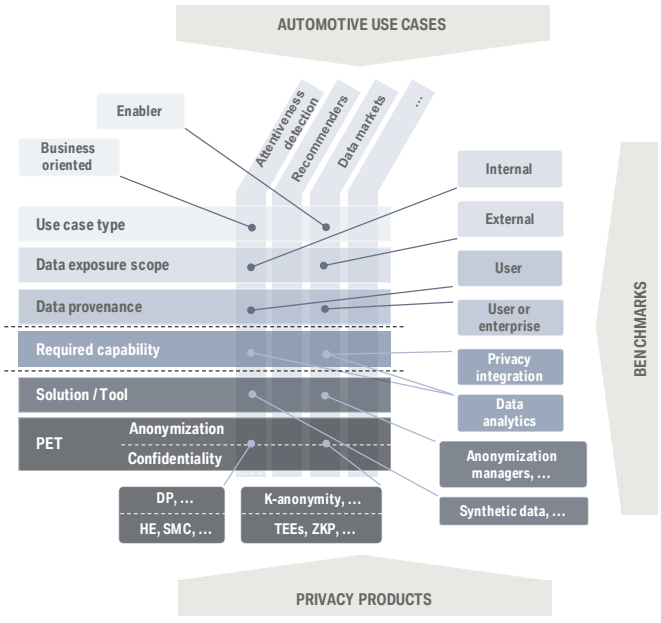

Fig. 1. Framework to map use cases and products. We selected attentiveness detection and data markets as examples. Legend: DP = Differential privacy; HE = Homomorphic encryption; SMC = Secure multiparty computation; TEE = Trusted execution environment, ZKP = Zero-knowledge proof

oriented. For example, creating a pipeline for automated risk and utility assessment and anonymization enables the use of data within an organization without measurable revenue.

Furthermore, the *data exposure scope* has a significant impact on the classification of the enabler-type products: internal (A) and external (B), because data markets and public dataset releases do not entirely apply internally. Without mastering (A), institutions would find (B) hard to tackle because (B) poses more restrictions, despite requiring similar products. Although the *data exposure scope* also applies to the business-oriented use cases (C) and (D), this dimension has more impact classifying their level of privacy risk than providing a sharp classification of business-oriented use cases. For example, attentiveness detection should require prior anonymization and confidentiality for user data, irrespective of whether the analysis is performed with data exposed internally or externally. However, executing analysis with attention data exposed externally (e.g., in a data market) would require a higher degree of perturbation and confidentiality because the probability of re-identification is higher, i.e., more potential attackers.

On the other hand, *data provenance* provides a sharper classification for business-oriented use cases because, for example, with enterprise data (D) there are often no customers (C) to protect under GDPR. Moreover, (D) often requires more authenticity to comply with a business need (e.g., fault analysis

TABLE IV
USE CASES (ENABLERS) AND DESCRIPTIONS.

Use Case (Enablers)	Description	Required Capabilities
A) *Exposing data internally*		
1. Websites and apps *In-car apps*[*] *Connected car*[*]	Customer-facing applications should first, empower users regarding data sovereignty and privacy and, secondly, enable institutions to leverage user data to improve their products and services in a privacy-enhancing manner. Furthermore, if the app shares data with third parties (e.g., connected cars interfacing with apps to provide new services) then this category A.1 would also pertain to Category B: *exposing data externally*.	Because this use case is user-focus, the solution capabilities chiefly needed are *user governance*, *consumer-facing applications*, and *privacy integration*. Furthermore, specific solutions such as website scanning can help the app provider with *regulation compliance management*. For the anonymization and analytics side of the applications, one may leverage different tools from Table III depending on the app's architecture and the goal of the app provider — analytics is covered more in detail in Table V.
2. Automate privacy functions *Automated risk assessment* *Automated anonymization* *Automated breach detection* *Utility assessment tooling*	Creating, streamlining, or automating parts of existing data pipelines to anonymize data and assess privacy leakage risk increases productivity (shorter delay to access data) and reduces human-error. Furthermore, detecting breaches and acting timely reduce user harm and legal backlash [30]. Lastly, tools that measure data utility before and after anonymization can help tune privacy thresholds, e.g., ε in DP or k in k-anonymity.	As the focus is the data privacy compliance of an enterprise, the solutions should be capable of *data* and *regulation compliance management*, and *privacy integration*. *User governance* tools such as consent and data subject request managers are also necessary to dynamically adapt to new regulatory requirements imposed by users. Furthermore, anonymization tools can replace or support some components of the solutions.
3. Prolong data access/storage	Based on current regulations like GDPR, organizations cannot access or store data longer than a predefined period. However, privacy-enhancing products can prolong access and storage and therefore benefit analytics about past events, predictions, and prescriptions. In addition to deploying products to monitor privacy-related activities like consent management, as long as there are strong anonymization and confidentiality guarantees, data may be accessed, e.g., using DP to prevent singling-out individuals and FL to train models without storing customer data.	Similarly to A.2, the solutions require *data* and *regulation compliance management*, and *privacy integration*. Moreover, given that customer consent is necessary for storage, *user governance* capabilities are also needed. Tools that allow to work with data confidentially chiefly match this use case: SMC, HE, TEE, and FL.
B) *Exposing data externally*		
1. Public dataset release *Research* *Cross-enterprise projects*	Releasing datasets publicly or confidentially for research or collaborative projects between institutions can lead an adversary to re-identify individuals [12].	Because the identity of individuals is directly threatened, the main capability is *privacy integration* (anonymization managers and video/photo anonymization); thus, organization can also employ tools based on DP and k-anonymity.
2. Data markets *Health data markets* *Automotive data markets*[*] *Financial data markets* *<Industry> data markets* *Cross-industry markets*	Data have become products in the electronic marketplace [38] in which institutions plan to exchange or trade data for profit. Initiatives such as GAIA-X or the automotive-related venture Catena-X endorse this type of ecosystems in Europe.	Data markets require a confluence of privacy products, as these ecosystems need privacy in different domains (e.g., storage, processing, verification, communication, among others) and involve different participants (users, companies, brokers, among others) [16]. Thus, data markets can benefit from many of the solutions from Table II and tools from Table III. Nonetheless, the primary capabilities that should be fulfilled are *privacy integration* and *data analytics*; without them, a data market would not be possible or alluring, respectively.

[*]Flags use cases prevalent in the *automotive* industry.

Legend: DP = Differential privacy; SMC = Secure multiparty computation; ZKP = Zero-knowledge proof; HE = Homomorphic encryption; TEE = Trusted execution environments; FL = Federated learning.

or set intersection). Thus, enterprise-data-driven use cases rely more often on secure and outsourced computation than on anonymization PETs. Examples of user data are geoposition or driving behavior, and instances of enterprise data are business secrets, performance metrics, or suppliers.

Each use case calls for a set of *required capabilities* (introduced in Section IV), which helped to narrow down the products necessary to enhance privacy in a particular use case. For example, prolonging data access/storage chiefly requires *regulation compliance management* because storing user data is strongly subject to data privacy law. Furthermore, to support such compliance, one must know what data are stored and how they will be handled and execute the appropriate privacy-enhancing software, thus, the solution should also incorporate *data management* and *privacy integration* capabilities. Lastly, because we require to monitor the type of consent given by the user, *user governance* capabilities should be present.

Finally, while the solutions provided by the startups already

TABLE V
BUSINESS-ORIENTED USE CASES AND DESCRIPTIONS

Use Case (Business-Oriented)	Description	Required Capabilities
C) *User data analytics*		
1. Recommender Systems *Driving behaviour*[*] *Eco-friendly driving*[*] *Entertainment systems*	ML models capable of training and recommending behavioural improvements (e.g., to reduce aggressive driving), provide scores (e.g., level of eco-friendly driving), or personalized options (e.g., which film to play) in a privacy-enhancing manner.	The main capabilities required are *data analytics*, and *user governance*. The open-source tools can be based on FL, DP, SMC, HE, or TEEs.
2. Across silos *Cross-app learning*	Analytics or ML models trained from different applications with different features can bring new insights that otherwise would not be possible, e.g., how does a home-entertainment system affect driving behavior or sleep.	The main capabilities are the same as in the case of C.1; furthermore, the added complexity of interacting with multiple companies requires stronger *privacy integration* capabilities.
3. Private search *Tailored Car Insurance*[*] *Internet browsing*	Finding assets on datasets without disclosing information from the query or accessing the dataset in plain text (e.g., comparing insurance prices without being tracked or disclosing one's driving history).	Solutions proposed with *data analytics* capabilities focused on secure and outsourced computation. Tools based on ZKP, SMC, HE, or TEEs.
4. Aggregate statistics *Users' statistics* *Employees' statistics* Car dealerships[*] Vehicle tracing[*]	Statistics help to understand how customers use products or services, how employees manage their hours or employ the facilities, or, more specifically, what customers care about and how they buy a car. However, most importantly, their data must be analyzed in a secure and privacy-enhancing manner (e.g., embedding DP in SQL queries).	*Privacy integration* capabilities, namely deployed with anonymization managers and some *data analytics* focused on anonymization. Tools capable of aggregating statistics can be based on DP or k-anonymity.
5. Attentiveness/mood detection *In-car environment adaption*[*] *Driving assistance*[*]	Assessing the attention of the driver can improve safety and facial micro-expressions can enhance the experience of a product or service (e.g., by adapting the lighting, music, or aroma of a car's interior). However, these expressions are highly sensitive information.	The capabilities and useful tools are equivalent to the ones in C.1.
6. Advertisement *Billboard advertisement*	Advertising in a privacy-enhancing manner can avoid price discrimination or unwarranted tracking. Furthermore, automakers can leverage demographics and geo-positioning data to propose advertisers' billboard locations without disclosing personal information.	The capabilities and useful tools are equivalent to the ones in C.1, in addition to k-anonymity.
D) *Enterprise data analytics*		
1. Set intersection *Merging logistics' data*[*]	Finding a common set of elements (e.g., vehicles' parts IDs) in private datasets across different entities without revealing any other element outside of the set (e.g., different automakers could order identical parts together to save costs).	The capabilities and useful tools are equivalent to the ones in C.3.
2. Across organizations *Cross-enterprise learning*	Leveraging databases across enterprises increases the training data of each enterprise (e.g., autonomous driving).	The capabilities and useful tools are equivalent to the ones in C.2.
3. KPI comparison *Supplier ranking*[*] *Industry benchmarking*	Compare KPIs from a set of entities without revealing the values, for instance, revealing the best offer without disclosing prices among suppliers or ranking industry competitors based on internal authenticated metrics.	The capabilities and useful tools are equivalent to the ones in C.3.
4. Computer Vision *Autonomous driving*[*] *Quality checks in production lines*[*]	De-identification of video material that contains people and training in a privacy-enhancing manner.	The main capability is *privacy integration*, namely video/photo material anonymization, and some *data analytics* capable solutions based on FL. Tools based on FL and DP can support or replace aspects of the solutions.
5. Fault analysis *Fault reporting*[*] *Predictive maintenance*[*]	Automakers could share information of parts common in their products if faults are detected in their supply chain or in vehicles to accelerate countermeasures. Moreover, automakers can use behavioral indicators of, for example, mechanical or electronic components, to predict performance deficiencies or failures. However, these indicators can sometimes be manufacturers' proprietary information.	The capabilities and useful tools are equivalent to the ones in C.3.

[*]Flags use cases prevalent in the *automotive* industry.

Legend: DP = Differential privacy; SMC = Secure multiparty computation; ML = Machine learning; ZKP = Zero-knowledge proof; HE = Homomorphic encryption; TEE = Trusted execution environments; FL = Federated learning.

include either the appropriate PETs or other software (e.g. data managers), institutions that desire to rely on open-source tools should enhance anonymity and confidentiality on any use case concerning analytics. The level of protection would depend on the privacy risk of the use case, which depends heavily on the *data exposure scope*.

The following is an end-to-end example of how we classified and mapped a use case to products; we selected the

public dataset release use case. Such release is an enabler because it does not directly leverage a business opportunity but facilitates data scientists to analyze the material. The *data exposure scope* is external, and the *data provenance* could be either from users or enterprise-specific. Assuming user data, we must enhance their anonymity (*privacy integration* capability), and, therefore, we need anonymization managers to devoid any links with the underlying users and video/photo anonymization in case the data are of such nature. Lastly, as this use case is focused on anonymization, tools employing DP or k-anonymity are most suitable, and because the *data exposure scope* is external, the values of ε for DP and k should be more restrictive. Lastly, note that as the data consumers are the public, PETs ensuring confidentiality are not necessarily required.

VI. Discussion

Based on our interviews and gray literature review, we have distilled the following key findings:

- There are products and use cases that enable others without directly enhancing privacy or producing measurable revenue, respectively.
- Business-oriented use cases are analytics-driven.
- There is no "one-size-fits-all" privacy-enhancing product.
- For the most part, enterprise data analytics use cases rely on secure and outsourced computation rather than on anonymization because these uses cases do not contain user data and require accuracy in the analysis, which is deteriorated with anonymization.
- Use cases requiring strong user sovereignty rely more heavily on secure and outsourced computation, as the data consumers should not host the data.
- Because exposing data internally (A) or externally (B) are "enabler" use cases and not business-specific, one may transfer most use cases in Table IV beyond the automotive domain.
- Because exposing data externally requires tighter privacy measures than exposing them within an organization, to engage in use cases such as data markets, institutions should consider first to be internally agile and confident with user data management and analytics.

VII. Related work

Most research focuses on applying specific PETs to address a particular use case, or investigate the use cases that a single PET can address. Examples include applying SMC for privacy-preserving deep learning [4], implementing DP in the context of sensitive health data [9], or identifying applications for which practitioners can employ TEEs [3]. However, these publications do not provide an overview of privacy use cases for different PETs or products thereof.

Other publications have surveyed how PETs fulfill privacy requirements in general [20] or from a particular context like data exchanges [35], or they highlight market opportunities for PETs to solve business problems (e.g., build trust or establish a competitive advantage [24]. However, mapping PETs with requirements or opportunities does not provide immediate insights regarding privacy use cases. The last set of publications we identified proposes industry use cases without explicitly mapping them to a list of privacy-enhancing products. Another set of publications proposes industry use cases without explicitly mapping them to a list of privacy-enhancing products. Examples range from outlining privacy use cases in the supply chain [19], the role of PETs in predictive maintenance in the automotive industry [39], or the use of PETs in the context of IoT [34] or smart cities [**curzon2019**].

We identified few publications that survey applications of PETs. There is a repository of implemented PET use cases [7] from different sectors (e.g., heath, transport, finance) and a list of case studies that used PETs to reach their objectives [13] in the financial sector. However, these surveys do not focus on production and industry use cases.

While the publications covering the domain of privacy and use cases are varied, to the best of our knowledge, they do not (i) identify patterns of PETs to classify and map privacy use cases to corresponding privacy-enhancing products, (ii) present a list of open-source tools for PETs, or (iii) present actionable use cases in the production industry.

VIII. Conclusion

We have conducted expert interviews and a gray literature review to uncover, describe, classify, and map use cases to their corresponding privacy-enhancing products, consisting of solutions from startups and open-source tooling. We encourage practitioners the use of the three-pronged approach depicted in our framework of Fig 1 to map privacy-enhancing products with use cases and consider performing benchmarks among the potentially suitable products. Furthermore, through the mappings of Tables II, III, IV, and V, we help practitioners understand the landscape of privacy from an industry perspective and swiftly select privacy-enhancing products based on a use case. Moreover, we provide a list of key findings beyond the gathered solutions and use cases. From these findings we underline the importance of engaging in "enabler" use cases first, otherwise analytics will be hindered and potentially subject to regulation infringements, and the elegant fit of secure and outsourced computation with enterprise-data analytics and strong user governance use cases. Given the landscape of privacy in the industry, we conclude that privacy-enhancing efforts should increase and that institutions should consider leveraging privacy products to enhance their competitive advantage and unlock new business models.

Future work. Overall, we recommend that practitioners expand the classification of use cases and privacy solutions with other categories, capabilities, industry domains, and examples. Additionally, practitioners can create more fine-grained decision trees based on this study to decide which PET to use according to a use case. Finally, using the existing tooling and privacy solutions, institutions can implement the use cases herein.

References

[1] Regulation (EU) 2016/679. *Protection of natural persons with regard to the processing of personal data and on the free movement of Such data, and repealing directive 95/46/EC. European Parliament and Council*. Luxembourg: Office for Official Publications of the European Communities. Apr. 2016.

[2] Peter Buxmann Anne Zöll Christian M. Olt. "Privacy-sensitive business models: barriers of organizational adoption of privacy-enhancing technologies". In: (2021), p. 22. URL: https://aisel.aisnet.org/ecis2021_rp/34/.

[3] Ghada Arfaoui, Saïd Gharout, and Jacques Traoré. "Trusted execution environments: A look under the hood". In: *2014 2nd IEEE International Conference on Mobile Cloud Computing, Services, and Engineering*. 2014, pp. 259–266. DOI: 10.1109/MobileCloud.2014.47.

[4] Kyle Bittner, Martine De Cock, and Rafael Dowsley. *Private speech classification with secure multiparty computation*. 2021. arXiv: 2007.00253 [cs.CR].

[5] Antonio Brogi and Stefano Forti. "QoS-aware deployment of IoT applications through the fog". In: *IEEE Internet of Things Journal* 4.5 (Oct. 2017), pp. 1185–1192. ISSN: 2327-4662. DOI: 10.1109/JIOT.2017.2701408. URL: http://ieeexplore.ieee.org/document/7919155/ (visited on 08/07/2021).

[6] Ann Cavoukian. "The 7 foundational principles". en. In: (2011), p. 2. URL: https://sites.psu.edu/digitalshred/2020/11/13/privacy-by-design-pbd-the-7-foundational-principles-cavoukian/ (visited on 06/13/2021).

[7] CDEI. "Privacy Enhancing Technologies Adoption Guide". In: (2021). URL: https://cdeiuk.github.io/pets-adoption-guide/ (visited on 07/15/2021).

[8] Pratibha Chaudhary et al. "Analysis and comparison of various Fully homomorphic encryption techniques". In: *2019 International Conference on Computing, Power and Communication Technologies, GUCON 2019* (2019), pp. 58–62.

[9] Olivia Choudhury et al. *Differential privacy-enabled federated learning for sensitive health data*. 2020. arXiv: 1910.02578 [cs.LG].

[10] C Dwork et al. "Our data, ourselves: privacy via distributed noise generation". In: *International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)* (2006).

[11] Cynthia Dwork and Aaron Roth. "The algorithmic foundations of differential privacy". en. In: *Foundations and Trends® in Theoretical Computer Science* 9.3-4 (2013), pp. 211–407. ISSN: 1551-305X, 1551-3068. DOI: 10.1561/0400000042. URL: http://www.nowpublishers.com/articles/foundations-and-trends-in-theoretical-computer-science/TCS-042 (visited on 04/07/2021).

[12] K. El Emam and L. Arbuckle. *Anonymizing health data*. O'Reilly Media, Inc., 2013.

[13] FFIS. "Case studies of the use of privacy preserving analysis to tackle financial crime". In: (2020). URL: https://www.gcffc.org/wp-content/uploads/2020/06/FFIS-Innovation-and-discussion-paper-Case-studies-of-the-use-of-privacy-preserving-analysis.pdf (visited on 11/06/2021).

[14] Xianyi Gao et al. "Elastic pathing: your speed is enough to track you". en. In: *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. Seattle Washington: ACM, Sept. 2014, pp. 975–986. ISBN: 978-1-4503-2968-2. DOI: 10.1145/2632048.2632077. URL: https://dl.acm.org/doi/10.1145/2632048.2632077 (visited on 06/12/2021).

[15] Rod Garratt and Maarten R.C. van Oordt. "Privacy as a public good: a case for electronic cash". In: *Journal of Political Economy* (2018). DOI: 10.1086/714133.

[16] Gonzalo Munilla Garrido et al. "Revealing the landscape of privacy-enhancing technologies in the context of data markets for the IoT: a systematic literature review". In: *arXiv:2107.11905 [cs]* (July 2021). URL: http://arxiv.org/abs/2107.11905 (visited on 09/02/2021).

[17] Oded Goldreich and Yair Oren. "Definitions and properties of zero-knowledge proof systems". en. In: *Journal of Cryptology* 7.1 (Dec. 1994), pp. 1–32. ISSN: 0933-2790, 1432-1378. DOI: 10.1007/BF00195207. URL: http://link.springer.com/10.1007/BF00195207 (visited on 02/11/2021).

[18] S. Goldwasser, S. Micali, and C. Rackoff. "The knowledge complexity of interactive proof systems". In: *SIAM J. Comput.* 18.1 (Feb. 1989), 186–208. ISSN: 0097-5397. DOI: 10.1137/0218012. URL: https://doi.org/10.1137/0218012.

[19] Peter Gonczol et al. "Blockchain implementations and use cases for supply chains – A survey". In: *IEEE Access* 8 (2020), pp. 11856–11871. DOI: 10.1109/ACCESS.2020.2964880.

[20] Johannes Heurix et al. "A taxonomy for privacy enhancing technologies". In: *Computers & Security* 53 (2015), pp. 1–17. ISSN: 0167-4048. DOI: 10.1016/j.cose.2015.05.002.

[21] Sally Hopewell, Mike Clarke, and Sue Mallett. "Grey literature and systematic reviews". In: *Publication bias in meta-analysis*. Chichester, UK: John Wiley & Sons, Ltd, Mar. 2006, pp. 49–72. DOI: 10.1002/0470870168.ch4.

[22] Patrik Hummel, Matthias Braun, and Peter Dabrock. "Own data? Ethical reflections on data ownership". In: *Philosophy & Technology* (2020). ISSN: 2210-5441. DOI: 10.1007/s13347-020-00404-9.

[23] IBM Security and Ponemon Institue LLC. *2020 cost of a data breach study*. 2020. URL: https://www.ibm.com/security/data-breach.

[24] Martin Jaatun et al. "Privacy enhancing technologies for information control". In: 2012, pp. 1–31. ISBN: 9781613505021. DOI: 10.4018/978-1-61350-501-4.ch001.

[25] Nesrine Kaaniche and Maryline Laurent. “Attribute-based signatures for supporting anonymous certification”. In: *Computer Security – ESORICS 2016*. Ed. by Ioannis Askoxylakis et al. Cham: Springer International Publishing, 2016, pp. 279–300. ISBN: 978-3-319-45744-4.

[26] Daniel Kondor et al. “Towards matching user mobility traces in large-scale datasets”. en. In: *IEEE Transactions on Big Data* 6.4 (Dec. 2020), pp. 714–726. ISSN: 2332-7790, 2372-2096. DOI: 10.1109/TBDATA.2018.2871693. URL: https://ieeexplore.ieee.org/document/8470173/ (visited on 06/12/2021).

[27] Jakub Konečný, Brendan McMahan, and Daniel Ramage. “Federated optimization:distributed optimization beyond the datacenter”. en. In: *arXiv:1511.03575 [cs, math]* (Nov. 2015). arXiv: 1511.03575. URL: http://arxiv.org/abs/1511.03575 (visited on 04/07/2021).

[28] T. Li et al. “Federated learning: challenges, methods, and future directions”. In: *IEEE Signal Processing Magazine* 37.3 (2020), pp. 50–60. DOI: 10.1109/MSP.2020.2975749.

[29] David López and Bilal Farooq. “A multi-layered blockchain framework for smart mobility data-markets”. In: *Transportation Research Part C: Emerging Technologies* 111.June 2019 (2020), pp. 588–615. ISSN: 0968090X. DOI: 10.1016/j.trc.2020.01.002.

[30] McKinsey & Company. *Four ways to accelerate the creation of data ecosystems*. URL: https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/four-ways-to-accelerate-the-creation-of-data-ecosystems (visited on 05/28/2021).

[31] Ahmadreza Montazerolghaem and Mohammad Hossein Yaghmaee. “Load-balanced and QoS-aware software-defined Internet of Things”. In: *IEEE Internet of Things Journal* 7.4 (Apr. 2020), pp. 3323–3337. ISSN: 2327-4662, 2372-2541. DOI: 10.1109/JIOT.2020.2967081. URL: https://ieeexplore.ieee.org/document/8962313/ (visited on 08/07/2021).

[32] Arvind Narayanan and Vitaly Shmatikov. “Robust de-anonymization of large sparse datasets”. en. In: *2008 IEEE symposium on security and privacy (sp 2008)*. ISSN: 1081-6011. Oakland, CA, USA: IEEE, May 2008, pp. 111–125. ISBN: 978-0-7695-3168-7. DOI: 10.1109/SP.2008.33. (Visited on 11/28/2020).

[33] OMTP. *Advanced trusted environment: OMTP TR1*. May 2009. URL: http://www.gsma.com/newsroom/wp-content/uploads/2012/03/omtpadvancedtrustedenvironmentomtptr1v11.pdf.

[34] Jan Pennekamp et al. “Dataflow challenges in an internet of production”. In: *ACMWorkshop on Cyber-Physical Systems Security & Privacy (CPS-SPC’19), November 11, 2019, London, United Kingdom. ACM*. 2019, pp. 27–38. ISBN: 9781450368315. DOI: 10.1145/3338499.3357357.

[35] Jan Pennekamp et al. “Dataflow challenges in an Internet of production: A aecurity & privacy perspective”. In: *Proceedings of the ACM Workshop on Cyber-Physical Systems Security & Privacy*. CPS-SPC’19. Association for Computing Machinery, 2019, 27–38. ISBN: 9781450368315. DOI: 10.1145/3338499.3357357.

[36] Per Runeson and Martin Höst. “Guidelines for conducting and reporting case study research in software engineering”. In: *Empirical software engineering* 14.2 (2009), pp. 131–164.

[37] Pierangela Samarati and Latanya Sweeney. “Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression”. en. In: (), p. 19.

[38] Florian Stahl et al. “A classification framework for data marketplaces”. en. In: *Vietnam Journal of Computer Science* 3.3 (Aug. 2016), pp. 137–143. ISSN: 2196-8888, 2196-8896. DOI: 10.1007/s40595-016-0064-2. (Visited on 11/13/2020).

[39] Andreas Theissler et al. “Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry”. In: *Reliability Engineering & System Safety* (2021), p. 107864. ISSN: 0951-8320. DOI: https://doi.org/10.1016/j.ress.2021.107864.

[40] Andrew Trask et al. *Beyond privacy trade-offs with dtructured transparency*. 2020. arXiv: 2012.08347 [cs.CR]. URL: https://www.researchgate.net/publication/347300876_Beyond_Privacy_Trade-offs_with_Structured_Transparency.

[41] Jan Vom Brocke et al. “Reconstructing the giant: on the importance of rigour in documenting the literature search process”. In: *Proceedings of the 17th European Conference on Information Systems*. Vol. 161. Verona, Italy, 2009. ISBN: 9788861293915.

[42] Mark A. Will and Ryan K.L. Ko. *A guide to homomorphic encryption*. Elsevier Inc., 2015, p. 101. ISBN: 9780128017807. DOI: 10.1016/B978-0-12-801595-7.00005-7. URL: http://dx.doi.org/10.1016/B978-0-12-801595-7.00005-7.

[43] A. C. Yao. “Protocols for secure computations”. In: *23rd annual symposium on foundations of computer science (sfcs 1982)*. 1982, pp. 160–164. DOI: 10.1109/SFCS.1982.38.

[44] Lian Yuming. “Data Sovereignty Theory”. In: *Sovereignty blockchain 1.0: Orderly Internet and community with a shared future for humanity*. Singapore: Springer Singapore, 2021, pp. 37–77. DOI: 10.1007/978-981-16-0757-8_2.

[45] Anne Zöll, Christian M Olt, and Peter Buxmann. “Privacy-sensitive business models: Barriers of organizational adoption of privacy-enhancing technologies”. In: *Proceedings of the 29th European Conference on Information Systems*. 2021.

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

journal homepage: www.elsevier.com/locate/jnca

Review

Revealing the landscape of privacy-enhancing technologies in the context of data markets for the IoT: A systematic literature review

Gonzalo Munilla Garrido [a,d,*], Johannes Sedlmeir [b], Ömer Uludağ [a], Ilias Soto Alaoui [a], Andre Luckow [c,d], Florian Matthes [a]

[a] Technical University of Munich, Department of Informatics, Munich, Germany
[b] Project Group Business & Information Systems Engineering of the Fraunhofer FIT, Bayreuth, Germany
[c] Ludwig Maximilian University of Munich, Department of Computer Science, Munich, Germany
[d] Group IT, BMW Group

ARTICLE INFO

Keywords:
Anonymization
Big data
Copy problem
Data exchange
Marketplace
Platform
Secure computation

ABSTRACT

IoT data markets in public and private institutions have become increasingly relevant in recent years because of their potential to improve data availability and unlock new business models. However, exchanging data in markets bears considerable challenges related to disclosing sensitive information. Despite considerable research focused on different aspects of privacy-enhancing data markets for the IoT, none of the solutions proposed so far seems to find a practical adoption. Thus, this study aims to organize the state-of-the-art solutions, analyze and scope the technologies that have been suggested in this context, and structure the remaining challenges to determine areas where future research is required. To accomplish this goal, we conducted a systematic literature review on privacy enhancement in data markets for the IoT, covering 50 publications dated up to July 2020, and provided updates with 24 publications dated up to May 2022. Our results indicate that most research in this area has emerged only recently, and no IoT data market architecture has established itself as canonical. Existing solutions frequently lack the required combination of anonymization and secure computation technologies. Furthermore, there is no consensus on the appropriate use of blockchain technology for IoT data markets and a low degree of leveraging existing libraries or reusing generic data market architectures. We also identified significant challenges remaining, such as the copy problem and the recursive enforcement problem that – while solutions have been suggested to some extent – are often not sufficiently addressed in proposed designs. We conclude that privacy-enhancing technologies need further improvements to positively impact data markets so that, ultimately, the value of data is preserved through data scarcity and users' privacy and businesses-critical information are protected.

1. Introduction

IoT devices have been improved, mass-produced, and deployed in the past few decades through steady progress in information and communication technologies (ICTs) and motivated by a trend of data-driven decision-making, automation, and the opportunity for new business models. IoT devices' primary collective purpose is to interact with the physical world and enable the measurement and collection of events and interactions [S23]. These characteristics apply to IoT devices deployed in, for example, a factory or a powerline network and many devices employed by people, such as cell phones, laptops, or wearables. The *volume*, *velocity*, and *variety* of the information generated by the IoT is immense, which drove practitioners to coin the term *big data* and develop tools for their analysis [S4]. Public and private institutions use *big data* to promote the public good, innovations, and improve products and services. Big data has become the foundation of the emerging data economy, which in Europe was worth nearly 2% of its GDP in 2016, close to 300 billion Euros [1]. However, the generation, collection, storage, processing, distribution, and analysis of *big data* to realize such economic potential also come with challenges for enterprises and responsibilities toward society.

Big data needs to be accessible to institutions that can harness their potential and develop innovations, lest society fails to materialize their advantages. Unfortunately, a significant share of the world's data is siloed and exploited solely by the institutions that host them [2], consequently locking the untapped potential of the data economy and hindering progress in science, business, and society. To surmount

* Correspondence to: Boltzmannstrasse 3, 85748, Garching, Germany.
E-mail addresses: gonzalo.munilla-garrido@tum.de (G.M. Garrido), johannes.sedlmeir@fit.fraunhofer.de (J. Sedlmeir).

https://doi.org/10.1016/j.jnca.2022.103465
Received 26 July 2021; Received in revised form 23 May 2022; Accepted 7 July 2022

this obstacle, a paradigm shift toward openness emerged in the form of electronic data markets for the IoT, i.e., mediums for the trade of information across the Internet based on electronic infrastructure [3]. This paradigm brings potential benefits, such as increasing the efficiency of business processes, facilitating growth by unlocking new business models [4], and profiting from trading. Decision-makers in governments and businesses have recognized the economic potential of data markets and hence recently supported significant projects that provide a shared digital infrastructure for data-sharing initiatives such as GAIA-X [5] or the automotive-related Catena-X, which promote the collaboration of large enterprises in data markets.

Despite data markets' promise to benefit society by fostering innovation and collaboration across enterprises, these markets hold the risk of exposing individuals' and businesses' sensitive information [6,7]. Moreover, confidence in privacy protection is an essential driver of users' willingness to share their data [S8]. Similarly, businesses are unwilling to bear the risk of unintentionally leaking their customers' private or business-critical information. Consequently, the adoption of data markets is generally hampered. Additionally, while a corporation may have taken security measures to protect collected data from unauthorized access or unintended use, data buyers might not have the same standards. Hence, in this case, exchanging data entails an additional risk that the seller needs to mitigate *before* the data are shared. Furthermore, blockchains are expected to play an essential role in the ability of institutions to trade data in tokenized form [8], but their inherent transparency further increases the need to make data exchanged in markets less sensitive [S25]. Additional trends that aggravate the negative consequences of lacking data protection for institutions are recent privacy laws such as the GDPR in Europe or the CCPA in California with their increasingly expensive fines for data breaches [9].

As a first reaction to these risks, practitioners and corporations have increased their systems' security. However, if data are sold and, thus, replicated, confidentiality is not sufficient to protect privacy: Only managing or modifying the data in a way that enhances privacy while preserving as much utility as possible is effective [10]. Thus, institutions have started to allocate more resources to balance data utility and privacy, employing privacy-enhancing technologies (PETs) [S27]. The term PET was coined in 1995 in a report by the Dutch Data Protection Authority and the Ontario Information Commissioner [11] that explored a novel approach to privacy protection [186]. These technologies take the form of architectures built with privacy-by-design principles and policies [S39,S6], or data modifications based on heuristics or mathematical privacy guarantees. Prominent examples of PETs are differential privacy [12,13], syntactic anonymization definitions like k-anonymity [14], homomorphic encryption [15–17], trusted execution environments [18], secure multiparty computation [19], zero-knowledge proofs [20,21], and a set of conventional de-identification approaches such as masking, rounding, or hashing [22].

The relevance of PETs in data markets also increases with the growing adoption of IoT devices, such as in vehicles, wearables, smartphones, and the applications that stream data daily from millions of individuals' private lives to data marketplaces [23]. Despite their current relevance and growing attention [S27], researchers and institutions still find PETs challenging to understand, integrate, and deploy in IoT data markets because most PETs are technically complex and have a wide range of variations and combinations with different tradeoffs [24]. Regarding research addressing these challenges, primary studies are predominant, i.e., studies based on original designs developed or data collected by their authors, while secondary studies collecting and systematizing existing knowledge are less frequent. The applications proposed by primary studies range from funneling data from markets into machine learning (ML) algorithms [S2,S25,S43], crowdsourcing data into markets [S16,S22,S24,S47], adopting data markets for smart mobility ecosystems [S3], smart manufacturing [S19], smart homes [S11], and smart wearables in the health industry [S48].

On the other hand, 9 out of the 50 studies that we identified in our systematic literature review (SLR) are secondary, and out of these, four studies [S19,S23,S35,S48] cover *some* of the PETs available for data markets for the IoT, yet without giving a detailed comparison of their functionalities, benefits, and limitations. The other five secondary studies [S4,S8,S14,S27,S38] perform high-level surveys revolving around challenges, non-technical privacy strategies, and user-centric perspectives on data markets for the IoT. However, none of these secondary studies provided a rigorous, *systematic* review that collected and mapped PETs and challenges comprehensively. Moreover, as we discuss in Section 9, we noted a low level of re-using existing components to build a more holistic architecture for data markets in related work, which may indicate the need for systematically analyzing the current seminal components, strengths, and weaknesses of solutions proposed for privacy-enhancing IoT data markets.

Consequently, we tackle the research gaps mentioned above with a comprehensive and detailed SLR that aims to guide decision-makers, privacy officers, policymakers, and researchers in the challenge of employing PETs to build or participate in privacy-enhancing IoT data markets. We guide these stakeholders by identifying, classifying, and describing how PETs are leveraged in the current body of scientific knowledge (see Sections 6 and 7) and presenting key findings from our SLR (see Section 9). Moreover, for the benefit of the reader, we distill terminology from the extant literature to differentiate and navigate the concepts of PETs in the scope of this SLR (see Section 5). We also organize related work into a reference model for the use of PETs in IoT data markets in distinct categories (see Fig. 10 and Fig. C.11) and identify narrow and broad challenges that PETs can tackle or circumvent (see Fig. 8). Through mapping PETs to the distilled terminology and the identified narrow challenges, we want to support practitioners in making informed decisions about the appropriate PETs to employ in the context of IoT data markets (see Table 3).

The remainder of the paper is structured as follows. Section 2 introduces the main concepts of privacy, data markets, and the IoT. Section 3 portrays how we conducted our SLR on publications dated before July 2020, followed by a discussion of related work in Section 4 and a distillation of terminology in Section 5. Sections 6 and 7 present the main results from analyzing the content of the studies in our SLR, followed by a structured review of challenges in Section 8. Based on these results and the studies' metadata, we extract a set of key findings and artifacts in Section 9, where we also provide future work and discuss the limitations of our research. Finally, Section 10 updates our study by including new selected publications from May 2022 to July 2020, and Section 11 concludes with a summary of the results.

2. Background

2.1. Privacy

Given the increased attention and relevance of privacy during the past decades, practitioners have provided many acknowledged definitions. For example, [25] stated that "[Privacy is] *the claim of individuals* [...] *to determine for themselves when, how and to what extent information about them is communicated* [...]". Similar definitions have been given by other authors like [26] or [27]. Despite these efforts, D. J. Solove argued that any attempt to distill a unique, timeless definition is infeasible due to privacy's multifaceted concept [28]. However, in the field of computer science, a narrower definition may be possible by adopting an *attack model* perspective, as the concept of privacy would likely not have emerged if transgressors would not exist: attackers of one's secrets tacitly give meaning to privacy. Therefore, a helpful definition in the context of computer science may be F. T. Wu's [29]: *"[Privacy] is defined not by what it is, but by what it is not – it is the absence of a privacy breach that defines a state of privacy"*. F. T. Wu hence defined privacy as a product of a threat model, the one from [30], in which a practitioner

needs to determine what information to hide, from whom, and what harms should be prevented before defining legal and technical tools.

Once IT architectures and tools enhance privacy appropriately, other advantages emerge. For example, from an economic perspective, privacy enables data utilization across organizations and applications to create new fair products and services and prevent price discrimination [31]. Furthermore, employing PETs may increase the number of sources and data harvested by institutions because PETs help to overcome regulatory barriers [32], in addition to mitigating the risk of fines and differentiating and appreciating a brand [4]. Moreover, political freedom and stability may only be achieved by unobtrusive forms of governments [33], privacy-enhancing journalism, and less pervasive forms of digital products such as social media that can enable malicious social engineering [S38]. Moreover, research indicates that compromising privacy can result in negative long-term economic effects [34].

Despite these benefits, and while consumers emphasize that privacy is important to them, they are typically not willing to make small additional efforts or pay for privacy [35], the so-called *privacy paradox*. Thus, in the past decades, governments have enacted rules and laws to protect consumers against violations of their privacy, specifically for data captured through advanced ICTs such as personal computers, the world wide web, or smartphones. Examples include the European Data Protection Directive in 1995, the HIPAA Privacy and Security Rule of 1996, the APEC cross-border privacy rules of 2011, the GDPR of 2016, and the Consumer Privacy Act of 2020 in the USA, which comprises Acts such as the CCPA of 2018 in California.

2.2. Data markets

According to [3], there is a misconception in everyday language between the terms "market" and "marketplace": A marketplace is the implementation of a market in terms of *infrastructure*, time, and location (virtual or physical) where the participants transact. Markets are the *environments* where buyers and sellers set the price and quantity of a particular good. Marketplaces have evolved over millennia; however, the most drastic changes arguably have happened in the past few decades. ICT has driven the costs of instant and ubiquitous communication to an often negligible amount, which has led to the digitization of many existing transaction-based ecosystems, including marketplaces [3]. Moreover, ICT has enabled the creation of virtual marketplaces that did not exist before [36] – the most prominent example being e-commerce. In this context, data have become goods themselves [3]. Data markets incentivize institutions to collect more data and to profit from trading, and, in turn, the resulting improvements and innovations benefit the public good [37].

Beyond the above formal definition, from the selected studies, we can carve out several characteristics of data marketplaces: [S1] indicated that most of the data marketplaces in operation are centralized, where the platform is run by either a trusted third party (a broker) that coordinates buyers and sellers or by the data owner (e.g., a large institution) who is also selling the data. Another 15 selected studies also proposed decentralized architectures employing distributed ledger technology to counter the drawbacks of centralized systems (e.g., single point of failure or trust on a potentially malicious entity). On the other hand, [S46] took another perspective, characterizing data trading platforms depending on the number and type of data domains: general platforms include data from any source type, while specialized platforms focus on one domain, e.g., financial, healthcare, or social media. [S27] identified two categories for data markets based on the type of participant: companies or private individual customers, e.g., owners of a smart home. Altogether, we distilled three dimensions for characterizing data markets: (i) the degree of centralization, (ii) the types and number of data domains, and (iii) the types of sellers and consumers. These dimensions permeate most of the identified solutions in this study, and all exhibit individual privacy trade-offs of which practitioners need to be aware (see Section 9).

2.3. The internet of things

The IoT is considered a network of physical devices that leverages sensors to measure and collect information from the real world and support the access and exchange of data via the Internet instantly and ubiquitously [38]. IoT devices are considered essential for gathering big data [39], which in turn brings new opportunities such as targeted advertisement, predictive maintenance, and quality improvements. Consequently, many companies have introduced the IoT in their strategy for participating in the data economy [38] and make substantial investments in the technologies that make them possible: sensors, wireless networks, and cloud computing infrastructure [39]. In this SLR, the definition of the IoT includes any device with a CPU connected to the Internet, including sensors in factories, supply chains, or vehicles, and devices such as smartphones, wearables, and computers that people use daily. These devices act as data collectors and as the gateway to a plethora of applications that collect users' actions and behavior, such as browsers, social media, e-commerce, or media entertainment, as well as sensor data generated in business processes like manufacturing and predictive maintenance, and use them for analyzes and predictions.

The design and implementation of data markets are dependent on the IoT. The ubiquity of IoT devices generates many constellations for different degrees of decentralization, with a myriad of possible sources and prosumer types. Furthermore, while such ubiquity will likely boost the data economy and its products and services, IoT devices also permeate many aspects of an individual's life, e.g., dealing with highly sensitive healthcare data or capturing sensitive information from a business perspective. Hence, the sensitivity of the data gathered from IoT devices calls for the implementation of PETs.

3. Research process

3.1. Goal and research questions

We employed the Goal-Question-Metric paradigm [40] to formulate the focus of this study as follows: we systematically analyze peer-reviewed literature to provide an overview of the state-of-the-art concerning available research and trade-offs on privacy-enhancing data markets for the IoT as well as potential research gaps from the point of view of both scholars and practitioners. Based on this paradigm, the research questions (RQs) we pursued were:

RQ1. *What relevant PETs enable IoT data markets?*

By answering this RQ, we aim to reveal, describe, and classify PETs in the context of data markets for the IoT based on their fundamentals and applications to give researchers and practitioners an overview of the PETs researched and employed so far.

RQ2. *What challenges and trade-offs hinder privacy-enhancing IoT data markets?*

Through answering this RQ, we account for explicit and implicit challenges depicted and tackled in existing work so that researchers may quickly identify pain points in the field and focus their research.

3.2. SLR execution

We conducted a SLR based on the guidelines of [41]. SLRs aim to collect, structure, and summarize the existing evidence and gaps in a particular research field to pave the way for future research. Furthermore, SLRs need to provide a rigorous and auditable methodology that can be reviewed and replicated [42]. SLRs define research questions, and a set of predefined inclusion and exclusion criteria that assess potentially relevant primary studies to answer them [43,44]. Table F.10 of the Appendix F contains the criteria for this SLR related to focus, quality, and accessibility.

To conduct the *study search*, we identified the most relevant publications in the field of privacy-enhancing data markets for the IoT to

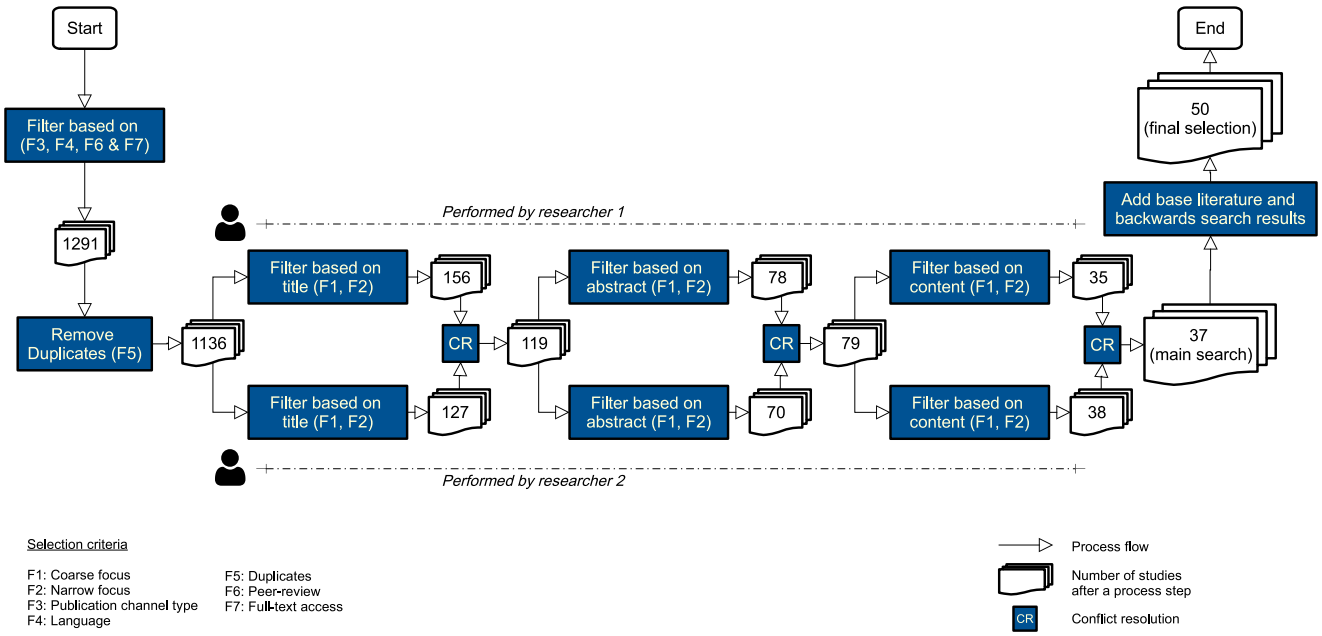

Fig. 1. Study selection process.

answer our research questions [45,46]. To obtain a corpus of high-quality publications, we defined a search strategy based on the work of [46]. Accordingly, our strategy consisted of three phases:

(i) A preliminary search of the base literature. The base literature includes representative papers (8) in the field of privacy-enhancing data markets for the IoT known to the researchers before the SLR, and some other publications found manually in the digital library of the researchers' university, which also complied with the criteria described in Table F.10. We created preliminary search strings based on identified keywords and synonyms that we found in the base literature and research questions. Afterward, we parsed our base literature with a tool to analyze frequent phrases and keywords. Using the results of this analysis, we refined our search terms [47]. Finally, we clustered the search terms into three strings based on the field of this SLR: Privacy, data markets, and IoT. Altogether, we composed the following search strings:

C1: *privacy OR private OR encryption OR encrypted OR encrypt OR data protection*

C2: *data market OR data marketplace OR data trading OR data broker OR data trader OR data auction*

C3: *Internet of things OR Internet of everything OR IoT OR sensor OR connected devices OR networked devices OR smart devices OR controller OR edge computing OR cloud infrastructure OR machine to machine OR M2M OR web-of-things OR WoT OR mobility OR automotive OR vehicle OR car OR automobile OR industry 4.0 OR smart grids OR V2V OR IIoT OR machine learning OR mobile OR cyber–physical OR microservice OR microcontroller OR micro-service OR micro-controller OR blockchain OR neural network OR smart learning OR automated driving OR autonomous driving OR smart city OR smart factory*.

Consequently, we defined our final search string as **C1** AND **C2** AND **C3**.

(ii) The main search. Since no single source may contain all the high-quality, relevant publications [48,49], we selected seven electronic databases (see Table F.9 that focus on computer science or software engineering and, according to L. Chen et al. cover the most relevant databases in these fields [50]. The time frame that we specified covers any publication included in the selected digital libraries before the 13th of July 2020. With the defined search string, time interval, databases, and following the process Fig. 1 depicts, we collected 1291 studies (1136 after duplicates removal), which two researchers filtered independently and redundantly by title (119 selected out of 1136), abstract (79 selected out of 119), and body (37 selected out of 79) following the predefined inclusion and exclusion criteria of Table F.10 to reduce bias. After each of the three filtering phases, both researchers resolved conflicts in an informed discussion and attended to the criteria.

(iii) A backward search of the references of the 37 studies resulting from the main search. After filtering by title, abstract, and body, considering our inclusion and exclusion criteria, we included another 11 studies in our corpus. The process resulted in a total of 50 studies from which we subsequently extracted and synthesized data. Hence, the SLR yielded a considerable but not excessive number of results. Furthermore, thanks to the multiple synonyms in the search string, the 37 studies only missed two studies from the base literature. Moreover, the backward search only added a modest number of new works (11). Thus, the process suggests that the choice of search terms was suitable.

To answer the research questions, we performed a *data extraction* of key information from the 50 publications in a structured manner [51]. To reduce the degree of bias, two scientists defined and independently followed an extraction card, which contained the following twenty fields: Authors, cite count, year, country, publication channel, publication type, publication source, research type, research approach, contribution type, tags, topic, subtopic, sub-subtopic, research goal, research questions, study findings, privacy-enhancing architecture or technologies, challenges, and future work. After the two scientists completed the data extraction, they held an informed discussion to resolve any possible conflicts on the extracted information. For the *data synthesis* necessary to answer RQ1 and RQ2, we adapted the "narrative synthesis" method described by [41] and performed the following synthesis procedure: (i) we developed a preliminary synthesis of the findings, followed by (ii) exploring relationships in the data and (iii) refining the preliminary synthesis with the newly acquired knowledge. After the refinement, we returned to the second step until we deemed the RQs answered.

Finally, with the goal of including significant updates and reaffirm the findings extracted from our initial research process, we conducted the same systematic search process for publications dated between July 2020 and May 2022 (24 new publications) and included the findings in Section 10.

Table 1
Short description of the involved layers.

Layer	Description
Storage	Persists data for future use.
Processing	Uses data from storage and runs algorithms on it, typically to extract information.
Communication	Exchange of data with other devices.
Verification	Checking via processing whether the data received and the identities involved are authentic.
Sovereignty	Ability to govern which (sensitive) information the communication with others exposes.
Consensus	A special case of communication and verification in which data is compared and synchronized with other devices' data.

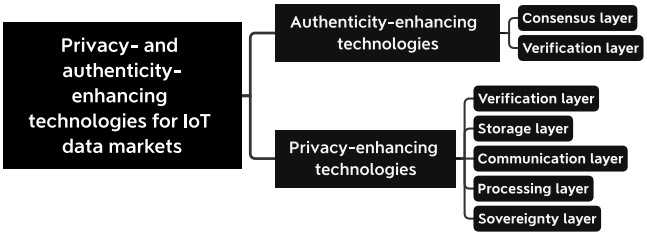

Fig. 2. Overview of the categories of our classification of the identified technologies among the selected studies. Note that some PETs also enhance authenticity.

4. Related work

In our SLR, we found nine secondary studies, i.e., studies that systematize and organize existing knowledge, conducted in the context of privacy enhancements for IoT data markets [S4,S8,S14,S19, S23,S27,S35,S38,S48], which Table B.4 summarizes. Some of these studies focused on privacy-related challenges in data markets for the IoT [S14,S19,S27,S38], while others delved into PETs from a technical perspective and discussed their challenges and opportunities [S23,S48, S35]. Lastly, [S8] analyzed users' preferences in privacy-enhancing data markets, and [S4] listed technical design choices for data markets for the IoT.

These secondary studies provided valuable contributions and built the foundation of our work; however, they focused on different aspects of IoT data markets and, therefore, lacked depth in the concepts we present in this study. For example, [S23,S35] and [S48] briefly discussed, altogether, blockchain technology, secure and outsourced computation, k-anonymity, and differential privacy and sketched their implications without considering other PETs that we identified in our SLR. Furthermore, although [S19] provided an overview of the available technologies and their challenges, the authors did not discuss PETs in detail, e.g., the study mentioned anonymization but did not delve into k-anonymity or differential privacy. Additionally, the authors only discussed three out of the six challenges we found: the recursive enforcement problem, the utility and privacy trade-off, and attacks on privacy. Moreover, [S19] based their results on observations from exemplary use cases and, therefore, cannot provide the scientific rigor and comprehensiveness of a SLR. The remaining secondary studies focused on privacy strategies instead of technology, discussed digital rights, described challenges at a high level, or provided a user-centric view of data markets.

Regarding PETs classification, some selected papers included in our SLR provided a framework. Notably, [S48] considered two categories for classification (*outsourced computations* and *information sharing*), whereas [S19] provided a more involved classification than [S48] with five layers: *data security*, *data processing*, *proving support*, *platform capabilities*, and *external measures*. Furthermore, outside of our SLR, a notable framework developed by [10], which was heavily inspired by H. Nissenbaum's work on contextual integrity [52], dissected an information flow into input, computation, and output, and assessed privacy and verifiability in each step. They also wrapped their framework with flow governance, i.e., the information flow rules upon which participants agree. To provide an improved classification of PETs, we inspired some components of our classification from [10] and [S19] and distilled from the 50 selected papers the set of layers necessary to build a privacy-enhancing IoT data market: *verification*, *storage*, *communication*, *processing*, and *sovereignty*, which Table 1 describes succinctly. Moreover, we also considered important layers necessary for a functional data market that do not require PETs (see Fig. 10).

Furthermore, not all of the technologies included in [S19] enhance privacy, e.g., version control and most distributed ledger technologies. Therefore, unlike in [S19], we have introduced another branch for technologies focused on authenticity, which we call authenticity-enhancing technologies (AET). Note that some PETs accomplish data authenticity or integrity while enhancing privacy or confidentiality, e.g., zero-knowledge proofs, homomorphic encryption, or some digital signatures (see terminology in Section 5). Specifically, some PETs can also be AETs, but AETs are not always PETs. Lastly, we classified the identified AETs into the *verification* and *consensus* layers, which are strongly associated with distributed ledger technology, as they coordinate entities and provide verification guarantees. We display our classification framework in Fig. 2. Accordingly, we structure our key results into privacy-enhancing technologies (PETs, Section 6), and authenticity-enhancing technologies (AETs, Section 7). The authors of the selected 41 primary studies jointly employed the technologies included in our classification to create holistic or parts of data market architectures for the IoT; Tables D.5, D.6, and D.7 describe the most salient architectures.

Overall, none of the related work conducted a SLR to create a holistic view of the body of scientific knowledge in privacy-enhancing IoT data markets, which, therefore, indicates the lack of an academically rigorous secondary study in this field [41]. Furthermore, despite the efforts in [S48,S19] and [10], there is not yet a *comprehensive classification and fine-grained analysis of technologies and challenges* that researchers have studied in the context of privacy-enhancing IoT data markets (see Sections Section 6, 7, and 8). Lastly, unlike other secondary studies, we also provide a detailed mapping of technologies, IoT data market layers, and challenges in Table 3.

5. Terminology

To help the reader follow our SLR, we first provide some terminology. These definitions are the distillation of the concepts found in the 50 selected studies and other seminal studies regarding utility and integrity [53], and confidentiality and privacy [54]. When we use the word *assure*, a technology fully guarantees the quality of the data or computation. In contrast, when we use the term *enhance*, a technology improves the quality of the data or computation to some extent. These qualities concern with *authenticity*, *integrity*, *confidentiality*, *privacy*, and *utility*. In line with the definition of privacy of Section 2.1 in the context of computer science, we define these qualities by the absence of an attack against them, if applicable.

Data authenticity is preserved when a malicious entity has not tampered with the truthfulness of the original data; truthfulness covers both *provenance* and *integrity*. In the context of PETs, the degree of authenticity of data can be reduced to enhance privacy. Correspondingly, *identity authenticity* is preserved when a malicious entity has not impersonated another entity. If identity authenticity is assured, then the *provenance* of the data is also assured. In the context of PETs, the identity of an entity can be concealed to enhance privacy. *Data integrity* is preserved if data that have been copied and stored or are in motion are equal to the original [53]. In practical scenarios where data are exchanged, if *integrity* is not preserved, then *data authenticity* is also inherently violated. *Computational integrity* is preserved when, even in the presence of malicious entities, the output of an algorithm that runs on data is

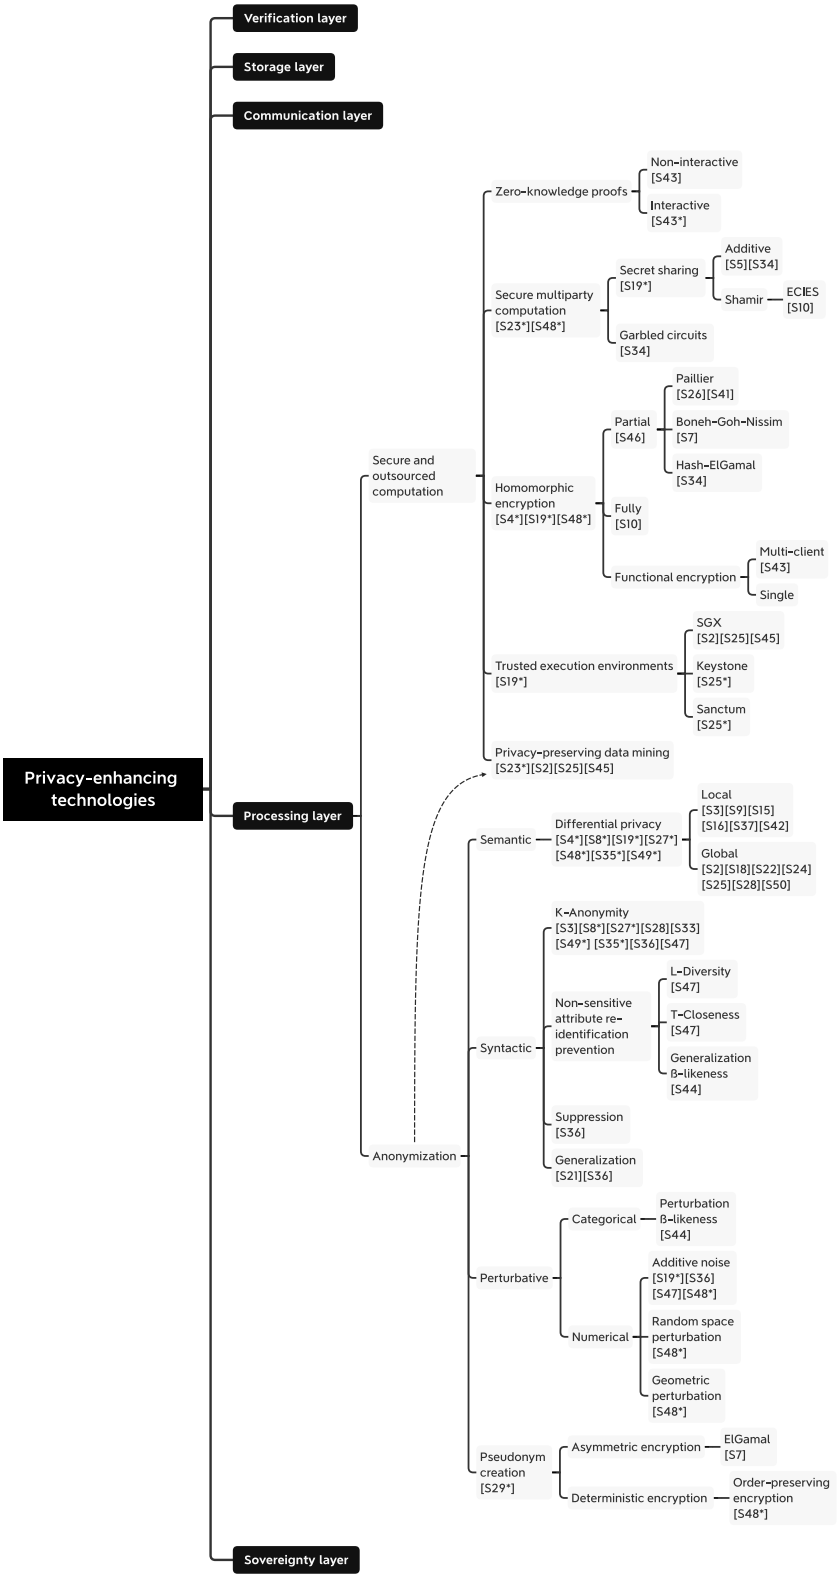

Fig. 3. Classification of PETs employed for data processing. Any other privacy approach encountered in the SLR without explicit inclusion of the underlying technology was either not included in a leaf node but in a parent node or completely dismissed if too vague. *The publication reviews or briefly comments on the technology without delving in-depth or using it as a building block of the architecture concept, e.g., included for future work.

computed correctly. In the context of PETs, the computation can be concealed to enhance privacy. Furthermore, some technologies *enhance confidentiality*, i.e., ensure that data or specific properties thereof are only shared only with the intended parties. Furthermore, we refer to *utility* as a measure of the usefulness of data for the successful completion of a task; it is high when the data is *authentic*. While PETs reduce *authenticity*, they are helpful in the balancing act between *utility* and *privacy*, as PETs may help to facilitate the sharing of data, which otherwise would not have been revealed (zero utility).

We briefly give some illustrative examples of the interplay of concepts: A digital signature can *assure data integrity* provided the corresponding private key is not accessible to an adversary, and *identity authenticity* if the signature includes a digital certificate that a trusted third party issued; otherwise, digital signatures *cannot* assure *identity authenticity*. Distributed ledgers can *assure data and computational integrity* by replicated storage and computation [55], but these ledgers cannot enhance *data authenticity*; additionally, replication is often problematic regarding *confidentiality* and, hence, *privacy* [56]. Furthermore, zero-knowledge proofs can provide evidence for *data and identity authenticity and computational integrity* without violating *privacy*, and truth discovery can *enhance* these qualities independent of *privacy* considerations. Moreover, privacy-preserving data mining can *enhance* privacy; however, if the right PETs are not employed, qualities such as *computational integrity* may not be *enhanced* or *assured*. As last examples, technologies such as differential privacy or k-anonymization *enhance* privacy by reducing *data authenticity*, and onion routing or ring signatures *enhance* privacy by forgoing or reducing *identity authenticity*, respectively. These latter technologies consequently reduce *data utility* in exchange for *privacy*.

Lastly, we are mindful of the term *tackling*, which refers to a technology that *directly* and fully or partially solves a current challenge in the context of privacy enhancement, e.g., the copy problem or the recursive enforcement problem (REP) (see challenges in Section 8). We use the term *circumvent* when a technology bypasses a problem, i.e., the technology does not directly address the issue. However, the entities that leverage the circumventing technology are still not affected by the problem. For example, obscuring the data and computation in a third-party server with homomorphic encryption (HE) does *not tackle* the REP; instead, HE *circumvents* such problem because the third-party server cannot see the contents. On the other hand, distributed ledger technology *tackles* the REP with a redundant and hence tamper-evident storage and execution.

The following Sections 6 and 7 describe the technologies that we identified in our SLR. We provide a new categorization of these technologies based on the characteristics emphasized in the corresponding selected publications and the technical properties described in this Section.

6. Privacy-enhancing technologies

6.1. Processing layer

The PETs we included in data processing aim to enhance the privacy of either data inputs, outputs, the intermediate steps of a computation, or a combination thereof while maintaining a high degree of utility. This Section follows the structure of Fig. 3.

6.1.1. Secure and outsourced computation

Secure and outsourced computation comprises PETs that enhance privacy through confidentiality. Furthermore, if the PET also employs digital signatures and their cryptography primitives, then the PET can also assure the integrity of the data and computation and identity authenticity in the presence of a digital certificate.

Zero-knowledge proofs (ZKPs). With ZKPs, a technology firstly conceived in the 1980s by [20], a *verifier* can verify the authenticity of the data and the integrity of a computation conducted by a *prover* without the need to access the data or replicate the computation itself [21]. If the statement that is proven is about claims attested in a digital certificate signed by a trusted entity (e.g., age over 18), ZKPs can verify identity authenticity while keeping the information leaked about the identity minimal.

Specifically, ZKPs exhibit (i) zero-knowledgeness, i.e., the *verifier* learns nothing new from the *prover* beyond the correctness of their statement, (ii) completeness, i.e., the prover can convince the verifier of a correct statement with high probability, and (iii) soundness, i.e., the *prover* cannot convince the *verifier* of a wrong statement with high probability [57,S43]. Furthermore, there are interactive and non-interactive ZKP protocols. With the latter, there is no need to engage in sequential message exchange, and the prover can convince multiple parties of a claim with a single, potentially short, message [57]. These characteristics make non-interactive ZKPs highly attractive for use in blockchains [56]. ZKPs are also the building block of many anonymous credentials, which are also known as privacy-preserving attributed-based credentials [58]. They allow the verification of information in a digital certificate without disclosing any unnecessary data, including the highly correlating value of the signature. Anonymous credentials were initially proposed in 1985 by [59], and developed further with ZKPs and blind signatures [60] chiefly by [61,62] and by [63]. Lately, anonymous credentials have seen renewed interest also in the context of digital wallets for end users' identity management [64,65].

Within our SLR in IoT data markets, [S43] employ non-interactive ZKPs to verify the correct computation of outputs, which, in turn, unlocks the payment from a smart contract in the Agora blockchain, eliminating a third-party verifier. While ZKPs have their limitations due to computational complexity, typically for the prover, and there is still a considerable gap between cryptographers and software engineers [66], we expect to see more publications such as [S43]. This projection is justified by the significant improvements in ZKPs' performance, and ease of use in recent years [67,68] and the availability of an increasing variety of domain-specific programming languages to implement ZKPs, such as bellman or circom in combination with snarkjs. Recently, first research has emerged that uses ZKPs to prove that a machine learning model was trained correctly on specific data [69], and there are many opportunities to leverage them in data markets, such as demonstrating that the input data of a computation was signed by a sensor that received a certificate from a trusted third party without revealing the sensor's identity or the data. In this case, the digital signature and certificate can be regarded as AETs, while their verification inside a ZKP enhances privacy and, hence, qualifies ZKPs as a PET.

Secure multiparty computation (MPC). In broad terms, MPC enables multiple parties to exchange information obliviously and jointly compute a function without revealing individual inputs to each other [19,70]. The MPC implementations that we observed in our SLR employ either secret sharing [S5,S10,S34] or garbled circuits [S34]. In secret-sharing-based MPC, each party first obfuscates the input by splitting it into shares. Secondly, this party distributes the shares among the other computing parties. Afterward, each party executes arithmetic operations independently on these shares, and finally, all parties share the outputs to reconstruct the result [19].

In Shamir's scheme [71], one can specify a minimum of shares that the recipient needs to reconstruct the output, and any combination of fewer shares does not reveal anything about the secret to the receiving entity [S19,187,188,S48]. On the other hand, in additive secret sharing, all the shares are needed. Outside MPC, Shamir's scheme has been commonly used for key management schemata for cryptographic systems so that if some shares that represent a private key are lost, one can still reconstruct the key with the remaining shares [71]. On the other hand, MPC can also be implemented by garbled circuits [19], for only two [72] or multiple [73] parties. Garbled circuits are protocols that enable secure computation by using functions translated into Boolean circuits, i.e., a sequence of basic logic gates such as AND, XOR, and OR that may be combined to construct any function [19,74,S48]. Garbled circuits make use of oblivious transfer [75], which in turn utilizes asymmetric encryption, and symmetric encryption for encrypting and decrypting each gate's truth table. Lastly, there are MPC hybrids that combine these approaches [76].

MPC allows computing functions without revealing the inputs to other participating parties. MPC protects inputs against brute force attacks and it is to date considered less computationally expensive than alternatives such as fully homomorphic encryption [77]. Drawbacks of

MPC include its high processing and communication costs [58] and sensitivity to network latency, which can considerably decrease the performance [S5,S9,S48]. Additionally, MPC protocols often need to be supplemented by mechanisms that prevent collusion [S5]. Moreover, since the individually provided inputs are only locally available, one cannot stop malicious entities from jeopardizing the authenticity of the input with false inputs. MPC can only prevent curious entities from learning information. A countermeasure for this reduction in accountability is zero-knowledge proofs to enforce the authenticity of participants' local computations while maintaining them confidential [78].

Three of the papers in our SLR implement MPC in their architecture: [S5] uses additive secret sharing, [S10] employs Shamir's secret sharing, and [S34] leverages a combination of garbled circuits an additive secret sharing. Additionally, other publications acknowledge the importance of MPC schemata by including them in their review [S19,S23,S48]. We provide additional details in Table D.5. While several *frameworks for MPC* are available, the MPC solutions employed by these three publications were handcrafted. This may indicate that the integration of MPC into existing systems requires features that are not yet available with generic tools, such as performance aspects.

Homomorphic encryption (HE). HE allows performing operations on encrypted data (ciphertext) as if they were not encrypted. After the computation, the entities with the corresponding secret key can decrypt the output [189]. There are variations of HE depending on the diversity of operations it can perform [15,16]: Fully homomorphic encryption (FHE) schemata support addition and multiplication, while partially homomorphic encryption (PHE) schemata allow for only one of these alternatives — typically in exchange for drastically improved performance. Any other schema in-between is called somewhat homomorphic encryption [S48].

Five out of the six studies that use HE in our SLR use a PHE variation [S41,S26,S7,S34,S46]. Each of the former four specifies the name of the employed schema, namely the Paillier cryptosystem in the first two [17], Boneh–Goh–Nissim [79], and Hash–ElGamal [80]. The latter study only briefly mentions the additive homomorphic property of their handled data. On the other hand, [S10] uses FHE with a schema called fully homomorphic non-interactive verifiable secret sharing [81]. Several other articles in our SLR underline the importance of HE [S4,S19,S48]. On the other hand, [S43] suggested the use of multi-client functional encryption [82,83] instead of HE so that the scheme combines data from some individuals with others, and, in turn, malicious entities cannot trace back the output of the computation to a single user, as it may happen in HE. Furthermore, there is a related scheme called functional encryption [S43] that allows to retrieve a prespecified function executed on a set of cyphertexts [82], e.g., decrypt only the mean of a set of encrypted numbers by deriving a function-specific decryption key from the secret key that was used for encrypting the data. A summary of papers from our SLR that mention or use HE is given in Table D.5.

The major limitation of FHE is its high computational complexity and the comparatively large storage needs of its cyphertext, which poses a significant challenge for its use and is aggravated in the context of IoT devices' limitations [S41,S7,S48,S24]. Therefore, the approach adopted by most authors is the use of PHE instead of FHE [84], which, while still not as efficient as other PETs, consumes significantly more computing resources than PHE [85].

As observed for the case of MPC, while there exist generic frameworks such as SEAL, HElib, or TFHE, the authors of the publications in our SLR utilized handcrafted solutions, which may indicate the lack of framework versatility or performance. Overall, practitioners and companies may use HE to perform lightweight functions on data privately on non-local resources, e.g., computing in the cloud, which otherwise would be too expensive to maintain in-house. MPC would usually be preferred over HE when the inputs to the function belong to multiple parties. Nonetheless, some selected publications also employ HE in these cases, e.g., when data brokers determine the winner of an auction [S26,S34,S41].

Trusted execution environments (TEE). TEEs were first defined in 2009 by the Open Mobile Terminal Platform as "hardware and software components providing facilities necessary to support applications" that are secure against attacks that aim to retrieve cryptographic key material or other sensitive information. These features include defense against more sophisticated hardware attacks such as probing external memory [18] or measuring execution times and energy consumption. Moreover, TEEs defend against adversaries who are legitimate owners of the hardware or remote access to the operating system that can run the code themselves. TEEs allow a user to define secure areas of memory ("enclaves") that enhance confidentiality and assure data and computation integrity of the code and data loaded in the TEE [S2], i.e., any other program outside the enclave cannot act on the data. Specifically, TEEs associate unique encryption keys to computer hardware, making software tampering at least as hard as hardware tampering and certifying the computation results within the TEE. The reason is that the only way to hacking a TEE is physical access to the hardware and, consequently, performing manipulations so that the hardware provides false certifications to bypass remote attestation and sealed storage [S45]. Seal-stored data may not be accessed unless the user employs the correct hardware and software, and remote attestation is a process whereby a trusted third-party assures that the execution of a program in a specific piece of hardware is correct [S45].

Four of the selected papers in our SLR leverage TEEs [S2,S12,S25,S45], and a review mentions their importance [S19]. [S2,S25] and [S45] proposed TEEs to confidentially train and evaluate machine learning models on data available through a data market. While the role of TEEs in data markets overlaps with the use of HE and MPC, authors have preferred the latter technologies to enhance confidentiality in auctions and data processing, which may be due to the limited memory TEEs offered at the time. The reviewed four studies used Intel's Software Guard Extension (SGX) [86], where Intel is the trusted third party, and, therefore, the single point of failure. However, practitioners should be mindful of the numerous vulnerabilities present in TEEs [87–89], and Intel's SGX deprecation in 2022 [90], which affects many of the designs found in this SLR dated before July 2020. Therefore, we suggest practitioners to explore Sanctum [91], Keystone [92] and AWS Nitro [93]. Specifically, [S2,S25] and [S45] used SGX for confidential computing, and [S12], employed SGX for their blockchain architecture to perform "Proof of Useful Work". In this type of consensus mechanism, nodes perform useful computations instead of computing hashes like in Bitcoin or Ethereum mining. Moreover, [S2] decided to use TEEs to enhance data and computation confidentiality for machine learning algorithms because of the low performance of MPC and HE on machine learning [94].

On the other hand, we noted that TEEs designed for resource-constrained devices – potentially at the cost of offering less functionality – were not prominently discussed in the selected papers. This includes, for instance, ARM TrustZone [95], which is relevant as many IoT devices run on ARM processors, and trusted execution modules [96].

Privacy-preserving data mining (PPDM). [S23] describe PPDM as a means to enhance privacy while extracting useful information from data mining. Data mining includes ML and conventional statistical analyses such as aggregations (e.g., mean or quantiles). PPDM is achieved by performing the computation where the data reside, protecting the computation with cryptographic or data perturbation means, or a combination thereof. As a comprehensive example, suppose the clients' local data and computation are cryptographically protected and the clients have the capability to perturb data. In that case, the computation can run anywhere, which is accomplished by deploying a ML model and input data in a trusted execution environment (TEE) or implementing a ML model using MPC or HE. With input or computation perturbation, the clients also enhance the privacy of the outputs.

A popular tool for PPDM is federated learning (FL) [97–99], as it avoids collecting users' data. Specifically, FL collaboratively trains

a seed ML model across multiple clients' local data, after which a server aggregates the resulting weights to form a unique model (process repeated across rounds). Researchers have increased the privacy of FL by sharing weights with secure aggregation protocols [100] (MPC, Shamir's secret sharing), and protected the privacy in client selection [101] and update parameter sharing [102] with additive HE (i.e., partial HE). Alternatively, split learning approaches [103,104] decompose neural networks' layers into elements and, thus, the input data and labels do not need to be within the same machine. Split learning presents advantages over FL when the local hardware for computations belongs to different network speeds or hardware configurations [105]. Additionally, gossip learning [106,107] proposes a framework whereby multiple models perform a random walk over clients, where they are trained and merged with other models they encounter.

Another way to achieve PPDM is by perturbing input data or weights of the ML model with anonymization techniques such as differential privacy (DP), resulting in privacy-enhancing optimization schemata like DP-stochastic-gradient-descent (DP-SGD) [108]. DP-SGD perturbs the weights' updates with noise and, therefore, one may not reconstruct the inputs based on the outputs, which may happen in ML or stand-alone FL [94]. Practitioners can plug in weight DP perturbation in central ML, FL, gossip learning, or split learning, in combination with MPC as well. We depict such leverage of anonymization technologies for PPDM with the dashed line connecting both elements in Fig. 3.

With PPDM, individuals may enjoy a higher degree of privacy than outsourcing the computation transparently to a trusted third party. Data markets can offer an infrastructure leveraged by PPDM, where data prosumers and consumers only need to provide the input data and ML models, much like the studies in our SLR propose [S2,S25,S45] using TEEs to train models with DP. Like data, trained ML models could also be exchanged in markets.

6.1.2. Anonymization

While the previously presented PETs hide sensitive data from unsolicited parties and, thus, provide confidentiality while enhancing or assuring data and computation integrity, the authorized receiver of the plaintext may reverse engineer the output and correlate data records with individuals (re-identification attack). Consequently, employing only secure and outsourced computation PETs is insufficient to provide the required degree of privacy in cases where the recipient may not be fully trusted. Anonymization technologies can help in these situations by protecting non-explicit identifiers and sensitive attributes [S47,S44]. The cost of this protection is forgoing data authenticity and thus decreasing utility. Given the frequency of re-identification attacks, anonymization should be a critical element of any survey or modern online application, and, in particular, IoT data markets [S8]. One may observe that anonymization technologies rely on statistics, probability theory, and heuristics, while secure and outsourced computation usually employs cryptography and trusted hardware.

This sub-section describes our findings for the most employed anonymization technologies identified in our SLR. We categorize most of them into two groups [190]. *Syntactic* technologies provide a numerical value to the degree of individuals' protection in a dataset, resulting in a perceptible perturbation of data, e.g., generalizing the values 42, 44, and 45 to the interval $[40, 45]$ such that it is harder for an attacker to distinguish between the three individuals. *Semantic* technologies enforce a privacy definition to a learning mechanism executed over a dataset, namely differential privacy, whereby the output distribution of the mechanism should be insensitive to the removal or addition of an individual in the dataset. Typically, the property is fulfilled by adding calibrated noise to the output of a mechanism, yielding a result that is not syntactically different from the original value, e.g., 42 could become 45 after noise addition.

Semantic technologies have an advantage over syntactic technologies, as they provide a mathematical guarantee of privacy agnostic to background information, i.e., an attacker cannot use related information to re-identify an individual in the dataset. Additionally, we discuss other anonymization technologies not covered in these two groups, namely noise perturbation and pseudonym creation. Perturbation, in this context, is not classified as semantic because its process does not necessarily provide a formal semantic privacy guarantee (such as in differential privacy) and, simultaneously, the outputs are not syntactically modified. In essence, anonymization techniques obfuscate information by perturbing the data during measurement or processing [58]. From this perspective, anonymization can be understood as a kind of statistical disclosure control [109], and, thus, also encompasses semantic techniques such as differential privacy.

Syntactic technologies. We identify the implementation of the privacy definitions of k-anonymity and its variations l-diversity and t-closeness, a newly proposed model called β-likeness, and also their building-blocks: generalization, and suppression. The most frequently utilized model for syntactic anonymization in our SLR is k-anonymity [S3,S28,S33,S36,S47], which was also reviewed or highlighted by [S8,S27,S49] and [S35]. K-anonymity is a privacy model that guarantees any individual in a dataset to be indistinguishable from at least $k-1$ others. K-anonymization, i.e., altering a dataset to fulfill k-anonymity, clusters a set of sensitive attribute values into equivalence classes of size k. However, finding an optimal value of k for minimum information loss is NP-hard. Thus, researchers have proposed alternative heuristics [191]. Nonetheless, some of the selected studies used the building blocks of k-anonymization (transformations): generalization [S21,S36] and suppression [S36]. Suppression deletes selected data points, while generalization substitutes data points for others that belong to a higher level in a manually pre-defined hierarchy, e.g., substituting a city by a country to make the location less detailed.

The selected studies [S28,S33] applied k-anonymization to aggregate data from a set of entities. [S47] innovated upon [S28] and [S33] by also employing the l-diversity model to ensure at least l different values in sensitive attributes, and t-closeness so that the distribution of the sensitive attributes within each equivalence class was at most at a distance t from the overall dataset distribution of that attribute. These two models have their own limitations, they prevent homogeneity and external knowledge attacks (l-diversity) and skewness and similarity attacks (t-closeness) [190], to which k-anonymity is vulnerable. Furthermore, [S36] tailored the use of k-anonymity based on record history, privacy policies, and disclosure context. Their new approach prevented a significant decrease in the data utility compared to homogeneously applying k-anonymity to all individuals' records equally.

Nonetheless, there are detractors of syntactic technologies in data markets because of the need for a centralized intermediary that sees and aggregates the data in a, e.g., k-anonymous fashion [S18]. Moreover, [S44] stated that these conventional syntactic approaches are not sufficient because they lack an attacker perspective in the model. For this reason, they designed a novel model called β-likeness that explicitly bounds the additional knowledge that an adversary gains from seeing the released data.

In the context of this SLR, k-anonymization is mainly employed before sharing data in an IoT data market. However, researchers also employ k-anonymity for privacy-enhancing location-based services that exchange location data in IoT data markets, whereby similar fake locations hide the real ones. This type of approach fits well with IoT devices embedded in phones, vehicles, and laptops, among other mobile *things*. Some of the approaches named by [S3] were *cloaking*, which consists of sending a more extensive region that encompasses the real one, and *geomasking*, whereby the real location is randomly displaced outside of an inner circle but within an outer one. [S3] adopted *geomasking* for situations where low accuracy is sufficient, and a high degree of privacy is required. A modern alternative to releasing anonymized data is synthetic data generation, which creates data by randomly sampling

from a distribution representative of the real data. Practitioners can employ generative adversarial networks (GANs) [110], or GANs with differential privacy for a higher protection [111] to synthesize data. Synthetic data could be helpful in some contexts as they "look" similar to the real data (unlike fulfilling k-anonymity), e.g., developing applications before testing them with the real data.

Altogether, anonymization technologies and PPDM compose the building blocks of *statistical disclosure control* [112], which organizations may leverage *internally* in their privacy-preserving data management and analysis solutions, or *externally*, by using privacy-enhancing publishing solutions [113] in, e.g., data markets. Among past surveys focused on the latter solutions (namely k-anonymity and related models, in addition to a few cryptographic primitives), the reader may refer to [113–116] and [117] for further specialized reading. Most notably, [115] provide a comprehensive survey of syntactic models and differential privacy and their *attacks*, and [117] compile a helpful *mapping* of data types to the appropriate syntactic and semantic techniques for anonymization.

Semantic technologies. Introduced by [12] in 2006, differential privacy (DP) proposes a formal guarantee of privacy that has become the golden standard for researchers [S37]. DP appears in its pure form or one of its flavors in 13 of the 35 studies that propose a solution in our SLR. Furthermore, another seven studies refer to DP to underline its importance or drawbacks. The potential reasons behind the high number of references and use of DP are multifaceted. While HE or MPC may protect the inputs' and computations' confidentiality, they do not prevent reverse-engineering the outputs (re-identification attacks). Moreover, syntactic technologies or other conventional anonymization technologies, e.g., additive noise, lack a mathematical guarantee of privacy and are subject to background knowledge attacks. DP, however, tackles these issues.

In broad terms, DP guarantees that the output distribution of an analysis (a statistical query or a ML model) over a dataset is "essentially" identical, irrespective of the presence or absence of an individual in the dataset. Additionally, DP is agnostic to auxiliary information available in the present or the future. DP is *typically* achieved by adding random noise sampled from a probability distribution such as the Laplacian or the Gaussian. Specifically, the noise limits the output distribution *difference* of an analysis executed over two datasets (one *with* and one *without* an individual) to be no greater than an upper bound, making the outputs "differentially" indistinguishable. Overall, DP bounds the amount of new information gained by an attacker after observing the output of an analysis.

The set of selected studies of our SLR that used DP in their proposed solutions are [S2,S3,S9,S15,S16,S18,S22,S24,S25,S28,S37,S42,S50]. Tables D.6 and D.7 summarize their proposed architectures. Six of these studies employ DP locally [S3,S9,S15,S16,S37,S42], i.e., the noise is added to the data of an individual on the client-side. In contrast, the rest of the studies apply DP centrally, i.e., on aggregated data on the server-side. Furthermore, we can cluster the studies into those that focus on a data trading design for data markets [S22,S37,S15,S50], crowdsensing data markets [S42,S9,S16,S24], and architectures that host a data market in an attempt to achieve end-to-end privacy [S2,S3,S18,S25,S28].

While DP offers a mathematical privacy guarantee, DP is not a panacea. DP still holds flaws in its real-world implementations [118] that the research community and practitioners should address. Moreover, DP's combination with ML needs further improvements regarding balancing privacy and accuracy [109]. In our SLR, [S48] and [S36] identify two specific problems with DP: firstly, DP cannot be used when a high level of accuracy is required [S48], e.g., analyzing data from the brakes of vehicles to improve safety. Secondly, releasing an entire dataset with current DP approaches is troublesome. Despite these challenges, the authors of [S2] and [S50] argue that the benefits of DP predominate, as DP can adapt to many use-cases and allows a practitioner to fine-tune the added noise to enhance privacy.

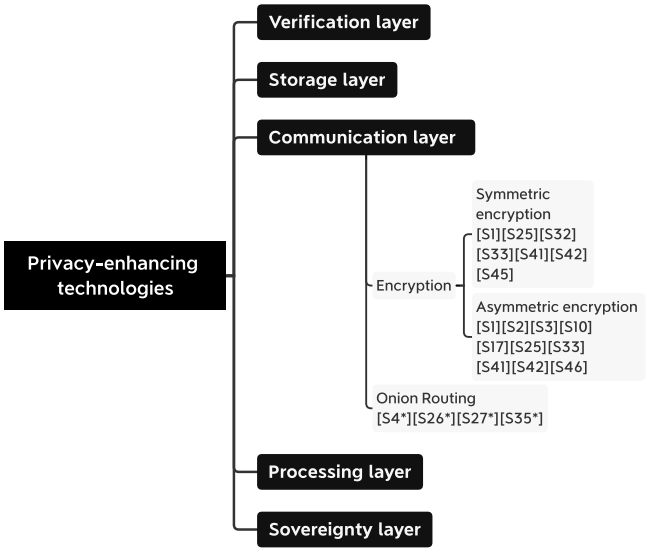

Fig. 4. Classification of PETs employed for communication.

As we already noted with ZKP, MPC and HE, the authors of the selected publications that used DP did not employ open-source DP libraries such as OpenDP, Google-DP, diffprivlib, TensorFlowprivacy, or Chorus. Instead, they used handcrafted implementations of DP. Aside from syntactic and semantic technologies, other anonymization technologies are simpler to implement, e.g., sampled data release, character masking, truncation, rounding, top and bottom coding, data swapping, randomization, creating pseudonyms, character scrambling, microaggregation, or noise perturbation [22,119]. Moreover, [120] designed an algorithm that combines the syntactic definition of k-anonymity with DP. The two other anonymization technologies employed by another three studies were noise perturbation [S36,S47] and pseudonym creation [S7].

Perturbative anonymization. Perturbation relies on the use of noise to obfuscate sensitive information. One of the simplest forms of perturbation is additive noise, employed in [S36] and [S47]. Additive noise consists of adding to a deterministic value a random value sampled from a uniform distribution whose bounds are set by a specific percentage of the deterministic value. Furthermore, [S48] reviews two novel perturbative technologies: First, random space perturbation [121], which strives to protect the privacy of cloud-stored data by utilizing a confluence of order-preserving encryption, dimensionality expansion, random noise injection, and projection. Second, geometric perturbation [122], which is motivated by the idea of protecting the geometric transformations that a machine learning model may perform on a dataset rather than the data itself. While perturbative technologies aim to tackle the same problems, unlike DP, they do not provide mathematical guarantees of privacy, even though some are also based on noise addition.

Pseudonym creation. Pseudonym creation is applied to direct identifiers, e.g., names or social security numbers, to enhance privacy while uniquely identifying each record. Practitioners create pseudonyms by hashing or deterministically encrypting an identifier, e.g., using order-preserving encryption [S48], or by applying asymmetric key encryption like ElGamal [S7]. However, researchers have demonstrated that pseudo-anonymization falters against some attacks like profiling, task tracing, or re-identification [S35].

6.2. Communication layer

The PETs included in this Section enhance the confidentiality of data in transit or of the sender's identity (see Fig. 4). These PETs rely on cryptography.

Encryption. Encryption is one of the most fundamental technologies to enhance confidentiality [S19] because after encrypting a piece of data (cipher), only the anointed holders of a decryption key can decipher such data. We underline that encryption cannot guarantee privacy because nothing stops an intended receiver from publicly sharing the decrypted message; this also emphasizes employing anonymization PETs. Encryption may be symmetric (one key to both encrypt and decrypt data) or asymmetric, known as public-key cryptography (two keys, a public key to encrypt and a private key to decrypt, or vice versa). Encryption is the building block of virtually every secure communication established through a network and takes a key role in digital signatures.

While some publications from our SLR employed asymmetric encryption for the confidential communication of data [S1,S2,S3,S10, S17,S25,S33,S41,S42,S46] (most of these publications also employed asymmetric encryption for digital signatures, hence the high frequency of digital signatures in Fig. G.14), other publications such as [S1, S25,S32,S33,S41] and [S45] employed symmetric encryption to also confidentially store data. Naturally, encryption is also a building block for digital signatures (verification layer) and for the storage layer to maintain data at rest confidential. We depict this relationship with the dashed lines connecting these elements in Fig. 5.

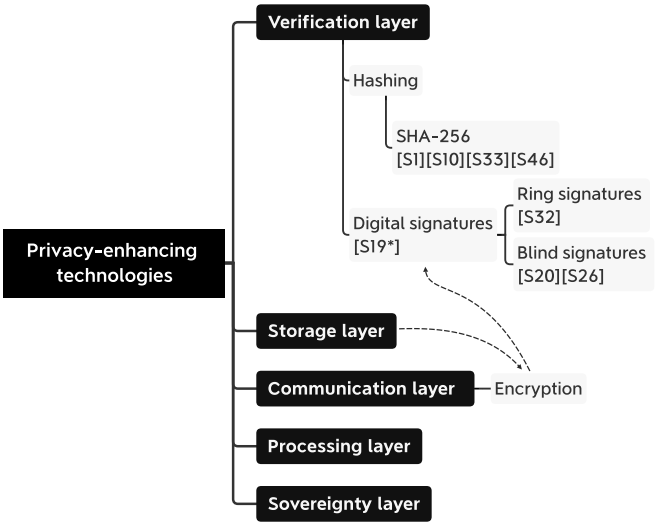

Fig. 5. Classification of PETs employed for verification.

Onion routing. Onion routing, the backbone of the P2P network resulting from the Tor project [123], consists of a series of re-transmission steps through the network's nodes. A sender's message is encrypted once for each step. The intermediaries decrypt only their appointed encryption layer. Thus, the node only knows the immediate sender and receiver but not the origin of the chain of messages. As the messages are encrypted, the nodes cannot see the contents either. Overall, onion routing renders one's messages unreadable and untraceable. The paper that suggests employing onion routing in a data market context is [S26], which some of the identified reviews equally appreciate [S4,S27, S35].

However, some drawbacks exist. Implementations backed by Tor have high-latency and redundant communication that challenges bandwidth, which can be hard to align with high transactional environments, such as IoT data markets. Moreover, if an architecture decides to use Tor, the network is often blocked by IT departments within organizations or even subject to state-level censorship by some governments [123]. Therefore, practitioners can use alternative technologies such as a VPN to enhance entities' privacy in a network in these contexts. However, these typically centralized alternatives generally offer lower anonymity guarantees, e.g., a VPN provider can identify a user [58].

Because there is no central authority to set privacy policies unilaterally, one must remember that onion routing enhances privacy by preventing malicious entities from collecting IP addresses to identify users. Onion routing will not help if the data that users submit to the network is intrinsically sensitive or correlating. To tackle these limitations, practitioners may use onion routing as a building block of a more extensive privacy-enhancing system that leverages other PETs [S26].

Notably, the publications surveyed in our SLR did not include many other untraceability protocols, which would include mixnet-based alternatives to onion routing (e.g., anonymous remailers, Chaum's mixes [124]), DC-nets [125], or peer-to-peer anonymous communication systems. An extensive overview of these systems is given by [126].

6.3. Storage layer

The authors of the selected papers that propose confidential storage functionality in their architecture leverage symmetric encryption, mostly AES [S1,S25,S32,S33,S41,S45] (encryption is described in the *communication layer*). Furthermore, researchers could leverage Inter-Planetary File Systems (IPFS) [127] to compensate for the lack of storage capacity in blockchains to some extent. Specifically, IPFS is a peer-to-peer protocol for data storage and access in a distributed file system. Among the PETs in the processing layer, practitioners could employ homomorphic encryption [128] to encrypt and store certain types of data, so that data are readily available to compute confidential operations. Furthermore, unless strictly necessary, practitioners should store encrypted data that are, in turn, *anonymized* with syntactic or semantic technologies. In case of a breach that leaks the decryption key, anonymized data would reduce the likelihood of attackers re-identifying individuals.

6.4. Verification layer

Some of the PETs that support data processing cannot verify the authenticity of data, identities, or the integrity of data [S19] by themselves. The PETs we include in this Section accomplish these verifications with different levels of privacy enhancement. The data processing PETs that can assure identity authenticity and data integrity use the digital signatures of the *verification layer* and the encryption technologies of the *communication layer* as building blocks. Furthermore, the credibility associated with verifying the information exchanged, analysis outputs, and identities can increase the willingness of users to share data [S8]. To navigate this Section, we refer to Fig. 5.

Privacy-enhancing digital signatures (DSs). DS schemata assure data integrity and identity authenticity if accompanied by a digital certificate. As a consequence, DSs also provide non-repudiation [S1], i.e., actions that an entity cannot deny later. The steps that usually constitute a DS scheme are private and public key generation, encrypting a digest of data with a private key, and a signature verifier that employs the public key to check whether the sender signed the data with the private key.

DSs and the encryption primitives of the communication layer are so fundamental that one of the selected studies solely relies on HTTPS for their data market architecture [S17]. However, this architecture does not consider privacy beyond data in transit. Hence, most selected studies rely on multiple PETs. Moreover, although not all of the selected studies explicitly mention DSs, we can safely assume that since DSs are already a living part of virtually any enterprise IT system, most selected studies employ them in their architectures (hence the high frequency of DS utilization in Fig. G.14). Nonetheless, while DSs allow verifying the integrity of data or the authentic identity of the sender, users still need to trust the sender with the authenticity of the data.

So far, we have only described DS as an authenticity-enhancing technology. However, some of the studies selected in this SLR employed

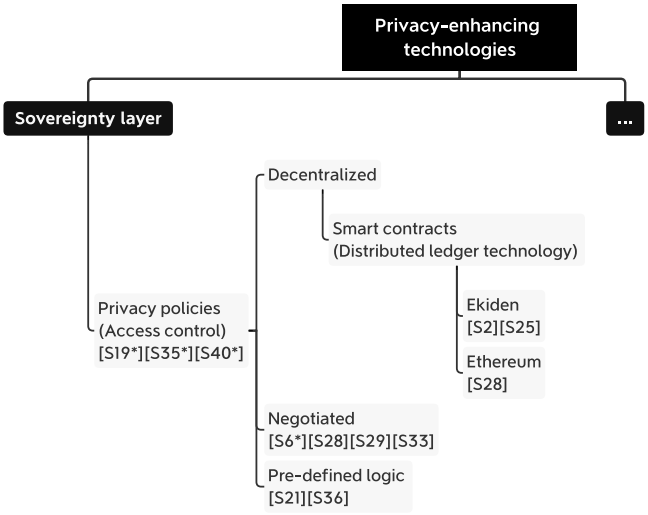

Fig. 6. Classification of PETs used for sovereignty purposes.

two DS schemata based on asymmetric encryption primitives that make DSs privacy-enhancing:

- **Ring signatures** [S32], whereby any party within a pre-defined set of parties could have been the signer of a message. Thus, the identity of the authentic signer is kept hidden [192,193].
- **Blind signatures** [S20,S26], whereby the signer does not have access to the content being signed [60]. It is possible to use blind signatures in combination with zero-knowledge proofs to convince the signer that the content to be signed has the expected properties. Also, one can make an entity sign multiple contents and allow for spot checks to detect fraud. The latter procedure has been employed in the first approaches toward privacy-enhancing payments [59].

Hashing. Hashing is a tool to deterministically map data of an arbitrary length to a fixed output length. In the context of privacy and verification, and aligned explicitly with some of the selected studies [S1,S10,S33,S46], hashing is used to verify the integrity of transferred data by hashing the data and making the hash public before transferring the data. In this manner, the recipient can verify the integrity of the confidentially transferred data by comparing the hash of the received data with the previously published hash. Provided the entropy of the data is sufficiently high, nobody except the intended recipient can determine the data from the published hashed value. Hence, hashing can be considered a form of version control with a privacy component.

The hash function employed by the publications mentioned above was SHA-256. Their authors commonly use the published hashed data on distributed ledger technologies to ensure immutability and availability. In this setting, hashing enhances the confidentiality of the sender's data while the parties (network nodes) ensuring the ledger's integrity (and inherently the persisted hash) cannot unveil the original data. The original data is only viewed by the intended receiver, which validates the integrity of the data received through another channel with the hash persisted in the ledger.

6.5. Sovereignty layer

The *sovereignty layer* deals with the concept of information control, the perceived ability to govern what is exposed from one's data [54]. Specifically, based on an entity's requirements, this layer defines the entity's rules and guidelines regarding ownership and management that can indirectly govern data processing, verification, and other IoT data market layers. Furthermore, to prevent entities' violation of privacy, practitioners should map these rules to the PETs capable of fulfilling them. For example, GAIA-X's high-level architecture contemplates privacy policies in their data *sovereignty layer* [5]. Privacy policies (closely related to access controls), have been predominant across publications, and, thus, dominate the sovereignty layer depicted in Fig. 6, which illustrates the three identified types.

Privacy policies and privacy by design. Privacy policies embody the requirements and guidelines of a data governance model and are meant to be part of any privacy-enhancing application. To define them, given the regulatory and human aspects of privacy policies, it is also helpful to adopt perspectives from definitions beyond computer science, such as the one underlined in Section 2. [129] indicates that the privacy requirements depend on the recipient of the information, e.g., an individual can have different reservations when disclosing information to a family member than to the government. Privacy policies should reflect this definition, which means that individuals should express the privacy policies they expect. Several studies in our SLR explicitly proposed policies as part of their solution [S6,S21,S25,S28,S29,S33,S36], while many others reviewed privacy policies [S19,S35,S40] or mentioned similar ideas. For example, [S8] did not provide a concrete implementation or explicitly named privacy policies. However, they mentioned that, in data sharing scenarios, the data owners should be able to control some fundamental aspects: data types to share, with whom to share, the required degree of trust in another party, the purpose of sharing, and for which benefit.

Among the publications discussing privacy policies, there is a discernible classification. Four of these publications [S6,S28,S29,S33] considered privacy policies as a negotiation between the user and a third party, such as the data consumer or data broker. [S6] provided a set of legal requirements and high-level technical solutions that facilitated the introduction of policies in international data markets, e.g., writing policies in a standard language. On the other hand, [S29] presented an approach where the data owner could choose among a set of four privacy policies, which included how data was aggregated, and [S33] relied on a *privacy policy manager* that acted as a gatekeeper and managed the privacy settings from a set of users. Furthermore, another set employs a pre-defined logic to execute PETs based on the desires and track record of the data shared by the individual [S21,S36], while the last set relied on smart contracts for decentralized pre-defined [S2,S25] or negotiated [S28] policies.

Nonetheless, the implementation of policies faces challenges. Firstly, there may be multiple colliding policies, i.e., applications must prioritize policies depending on the context [S21]. Secondly, there is no uniformly accepted global standard for electronic privacy policies [S6]. [S40] investigate how practitioners model privacy policies in different domains, focusing on IoT applications. They point to the lack of a uniform standard and propose to utilize ontology-based privacy-knowledge modeling. Thirdly, policy enforcement also causes overhead and an increase in latency due to the need for compliance checks and a lack of automation [S21]. Furthermore, conventional users should include their privacy preferences with minimal manual effort, as they could be overwhelmed otherwise. [S40] suggested using recommender systems based on similar users' data to address this issue. However, this solution may incur a biased recommendation. Moreover, data acquisition expenditure for privacy policies should not incur costly computational resources as they scale to a growing number of transactions [S40].

Privacy policies are crucial to protect users' privacy; however, they are not enough. Organizations must consider privacy issues at each stage of the data pipeline (i.e., processing data end-to-end with the extract-transform-load framework [130]), contemplating aspects that escape user-defined or mutually-agreed policies, and taking into account that typically, neither users nor data brokers will be privacy experts. If a user does not know the potential harms of sharing sensitive information such as DNA data, a data consumer may take advantage

of the user. Therefore, while privacy policies are a stepping stone toward end-to-end privacy, practitioners must develop systems with a privacy-by-design philosophy [S36,S4].

Privacy by design is a term coined in the '90s by the former information and privacy commissioner for the Canadian province of Ontario, A. Cavoukian, who created seven principles [131]. Privacy by design claims that privacy goes beyond current regulations and must be an ever-present concern in the minds of organizations [131]. Following privacy-by-design principles entails, for example, preventing sensitive information extraction by default [S36], minimizing the amount of shared data at each exchange (proportionality) [194], and increasing the price of large data packages [194], among others. However, adopting these design principles comes with effort, forcing developers to adapt their design patterns. For example, current homomorphic encryption techniques force data scientists to express their analysis in terms of additions and multiplications, and differential privacy requires new software engineering design patterns that track the privacy budget of individuals or data scientists.

Smart contracts (SCs). A SC alone is mainly equivalent to conventional scripts. Nevertheless, because SCs are executed in distributed-ledger-technology-based architectures (DLT) (see Section 7.1), SCs inherit from DLT their enhanced availability and integrity guarantees [55]. DLTs execute SCs synchronously on every node of a P2P network if an arbitrary transaction demands a function's execution. Once deployed, no one can change the script, not even the creators (unless there is an intended call of the script that enables modification), and the script will remain in the network as long as the network exists unless specified differently (e.g., through a self-destruct call). This inherited integrity property of SCs makes them a unique tool to specify and enforce policies between parties or any other process where no trusted third party is available.

Within this review, all the studies that used the Ethereum, Quorum, Hyperledger Iroha, or Ekiden blockchains relied on SCs to declare privacy policies [S2,S25] (Ekiden) [S28] (Ethereum), fair auctions [S3] (Hyperledger Iroha), or payments or incentives [S32] (Ethereum) [S43] (Agora) [S3] (Hyperledger Iroha). However, while SCs ease verification and enable democratic proposals of privacy policies, SCs also inherit the privacy flaws of DLT, i.e., SCs by default imply the disclosure of data and computations to all DLT network nodes [56,132]. For example, the architecture from [S25] employed SCs to set user-defined policies, yet it relied on trusted execution environments to enforce them. *SCs alone cannot enforce privacy policies without relying on other PETs.* The only privacy-related feature that a SC can offer to an IoT data market is declaring privacy policies.

Data access control. Data access control refers to allowing an organization or an individual to choose *who* has access to *which* data. Access control represents a subset of privacy policies in data markets and may utilize different PETs to enforce access rights. While access control is a long-established approach, [S25] propose a novel method, using a key-rotation system [195] in combination with a key manager. Thus, the potential impact of a leaked key is only temporal, with the downside of shorter access permissions.

7. Authenticity-enhancing technologies

The included authenticity-enhancing technologies (AET) focus on enhancing the authenticity of data and identities and also cover data integrity as described in Section 5. Some of the AETs that we describe incorporate privacy-enhancing features, while others do not address or even aggravate privacy protection issues and, thus, need to be combined with PETs.

7.1. Consensus layer

Distributed ledger technology (DLT). While DLT may take different forms, most architectures follow the blockchain design pattern, except for IOTA, which uses the so-called Tangle [196]. A blockchain is a tamper-proof distributed database whose state is stored, synchronized, and replicated by nodes in a P2P network following a consensus algorithm [55]. By its distributed nature, the shared ledger becomes a medium to verify claims, data, payments, or contracts, as once an entity writes something on the ledger, it is practically impossible to modify or erase this record in the future. This property makes blockchain a decentralized and highly reliable alternative to conventional auditing methods like version control [S19].

Benefits of DLT in IoT data markets are the ability to represent the governance, distribution, and roles of authorities on a technical basis [S3], and the enforcement or transparent storage of pre-defined rules by the architects of the respective platform [S25]. Other benefits include eliminating the need for a trusted third party, which removes a single point of failure, improves censorship resistance, and provides more robust data and computational integrity guarantees. DLTs also enable payments through their often built-in cryptocurrencies or other payment systems implemented via smart contracts [S32,S46].

However, some of the studies in our SLR also point at the challenges of current DLT designs: IOTA fails to deliver regarding throughput [S28], is still centralized [S20], and provably has security flaws [133]. Furthermore, blockchains exhibit low transaction throughput [S25], high latency [S11], limited storage [S1] and scalability [S11,S25], computational overhead [S25], high energy consumption [S12], and, most importantly, excessive information exposure that can entail a privacy violation [134]. However, some of these aspects can be mitigated. For example, the energy consumption issue only concerns proof-of-work blockchains [135], and performance can be improved to some extent by private permissioned blockchains that restrict participation in consensus and read access to a small number of nodes in a consortium [136].

Despite the possible operational improvements, employing a DLT for a privacy-enhancing IoT data market needs in-depth consideration. Firstly, through highly replicated storage, a DLT is not suitable for storing large amounts of data produced by IoT devices, not even in a privacy-compliant manner. Consequently, most architectures of the selected studies transfer data through interplanetary file systems [S28], employ a hashing verification approach as described in the *communication layer* [S1,S10,S33] or use Merkle trees [S46]. Secondly, while DLT allows for disintermediation and verification in a trust-less manner, it exposes to the network whatever information someone writes on the ledger for as long as the network exists, which may, among others, violate GDPR's Article 17 "*Right to be forgotten*" for personally identifiable information [137]. Lastly, even if an organization uses a DLT only for the matching and clearing steps of an auction, potentially sensitive business information such as turnover can become available to other network participants, which can conflict with antitrust regulation.

Despite the privacy and performance issues of DLT, 31% of the included papers implemented a DLT as the backbone of IoT data market architectures, employing the Ethereum blockchain [S1,S13,S20,S28,S31,S32,S45,S46], Quorum [S18], the Agora blockchain [S43], Hyperledger Iroha [S3], Hyperledger Fabric [S10,S33], IOTA [S20,S28,S30], Intel's TEE-based consensus Rem [S12], and Ekiden [S2,S25]. Other publications only considered them agnostically [S11,S49] or in a review [S19,S23]. The most salient architectures are described in Table D.7.

Some of the selected studies included privacy-enhancing features in their stack. For example, Quorum supports private transactions and private contracts through a public–private state separation and P2P encrypted message exchange for the direct transfer of private data [S18]. However, the interaction between the private and public ledgers is thus naturally limited and cannot be directly applied, for example, to an on-chain payment system. Another example is Ekiden, which offers a

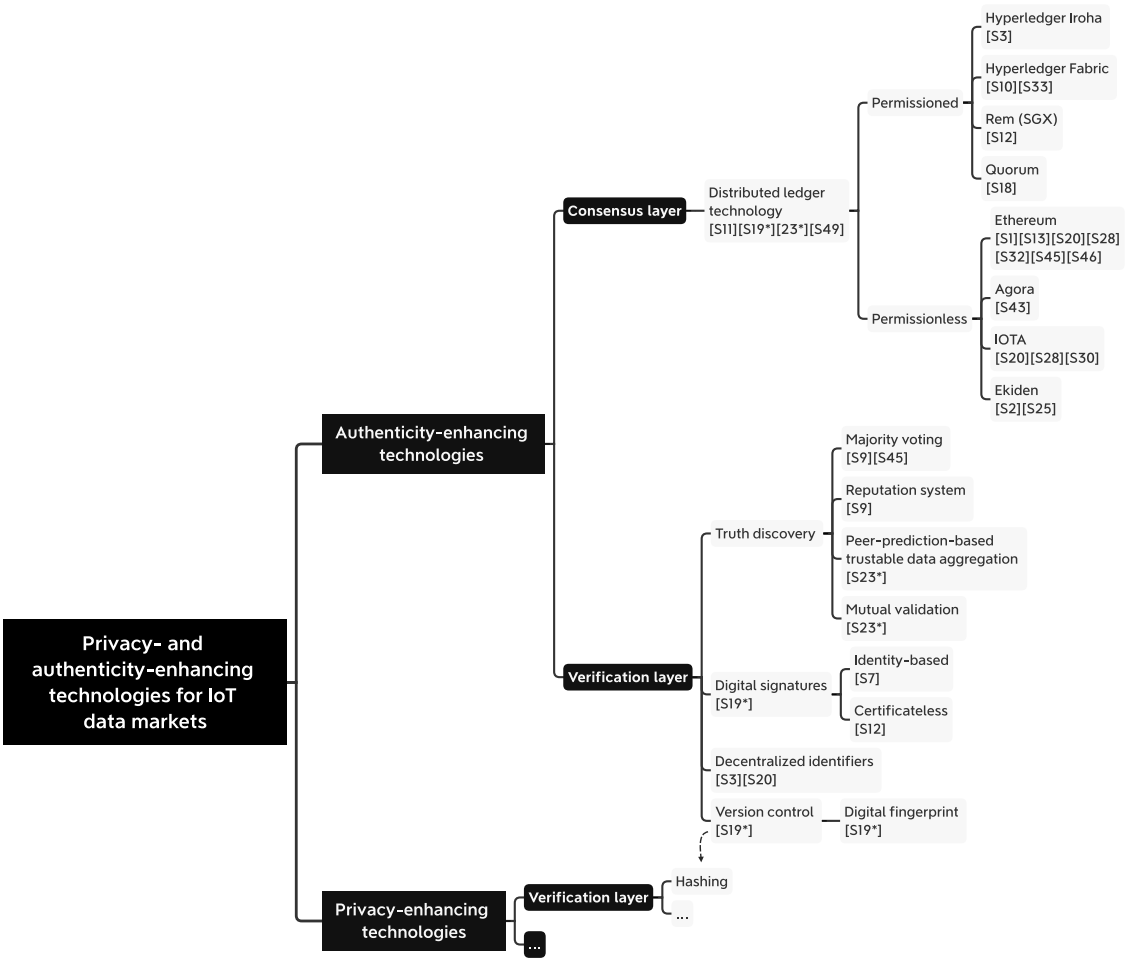

Fig. 7. Classification of the identified authenticity-enhancing technologies in the selected studies of this SLR.

horizontally scalable blockchain potentially capable of hosting end-to-end privacy-enhancing applications through key management protocols and Intel's TEEs [S25] (Note that [S12] uses these TEEs only for consensus, not privacy). Like the solution that [S45] presented, Ekiden allows for smart contracts to execute data analysis in TEEs. However, it is essential to note that these DLTs accomplish the described privacy and integrity functionalities not because of the DLT characteristics but by leveraging the PETs described throughout Section 6.

7.2. Verification layer

This verification layer corresponds to AETs that can be employed for the verification of data and identities. We structure this Section according to Fig. 7.

Truth discovery (TD). TD encompasses algorithms aiming to find the authentic value when different data sources provide conflicting information. As a consequence, TD enhances data and computation integrity and can also enhance identity authenticity, e.g., through reputation systems [S9]. In our SLR, we found that TD takes different forms. For example, the survey by [S23] mentioned a mechanism called peer-prediction-based trustable data aggregation [197], in which the system administrator rewards participants for predicting outcomes of arbitrary events based on other participants' data. This design created incentives for honest reports and therefore enhanced data correctness, resulting in almost all participants choosing to report their bids truthfully [S23]. Moreover, [S23] also proposed mutual validation in which an IoT device compares its data with that of other nearby IoT devices. However, this only applies to specific measurements that are positively correlated for neighboring devices, e.g., temperature, speed of a vehicle, or location in particular settings. It also seems challenging to establish generic handling of differences. Other TD approaches are majority voting, implemented by [S9] in their crowdsourcing architecture and by [S45] in their data processing-as-a-service model. Specifically, given the use of differential privacy in the former approach, they systematically discovered high-quality data with an *estimated* measure of utility that compares individual data points with an aggregate (the "majority"). The latter publication created a reputation system based on the quality of previously sold data.

While most TD approaches leverage transparency to enhance data and identity authenticity and data and computation integrity, TDs are flexible to include PETs such as ZKPs, MPC, HE, TEEs, and DP such as in [S9]. Furthermore, TD can also tackle the oracle problem of DLT, i.e., nodes within the network cannot assure the authenticity of data from outside the network, e.g., the price of a physical asset or the result of an election. For example, ChainLink [138] is an initiative that utilizes incentives to create a trusted oracle network and incorporates many of the principles of TD.

Digital signatures (DS). While DS[1] schemata are commonplace for authentication purposes in today's IT architectures, we have found in selected studies the use of two notable public key cryptography (PKC) schemata that offer some convenience-related advantages over conventional PKC systems:

[1] We introduced the fundamentals of DSs in the *verification layer* within the PETs branch.

- **Identity-based** [S7], where a key generation center (KGC) creates a secret key in a way that the entity's public key can be a publicly available unique string, e.g., the entity's email address. The KGC must be trusted because it holds the master secret key from which all parties' secret keys can be derived. This digital signature assures identity authenticity as the signature is digitally certified by the KGC.
- **Certificateless** [S12] DS schemata are a special form of identity-based PKC whereby an entity's private key is generated by both the entity and a KGC so that the KGC is not aware of the private key of the entity. However, the entity can prove that the KGC was involved in the key generation [139]. This approach assures the authenticity of an entity while tackling the single point of failure of the KGC.

Decentralized identifiers (DIDs). Identifiers can link an entity electronically across multiple IT systems, such as mobile phone numbers, ID cards, user names, or emails. These links are sometimes but not always unique and are facilitated by identity providers that centrally host registries of these identifiers' [140]. In contrast, DIDs are globally unique (with certainty through publishing them on a DLT or probabilistically through randomized generation) identifiers decoupled from centralized registries. DIDs essentially correspond to URLs linked to a file containing one or several public keys and associated metadata that specifies the policies of controlling or interacting with the associated identity. There are two studies in our SLR that employed DIDs in combination with DLT in their conceptual frameworks [S3,S20], described in Table D.7.

Digital fingerprints (DF). DFs are unique physical identifiers that can be attached to or are inherent of items, and thus, one can be sure to interact with, e.g., the right IoT device [S19]. DFs can be seen as a form of version control at a high level. However, attaching an identifier securely to a physical object is difficult unless it has a unique property, e.g., unique metal patterns in the soldering of a chip. However, even if the attachment is relatively tamper-proof, e.g., with a crypto-chip, the same problem also pervades the items that interact with the digital fingerprinted item, e.g., tracking scanners. Therefore, despite the authenticity assurance of DFs, their authentication can only be as truthful as the honesty of the devices that scan the DF.

8. Privacy challenges in IoT data markets

This Section aims to answer RQ2 by distilling the implicit and explicit challenges unveiled in our SLR and other seminal studies [10,24] concerning privacy in the context of IoT data markets. We further classify them into narrow and broad challenges depending on the scope of their definition. Fig. 8 summarizes and outlines the structure of this Section.

8.1. Narrow challenges

Aside from the inherent complexity and low maturity of some PETs and the compatibility issues with legacy systems [24], we identified another specific set of challenges tackled or circumvented by the selected studies.

8.1.1. The trade-off between utility and privacy

Practitioners working with personal data face the challenge of balancing the enhancement of individuals' privacy with the preservation of data's utility [S14]. This challenge is explicitly mentioned by some of the selected studies [S6,S9,S13] and implicitly tackled by others [S2,S24,S25,S35,S47]. This dichotomy is the underlying reason behind the tension between data owners and consumers: the former aim to maximize privacy while the latter intends to maximize utility, which, in turn, is frequently determined by data authenticity (see terminology in Section 5). Furthermore, privacy officers should consider balancing this trade-off at each stage of an information flow [10]: input, computation, output, in transit, and at rest, which increases the complexity of the task. On the other hand, decision makers' or data scientists' quality of judgment depends on computational integrity and the authenticity of data and identities, which is affected by the privacy-utility trade-off.

While PETs from the secure and outsourced computation category seem to circumvent the utility-privacy trade-off by concealing inputs, computation, and outputs, the anointed recipients of these outputs can still perform a re-identification attack. Thus, anonymization PETs, such as differential privacy, should also be included in the stack as they lower the probability of successful re-identification attacks [10].

Data and identity authenticity and accountability bring another problem in the utility-privacy trade-off. Some PETs, namely anonymization technologies, increase plausible deniability at the expense of reducing authenticity and, therefore, accountability [S14,S43]. If the data is fuzzy, the data owner may claim that such a result is not resembling the truth, which is favorable for individual users. However, such protection is not beneficial for society in some contexts, e.g., in criminal contexts. Regarding authentic data from fuzzy identities, if an authority cannot trace data back to the origin, an individual could try to claim plausible deniability, which would hinder processes such as tracking COVID-19 patients to improve pandemic countermeasures. For practitioners to find a balance in these contexts, the requirements of any application or platform should first define the accountability of the involved entities to strike an optimal balance between utility and privacy.

Regarding accountability in data markets specifically, mechanisms to punish misbehavior, such as banning an entity for re-selling or not selling authentic data [S7], can be beneficial to enhance the utility of the market. While AETs such as truth discovery, e.g., majority voting [S9,S45] or reputation systems [S9], incentivize market participants to report honestly about the exchanged data, their identity, and computation integrity, among others, the privacy of the entities is not necessarily enhanced. Additional related issues that may arise are verifying the purchased data's authenticity without violating the individuals' privacy [S7]. For example, if a data broker sells analysis outputs (insights) and not the (privacy-enhanced) original data, the original data owner's digital signature is not valid to authenticate the insights [S7]. Nevertheless, systems could use zero-knowledge-proof-based authentication of data and computation, coupled with value deposits locked in smart contracts to hold participants accountable through enforcing reimbursement across intermediaries. Moreover, other nascent solutions exploit the ubiquity and proximity of IoT devices in specific contexts because the data gathered is likely to be correlated, which allows for mutual data authenticity verification among IoT devices [S23]. However, exploiting correlation can only be used for specific measurements, e.g., weather conditions, vehicle speed, and location, among others.

8.1.2. The recursive enforcement problem

Trust is an essential component in distributed systems that involve different stakeholders and, thus, trust is specifically relevant in the context of IoT data markets. Definitions of trust generally refer to a "*[...] directional relationship between two entities*" [S49] where one entity (the trustor) has subjective expectations on the behavior of the other entity (the trustee) based on previously observed behavior (reputation-based trust) [141] or the belief in competencies and corresponding actions — often within a specific context [142,143] and incentivized by joint interests [144]. Following this definition, we can consider users as trustors of application owners protecting their data when they engage in digital activity. However, the number of data breaches [9] and privacy scandals such as Cambridge Analytica indicate that this trust is not always deserved. The recursive enforcement problem (REP) encompasses the underlying problem of third-party trust with more nuance: Given a third-party authority (A), there ought to be another authority (B) to supervise A, so that A can be trusted. In turn, there should be yet another authority C to supervise B [10], and so forth.

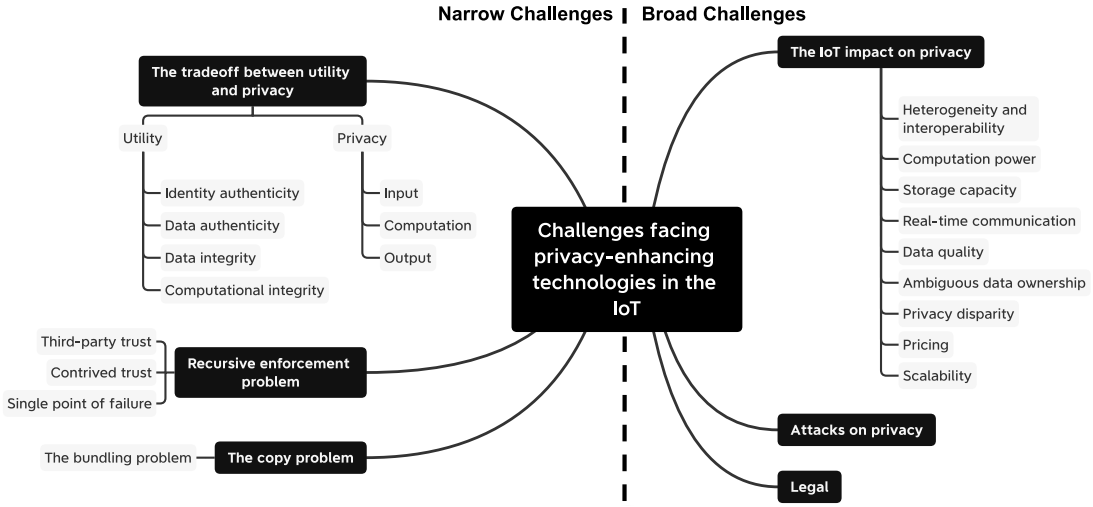

Fig. 8. Overview of the narrow and broad challenges facing privacy-enhancing IoT data markets.

The REP is a significant challenge that has been covered and tackled implicitly by some of the studies in our review [S1,S3,S4,S7,S10,S16, S25,S41]. Additionally, others tackle a sub-set of the REP, which is the single point of failure of trusting a unique third party [S6,S12]. According to [S8], the hesitation in trusting third parties is one of the main reasons for the slow adoption of IoT data markets. Indeed, it is hard to technically ensure and prove that the third party will not use one's data for purposes other than those agreed [S8]. Additionally, adoption is further slowed down because the use of third parties to supervise other parties incurs costs [S46]. Furthermore, users' daily interactions with "trusted" third parties can be regarded as a product of *contrived trust*, another form of the REP. For instance, applications from large service providers with negligible competitors push users to accept the sometimes poor privacy conditions, e.g., GPS apps. Note that *contrived trust* is different from the trust users have on cryptography, open-source code, or consensus mechanisms that the broad scientific community has audited over the years.

Tackling the REP requires reducing the power and the responsibility of the third party in a particular aspect of a specific service by, e.g., distributing such responsibility among other parties or distributing the power among multiple parties that enforce rules on each other. These measures can ease the hesitation to trust a single third party, tackle contrived trust, and reduce the single point of failure because the third party would be supervised and held accountable by other third parties in a flat hierarchy. Fortunately, the PETs included in this study can also circumvent – only onion routing can tackle – the REP, which, in turn, reduces the need for third-party trust and, therefore, reduce contrived trust and a single point of failure. Additionally, 5 of the 7 AETs included in this study can tackle the REP, primarily distributed ledger technology, whose architecture was purposefully built to tackle the byzantine generals' problem [145,146], a manifestation of the REP.

8.1.3. The copy problem

Once an entity releases data freely or for profit-seeking, the data is no longer under the original owner's control. Consequently, the recipients of such data can *copy* and, e.g., re-sell or use the entity's data for a non-agreed purpose without informing or acknowledging the original owner [10]. Beyond the privacy threats the copy problem (CP) entails for users of service providers under poor privacy conditions, the CP is a major obstacle for organizations to engage in data markets, which some of the selected studies implicitly tackle [S2,S25,S45]. The CP leads companies to either hoard or sell data as fast as organizations obtain the data, lest its value drops [10]. Nonetheless, secure and outsourced computation PETs such as trusted execution environments or homomorphic encryption can tackle the CP by allowing other entities to extract value without losing control over the input data beyond the specific information sold, such as an algorithm's evaluated output on this data. This paradigm is profound because tackling the CP makes data scarce (to some extent), as the original data is not shared, and the data owner would not allow a non-agreed computation. Thus, selling the *access* to data can be more attractive to companies, as data would preserve their value longer than releasing the data.

A subset of the CP is the *bundling problem* (BP) [10], which is an attack vector different from re-identification that occurs when an entity requests actively or passively more data than strictly needed to (i) prove a claim or (ii) perform an analysis. Harvesting more information than needed worsens data breaches' consequences for individuals and companies and indicates questionable business ethics. For instance, (i) to prove one's age with an ID card, the prover usually shares all the information in the ID instead of only the age and proof of the card's authenticity. The BP is a subset of the CP because if one tackles the CP, neither necessary nor additional information is released beyond the required computation. For example, tackling the CP by restricting verification and processing to a trusted execution environment also tackles the BP. In this setting, the data consumers cannot copy the necessary data or metadata for other unsolicited analyses, despite being able to process metadata to verify the authenticity of the data and the integrity of the computation and obtaining the desired outputs of the analysis. Additionally, (ii) anonymization-based PETs such as differential privacy or k-anonymity reduce data authenticity to tackle the BP. For instance, in a demographic analysis that only requires the first digits of the ZIP code to perform clustering, data curators can generalize the ZIP codes with k-anonymity, so only the strictly necessary information is revealed to data scientists. Nonetheless, anonymization can suffer from background-knowledge-based attacks [S48,156] and does not solve the CP because the data consumers can replicate the privacy-enhanced data.

8.2. Broad challenges

8.2.1. The IoT impact on privacy

The paradigm brought by the IoT brings significant amounts of data to markets. However, this paradigm also bears some of the shortcomings of IoT devices [198]. Table 2 contains an overview of these challenges and briefly discusses their impact on privacy. In summary, privacy is always affected by the context and employed technologies, which underlines the importance of adhering to privacy-by-design principles [131] and the need for practitioners in other fields such as software engineering, economics, law, and politics to tackle together the diverse issues that IoT entails for privacy.

Table 2
Overview of challenges brought by the IoT paradigm into data markets explicitly covered by some of the studies included in this SLR.

Challenge	Studies	Description	Impact on privacy
Heterogeneity and Interoperability	[S36] and [S39]	The IoT consists of billions of IoT devices from different manufacturers, running different software on different local networks and geographic regions, with different computation power and storage capacity [147]. Furthermore, different communications standards, connectivity and availability aggravate the interplay of IoT devices.	An IoT data market should be agnostic to these differences and minimize any additional requirements; however, it is unclear how global data markets should harmonize data coming from different jurisdictions with different privacy regulations and how an IoT device can interact with another whose, e.g., verification schemata are considered inadequate. In addition to these obstacles, a lack of interoperability may restrain PETs that involve the communication between many devices, e.g., MPC.
Computation power	[S5], [S11] and [S48]	Manufacturers produce many IoT devices designed to consume low energy and require minimal volume, limiting these IoT devices to the core functionalities of monitoring and communication [S11].	Any additional computation requires a higher investment in resources and manufacturing, and running some PETs becomes infeasible without this extra investment. Consequently, a set of PETs is excluded without more computation power, e.g., cryptography-based PETs such as HE, MPC, ZKP, some digital signatures, or consensus algorithms. This limitation, however, may only apply to contexts where it might not be possible to connect IoT devices acting as clients with proprietary or trusted third-party nodes where these PETs are executed.
Storage capacity and real-time communication	[S1], [S5], [S11], [S17], [S29] and [S48]	Minimizing the physical volume of an IoT device reduces their price but limits their storage capacity, forcing IoT devices to transmit the data to a data warehouse or a data market as quickly as possible. This tendency intensifies in some IoT applications where the time delay tolerance is low to enhance the utility of real-time information [S17].	Processing time constrains the number of usable PETs, excluding those that require long execution times, such as fully HE or creating a ZKP.
Data quality	[S14]	An unreliable IoT design may afflict thousands of IoT devices mass-produced by a manufacturer, which at deployment may lead to millions of unreliable data points. Furthermore, networks may also be unreliable, further worsening the quality [S14].	The impact may seem beneficial in terms of privacy; however, unreliable data leads to verification and secure computation schemata to fail and anonymization technologies to over-perturb the data as the underlying data is not entirely truthful.
Ambiguous data ownership	[S4], [S10] and [S14]	When purchasing a device or a cluster of IoT devices, e.g., a phone, consumers also expect to own the data they are generating. However, the phone manufacturer and service providers expect to receive parts of this data nowadays with meager consent. In addition to this clash of interests, there are scholars that ponder whether data belongs to anyone in the first place, like [148].	Having unclear data ownership leads to a misguided deployment of PETs, which may cause detrimental consequences if the privacy measures fall short. On the other hand, if the practitioner knows who has the right to the data and what the owner is reticent to share with a third party, then selected PETs and their privacy tuning can be optimized accordingly.
Privacy disparity	[S4] and [S40]	Depending on the IoT devices' deployment location, the degree of privacy measures should be higher or lower, e.g., sensors in vehicles, smart homes, phones, and wearables. Furthermore, IoT deployments should adapt the monitoring time to an adequate amount depending on the context [S40].	Some PETs, such as semantic and syntactic technologies, allow adjusting the degree of privacy; however, others are more rigid. Selecting and adapting a PET to the IoT devices' deployment context requires expertise.
Pricing	[S4], [S14], [S17] and [S50]	There are multiple variables imposing the price of data aside from supply and demand: the truthfulness, the source, either purchasing the data or the access [S17], and the privacy level. These factors add additional complexity to pricing, e.g., the sources have become disparate with the IoT, which drives pricing to a more granular task than before, when aggregated data could be sold as a unit [S17].	Aside from payment enforcement mechanisms, pricing involves negotiations, which frequently must ensure privacy. This adds an extra layer of complexity to the deployment of PETs. Furthermore, as data markets trade with more granular data points, PETs that need aggregation might be excluded in some contexts, e.g., syntactic technologies such as k-anonymity.
Scalability	[S25] and [S36]	The number of IoT devices and streamed data grow exponentially across industries [147,149], which extends data collection and improves analytics across different domains, e.g., health, insurance, or finance. To gain these benefits, there is a need to increase networks' communication and overall storage capacity as well as interoperability and security efforts.	As the IoT scales, analysts will access more datasets from different domains to create new products and services, e.g., linking driving behavior with insurance in pay-how-you-drive schemata [150]. Such innovations stem from the "*mosaic effect*" [151], where disparate datasets with limited information value can obtain significance when combined with other datasets. However, malicious entities can leverage such an effect to extract sensitive personal information not explicitly contained in a dataset [152].

8.2.2. Attacks on privacy

Adversaries can be malicious, actively trying to breach users' privacy through hacking, or honest but curious, passively gathering data from users to reveal hidden insights [S35]. Both of these entities can carry re-identification attacks with the collected information. Within the context of the IoT, the list of security and privacy attacks is extensive (sniffing, cache poisoning, DoS/DDoS, sinkhole attacks, replay attacks, among others) [153]. Furthermore, within our SLR, [S45] discuss some of the additional attack vectors these malicious or curious entities may execute in the context of IoT data markets to learn sensitive information from users. Notable ones include: *Data forwarding*, which is one way the *copy problem* materializes; *roles collision*, where data brokers and buyers may be the same or collaborating entities, and, therefore, the broker could rig the auction for its benefit and access the sold data; and *side channel attacks*, where attackers exploit the physical properties of the hardware or its power consumption to extract knowledge from the hidden computations (trusted execution environments suffer mainly from this attack).

Such attacks make the possession of data intrinsically risky because if attackers are successful, data re-identification is possible [S38], even if data have undergone some form of privacy enhancement [154]. There are common attacks used to re-identify data, e.g., reconstruction, tracing, or linkage attacks [118,155]. Some of the most famous re-identification *white-hat* attacks involve [156] who deanonymized the Netflix Prize dataset with IMDB's public dataset in 2008, [157] who performed the same feat with Amazon's public review data, and [6] (the inventor of k-anonymity) who re-identified participants within a genome sequence dataset in 2013. Furthermore, in 2014, [7] tracked drivers with home address and vehicle speed as inputs, and in 2020, [154] matched users with large-scale mobility datasets from a mobile network operator and transportation smart card usage.

Overall, IoT data markets will facilitate access to large quantities of data from different domains, including biometrics, which will increase the impact of these attacks and the potential harms to individuals, e.g., insurance, employment, or price discrimination. Therefore, IoT data markets require a more robust adoption of PETs and security standards.

8.2.3. Legal challenges

Progressively along the past decades, governmental institutions have released laws to protect the privacy of their citizens (see Section 2.1). These laws also refer to an individual's and businesses' right to exploit their data commercially, which provides leeway for data markets [S28] and aims to uncover the untapped potential of data for innovations.

Nonetheless, research points out the sometimes unrealistic expectation to monitor the entirety of the Internet for privacy violations [S45], and the dexterity of hackers to find novel deception methods [S3], and that laws are more reactive than preventative. Well-known networks of illegal proprietary digital asset exchanges, e.g., scientific works and how users of digital services give away data, tacitly provide testimony of the failure of data-related legal measures today, and the problems will likely increase with the accruing number of IoT devices [S38]. Moreover, privacy regulations can strangle free markets and innovations if they are too stringent [S45].

Aligned with these deficiencies, [S38] introduced privacy regulation pitfalls that the IoT unfolds in data markets in 2013. They note that (i) definitions of personally identifiable information will be deprecated as unprecedented amounts of data can be aggregated, easing re-identification, (ii) the development and audit of PETs is costly, which may limit business models and potentially make disregarding privacy regulation profitable [199], (iii) privacy violations result on small fines or remain unpunished, (iv) technology tends to outpace regulation, and (v) the ubiquity of IoT devices will yield more illegal secondary personal data markets. After almost a decade of further research, (i) seems valid, at least in some scenarios. The ambiguity of privacy regulation is a barrier in some cases, as practitioners may default to weaker forms of privacy if their architecture appears to comply. This leads to re-identification – an attack that is also more practical with the increasing number of IoT devices [156] – being more likely to succeed. However, in defense of these practitioners, while PETs have improved since 2013, some PETs that offer better privacy enhancements are still complex and not yet performant in 2021.

Based on the prior arguments, pitfall (ii) seems to hold; however, (iii) is no longer a strong pitfall. Since the enforcement of GDPR [137] in 2018, GDPR has punished multiple corporations with considerable fines ranging between €20 million and up to 4% of a corporation's annual worldwide turnover of the preceding financial year. As of the writing of this publication, GDPR has harvested considerable fines assigned to Google in France on two occasions [200], Amazon [158], H&M [159] or the telecommunications operator TIM [160]. These fines alone accumulate to €282 million. These statistics are a sign that PETs are not appropriately introduced in production applications even by *big* technology companies and that not complying with privacy regulations in an IoT data market has dire economic consequences. While these fines could indicate how profitable it still is to violate privacy regulation (iii), one can no longer vigorously defend (iii). Pitfall (iv) seems to materialize as long as the nature of law-making does not change. Lastly, pitfall (v) is concerning, given the existence of legal personal data markets that store up to 750 million user profiles and trade 75 million online auctions daily like BlueKai [23], whose data could leak to the increasing number of illegal *shadow markets* [S6].

9. Discussion

This Section presents a set of key findings (KF) distilled from the two research questions answered in Sections 6 and 7 as well as 8, the content and metadata of the 50 publications included in our SLR, and other seminal studies that we encountered throughout our SLR but which do not necessarily address IoT data markets directly. Lastly, we cover the limitations of this study and future work.

9.1. Key findings

(KF1) *The attention of scientists toward privacy-enhancing technologies in the field of data markets for iot devices has increased notably in recent years.* The selected publications are modern, as 49 of the 50 studies were published between 2012 and 2020, and 34 of them (68%) were published either in 2018, 2019, or during the first half of 2020. While the absolute number of publications in 2020 is lower than in 2019 because we captured only the first seven months of 2020, Fig. 9 illustrates the arguably accelerating trend of the cumulative curve of publications in the field of privacy-enhancing IoT data markets.

(KF2) *The most frequent research type (design and creation) and least common research contribution (lessons learned) suggest that privacy-oriented iot data markets are still maturing and have not faced many production-grade implementations yet.* According to Fig. G.13, around 76% of the publications use a *design and creation* research approach, while only 4% perform a case study. A further indication of field novelty is that only one out of the 50 publications had the contribution type *lessons learned* [S40]. Furthermore, while 35 studies (70%) were of research type *solution proposal*, to the best of the author's knowledge, only one solution appears to have an implemented system that is applied in production [S25].

(KF3) *The selected studies rarely leverage existing libraries that provide PETs and often only build upon architectures developed in previous work to a small degree. Therefore, to gain more practical relevance, it may be beneficial for researchers to improve and extend existing work instead of reinventing the wheel.* The research community and industry have developed many open-source libraries to employ zero-knowledge proofs,

Table 3
A mapping of privacy- and authenticity-enhancing technologies (PET and AET) to the narrow challenges using the terminology defined for this review. The extent of enhancement of privacy, utility, and characteristics of the different PETs and AETs varies from significantly increasing ++, over +, +−, − to significantly decreasing − −. na denotes *not applicable*. w/ denotes *with*. *Considering a digital certificate when using digital signatures, if applicable. The privacy column assumes data and identity are authentic.

Layer	Technology category	Technology	Narrow challenges: Privacy-utility tradeoff — Privacy [Confidentiality (Conf)] Input [Data (D)] [Identity (I)]	Privacy: Computation	Privacy: Output	Utility: Authenticity [Data (D)] [Identity (I)]	Utility: Integrity [Data (D)] [Computation (C)]	Recursive enforcement problem	Copy problem [Bundling problem (BP)]
		PET							
Processing	Secure and outsourced computation (SOC)	Zero knowledge proofs for computational integrity	++ D	++	++	++ D	++ D, C	Circumvents	Tackles (BP)
		ZKPs of anonymous credentials	++ I	++	++	++ I	++ D, C	Circumvents	Tackles (BP)
		Trusted execution environments	++ D	++	+− Conf	na	++ D, C	Circumvents	Tackles
		Partially homomorphic encryption	++ D	++	+− Conf	na	++ D, C	Circumvents	Tackles
		Fully homomorphic encryption	++ D	++	+− Conf	na	++ D, C	Circumvents	Tackles
		Secure multiparty computation	++ D	++	+− Conf	na	++ D, C	Circumvents	Tackles
		Privacy-preserving data mining — W/ SOC (w/ AN)	++ D (+ D)	++ (na)	+− , Conf (+)	na (+− D)	++ D, C (na)	Circumvents	Tackles
		Privacy-preserving data mining — Federated (w/ AN)				na (+− D)	na (na)		
	Anonymization (AN)	Differential privacy (DP)	+ D	na	+	+− D	na	na	Tackles (BP)
		K-anonymity	+− D	na	+−	− D	na	na	Tackles (BP)
		Perturbative	+− D	na	+−	− D	na	na	Tackles (BP)
		Pseudonym creation	+− I	na	+−	− I	na	na	Tackles (BP)
Storage / Communication		Encryption	Storage layer: ++ D, Conf; Communication layer: ++ D, Conf			na	++ *D	Circumvents	na
Storage / Communication		Onion routing	++ I +− D, Conf	na	na	na	++ D	Tackles	na
Verification		Hashing	+− D, Conf	na	na	na	++ D	Circumvents	na
Verification		Ring digital signatures	+ I	na	na	− I	++ D	na	Tackles (BP)
Verification		Blind digital signatures	−− *I + D	na	na	++ *I na D	++ D	na	Tackles (BP)
Sovereignty		Smart contracts (for privacy policies)	Characteristics and challenges are pegged to distributed ledger technology.						
Sovereignty		Privacy policies (Access control)	Characteristics and challenges are pegged to the selected PETs and AETs employed to fulfil the privacy requirements.						
		AET							
Consensus		Distributed ledger technology (permissioned)	− I +− D, Conf	+− Conf	+− Conf	+ I na D	+ D, C	Tackles	na
Consensus		Distributed ledger technology (permissionless)	+− I − D	−	−	+− I na D	++ D, C	Tackles	na
Verification		Truth discovery	− I − D	−	−	+ I + D	++ D, C	Tackles	na
Verification		Decentralized identifiers	− I	na	na	++ I	na	Tackles	na
Verification		Digital fingerprint (Version control)	−− I − D	na	na	++ I ++ D	++ D	na	na
Verification		Identity-based digital signatures	− *I	na	na	++ *I	++ D	na	na
Verification		Certificateless digital signatures	− *I	na	na	++ *I	++ D	Tackles	na

homomorphic encryption, secure multiparty computation, or differential privacy (see Section 6). However, none of the studies have indicated their use. Furthermore, studies often do not build upon each other, leading to overlapping further. For example, [S41] and [S34] both showcase an auction that obscures the bids by employing partially homomorphic encryption. However, W. Gao et al. only refers to the work from Z. Chen et al. in one line, noting that "[...] *there is only few literatures on designing privacy-preserving schemes in data market auctions*". Moreover, [S42] builds upon [S27] and [S2] upon [S25], but each of these two sets belongs to the same group of researchers. In conclusion, it may be beneficial for researchers to incorporate building blocks from previous data market architectures to advance privacy-oriented research.

Moreover, many studies included in Table D.7 aim to create an IoT data marketplace employing distributed ledger technology (DLT). However, there seems not to be a consensus about which DLT to use for IoT data markets, as the authors build upon Ethereum, IOTA, Hyperledger Iroha, Fabric, Agora, or Quorum, among others. Specifically, as an example, [S32] uses Ethereum smart contracts for payments while [S28] only uses these contracts for safelisting and employs IOTA for payments instead.

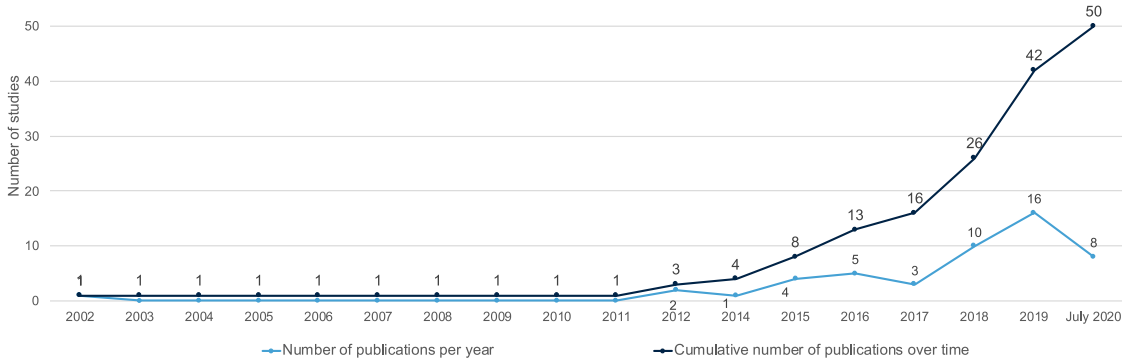

Fig. 9. Publications in the field of privacy-enhancing data markets for the IoT from January 2002 to July 2020.

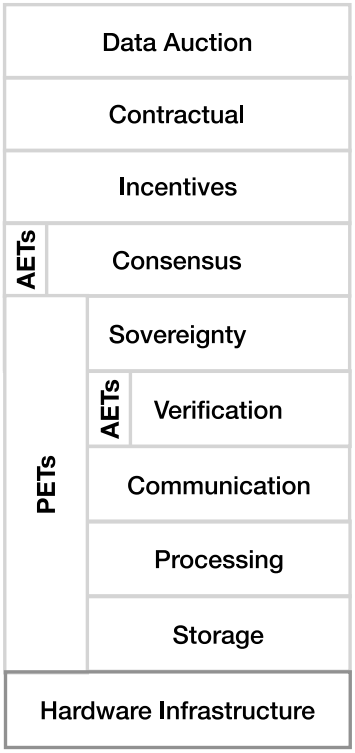

Fig. 10. Reference model for the layers of a privacy-enhancing IoT data market.

(KF4) *The content of the selected studies can be categorized into two main orthogonal research streams within the context of privacy-enhancing iot data markets: architectures and data trading schemata.* The first research stream is dedicated to the design of privacy-enhancing architectures for the exchange of data in IoT data markets (25 studies, 50%), and the second one focuses on the design of privacy-enhancing data trading such as auctions (12 studies, 24%). The remaining studies can be associated with domains like legal [S6], user preferences [S8], or IoT data market challenges [S14,S19]. The selected studies, and also international initiatives such as the European GAIA-X [5], hence envision data markets beyond matchmaking and auction capabilities. Specifically, the studies that we analyzed structure the software, hardware, abstract entities, and their coordination, data processing, storage, communication, and the offered services to build a holistic or part of a privacy-enhancing IoT data market that includes PETs to tackle some of the challenges described in Section 8.

(KF5) *Despite the acknowledged need for combining anonymization and secure and outsourced computation techniques, none of the researchers behind the 12 studies proposing data trading schemata, and only two publications out of the 25 designing data market architectures employ both PET categories in combination.* Although PETs such as homomorphic encryption (HE) or secure multiparty computation conceal inputs and computation, the outputs can leak information about the underlying data and hence may be exposed to re-identification attacks [10]. Combining secure and outsourced computation techniques with anonymization-based PETs like differential privacy (DP) can help to make the outputs less sensitive. Moreover, leveraging only anonymization PETs does not sufficiently address the copy problem.

Within the 12 selected studies focused on data trading, DP is the most frequently used PET in auctions to enhance the privacy of the exchanged data. [S16,S22,S24] and [S37] employ DP in various forms to set the privacy levels and, subsequently, the price of the traded IoT data. Researchers might choose DP over other anonymization technologies because DP is the only PET with a mathematical guarantee of privacy [13]. At the same time, partially HE (PHE) is the PET of choice to enhance the privacy of the bidding process. A group of authors [S26,S34,S41] chose PHE primarily for hiding the bids, confidentially computing the winner, and only revealing the output to the auction's winner. Researchers might decide to use PHE over other forms of HE, secure multiparty computation, or trusted execution environments (TEEs) despite PHE's significantly less general scope because PHE has relatively high performance and is conceptually simple.

Together, DP and PHE can holistically enhance the privacy of auctions, which is a contribution we have not found in this review. [S48] emphasize that some HE schemata, such as Paillier's, must complement other methods to guarantee more protection. Moreover, while HE protects the input and the computation itself, if the intended recipients of the decrypted output are malicious, they may reverse engineer the output to learn properties about the input. An additional modification employing, e.g., differential privacy, of inputs or decrypted outputs before sharing may help prevent this attack in exchange for accuracy and thus utility. The same argument applies to other secure and outsourced computation methods when used in isolation. We consequently point to a lack of combination in the research stream of data market architectures, except for two publications from the same group of researchers [S2,S25], which use TEEs to train machine learning models with DP.

(KF6) *The selected studies employ three dimensions to characterize data markets that entail privacy concerns: the degree of decentralization, the types and number of data domains, and the types of sellers and consumers.* Each of these dimensions, for example, characterized by [S1,S27] and [S46] respectively, brings privacy concerns. Data may be stored by the seller, the platform provider, or a decentralized platform using, e.g., a combination of commercial cloud storage, interplanetary file systems, or blockchains. Depending on the degree of decentralization and replication, practitioners need to consider different leakage risks. In particular, if the architecture relies on a blockchain, PETs are particularly important [134].

An increase in the number and types of data domains opens additional attack vectors and more possibilities for malicious entities to link an individual's data across databases. This hyper-connectivity between datasets can render the definitions of de-identified data, such as HIPAA's, obsolete and suggests that privacy enhancements in the data economy should be defined globally and not locally.

The degree of privacy enhancement should depend on the type of seller and consumer, e.g., consumers may expect higher privacy guarantees when a health insurance company gathers their data than when the collector is a renowned health research institution.

(KF7) *Based on our classifications in Sections 6 and 7 and inspired by a set of seminal selected studies, we have created a reference model for the design of iot data markets in Fig. 10, and detailed in Table 3.* Most of the studies included in this SLR proposed solutions without following a reference model, except for [S3] and [S7], who developed their own without a systematic research. [S7] condense their architecture into two layers: data acquisition and trading. On the other hand, [S3] present a more holistic view of privacy-enhancing IoT data markets with six layers (*identification, privacy, contractual, communication, consensus, and incentive*) inspired by the Open System Interconnection model and heavily conditioned by the use of blockchain technology. This model, however, lacks essential steps of an IoT data market that several publications in our SLR focused on, namely storage [S1,S12,S25,S28] and processing [S7,S19,S20,S45]. Furthermore, the *identification layer* [S3] can be regarded as a subset of verification, which also includes data verification. Other studies base their market design on the type of participants [S15,S24,S27,S33,S43], e.g., sellers, aggregators, brokers, among others, and the type of data domain [S46], e.g., health, financial, or a combination. However, these categories cannot be transferred to other contexts as easily as a reference model agnostic to entity and data domain types.

Our reference model hence combines and generalizes some of the layers from [S3,S7] and complements them with additional layers such as the data auction, storage, verification, processing, and sovereignty layer (see Fig. 10). Most of these layers need multiple PETs, as there is no "one-size-fits-all" technology to enhance privacy. To navigate these layers in detail, refer to Table 3 amd Fig. C.11. Furthermore, we distinguish between a contractual and sovereignty-related design to separate formal agreements from privacy and ownership policies. Furthermore, given the distinct purpose and implementation that auction schemata play in a data market, they should be respected by a unique IoT data market layer (auction dedicated studies: S16,S22,S24,S37, among others). Lastly, incentives are necessary to encourage behavior that preserves the pre-defined qualities of the IoT data market, e.g., optimized prices [S3,S16], data authenticity [S16,S23], or maintaining the infrastructure like a permissionless DLT.

(KF8) *Aside from the ubiquity of digital signatures in IT systems, in this slr, distributed ledger technology (DLT) is most frequently employed as the backbone of iot data market design (see fig. g.14), despite the lack of consensus on its use and DLT-based applications in production.* Although centralized systems seem more efficient and easier to deploy, and despite the seemingly few industrial applications running on blockchain today, many researchers in this SLR still advocate for distributed systems using DLT. Within the 35 solution proposals, around 31% chose permissionless DLT, 14% consortium DLTs, and the authors of the remaining 55% either reviewed DLT, implemented a centralized solution, or focused on designing narrow features. However, we noted that within the 45% of DLT-based designs, many authors still relied on single entities for data processing or storage. Specifically, only one of the 50 studies [S25] has a public blockchain-based ecosystem in production, yet without a real-world use case running. These statistics indicate a lack of adoption despite substantial research efforts.

Furthermore, while blockchains enhance authenticity, assure integrity, and enable payments without the need for a trusted third party, blockchains are limited in storage capacity [S1], computation power [S25], and can exacerbate privacy issues because of their tamper proof-quality and inherent data and computation replication [56,132,136,161]. Consequently, almost all studies that include blockchain technology to support an IoT data market require PETs to protect users' data and identities. These studies go as far as creating innovative privacy-enhancing blockchain architectures with other PETs as building blocks, e.g., trusted execution environments [S25,S45], or adding a privacy layer to their market design based on differential privacy [S3]. However, within the literature, there are also questionable statements such as "[...] *researchers and technologists have found that blockchain can be a potential solution to the privacy problem by decentralizing information* [...] *Blockchain can be used to securely share private information* [...]" [S3], "*Blockchain-based approaches provide decentralized security and privacy* [...]" [S11], or "*Blockchain has been proven to possess security, immutability, and privacy properties, which has caused a lot of researchers to introduce it into the privacy and security concerned IoT*" [S23]. These statements, coupled with the current excitement around blockchain, can lead practitioners in the industry to wrongfully push blockchain for "privacy". Therefore, the community would benefit from clear explanations of why authors employ blockchain and clearly state the need for other technologies to enhance privacy.

9.2. Limitations

Even though we have adopted a rigorous research design and paid particular attention to the selection and analysis of published studies, SLRs have limitations that may have undermined our effectiveness. These threats include (i) incompleteness of study search, (ii) bias in study selection and (iii) inaccuracy of data extraction.

(i) Some relevant publications might be absent. To mitigate this limitation, we searched in several highly reputed digital libraries, performed a preliminary search to determine suitable search strings, conducted a backward search to identify additional related work, and included studies in advance that met the standards and filters of this SLR. These measures reduce the probability of missing relevant publications. (ii) The experience and knowledge of the researchers may drive the study selection with an inherent bias. Nonetheless, following Kitchenham [43], we aimed to create a set of explicit inclusion and exclusion criteria to maximize the degree of objectivity. To mitigate different appreciations of these criteria, we conducted a preliminary search to ensure researchers have a consistent understanding of the requirements. Furthermore, two researchers conducted the selection process independently and resolved the conflicts between their decisions interactively. (iii) There might be a bias in selecting the extracted data, which may affect the classification results of the selected studies. To mitigate this potential limitation, the two researchers specified a set of data extraction cards (see Section 3.2) to eliminate any misalignment in the data extraction process results.

9.3. Future work

The opportunities and need for future work in the context of privacy and data markets for the IoT highlighted by the selected studies resonate with the challenges covered in Section 8. Most notably, there is a need to solve the copy problem [10,S4,S17] and to lessen IoT devices' limitations regarding computation [S5,S11], storage and capacity [S29,S48] to tackle or circumvent the constraints PETs may induce. Moreover, to decrease the probability of re-identification attacks, further work is needed to advance the maturity of PETs and combine them, e.g., bringing together differential privacy and secure and outsourced computation efficiently. Additional research is also necessary to create standards for data markets, such as a language to describe privacy requirements, universal APIs to interact between different IoT devices with various degrees and techniques for privacy protection, and machine-readable definitions of privacy, e.g., using ontologies [S40]. In this context, a more detailed description and

classification of the layers that we found relevant for classifying privacy and authenticity enhancing technologies (see Fig. 10) constitutes a promising and relevant avenue for future research. Nonetheless, we want to emphasize that privacy is not the only challenge that needs to be addressed, as future research must also consider, for instance, scalability.

If society considers privacy a necessity, it should be enhanced by default and optimally in any system without attaching price tags to one's privacy, as some of the selected studies pursued suggest [S16,S15, S22]. We find this posture a worthy research endeavor and encourage researchers to ponder whether monetizing privacy in a competitive market ultimately benefits society. Furthermore, legal practitioners have ample ground to develop legislation specifically around privacy in IoT data markets and for economists to delve into data pricing and decentralized market interactions. Legal researchers could investigate how stringent privacy regulations should be, as heavy regulation may strangle free markets and innovations [S45]. Additionally, the legal, pricing and privacy aspects hinge around data sovereignty. As long as ownership is ambiguous, researchers' efforts will struggle to maximize impact. Furthermore, the relevance of our results may reach beyond IoT data markets, as the analysis of PETs and derived insights, e.g., how IoT impacts privacy, can permeate other research areas such as privacy-by-design software engineering, policy-making, and data governance, politics, and economics. Moreover, most PETs have specific performance-, complexity- or utility-related shortcomings (which we describe in Section 6) that researchers can address.

Lastly, we recommend that researchers derive decision trees based on Table 3 to enhance the decision-making of privacy officers beyond our work. Moreover, we could not find any formulation of an information-theoretic quantification of the data leaked from a data market. We also encourage social scientists to focus on questions related to data sovereignty. To realize a vision of data markets that benefit society, we suggest researchers concentrate on roadblocks such as the copy problem. Finally, institutions should consider updating their privacy-enhancing processes to effectively participate in IoT data markets.

10. Reassessment of the results

This section provides and discusses new key publications since the research process ended. Accordingly, we conducted a research process as per Section 3 for studies dated between July 2020 and May 2022 and, among them, picked for discussion the ones providing the most significant updates to our systematic literature review or, on the contrary, underlining our previous findings. Note that the references included in this section correspond *only* to the newly found publications.

In our new search, we again selected primary and secondary studies. Overall, our new search resulted in 24 publications: 3 more from 2020, 14 from 2021, and – as of May – 7 from 2022. These statistics indicate that the trend depicted in Fig. 9 (consolidating KF1) has not reversed. Notably, we could still not find publications discussing production-ready deployments of privacy-enhancing architectures or auction schemata and no reference to open-source tooling despite PETs being more mature since July 2020 (underlining KF2 and KF3).

Among the *secondary studies*, [162] explored the concept of privacy in the digital economy more broadly and pointed out the need for interdisciplinary research to supplement the purely technical PET constructions with the economic (tradeoff between accuracy and privacy) and governance perspectives (privacy policies) that we elaborate on in our paper. [163] systematically investigated privacy-oriented identity management in the context of the IoT, such as anonymous credentials and other techniques that our review covers. Moreover, [164] conducted a less systematic survey of standards and future challenges, including discussions regarding authentication and access control, and highlighted a subset of the privacy challenges of IoT that we present in Table 2. Additionally, [165] presented a recent systematic literature review on designing data markets. However, their work did not focus on privacy.

More secondary studies, such as from [166] considered blockchain and smart contracts beneficial for privacy. Furthermore, [167] reviewed PETs in the context of crowdsensing and emphasized the privacy issues with smart contracts, along with practical challenges in security and feed-in of reliable data. They suggested a subset of the anonymization techniques that we present in Section 6, such as privacy-oriented digital signatures, anonymous networking, k-anonymity, l-diversity, t-closeness, and differential privacy (DP), and some more specific ones in the context of location. While they also mentioned ZKPs, there is no detailed discussion of the secure computing techniques we survey. [168] considered privacy mechanisms in data sharing for collaborative forecasting and discussed the tradeoff between privacy and accuracy. They distinguished between perturbative techniques ("data transformation"), MPC-based protocols, and distributed or federated approaches ("decomposition") combined with DP. Lastly, [169] provided a survey that examined the privacy risks that machine learning poses on IoT data markets supported by blockchains. Hence, our SLR, with its comprehensive focus on privacy, still fills the gaps that we discussed in Section 4.

The *primary studies* followed a similar pattern to the previously collected studies. Above all, many still employed blockchains and often did not provide clear explanations the corresponding benefits and acknowledgments of the corresponding challenges, specifically regarding privacy (reaffirming KF8). We again encountered questionable claims such as "[...] a decentralized approach based on distributed ledger technologies (DLT) enables data trading while ensuring trust, security, and privacy" [170], without discussing why DLT enhances privacy in the rest of the publication about benchmarking IoT data trading protocols in blockchains. Others followed suit on the use of blockchain to support electric vehicle trading marketplaces with IPFS and a scheme to hide payment sources [171] and cloaking location with k-anonymity [172] or proposing a new blockchain architecture with permissioned domains to enhance privacy for data market places [173]. Another presented several building blocks (blockchain, trusted execution environments, gossip learning) without an evaluation of the proposal [174].

A notable exception is the comprehensive details provided by [175] in their blockchain architecture. Their architecture stores encrypted sensor data in cloud storage, and smart contracts support sensor registration, data auctioning, and payments. While the smart contract emits notifications and displays the endpoint for retrieving proxy re-encrypted data, the data are exchanged off-chain confidentially via proxy re-encryption. This construction addresses transparency and scalability issues regarding sensor data. Nevertheless, bidding and payment processes may still reveal sensitive information and require future research by combining this approach with some of the PETs we surveyed. [176] also acknowledged the aggravation of privacy issues on blockchains and combined federated learning with DP to obfuscate clients' weights and use ZKPs to prove the integrity of the training and evaluation process, which they required for providing fair incentives managed by a smart contract. Another related publication by [177] presented a blockchain-based solution for tracking IoT sensor data across marketplaces and, thus, only detecting but not preventing illegitimate replication and resale.

These publications fall into the category of *architectures* identified in KF4. We also found papers in the *data trading schemata* category (or related): two new data auction schemata enhanced with DP [178, 179], a task assignment scheme in crowdsensing that hides the tasks' content with homomorphic encryption (HE) for crowdsensing [180], and another where they employ DP on billing data [181]. The latter publication, however, does not discuss *fairness*, which is critical in monetary use cases as a noisy bill can make data prosumers profit less from their data on some occasions. Furthermore, [182] and [183] focused on determining fair prices for end users' datasets that are

anonymized with DP according to their accuracy and, correspondingly, risks of revealing sensitive information.

One interesting development that explored the paradigm in ML markets comes from [184]. They developed a privacy-enhanced framework to evaluate the quality of ML models and data for sale with functional encryption, achieving improvements over similar schemata implemented with HE. Another novel concept by [185] strives to empower users with tools to help them determine the risk of sharing their information and, accordingly, make an informed decision about their data framework. It is thus closely related to enforcing privacy policies in the sovereignty layer. Except for [176], who focused narrowly on federated learning, the rest of the new (notable) primary studies did not leverage the combination of anonymization and outsourced computation technologies, underlining KF5. Notably, the new publications have not altered the data market characterization of KF6 or the reference model of KF7.

11. Conclusion

With this review, we reveal the landscape of PETs in data markets for the IoT. We have conducted a systematic literature review (SLR) to identify and filter the studies aiming to solve this landscape's challenges. Consecutively, we formulated terminology to dissect the selected studies' architectures and findings and identified the PETs that related work employed and which specific challenges they addressed.

The authors of the selected studies in this SLR have devised proposals for privacy-enhancing IoT data marketplaces to comply with privacy requirements while maintaining utility, profitability, and fair and seamless data exchange. Since this is a relatively new, multidisciplinary research field, the optimal combination of technologies and theoretical foundations employed in these proposals is still in the development phase. Therefore, no proposal has established itself as canonical yet. Moreover, we observed that the research community needs to further explore the balancing act of utility and privacy before data markets flourish. We conclude that the practicality of PETs needs to advance further to positively impact data markets for the IoT. Additionally, we suggest researchers solve the copy problem and improve privacy-enhancing verification as their absence discourages data markets from forming. We also discovered that research on privacy-oriented data markets could benefit from increased reuse of components from previous articles and existing open-source libraries and a more explicit description of critical objectives. For example, the benefits of utilizing distributed ledger technology (DLT) in data markets for IoT architectures often remain unclear, and authors do not sufficiently consider DLT's lack of maturity and inherent privacy challenges.

The IoT's particular characteristics bring new challenges for privacy enhancement, most notably, the consequences of a lack of interoperability, computation and storage constraints, and the privacy disparity across jurisdictions. We have also observed the importance of *first* determining the sovereignty layer in data market design, as the participants' ownership and management rules impact the PETs in the rest of the layers. We also must underline that there is no "one-size-fits-all" PET. Only a combination may tackle the various privacy challenges facing data markets for the IoT. Lastly, we recommend that institutions invest resources in the research and adoption of PETs to remain competitive in the advent of a more privacy-enhancing IoT.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We would like to thank the Bayerisches Forschungsinstitut für Digitale Transformation for supporting our research on differential privacy, and the Bavarian Ministry of Economic Affairs, Regional Development and Energy for their funding of the project "Fraunhofer Blockchain Center (20-3066-2-6-14)" that made this paper possible, and the BMW Group for generously funding this project.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jnca.2022.103465.

Selected studies

[S1] Y.N. Li, X. Feng, J. Xie, H. Feng, Z. Guan, Q. Wu, A decentralized and secure blockchain platform for open fair data trading, Concurr. Comput. 32 (7) (2019a) 1–11, http://dx.doi.org/10.1002/cpe.5578.

[S2] N. Hynes, D. Dao, D. Yan, R. Cheng, D. Song, A demonstration of sterling: A privacy-preserving data marketplace, Proc. VLDB Endow. 11 (12) (2018) 2086–2089, http://dx.doi.org/10.14778/3229863.3236266.

[S3] D. López, B. Farooq, A multi-layered blockchain framework for smart mobility data-markets, Transp. Res. C 111 (June 2019) (2020) 588–615, http://dx.doi.org/10.1016/j.trc.2020.01.002.

[S4] F. Liang, W. Yu, D. An, Q. Yang, X. Fu, W. Zhao, A survey on big data market: Pricing, trading and protection, IEEE Access 6 (May) (2018) 15132–15154, http://dx.doi.org/10.1109/ACCESS.2018.2806881.

[S5] D. Bogdanov, R. Jagomägis, S. Laur, A universal toolkit for cryptographically secure privacy-preserving data mining, in: LNCS 7299 - Intelligence and Security Informatics, Vol. 7299, 2012, URL https://link.springer.com/content/pdf/10.1007{%}2F978-3-642-30428-6.pdf.

[S6] S. Spiekermann, A. Novotny, A vision for global privacy bridges: Technical and legal measures for international data markets, Comput. Law Secur. Rev. 31 (2) (2015) 181–200, http://dx.doi.org/10.1016/j.clsr.2015.01.009.

[S7] C. Niu, Z. Zheng, F. Wu, X. Gao, G. Chen, Achieving data truthfulness and privacy preservation in data markets, IEEE Trans. Knowl. Data Eng. 31 (1) (2019) 105–119, http://dx.doi.org/10.1109/TKDE.2018.2822727, arXiv:1812.03280.

[S8] E.M. Schomakers, C. Lidynia, M. Ziefle, All of me? Users' preferences for privacy-preserving data markets and the importance of anonymity, Electron. Mark. (2020) http://dx.doi.org/10.1007/s12525-020-00404-9.

[S9] Y. Li, C. Miao, L. Su, J. Gao, Q. Li, B. Ding, Z. Qin, K. Ren, An efficient two-layer mechanism for privacy-preserving truth discovery, in: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2018, pp. 1705–1714, http://dx.doi.org/10.1145/3219819.3219998.

[S10] L. Zhou, L. Wang, T. Ai, Y. Sun, BeeKeeper 2.0: Confidential blockchain-enabled IoT system with fully homomorphic computation, Sensors (Switzerland) 18 (11) (2018a) http://dx.doi.org/10.3390/s18113785.

[S11] A. Dorri, S.S. Kanhere, R. Jurdak, P. Gauravaram, Blockchain for IoT security and privacy: The case study of a smart home, in: 2017 IEEE International Conference on Pervasive Computing and Communications Workshops, PerCom Workshops 2017, IEEE, 2017, pp. 618–623, http://dx.doi.org/10.1109/PERCOMW.2017.7917634.

[S12] R. Li, T. Song, B. Mei, H. Li, X. Cheng, L. Sun, Blockchain for large-scale internet of things data storage and protection, IEEE Trans. Serv. Comput. 12 (5) (2019b) 762–771, http://dx.doi.org/10.1109/TSC.2018.2853167.

[S13] J. Wei, M. Sabonuchi, R. Roche, Blockchain-enabled peer-to-peer data trading mechanism, in: 2018 IEEE International Conference on Internet of Things (IThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData), 2018, pp. 1349–1354, http://dx.doi.org/10.1109/Cybermatics.

[S14] Z. Zheng, W. Mao, F. Wu, G. Chen, Challenges and opportunities in IoT data markets, in: SocialSense 2019 - Proceedings of the 2019 4th International Workshop on Social Sensing, 2019, pp. 1–2, http://dx.doi.org/10.1145/3313294.3313378.

[S15] M.M. Khalili, X. Zhang, M. Liu, Contract design for purchasing private data using a biased differentially private algorithm, in: Proceedings of NetEcon 2019: 14th Workshop on the Economics of Networks, Systems and Computation - in Conjunction with ACM EC 2019 and ACM SIGMETRICS 2019, 2019, http://dx.doi.org/10.1145/3338506.3340273.

[S16] L. Yang, M. Zhang, S. He, M. Li, J. Zhang, Crowd-empowered privacy-preserving data aggregation for mobile crowdsensing, in: Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc), 2018, pp. 151–160, http://dx.doi.org/10.1145/3209582.3209598.

[S17] K. Mišura, M. Žagar, Data marketplace for internet of things, in: Proceedings of 2016 International Conference on Smart Systems and Technologies, SST 2016, IEEE, 2016, pp. 255–260, http://dx.doi.org/10.1109/SST.2016.7765669.
[S18] X. Zheng, Data trading with differential privacy in data market, ACM Int. Conf. Proc. Ser. (8) (2020) 112–115, http://dx.doi.org/10.1145/3379247.3379271.
[S19] J. Pennekamp, M. Henze, S. Schmidt, P. Niemietz, M. Fey, D. Trauth, T. Bergs, C. Brecher, K. Wehrle, Dataflow challenges in an internet of production, in: ACM Workshop on Cyber–Physical Systems Security & Privacy (CPS-SPC'19), November 11, 2019, ACM, London, United Kingdom, 2019, pp. 27–38, http://dx.doi.org/10.1145/3338499.3357357.
[S20] Z.J. Wang, C.H.V. Lin, Y.H. Yuan, C.C.J. Huang, Decentralized data marketplace to enable trusted machine economy, in: 2019 IEEE Eurasia Conference on IOT, Communication and Engineering, ECICE 2019, IEEE, 2019a, pp. 246–250, http://dx.doi.org/10.1109/ECICE47484.2019.8942729.
[S21] M. Guerriero, D.A. Tamburri, E. Di Nitto, Defining, enforcing and checking privacy policies in data-intensive applications, in: Proceedings - International Conference on Software Engineering, 2018, pp. 172–182, http://dx.doi.org/10.1145/3194133.3194140.
[S22] M. Shi, Y. Qiao, X. Wang, Differentially private auctions for private data crowdsourcing, in: Proceedings - 2019 IEEE Intl Conf on Parallel and Distributed Processing with Applications, Big Data and Cloud Computing, Sustainable Computing and Communications, Social Computing and Networking, ISPA/BDCloud/SustainCom/SocialCom 2019, IEEE, 2019, pp. 1–8, http://dx.doi.org/10.1109/ISPA-BDCloud-SustainCom-SocialCom48970.2019.00013.
[S23] J. Du, C. Jiang, E. Gelenbe, L. Xu, J. Li, Y. Ren, Distributed data privacy preservation in IoT applications, IEEE Wirel. Commun. 25 (December) (2018) 68–76, http://dx.doi.org/10.1109/MWC.2017.1800094.
[S24] G. Gao, M. Xiao, J. Wu, S. Zhang, L. Huang, G. Xiao, DPDT: A differentially private crowd-sensed data trading mechanism, IEEE Internet Things J. 7 (1) (2020a) 751–762, http://dx.doi.org/10.1109/JIOT.2019.2944107.
[S25] R. Cheng, F. Zhang, J. Kos, W. He, N. Hynes, N. Johnson, A. Juels, A. Miller, D. Song, Ekiden: A platform for confidentiality-preserving, trustworthy, and performant smart contracts, in: Proceedings - 4th IEEE European Symposium on Security and Privacy, EURO S and P 2019, 2019, pp. 185–200, http://dx.doi.org/10.1109/EuroSP.2019.00023.
[S26] T. Jung, X.Y. Li, Enabling privacy-preserving auctions in big data, in: Proceedings - IEEE INFOCOM, Vol. 2015-Augus (BigSecurity), IEEE, 2015, pp. 173–178, http://dx.doi.org/10.1109/INFCOMW.2015.7179380, arXiv:1308.6202.
[S27] C. Perera, R. Ranjan, L. Wang, End-to-end privacy for open big data markets, IEEE Cloud Comput. 2 (4) (2015b) 44–53, http://dx.doi.org/10.1109/MCC.2015.78.
[S28] M. Zichichi, M. Contu, S. Ferretti, V. Rodríguez-Doncel, Ensuring personal data anonymity in data marketplaces through sensing-as-a-service and distributed ledger technologies, in: CEUR Workshop Proceedings, Vol. 2580, 2020, URL https://www.researchgate.net/publication/340183476_Ensuring_Personal_Data_Anonymity_in_Data_Marketplaces_through_Sensing-as-a-Service_and_Distributed_Ledger.
[S29] S. Duri, M. Gruteser, X. Liu, P. Moskowitz, R. Perez, M. Singh, J.M. Tang, Framework for security and privacy in automotive telematics, in: Proceedings of the ACM International Workshop on Mobile Commerce, 2002, pp. 25–32, http://dx.doi.org/10.1145/570709.570711.
[S30] P. Tzianos, G. Pipelidis, N. Tsiamitros, Hermes: An open and transparent marketplace for iot sensor data over distributed ledgers, in: ICBC 2019 - IEEE International Conference on Blockchain and Cryptocurrency, IEEE, 2019, pp. 167–170, http://dx.doi.org/10.1109/BLOC.2019.8751331.
[S31] K. Li, L. Tian, W. Li, G. Luo, Z. Cai, Incorporating social interaction into three-party game towards privacy protection in iot, Computer Networks 150 (2019) 90–101, http://dx.doi.org/10.1016/j.comnet.2018.11.036.
[S32] Y. Zhao, Y. Yu, Y. Li, G. Han, X. Du, Machine learning based privacy-preserving fair data trading in big data market, Inform. Sci. 478 (2019) 449–460, http://dx.doi.org/10.1016/j.ins.2018.11.028.
[S33] S. Kiyomoto, M.S. Rahman, A. Basu, On blockchain-based anonymized dataset distribution platform, in: 2017 IEEE 15th International Conference on Software Engineering Research, Management and Applications (SERA), IEEE, 2017, pp. 85–92, URL https://ieeexplore.ieee.org/document/7965711.
[S34] Z. Chen, L. Chen, L. Huang, H. Zhong, On privacy-preserving cloud auction, in: Proceedings of the IEEE Symposium on Reliable Distributed Systems, IEEE, 2016, pp. 279–288, http://dx.doi.org/10.1109/SRDS.2016.045.
[S35] L. Pournajaf, D.A. Garcia-Ulloa, L. Xiong, V. Sunderam, Participant privacy in mobile crowd sensing task management, ACM SIGMOD Rec. 44 (4) (2016) 23–34, http://dx.doi.org/10.1145/2935694.2935700.
[S36] D. Sánchez, A. Viejo, Personalized privacy in open data sharing scenarios, Online Inf. Rev. 41 (3) (2017) 298–310, http://dx.doi.org/10.1108/OIR-01-2016-0011.
[S37] K. Jung, S. Park, Privacy bargaining with fairness: Privacy-price negotiation system for applying differential privacy in data market environments, in: 2019 IEEE International Conference on Big Data, IEEE, 2019, pp. 1389–1394, http://dx.doi.org/10.1109/BigData47090.2019.9006101.
[S38] J.H. Ziegeldorf, O.G. Morchon, K. Wehrle, Privacy in the internet of things: Threats and challenges, Secur. Commun. Netw. 7 (12) (2014) 2728–2742, http://dx.doi.org/10.1002/sec.795.
[S39] C. Perera, C. McCormick, A.K. Bandara, B.A. Price, B. Nuseibeh, Privacy-by-design framework for assessing internet of things applications and platforms, ACM Int. Conf. Proc. Ser. 07-09-Nove (2016b) 83–92, http://dx.doi.org/10.1145/2991561.2991566.
[S40] C. Perera, C. Liu, R. Ranjan, L. Wang, A. Zomaya, Privacy-knowledge modeling for the internet of things: A look back, Computer 49 (12) (2016a) 60–68, http://dx.doi.org/10.1109/MC.2016.366.
[S41] W. Gao, W. Yu, F. Liang, W.G. Hatcher, C. Lu, Privacy-preserving auction for big data trading using homomorphic encryption, IEEE Trans. Netw. Sci. Eng. 7 (2) (2020b) 776–791, http://dx.doi.org/10.1109/TNSE.2018.2846736.
[S42] S. Park, K. Park, J. Lee, K. Jung, PRIVATA: Differentially private data market framework using negotiation-based pricing mechanism, in: Proceedings of ACM CIKM Conference (CIKM'19), November 3–7, 2019, Beijing, China, 2019, pp. 156–157, http://dx.doi.org/10.1007/978-3-663-10915-0_47.
[S43] V. Koutsos, D. Papadopoulos, D. Chatzopoulos, S. Tarkoma, P. Hui, Agora: A privacy-aware data marketplace, 2020, p. 13, URL https://eprint.iacr.org/2020/865.pdf.
[S44] J. Cao, P. Karras, Publishing microdata with a robust privacy guarantee, Proc. VLDB Endow. 5 (11) (2012) 1388–1399, http://dx.doi.org/10.14778/2350229.2350255, arXiv:1208.0220.
[S45] W. Dai, C. Dai, K.K.R. Choo, C. Cui, D. Zou, H. Jin, SDTE: A secure blockchain-based data trading ecosystem, IEEE Trans. Inf. Forensics Secur. 15 (2020) 725–737, http://dx.doi.org/10.1109/TIFS.2019.2928256.
[S46] Z. Guan, X. Shao, Z. Wan, Secure, fair and efficient data trading without third party using blockchain, in: 2018 IEEE International Conference on Internet of Things (IThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData), IEEE, 2018, pp. 1349–1354, doi:10.1109/Cybermatics.
[S47] M.A. Alsheikh, Y. Jiao, D. Niyato, P. Wang, D. Leong, Z. Han, The accuracy-privacy trade-off of mobile crowdsensing, IEEE Commun. Mag. 55 (6) (2017) 132–139, http://dx.doi.org/10.1109/MCOM.2017.1600737, arXiv:1702.04565.
[S48] S. Sharma, K. Chen, A. Sheth, Toward practical privacy-preserving analytics for IoT and cloud-based healthcare systems, IEEE Internet Comput. 22 (2) (2018) 42–51, http://dx.doi.org/10.1109/MIC.2018.112102519.
[S49] A. Colman, M.J.M. Chowdhury, M. Baruwal Chhetri, Toward a trusted marketplace for wearable data, in: Proceedings - 2019 IEEE 5th International Conference on Collaboration and Internet Computing, CIC 2019 (Cic), 2019, pp. 314–321, http://dx.doi.org/10.1109/CIC48465.2019.00044.
[S50] Z. Cai, Z. He, Trading private range counting over big IoT data, in: Proceedings - International Conference on Distributed Computing Systems, Vol. 2019-July, IEEE, 2019, pp. 144–153, http://dx.doi.org/10.1109/ICDCS.2019.00023.

References

[1] IDC, Open Evidence, European data market SMART, 2017, URL https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=44400.
[2] A.R. Miller, C. Tucker, Health information exchange, system size and information silos, 2013, p. 29, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1457719.
[3] F. Stahl, F. Schomm, G. Vossen, L. Vomfell, A classification framework for data marketplaces, Vietnam J. Comput. Sci. 3 (3) (2016) 137–143, http://dx.doi.org/10.1007/s40595-016-0064-2.
[4] McKinsey & Company, Four ways to accelerate the creation of data ecosystems, 2020, https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/four-ways-to-accelerate-the-creation-of-data-ecosystems.
[5] G. Eggers, B. Fondermann, B. Maier, K. Ottradovetz, J. Pformmer, R. Reinhardt, H. Rollin, A. Schmieg, S. Steinbuß, P. Trinius, A. Weis, C. Weiss, S. Wilfling, GAIA-X: Technical architecture, 2020, URL https://www.data-infrastructure.eu/GAIAX/Redaktion/EN/Publications/gaia-x-technical-architecture.pdf?_blob=publicationFile&v=5.
[6] L. Sweeney, A. Abu, J. Winn, Identifying participants in the personal genome project by name, SSRN Electron. J. (2013) http://dx.doi.org/10.2139/ssrn.2257732.
[7] X. Gao, B. Firner, S. Sugrim, V. Kaiser-Pendergrast, Y. Yang, J. Lindqvist, Elastic pathing: your speed is enough to track you, in: Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '14 Adjunct, ACM Press, Seattle, Washington, 2014, pp. 975–986, http://dx.doi.org/10.1145/2632048.2632077.
[8] A. Sunyaev, N. Kannengießer, R. Beck, H. Treiblmaier, M. Lacity, J. Kranz, G. Fridgen, U. Spankowski, A. Luckow, Token economy, Bus. Inf. Syst. Eng. (2021) URL https://link.springer.com/article/10.1007/s12599-021-00684-1.
[9] IBM Security and Ponemon Institute LLC, 2018 Cost of a data breach study: Global overview, 2018, p. 47, URL https://www.intlxsolutions.com/hubfs/2018_Global_Cost_of_a_Data_Breach_Report.pdf.
[10] A. Trask, E. Bluemke, B. Garfinkel, C.G. Cuervas-Mons, A. Dafoe, Beyond privacy trade-offs with structured transparency, 2020, arXiv:2012.08347. URL https://www.researchgate.net/publication/347300876_Beyond_Privacy_Trade-offs_with_Structured_Transparency.

[11] R. Hes, J.J. Borking, Netherlands, I.a.P. Commissioner/Ontario (Eds.), Privacy-enhancing Technologies: the Path to Anonymity, rev. ed., in: Achtergrondstudies en Verkenningen, (11) Registratiekamer, The Hague, 1998, URL https://www.researchgate.net/publication/243777645_Privacy-Enhancing_Technologies_The_Path_to_Anonymity.

[12] C. Dwork, F. McSherry, K. Nissim, A. Smith, Calibrating noise to sensitivity in private data analysis, in: S. Halevi, T. Rabin (Eds.), Theory of Cryptography, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 265–284, URL https://link.springer.com/chapter/10.1007/11681878_14. Online; accessed 30 December 2021.

[13] C. Dwork, A. Roth, The algorithmic foundations of differential privacy, Found. Trends® Theor. Comput. Sci. 9 (3–4) (2013) 211–407, http://dx.doi.org/10.1561/0400000042.

[14] P. Samarati, L. Sweeney, 1998. Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. 19. URL https://epic.org/privacy/reidentification/Samarati_Sweeney_paper.pdf.

[15] M.A. Will, R.K. Ko, A Guide to Homomorphic Encryption, Elsevier Inc., 2015, p. 101, http://dx.doi.org/10.1016/B978-0-12-801595-7.00005-7.

[16] P. Chaudhary, R. Gupta, A. Singh, P. Majumder, Analysis and comparison of various fully homomorphic encryption techniques, in: 2019 International Conference on Computing, Power and Communication Technologies, GUCON 2019, Galgotias University, 2019, pp. 58–62, URL https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8940577.

[17] P. Paillier, Public-key cryptosystems based on composite degree residuosity classes, Eurocrypt (1999) http://dx.doi.org/10.1007/3-540-48910-X_9.

[18] OMTP, Advanced trusted environment: OMTP TR1, 2009, p. 204, URL http://www.gsma.com/newsroom/wp-content/uploads/2012/03/omtpadvancedtrustedenvironmentomtptr1v11.pdf.

[19] A.C. Yao, Protocols for secure computations, in: 23rd Annual Symposium on Foundations of Computer Science (Sfcs 1982), IEEE, Chicago, IL, USA, 1982, pp. 160–164, http://dx.doi.org/10.1109/SFCS.1982.38.

[20] S. Goldwasser, S. Micali, C. Rackoff, The knowledge complexity of interactive proof systems, SIAM J. Comput. 18 (1) (1989) 186–208, http://dx.doi.org/10.1137/0218012.

[21] O. Goldreich, Y. Oren, Definitions and properties of zero-knowledge proof systems, J. Cryptol. 7 (1) (1994) 1–32, http://dx.doi.org/10.1007/BF00195207.

[22] G. Bondel, G.M. Garrido, K. Baumer, F. Matthes, 2020. Towards a privacy-enhancing tool based on de- identification methods. 8. URL https://aisel.aisnet.org/pacis2020/157/.

[23] S. Spiekermann, R. Böhme, A. Acquisti, K.-L. Hui, Personal data markets, Electron. Mark. 25 (2) (2015) 91–93, http://dx.doi.org/10.1007/s12525-015-0190-1.

[24] P.B. Anne Zöll, C.M. Olt, Privacy-sensitive business models: Barriers of organizational adoption of privacy-enhancing technologies, 2021, p. 22, URL https://aisel.aisnet.org/ecis2021_rp/34/.

[25] A.F. Westin, Privacy and Freedom, IG Publishing, New York, 1967, URL https://scholarlycommons.law.wlu.edu/wlulr/vol25/iss1/20/.

[26] G.A. Fink, H. Song, S. Jeschke (Eds.), Security and Privacy in Cyber-Physical Systems: Foundations, Principles, and Applications, first ed., Wiley IEEE Press, Hoboken, NJ, 2018, URL https://ieeexplore.ieee.org/servlet/opac?bknumber=8068866.

[27] K. Renaud, D. Galvez-Cruz, Privacy: Aspects, definitions and a multi-faceted privacy preservation approach, in: Proceedings of the 2010 Information Security for South Africa Conference, ISSA 2010, 2010, pp. 1–8, http://dx.doi.org/10.1109/ISSA.2010.5588297.

[28] D.J. Solove, The meaning and value of privacy, in: B. Roessler, D. Mokrosinska (Eds.), Social Dimensions of Privacy, Cambridge University Press, Cambridge, 2015, pp. 71–82, http://dx.doi.org/10.1017/CBO9781107280557.005.

[29] F.T. Wu, Defining privacy and utility in data sets, in: 84 University of Colorado Law Review 1117 (2013); 2012 TRPC, 2012, pp. 1117–1177, http://dx.doi.org/10.2139/ssrn.2031808.

[30] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, W. Joosen, A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements, Requir. Eng. 16 (1) (2011) 3–32, http://dx.doi.org/10.1007/s00766-010-0115-7.

[31] R. Garratt, M.R.v. Oordt, Privacy as a public good: A case for electronic cash, J. Polit. Econ. (2018) http://dx.doi.org/10.1086/714133.

[32] N. Kaaniche, M. Laurent, Attribute-based signatures for supporting anonymous certification, in: I. Askoxylakis, S. Ioannidis, S. Katsikas, C. Meadows (Eds.), Computer Security – ESORICS 2016, Springer International Publishing, Cham, 2016, pp. 279–300, URL https://www.semanticscholar.org/paper/Attribute-Based-Signatures-for-Supporting-Anonymous-Kaaniche-Laurent-Maknavicius/3b0624ff32b9258ca2351c894d320d83a546fcd6.

[33] J.E. Campbell, M. Carlson, Panopticon.com: Online surveillance and the commodification of privacy, J. Broadcast. Electron. Media 46 (4) (2002) 586–606, http://dx.doi.org/10.1207/s15506878jobem4604_6.

[34] A. Lichter, M. Löffler, S. Siegloch, The long-term costs of government surveillance: Insights from stasi spying in east Germany, J. Eur. Econom. Assoc. 19 (2) (2020) 741–789, http://dx.doi.org/10.1093/jeea/jvaa009, arXiv:https://academic.oup.com/jeea/article-pdf/19/2/741/37108669/jvaa009.pdf.

[35] S. Kokolakis, Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon, Comput. Secur. 64 (2017) 122–134, http://dx.doi.org/10.1016/j.cose.2015.07.002.

[36] J. Coppel, E-Commerce: Impacts and Policy Challenges, OECD Economics Department Working Papers 252, 2000, http://dx.doi.org/10.1787/801315684632, Series: OECD Economics Department Working Papers Volume: 252.

[37] J. Kennedy, Big data's economic impact, 2021, [Online]. Available: https://www.ced.org/blog/entry/big-datas-economic-impact, [Accessed on 04 Jul. 2021].

[38] A.M. Oberländer, M. Röglinger, M. Rosemann, A. Kees, Conceptualizing business-to-thing interactions – a sociomaterial perspective on the internet of things, Eur. J. Inf. Syst. 27 (4) (2018) 486–502, http://dx.doi.org/10.1080/0960085X.2017.1387714.

[39] I. Lee, K. Lee, The internet of things (IoT): Applications, investments, and challenges for enterprises, Bus. Horiz. 58 (4) (2015) 431–440, http://dx.doi.org/10.1016/j.bushor.2015.03.008.

[40] V. Basili, G. Caldiera, D. Rombach, The goal question metric approach, Encycl. Softw. Eng. (1994) 528–532, URL http://www.cs.toronto.edu/~sme/CSC444F/handouts/GQM-paper.pdf.

[41] B.A. Kitchenham, D. Budgen, Evidence-Based Software Engineering and Systematic Reviews, Chapman and Hall/CRC, 2015, URL https://dl.acm.org/doi/book/10.5555/2994449.

[42] B. Kitchenham, Procedures for Performing Systematic Reviews, Joint Technical Report, 2004, http://dx.doi.org/10.5144/0256-4947.2017.79.

[43] D.C.B. Mariano, C. Leite, L.H.S. Santos, R.E.O. Rocha, R.C. de Melo-Minardi, A guide to performing systematic literature reviews in bioinformatics, 2017, arXiv:1707.05813.

[44] T. Dybå, T. Dingsøyr, G. Hanssen, Applying systematic reviews to diverse study types: An experience report, in: Proceedings - 1st International Symposium on Empirical Software Engineering and Measurement, ESEM 2007 (7465), 2007, pp. 126–135, http://dx.doi.org/10.1109/ESEM.2007.59.

[45] O. Dieste, A. Grimán, N. Juristo, Developing search strategies for detecting relevant experiments, Empir. Softw. Eng. 14 (5) (2009) 513–539, http://dx.doi.org/10.1007/s10664-008-9091-7.

[46] H. Zhang, M.A. Babar, P. Tell, Identifying relevant studies in software engineering, Inf. Softw. Technol. 53 (6) (2011) 625–637, http://dx.doi.org/10.1016/j.infsof.2010.12.010.

[47] A. Kilgarriff, V. Baisa, J. Bušta, M. Jakubíček, V. Kovář, J. Michelfeit, P. Rychlý, V. Suchomel, The sketch engine: ten years on, Lexicography (2014) URL https://www.researchgate.net/publication/271848017_The_Sketch_Engine_Ten_Years_On.

[48] O.P. Brereton, B.A. Kitchenham, D. Budgen, M. Turner, M. Khalil, Lessons from applying the systematic literature review process within the software engineering domain, J. Syst. Softw. 80 (4) (2007) 571–583, http://dx.doi.org/10.1016/j.jss.2006.07.009.

[49] B.A. Kitchenham, O.P. Brereton, A systematic review of systematic review process research in software engineering, Inf. Softw. Technol. 55 (12) (2013) 2049–2075, http://dx.doi.org/10.1016/j.infsof.2013.07.010.

[50] L. Chen, M.A. Babar, H. Zhang, Towards an evidence-based understanding of electronic data sources (january 2015), 2010, http://dx.doi.org/10.14236/ewic/ease2010.17.

[51] C. Wohlin, P. Runeson, M. Höst, M.C. Ohlsson, B. Regnell, A. Wesslén, Experimentation in Software Engineering, Springer Science & Business Media, 2012, http://dx.doi.org/10.1007/978-3-642-29044-2.

[52] H. Nissenbaum, Privacy in Context: Technology, Policy, and the Integrity of Social Life, Stanford University Press, 2009, URL https://www.sup.org/books/title/?id=8862.

[53] R.Y. Wang, D.M. Strong, Beyond accuracy: What data quality means to data consumers, J. Manage. Inf. Syst. 12 (4) (1996) 5–33, http://dx.doi.org/10.1080/07421222.1996.11518099.

[54] T. Dinev, H. Xu, J.H. Smith, P. Hart, Information privacy and correlates: an empirical attempt to bridge and distinguish privacy-related concepts, Eur. J. Inf. Syst. 22 (3) (2013) 295–316, http://dx.doi.org/10.1057/ejis.2012.23.

[55] B.-J. Butijn, D.A. Tamburri, W.-J.v.d. Heuvel, Blockchains: a systematic multivocal literature review, ACM Comput. Surv. 53 (3) (2020) 1–37, URL https://dl.acm.org/doi/abs/10.1145/3369052.

[56] R. Zhang, R. Xue, L. Liu, Security and privacy on blockchain, ACM Comput. Surv. 52 (3) (2019) 1–34, URL https://dl.acm.org/doi/10.1145/3316481.

[57] G.I. Simari, A primer on zero knowledge protocols, 2002, p. 12, URL http://cs.uns.edu.ar/~gis/publications/zkp-simari2002.pdf.

[58] N. Kaaniche, M. Laurent, S. Belguith, Privacy enhancing technologies for solving the privacy-personalization paradox: Taxonomy and survey, J. Netw. Comput. Appl. (2020).

[59] D. Chaum, Security without identification: transaction systems to make big brother obsolete, Commun. ACM 28 (10) (1985) 1030–1044, http://dx.doi.org/10.1145/4372.4373.

[60] J.L. Camenisch, J.-M. Piveteau, M.A. Stadler, Blind signatures based on the discrete logarithm problem, in: A. De Santis (Ed.), Advances in Cryptology — EUROCRYPT'94, Springer Berlin Heidelberg, Berlin, Heidelberg, 1995, pp. 428–432, URL https://link.springer.com/chapter/10.1007/BFb0053458.

[61] J. Camenisch, A. Lysyanskaya, Dynamic accumulators and application to efficient revocation of anonymous credentials, in: G. Goos, J. Hartmanis, J. van Leeuwen, M. Yung (Eds.), Advances in Cryptology — CRYPTO 2002, Vol. 2442, in: Lecture Notes in Computer Science, Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 61–76, http://dx.doi.org/10.1007/3-540-45708-9_5.

[62] J. Camenisch, T. Groß, Efficient attributes for anonymous credentials, 2010, p. 29, URL https://eprint.iacr.org/2010/496.pdf.

[63] S.A. Brands, Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy, MIT Press, Cambridge, MA, USA, 2000, URL https://direct.mit.edu/books/book/1912/Rethinking-Public-Key-Infrastructures-and-Digital.

[64] J. Sedlmeir, R. Smethurst, A. Rieger, G. Fridgen, Digital identities and verifiable credentials, Bus. Inf. Syst. Eng. 63 (5) (2021) 603–613.

[65] V. Schlatt, J. Sedlmeir, S. Feulner, N. Urbach, Designing a framework for digital KYC processes built on blockchain-based self-sovereign identity, Inf. Manage. (2021) 103553.

[66] E. Bangerter, S. Barzan, S. Krenn, A.-R. Sadeghi, T. Schneider, J.-K. Tsay, Bringing zero-knowledge proofs of knowledge to practice, 2009, p. 12, URL https://eprint.iacr.org/2009/211.pdf.

[67] M. Hoffmann, M. Klooß, A. Rupp, Efficient zero-knowledge arguments in the discrete log setting, revisited, in: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, ACM, London United Kingdom, 2019, pp. 2093–2110, http://dx.doi.org/10.1145/3319535.3354251.

[68] T. Nakanishi, H. Yoshino, T. Murakami, G.-V. Policharla, Efficient zero-knowledge proofs of graph signature for connectivity and isolation using bilinear-map accumulator, in: Proceedings of the 7th ACM Workshop on ASIA Public-Key Cryptography, ACM, Taipei Taiwan, 2020, pp. 9–18, http://dx.doi.org/10.1145/3384940.3388959.

[69] Y. Zhang, Zero-knowledge proofs for machine learning, in: Proceedings of the 2020 Workshop on Privacy-Preserving Machine Learning in Practice, Association for Computing Machinery, 2020, p. 7, URL https://doi.org/10.1145/3411501.3418608.

[70] M. Stadler, Publicly verifiable secret sharing, in: Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), Vol. 1070, 1996, pp. 190–199, http://dx.doi.org/10.1007/3-540-68339-9_17.

[71] A. Shamir, How to share a secret, Commun. ACM 22 (11) (1979a) 612–613, http://dx.doi.org/10.1145/359168.359176.

[72] Y. Lindell, B. Pinkas, A proof of security of Yao's protocol for two-party computation, J. Cryptol. 22 (2) (2009) 161–188, http://dx.doi.org/10.1007/s00145-008-9036-8.

[73] A. Ben-David, N. Nisan, B. Pinkas, Fairplaymp: a system for secure multi-party computation, in: Proceedings of the 15th ACM Conference on Computer and Communications Security - CCS '08, ACM Press, Alexandria, Virginia, USA, 2008, p. 257, http://dx.doi.org/10.1145/1455770.1455804.

[74] S. Yakoubov, A gentle introduction to Yao' s garbled circuits, 2017, URL https://web.mit.edu/sonka89/www/papers/2017ygc.pdf.

[75] Z.A. Genç, V. Iovino, A. Rial, The simplest protocol for oblivious transfer, Inform. Process. Lett. 161 (2020) 1–12, http://dx.doi.org/10.1016/j.ipl.2020.105975.

[76] P. Pullonen, S. Siim, Combining secret sharing and garbled circuits for efficient private IEEE 754 floating-point computations, in: M. Brenner, N. Christin, B. Johnson, K. Rohloff (Eds.), Financial Cryptography and Data Security, Vol. vol. 8976, in: Lecture Notes in Computer Science, Springer Berlin Heidelberg, Berlin, Heidelberg, 2015, pp. 172–183, http://dx.doi.org/10.1007/978-3-662-48051-9_13.

[77] Y. Yang, X. Huang, X. Liu, H. Cheng, J. Weng, X. Luo, V. Chang, A comprehensive survey on secure outsourced computation and its applications, IEEE Access 7 (2019a) 159426–159465, URL https://ieeexplore.ieee.org/document/8884162/.

[78] E. Boyle, N. Gilboa, Y. Ishai, A. Nof, Practical fully secure three-party computation via sublinear distributed zero-knowledge proofs, in: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, ACM, London United Kingdom, 2019, pp. 869–886, URL https://dl.acm.org/doi/10.1145/3319535.3363227.

[79] D. Boneh, E.-J. Goh, K. Nissim, Evaluating 2-DNF formulas on ciphertexts, 3378, 2005, pp. 325–341, URL https://www.researchgate.net/publication/221354138_Evaluating_2-DNF_Formulas_on_Ciphertexts,

[80] V. Nikolaenko, S. Ioannidis, U. Weinsberg, M. Joye, N. Taft, D. Boneh, Privacy-preserving matrix factorization, in: Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security, Association for Computing Machinery, 2013, pp. 801–812, URL https://doi.org/10.1145/2508859.2516751.

[81] L. Zhou, L. Wang, Y. Sun, T. Ai, AntNest: Fully non-interactive secure multi-party computation, IEEE Access 6 (2018b) 75639–75649, URL https://ieeexplore.ieee.org/document/8550709/.

[82] D. Boneh, A. Sahai, B. Waters, Functional encryption: Definitions and challenges, in: Y. Ishai (Ed.), Theory of Cryptography, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 253–273, URL https://eprint.iacr.org/2010/543.pdf.

[83] J. Chotard, E. Dufour Sans, R. Gay, D.H. Phan, D. Pointcheval, Decentralized multi-client functional encryption for inner product, in: T. Peyrin, S. Galbraith (Eds.), Advances in Cryptology – ASIACRYPT 2018, Springer International Publishing, Cham, 2018, pp. 703–732, URL https://eprint.iacr.org/2017/989.pdf.

[84] Z. Brakerski, C. Gentry, V. Vaikuntanathan, (Leveled) fully homomorphic encryption without bootstrapping, ACM Trans. Comput. Theory 6 (3) (2014) http://dx.doi.org/10.1145/2633600.

[85] W. Wang, Y. Hu, L. Chen, X. Huang, B. Sunar, Exploring the feasibility of fully homomorphic encryption, IEEE Trans. Comput. 64 (3) (2015) 698–706, http://dx.doi.org/10.1109/TC.2013.154.

[86] I. Anati, S. Gueron, S.P. Johnson, V.R. Scarlata, Innovative technology for CPU based attestation and sealing, 2013, p. 7, URL https://software.intel.com/content/dam/develop/external/us/en/documents/hasp-2013-innovative-technology-for-attestation-and-sealing-413939.pdf.

[87] F. Khalid, A. Masood, Vulnerability analysis of qualcomm secure execution environment, Comput. Secur. 116 (2022) 102628, http://dx.doi.org/10.1016/j.cose.2022.102628, URL https://www.sciencedirect.com/science/article/pii/S016740482200027X.

[88] F. Alder, J. Van Bulck, J. Spielman, D. Oswald, F. Piessens, Faulty point unit: ABI poisoning attacks on trusted execution environments, Digit. Threats Res. Pract. 3 (2) (2022) http://dx.doi.org/10.1145/3491264, URL https://doi-org.eaccess.ub.tum.de/10.1145/3491264.

[89] D. Skarlatos, M. Yan, B. Gopireddy, R. Sprabery, J. Torrellas, C.W. Fletcher, MicroScope: Enabling microarchitectural replay attacks, in: Proceedings of the 46th International Symposium on Computer Architecture, ACM, 2019, pp. 318–331, http://dx.doi.org/10.1145/3307650.3322228.

[90] Intel, 12Th generation intel core processors, 2022, URL https://cdrdv2.intel.com/v1/dl/getContent/655258.

[91] V. Costan, I. Lebedev, S. Devadas, Sanctum: Minimal hardware extensions for strong software isolation, 2016, p. 19, URL https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/costan.

[92] D. Lee, D. Kohlbrenner, S. Shinde, K. Asanović, D. Song, Keystone: an open framework for architecting trusted execution environments, in: Proceedings of the Fifteenth European Conference on Computer Systems, ACM, Heraklion Greece, 2020, pp. 1–16, http://dx.doi.org/10.1145/3342195.3387532.

[93] AWS, Nitro system, 2022, Online; accessed 4 May 2022. URL https://aws.amazon.com/ec2/nitro/.

[94] W. Wei, L. Liu, M. Loper, K.-H. Chow, M.E. Gursoy, S. Truex, Y. Wu, A framework for evaluating gradient leakage attacks in federated learning, 2020, arXiv:2004.10397. URL https://www.semanticscholar.org/paper/A-Framework-for-Evaluating-Gradient-Leakage-Attacks-Wei-Liu/9853a348f61aec83b410f307ab905a4ae001fcd4.

[95] S. Pinto, N. Santos, Demystifying ARM TrustZone: A comprehensive survey, ACM Comput. Surv. 51 (6) (2019).

[96] V. Costan, L.F. Sarmenta, M.v. Dijk, S. Devadas, The trusted execution module: Commodity general-purpose trusted computing, in: International Conference on Smart Card Research and Advanced Applications, Springer, 2008, pp. 133–148.

[97] J. Konečný, B. McMahan, D. Ramage, Federated optimization:Distributed optimization beyond the datacenter, 2015, arXiv:1511.03575. URL https://docplayer.net/15450695-Federated-optimization-distributed-optimization-beyond-the-datacenter.html.

[98] T. Li, A.K. Sahu, A. Talwalkar, V. Smith, Federated learning: Challenges, methods, and future directions, IEEE Signal Process. Mag. 37 (3) (2020) 50–60, http://dx.doi.org/10.1109/MSP.2020.2975749.

[99] Q. Yang, Y. Liu, Y. Cheng, Y. Kang, T. Chen, H. Yu, Federated learning, Synth. Lect. Artif. Intell. Mach. Learn. 13 (3) (2019b) 1–207, http://dx.doi.org/10.2200/S00960ED2V01Y201910AIM043.

[100] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H.B. McMahan, S. Patel, D. Ramage, A. Segal, K. Seth, Practical secure aggregation for privacy-preserving machine learning, in: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, in: CCS '17, Association for Computing Machinery, New York, NY, USA, 2017, pp. 1175–1191, http://dx.doi.org/10.1145/3133956.3133982.

[101] S. Zhang, Z. Li, Q. Chen, W. Zheng, J. Leng, M. Guo, Dubhe: Towards data unbiasedness with homomorphic encryption in federated learning client selection, in: 50th International Conference on Parallel Processing, in: ICPP 2021, Association for Computing Machinery, New York, NY, USA, 2021a, http://dx.doi.org/10.1145/3472456.3473513, URL https://doi-org.eaccess.ub.tum.de/10.1145/3472456.3473513.

[102] W. Yang, B. Liu, C. Lu, N. Yu, Privacy Preserving on Updated Parameters in Federated Learning, ACM TURC'20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 27–31, http://dx.doi.org/10.1145/3393527.3393533.

[103] P. Vepakomma, O. Gupta, T. Swedish, R. Raskar, Split learning for health: Distributed deep learning without sharing raw patient data, 2018, URL https://aiforsocialgood.github.io/iclr2019/accepted/track1/pdfs/31_aisg_iclr2019.pdf.

[104] O. Gupta, R. Raskar, Distributed learning of deep neural network over multiple agents, J. Netw. Comput. Appl. 116 (2018) 1–8, http://dx.doi.org/10.1016/j.jnca.2018.05.003.

[105] M.G. Poirot, P. Vepakomma, K. Chang, J. Kalpathy-Cramer, R. Gupta, R. Raskar, 2019. Split learning for collaborative deep learning in healthcare. 9. URL https://arxiv.org/abs/1912.12115.
[106] L. Giaretta, S. Girdzijauskas, Gossip learning: Off the beaten path, in: 2019 IEEE International Conference on Big Data (Big Data), 2019, pp. 1117–1124, http://dx.doi.org/10.1109/BigData47090.2019.9006216.
[107] R. Ormándi, I. Hegedűs, M. Jelasity, 2012. Gossip learning with linear models on fully distributed data: Efficient p2p ensemble learning with linear models on fully distributed data. 25 (4), 556–571. http://dx.doi.org/10.1002/cpe.2858. URL https://onlinelibrary.wiley.com/doi/10.1002/cpe.2858.
[108] M. Abadi, A. Chu, I. Goodfellow, H.B. McMahan, I. Mironov, K. Talwar, L. Zhang, Deep learning with differential privacy, in: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, ACM, 2016, http://dx.doi.org/10.1145/2976749.2978318.
[109] J. Domingo-Ferrer, D. Sánchez, A. Blanco-Justicia, The limits of differential privacy (and its misuse in data release and machine learning), Commun. ACM 64 (7) (2021) 33–35, http://dx.doi.org/10.1145/3433638, URL https://doi-org.eaccess.ub.tum.de/10.1145/3433638.
[110] E. Dikici, L.M. Prevedello, M. Bigelow, R.D. White, B.S. Erdal, Constrained generative adversarial network ensembles for sharable synthetic data generation, 2020, arXiv:2003.00086. URL https://www.researchgate.net/publication/339642358_Constrained_Generative_Adversarial_Network_Ensembles_for_Sharable_Synthetic_Data_Generation.
[111] A. Torfi, E.A. Fox, C.K. Reddy, Differentially private synthetic medical data generation using convolutional GANs, 2020, URL https://www.researchgate.net/publication/347624671_Differentially_Private_Synthetic_Medical_Data_Generation_using_Convolutional_GANs.
[112] 2010. Nin, J., Herranz, J. (Eds.), Privacy and Anonymity in Information Management Systems. In: Advanced Information and Knowledge Processing, Springer, London. http://dx.doi.org/10.1007/978-1-84996-238-4. URL http://link.springer.com/10.1007/978-1-84996-238-4.
[113] V. Puri, S. Sachdeva, P. Kaur, Privacy preserving publication of relational and transaction data: Survey on the anonymization of patient data, Comput. Sci. Rev. 32 (C) (2019) 45–61, http://dx.doi.org/10.1016/j.cosrev.2019.02.001.
[114] B.C.M. Fung, K. Wang, R. Chen, P.S. Yu, Privacy-preserving data publishing: A survey of recent developments, ACM Comput. Surv. 42 (4) (2010) http://dx.doi.org/10.1145/1749603.1749605.
[115] C.C. Aggarwal, P.S. Yu (Eds.), Privacy-Preserving Data Mining - Models and Algorithms, in: Advances in Database Systems, vol. 34, Springer, 2008, http://dx.doi.org/10.1007/978-0-387-70992-5.
[116] P. Ram Mohan Rao, S. Murali Krishna, A.P. Siva Kumar, 2018. Privacy preservation techniques in big data analytics: a survey. 5 (1), 33. http://dx.doi.org/10.1186/s40537-018-0141-8. URL https://journalofbigdata.springeropen.com/articles/10.1186/s40537-018-0141-8.
[117] M. Cunha, R. Mendes, J. ao P. Vilela, A survey of privacy-preserving mechanisms for heterogeneous data types, Comp. Sci. Rev. 41 (2021) 100403, http://dx.doi.org/10.1016/j.cosrev.2021.100403, URL https://www.sciencedirect.com/science/article/pii/S1574013721000435.
[118] C. Dwork, A. Smith, T. Steinke, J. Ullman, Exposed! A survey of attacks on private data, Annu. Rev. Stat. Appl. 4 (1) (2017) 61–84, http://dx.doi.org/10.1146/annurev-statistics-060116-054123.
[119] ISO, Privacy enhancing data de-identification terminology and classification of techniques, 2018, URL https://www.iso.org/standard/69373.html.
[120] N. Li, W. Qardaji, D. Su, On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy, in: Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security, in: ASIACCS '12, Association for Computing Machinery, New York, NY, USA, 2012, pp. 32–33, http://dx.doi.org/10.1145/2414456.2414474.
[121] H. Xu, S. Guo, K. Chen, Building confidential and efficient query services in the cloud with RASP data perturbation, 2013, http://dx.doi.org/10.1109/TKDE.2012.251.
[122] K. Chen, L. Liu, Geometric data perturbation for privacy preserving outsourced data mining, Knowl. Inf. Syst. 29 (3) (2011) 657–695, URL http://link.springer.com/10.1007/s10115-010-0362-4.
[123] R. Henry, A. Herzberg, A. Kate, Blockchain access privacy: Challenges and directions, IEEE Secur. Priv. 16 (4) (2018) 38–45, http://dx.doi.org/10.1109/MSP.2018.3111245.
[124] D. Chaum, Untraceable electronic mail, return addresses, and digital pseudonyms, Commun. ACM 24 (2) (1981) 84–90.
[125] D. Chaum, The dining cryptographers problem: Unconditional sender and recipient untraceability, J. Cryptol. 1 (1) (1988) 65–75.
[126] J. Ren, J. Wu, Survey on anonymous communications in computer networks, Comput. Commun. 33 (4) (2010) 420–431.
[127] M.S. Ali, K. Dolui, F. Antonelli, Iot data privacy via blockchains and IPFS, in: Proceedings of the Seventh International Conference on the Internet of Things, in: IoT '17, Association for Computing Machinery, New York, NY, USA, 2017, http://dx.doi.org/10.1145/3131542.3131563.
[128] M. Kesarwani, A. Kaul, S. Braghin, N. Holohan, S. Antonatos, Secure k-anonymization over encrypted databases, in: 2021 IEEE 14th International Conference on Cloud Computing (CLOUD), 2021, pp. 20–30, http://dx.doi.org/10.1109/CLOUD53861.2021.00015.
[129] A.F. Westin, Privacy and freedom, 1970, URL https://www.worldcat.org/title/privacy-and-freedom/oclc/792862.
[130] A. Raj, J. Bosch, H.H. Olsson, T.J. Wang, Modelling data pipelines, in: 2020 46th Euromicro Conference on Software Engineering and Advanced Applications (SEAA), 2020, pp. 13–20, http://dx.doi.org/10.1109/SEAA51224.2020.00014.
[131] A. Cavoukian, The 7 foundational principles, 2011, p. 2, URL https://sites.psu.edu/digitalshred/2020/11/13/privacy-by-design-pbd-the-7-foundational-principles-cavoukian/.
[132] J. Sedlmeir, J. Lautenschlager, G. Fridgen, N. Urbach, The transparency challenge of blockchain in organizations, Electron. Mark. (2022) http://dx.doi.org/10.1007/s12525-022-00536-0.
[133] E. Heilman, N. Narula, G. Tanzer, J. Lovejoy, M. Colavita, M. Virza, T. Dryja, Cryptanalysis of curl-p and other attacks on the IOTA cryptocurrency, IACR Trans. Symm. Crypt. (2020) 367–391, http://dx.doi.org/10.46586/tosc.v2020.i3.367-391.
[134] D. Wang, J. Zhao, Y. Wang, A survey on privacy protection of blockchain: The technology and application, IEEE Access 8 (2020) 108766–108781, http://dx.doi.org/10.1109/ACCESS.2020.2994294, URL https://ieeexplore.ieee.org/document/9093015/.
[135] J. Sedlmeir, H.U. Buhl, G. Fridgen, R. Keller, The energy consumption of blockchain technology: beyond myth, Bus. Inf. Syst. Eng. 62 (6) (2020) 599–608, URL https://www.researchgate.net/publication/342313238_The_Energy_Consumption_of_Blockchain_Technology_Beyond_Myth.
[136] N. Kannengießer, S. Lins, T. Dehling, A. Sunyaev, Trade-offs between distributed ledger technology characteristics, ACM Comput. Surv. 53 (2) (2020) 1–37, URL https://dl.acm.org/doi/10.1145/3379463.
[137] European Parliament, European Parliament and Council of the European Union, Regulation (EU) 2016/679 directive 95/46/EC (general data protection regulation): General data protection regulation, Off. J. Eur. Union (L119) (2016) 1–88, URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679.
[138] S. Ellis, A. Juels, S. Nazarov, Chainlink a decentralized oracle network, 2021, URL https://link.smartcontract.com/whitepaper.
[139] S.S. Al-Riyami, K.G. Paterson, Certificateless public key cryptography, in: Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), Vol. 2894, 2003, pp. 452–473, URL https://eprint.iacr.org/2003/126.pdf.
[140] D. Reed, M. Sporny, D. Longley, C. Allen, A. Grant, M. Sabadello, Decentralized identifiers (dids) v1.0, 2021, URL https://w3c.github.io/did-core/.
[141] L. Mui, M. Mohtashemi, A. Halberstadt, A computational model of trust and reputation, in: Proceedings of the 35th Hawaii International Conference on System Sciences, IEEE, 2002, pp. 2431–2439.
[142] T. Grandison, M. Sloman, A survey of trust in internet applications, IEEE Commun. Surv. Tutor. 3 (4) (2000) 2–16.
[143] D. Artz, Y. Gil, A survey of trust in computer science and the semantic web, J. Web Semant. 5 (2) (2007) 58–71.
[144] K.S. Cook, R. Hardin, M. Levi, Cooperation Without Trust? Russell Sage Foundation, 2005.
[145] L. Ismail, H. Hameed, M. AlShamsi, M. AlHammadi, N. AlDhanhani, Towards a blockchain deployment at UAE university: Performance evaluation and blockchain taxonomy, in: Proceedings of the 2019 International Conference on Blockchain Technology, ACM, Honolulu HI USA, 2019, pp. 30–38, http://dx.doi.org/10.1145/3320154.3320156.
[146] L. Lamport, R. Shostak, M. Pease, The Byzantine generals problem, in: Concurrency: The Works of Leslie Lamport, Association for Computing Machinery, New York, NY, USA, 2019, pp. 203–226, http://dx.doi.org/10.1145/3335772.3335936.
[147] C. Perera, C.H. Liu, S. Jayawardena, The emerging internet of things marketplace from an industrial perspective: A survey, IEEE Trans. Emerg. Top. Comput. 3 (4) (2015a) 585–598, URL https://ieeexplore.ieee.org/document/7004800.
[148] L. Determann, No one owns data, UC Hast. Law 70 (2018) 44, http://dx.doi.org/10.2139/ssrn.3123957.
[149] J.H. Nord, A. Koohang, J. Paliszkiewicz, The internet of things: Review and theoretical framework, Expert Syst. Appl. 133 (2019) 97–108, http://dx.doi.org/10.1016/j.eswa.2019.05.014, URL https://www.sciencedirect.com/science/article/pii/S0957417419303331.
[150] S. Arumugam, R. Bhargavi, 2019. A survey on driving behavior analysis in usage based insurance using big data. 6 (1), 86. http://dx.doi.org/10.1186/s40537-019-0249-5. URL https://journalofbigdata.springeropen.com/articles/10.1186/s40537-019-0249-5.
[151] D.E. Pozen, 2005. The mosaic theory, national security, and the freedom of information act. 52.
[152] M. Archie, S. Gershon, A. Katcoff, A. Zeng, Who's watching? De-anonymization of netflix reviews using amazon reviews, 2018b, Online; accessed 30 December 2021. URL https://www.oasislabs.com/how-it-works.
[153] L.M. Zagi, B. Aziz, Privacy attack on IoT: A systematic literature review, in: 7th International Conference on ICT for Smart Society: AIoT for Smart Society, ICISS 2020 - Proceeding, Institute of Electrical and Electronics Engineers Inc., 2020, URL https://ieeexplore.ieee.org/document/9307568.

[154] D. Kondor, B. Hashemian, Y.-A. de Montjoye, C. Ratti, Towards matching user mobility traces in large-scale datasets, IEEE Trans. Big Data 6 (4) (2020) 714–726, http://dx.doi.org/10.1109/TBDATA.2018.2871693.

[155] A. Wood, M. Altman, A. Bembenek, M. Bun, M. Gaboardi, J. Honaker, K. Nissim, D. O'Brien, T. Steinke, S. Vadhan, Differential privacy: A primer for a non-technical audience, SSRN Electron. J. (2018) http://dx.doi.org/10.2139/ssrn.3338027.

[156] A. Narayanan, V. Shmatikov, Robust de-anonymization of large sparse datasets, in: 2008 IEEE Symposium on Security and Privacy (Sp 2008), IEEE, Oakland, CA, USA, (ISSN: 1081-6011) 2008, pp. 111–125, http://dx.doi.org/10.1109/SP.2008.33.

[157] M. Archie, S. Gershon, A. Katcoff, A. Zeng, De-anonymization of netflix reviews using amazon reviews, 2018a, p. 5, URL https://www.readkong.com/page/de-anonymization-of-netflix-reviews-using-amazon-reviews-1439089.

[158] N. Lomas, France fines google $120M and amazon $42M for dropping tracking cookies without consent, 2020, URL https://dataprotection.news/france-fines-google-120m-and-amazon-42m-for-dropping-tracking-cookies-without-consent/.

[159] BBC, H&M fined for breaking GDPR over employee surveillance - BBC news, in: BBC, 2020, URL https://www.bbc.com/news/technology-54418936.

[160] Marketing : the Italian SA fines TIM EUR27.8 million, 2020, URL https://edpb.europa.eu/news/national-news/2020/marketing-italian-sa-fines-tim-eur-278-million_en.

[161] Q. Feng, D. He, S. Zeadally, M.K. Khan, N. Kumar, A survey on privacy protection in blockchain system, J. Netw. Comput. Appl. 126 (2019) 45–58, http://dx.doi.org/10.1016/j.jnca.2018.10.020.

[162] C. Wang, N. Zhang, C. Wang, Managing privacy in the digital economy, Fund. Res. 1 (5) (2021) 543–551, http://dx.doi.org/10.1016/j.fmre.2021.08.009.

[163] M. Akil, L. Islami, S. Fischer-Hübner, L.A. Martucci, A. Zuccato, Privacy-preserving identifiers for IoT: A systematic literature review, IEEE Access 8 (2020) 168470–168485, http://dx.doi.org/10.1109/ACCESS.2020.3023659.

[164] T. Gebremichael, L.P.I. Ledwaba, M.H. Eldefrawy, G.P. Hancke, N. Pereira, M. Gidlund, J. Akerberg, Security and privacy in the industrial internet of things: Current standards and future challenges, IEEE Access 8 (2020) 152351–152366, http://dx.doi.org/10.1109/ACCESS.2020.3016937.

[165] S.W. Driessen, G. Monsieur, W.-J. Van Den Heuvel, Data market design: A systematic literature review, IEEE Access 10 (2022) 33123–33153, http://dx.doi.org/10.1109/ACCESS.2022.3161478.

[166] N. Deepa, Q.-V. Pham, D.C. Nguyen, S. Bhattacharya, B. Prabadevi, T.R. Gadekallu, P.K.R. Maddikunta, F. Fang, P.N. Pathirana, A survey on blockchain for big data: Approaches, opportunities, and future directions, Future Gener. Comput. Syst. 131 (2022) 209–226.

[167] A.J. Perez, S. Zeadally, Secure and privacy-preserving crowdsensing using smart contracts: Issues and solutions, Comp. Sci. Rev. 43 (2022) 100450, http://dx.doi.org/10.1016/j.cosrev.2021.100450.

[168] C. Gonçalves, R.J. Bessa, P. Pinson, A critical overview of privacy-preserving approaches for collaborative forecasting, Int. J. Forecast. 37 (1) (2021) 322–342, http://dx.doi.org/10.1016/j.ijforecast.2020.06.003.

[169] Y. Wu, Z. Wang, Y. Ma, V.C.M. Leung, Deep reinforcement learning for blockchain in industrial IoT: A survey, Comput. Netw. 191 (2021) 108004, http://dx.doi.org/10.1016/j.comnet.2021.108004, URL https://www.sciencedirect.com/science/article/pii/S1389128621001213.

[170] L.D. Nguyen, I. Leyva-Mayorga, A.N. Lewis, P. Popovski, Modeling and analysis of data trading on blockchain-based market in IoT networks, IEEE Internet Things J. 8 (8) (2021) 6487–6497, http://dx.doi.org/10.1109/JIOT.2021.3051923.

[171] A. Sadiq, M.U. Javed, R. Khalid, A. Almogren, M. Shafiq, N. Javaid, Blockchain based data and energy trading in internet of electric vehicles, IEEE Access 9 (2021) 7000–7020, http://dx.doi.org/10.1109/ACCESS.2020.3048169.

[172] Y. Long, Y. Chen, W. Ren, H. Dou, N.N. Xiong, DePET: A decentralized privacy-preserving energy trading scheme for vehicular energy network via blockchain and K-anonymity, IEEE Access 8 (2020) 192587–192596, http://dx.doi.org/10.1109/ACCESS.2020.3030241.

[173] R. Xu, Y. Chen, Fed-DDM: A federated ledgers based framework for hierarchical decentralized data marketplaces, in: 2021 International Conference on Computer Communications and Networks, 2021.

[174] L. Giaretta, I. Savvidis, T. Marchioro, S. Girdzijauskas, G. Pallis, M.D. Dikaiakos, E. Markatos, PDS2: A user-centered decentralized marketplace for privacy preserving data processing, in: 2021 IEEE 37th International Conference on Data Engineering Workshops (ICDEW), 2021, pp. 92–99, http://dx.doi.org/10.1109/ICDEW53142.2021.00024.

[175] A. Manzoor, A. Braeken, S.S. Kanhere, M. Ylianttila, M. Liyanage, Proxy re-encryption enabled secure and anonymous IoT data sharing platform based on blockchain, J. Netw. Comput. Appl. 176 (2021) 102917, http://dx.doi.org/10.1016/j.jnca.2020.102917.

[176] T. Rückel, J. Sedlmeir, P. Hofmann, Fairness, integrity, and privacy in a scalable blockchain-based federated learning system, Comput. Netw. 202 (2022) 108621, http://dx.doi.org/10.1016/j.comnet.2021.108621.

[177] P. Gupta, V. Dedeoglu, S.S. Kanhere, R. Jurdak, TrailChain: Traceability of data ownership across blockchain-enabled multiple marketplaces, J. Netw. Comput. Appl. (2022) 103389, http://dx.doi.org/10.1016/j.jnca.2022.103389.

[178] Y. Tian, B. Song, T. Ma, A. Al-Dhelaan, M. Al-Dhelaan, Bi-tier differential privacy for precise auction-based people-centric IoT service, IEEE Access 9 (2021) 55036–55044, http://dx.doi.org/10.1109/ACCESS.2021.3067138.

[179] M. Zhang, L. Yang, S. He, M. Li, J. Zhang, Privacy-preserving data aggregation for mobile crowdsensing with externality: An auction approach, IEEE/ACM Trans. Netw. 29 (3) (2021b) 1046–1059, http://dx.doi.org/10.1109/TNET.2021.3056490.

[180] X. Xu, Z. Yang, Y. Xian, ATM: Attribute-based privacy-preserving task assignment and incentive mechanism for crowdsensing, IEEE Access 9 (2021) 60923–60933, http://dx.doi.org/10.1109/ACCESS.2021.3074142.

[181] F. Kserawi, S. Al-Marri, Q. Malluhi, Privacy-preserving fog aggregation of smart grid data using dynamic differentially-private data perturbation, IEEE Access 10 (2022) 43159–43174, http://dx.doi.org/10.1109/ACCESS.2022.3167015.

[182] Y. Shen, B. Guo, Y. Shen, X. Duan, X. Dong, H. Zhang, C. Zhang, Y. Jiang, Personal big data pricing method based on differential privacy, Comput. Secur. 113 (2022) 102529, http://dx.doi.org/10.1016/j.cose.2021.102529.

[183] Y. Hu, C. Li, A. Hu, A. Hu, J. Zhao, Trading off data resource availability and privacy preservation in multi-layer network transaction, Phys. Commun. 46 (2021) 101317, http://dx.doi.org/10.1016/j.phycom.2021.101317.

[184] Q. Song, J. Cao, K. Sun, Q. Li, K. Xu, Try before you buy: Privacy-preserving data evaluation on cloud-based machine learning data marketplace, in: Annual Computer Security Applications Conference, ACM, 2021, pp. 260–272, http://dx.doi.org/10.1145/3485832.3485921.

[185] M.N. Alraja, H. Barhamgi, A. Rattrout, M. Barhamgi, 2021. An integrated framework for privacy protection in IoT — Applied to smart healthcare. 91, 107060. http://dx.doi.org/10.1016/j.compeleceng.2021.107060. URL https://www.sciencedirect.com/science/article/pii/S0045790621000744.

[186] R. Oppliger, Privacy-enhancing technologies for the world wide web, Comput. Commun. 28 (16) (2005) 1791–1797, http://dx.doi.org/10.1016/j.comcom.2005.02.003.

[187] T.P. Pedersen, Non-interactive and information-theoretic secure verifiable secret sharing, in: Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), in: LNCS, vol. 576, 1992, pp. 129–140, http://dx.doi.org/10.1007/3-540-46766-1_9.

[188] A. Shamir, How to share a secret, Publ. ACM (1979b) http://dx.doi.org/10.1007/978-3-642-15328-0_17.

[189] X. Chen, Introduction to Secure Outsourcing Computation, Morgan & Claypool publishers, 2016, p. 94, URL https://www.researchgate.net/publication/295681472_Introduction_to_Secure_Outsourcing_Computation.

[190] S. De Capitani Di Vimercati, S. Foresti, G. Livraga, P. Samarati, Data privacy: Definitions and techniques, Int. J. Uncertain. Fuzz. Knowl.-Based Syst. 20 (6) (2012) 793–817, http://dx.doi.org/10.1142/S0218488512400247, URL https://www.worldscientific.com/doi/abs/10.1142/S0218488512400247.

[191] A. Meyerson, R. Williams, On the complexity of optimal k-anonymity, in: Proceedings of the ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, 23, 2004, pp. 223–228, http://dx.doi.org/10.1145/1055558.1055591, URL https://dl.acm.org/doi/10.1145/1055558.1055591.

[192] R. Rivest, A. Shamir, Y. Tauman, How to leak a secret, in: Lecture Notes in Computer Science, Vol 2248, Springer, Berlin, Heidelberg, 2001, URL https://ieeexplore.ieee.org/document/6032224{%}0Ahttps://cryptoslate.com/ethereum-network-congestion-doubles-gas-fees-as-game.

[193] M. Bellare, D. Micciancio, B. Warinschi, Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions, Eurocrypt 2656 (2003) 1–27, URL https://cseweb.ucsd.edu/~mihir/papers/gs.pdf.

[194] S. Gürses, C. Troncoso, C. Diaz, Engineering: Privacy by design, Science 317 (5842) (2011) 1178–1179, URL https://www.esat.kuleuven.be/cosic/publications/article-1542.pdf.

[195] M. Yung, S. Jarecki, H. Krawczyk, A. Herzberg, Proactive secret sharing or : How to cope with perpetual leakage, Communication (1995) URL https://www.researchgate.net/publication/221355399_Proactive_Secret_Sharing_Or_How_to_Cope_With_Perpetual_Leakage.

[196] IOTA-Foundation, About the tangle, 2020, URL https://legacy.docs.iota.org/docs/getting-started/1.1/the-tangle/overview.

[197] A. Ghosh, K. Ligett, A. Roth, G. Schoenebeck, Buying private data without verification, in: EC 2014 - Proceedings of the 15th ACM Conference on Economics and Computation, 2014, pp. 931–948, http://dx.doi.org/10.1145/2600057.2602902, arXiv:1404.6003.

[198] C. Perera, R. Ranjan, L. Wang, S.U. Khan, A.Y. Zomaya, Big data privacy in the internet of things era, IT Prof. 17 (3) (2015c) 32–39, http://dx.doi.org/10.1109/MITP.2015.34.

[199] The Wall Street Journal, Google to pay $22.5 million in FTC settlement, 2012, URL https://www.wsj.com/articles/SB10000872396390443404004577579232818727246.

[200] J. Porter, Google fined €50 million for GDPR violation in France, 2019, URL https://www.theverge.com/2019/1/21/18191591/google-gdpr-fine-50-million-euros-data-consent-cnil.

Further reading

[1] G. Goos, J. Hartmanis, J. van Leeuwen, D. Hutchison, T. Kanade, J. Kittler, J.M. Kleinberg, F. Mattern, J.C. Mitchell, M. Naor, O. Nierstrasz, C.P. Rangan, B. Steffen, 1973. Lecture notes in computer science. 556. URL https://doi.org/10.1007/978-3-319-70139-4_56.

[2] P. Cramton, Y. Shoham, R. Steinberg, An overview of combinatorial auctions, ACM SIGECOM Exch. 7 (1) (2007) 3–14, URL https://dl.acm.org/doi/10.1145/1345037.1345039.

[3] R. Dingledine, N. Mathewson, P. Syverson, Tor: The Second-Generation Onion Router, Tech. Rep., Defense Technical Information Center, Fort Belvoir, VA, 2004, URL http://www.dtic.mil/docs/citations/ADA465464.

[4] D. Bogdanov, S. Laur, J. Willemson, 2008. Sharemind: a framework for fast privacy-preserving computations. 15. URL https://link.springer.com/chapter/10.1007/978-3-540-88313-5_13.

[5] N. Wang, X. Xiao, Y. Yang, J. Zhao, S.C. Hui, H. Shin, J. Shin, G. Yu, Collecting and analyzing multidimensional data with local differential privacy, 2019b, p. 13, URL https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8731512.

[6] R. Bassily, A. Smith, Local, private, efficient protocols for succinct histograms, in: Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, ACM, 2015, http://dx.doi.org/10.1145/2746539.2746632.

[7] A. Herzberg, S. Jarecki, H. Krawczyk, M. Yung, Proactive secret sharing or: How to cope with perpetual leakage, in: D. Coppersmith (Ed.), Advances in Cryptology — CRYPT0' 95, Springer Berlin Heidelberg, Berlin, Heidelberg, 1995, pp. 339–352, URL https://link.springer.com/chapter/10.1007/3-540-44750-4_27.

[8] A. Bellet, A. Habrard, M. Sebban, A survey on metric learning for feature vectors and structured data, 2014, URL https://arxiv.org/abs/1306.6709.

[9] B. Poettering, D. Stebila, Double-authentication-preventing signatures, Int. J. Inf. Secur. 16 (1) (2017) URL http://link.springer.com/10.1007/s10207-015-0307-8.

[10] S. Yu, C. Wang, K. Ren, W. Lou, Achieving secure, scalable, and fine-grained data access control in cloud computing, in: 2010 Proceedings IEEE INFOCOM, 2010, pp. 1–9, http://dx.doi.org/10.1109/INFCOM.2010.5462174.

[11] R. Wieringa, N. Maiden, N. Mead, C. Rolland, Requirements engineering paper classification and evaluation criteria: A proposal and a discussion, Requir. Eng. 11 (1) (2006) 102–107, http://dx.doi.org/10.1007/s00766-005-0021-6.

Gonzalo Munilla Garrido joined the chair for Software Engineering of Business Information Systems (Sebis) at the Technical University Munich in October 2019. In cooperation with The BMW Group, he currently investigates how privacy-preserving techniques enable new iterations of data marketplaces. He received a M.Sc. mult. in Mechanical Engineering from the TUM and the Polytechnique University of Madrid.

Johannes Sedlmeir is a researcher at the Project Group Business & Information Systems Engineering of the Fraunhofer FIT and currently pursuing his Ph.D. in Information Systems at the FIM Research Center, University of Bayreuth. In his research, Mr. Sedlmeir works on the energy consumption and performance benchmarking of different blockchains, decentralized digital identities, and the application of cryptographic methods such as Zero-Knowledge Proofs for scalability and privacy on blockchains. He received his M.Sc. in Theoretical and Mathematical Physics.

Ömer Uludağ joined the chair for Software Engineering of Business Information Systems at the Technical University Munich in April 2016. He holds a M.Sc. in Information Systems and currently works at Allianz SE.

Ilias Soto is currently a M.Sc. student in Information Systems at the Technical University of Munich and holds a B.Sc. in the same field.

Andre Luckow work currently focuses on applications of machine learning, quantum computing and blockchain. Mr. Luckow drives the application of computing technologies to problems in business and science bridging cross-functional gaps to create value via process improvements or the enablement of new products. Previously, Mr. Luckow held several functions at BMW Group IT in Munich, Germany. Furthermore, he teaches several courses at different universities, e.g. the Ludwig-Maximilians-University Munich, Germany, and Clemson University, SC, USA. Andre Luckow holds a Ph.D. in the field of distributed computing from the University of Potsdam.

Florian Matthes Since 2002 Florian Matthes holds the chair for Software Engineering for Business Information Systems at Technische Universität München.

The current focus of his research is on technologies driving the digital transformation of enterprises and societies: Enterprise architecture management, service platforms and their ecosystems, semantic analysis of legal texts and executable contracts on blockchains.

He is co-founder of CoreMedia, infoAsset and Tr8cy, co-founder and chair of Blockchain Bayern e.V. scientifc advisor of Noumena Digital, member of the advisory board of the Ernst Denert-Stiftung für Software Engineering, and initiator and organizer of international conferences and workshops in software and enterprise engineering

Earlier stations of his academic career are the Goethe-University Frankfurt (Diploma 1988) the University of Hamburg(Ph.D. 1992), the Digital Systems Research Center (now HP SRC Classic) in Palo Alto, USA (Researcher 1992–1993), and the Technical University Hamburg-Harburg (Associate Professor 1997–2002).

Until 2010 he served as dean of studies at the Faculty for Informatics and member of the teaching board of TU München. Since 2020 he serves as "venture ambassador" of TUM venture labs.

Appendix A. Acronyms

AET	Authenticity-enhancing technologies
BP	Bundling problem
CP	Copy problem
DF	Digital fingerprint
DID	Decentralized identifier
DLT	Distributed ledger technology
DP	Differential privacy
DS	Digital signature
FHE	Fully homomorphic encryption
FL	Federated learning
GAN	Generative adversarial networks
GDPR	General Data Protection Regulation
HE	Homomorphic encryption
ICT	Information and communication technology
IoT	Internet of things
KGC	Key generation center
ML	Machine learning
PET	Privacy-enhancing technology
PHE	Partially homomorphic encryption
PKI	Public key infrastructure
PPDM	Privacy-preserving data mining
REP	Recursive enforcement problem
RQ	Research question
SC	Smart contract
SLR	Systematic literature review
MPC	Secure multiparty computation
TD	Truth discovery
TEE	Trusted execution environment
ZKP	Zero-knowledge proof

Appendix B. Summaries of the selected secondary studies

Table B.4: Secondary studies on privacy in data markets.

Year	Study	Topic	Description
2020	[S8]	Privacy-enhancing design of data markets	Analyzes internet users' preferences for privacy in data sharing to uncover mental models of these preferences and their motives, barriers, and conditions for a privacy-enhancing data market. It provides a set of key findings, the two most notable ones being that the primary barrier to creating data markets is privacy and moral concerns and that the level of anonymization has the largest effect on the willingness to share.
2019	[S19]	Privacy and security data flow challenges in an internet of production	Introduces the internet of production and illustrates its inter-organizational data flows. It also identifies security and privacy demands and challenges within these data flows: authenticity, data access scope, and anonymity. Furthermore, it provides a survey of PETs to tackle these challenges: provide confidentiality, hide information during computation (data processing), verify the authenticity of information (providing support), deploy mechanisms that enforce rules (platform capabilities), and support approaches that focus on the security of data flows (external measures).
2019	[S14]	Challenges and research opportunities in data markets for the IoT	A short study that identifies three research opportunities in IoT data markets: Procurement, pricing, and privacy. Significant identified challenges are: ambiguity in data ownership that hinders trading, the difficulty to detect data piracy, and that privacy must be considered before trading.
2018	[S23]	Privacy enhancing in IoT applications	Introduces and surveys privacy-enhancing technologies in the processes of data aggregation, trading, and analysis; in particular, it discusses outsourced computation, data validation, and blockchain technology. Additionally, it describes types of privacy breaches and their countermeasures. Furthermore, it reviews relevant aspects of pricing procedures as well as game-theoretical approaches and auction schemes.
2018	[S4]	Pricing, trading, and protection in data markets for the IoT	Surveys the three fields of pricing models and strategies, design of platforms and data trading, and digital copyright mechanisms with a focus on privacy enhancement.
2018	[S48]	Privacy-enhancing analytics for IoT and cloud-based systems	Summarizes privacy-enhancing technologies in the specific use case of a health data collecting app in the health industry. More specifically, it separates privacy-enhancing technologies into two scenarios: Outsourced computation and information sharing.
2016	[S35]	Privacy enhancing in crowdsourcing task management	Surveys privacy-enhancing technologies and the challenges of crowdsourcing task management. The proposed technologies are anonymization, such as k-anonymity, spatio-temporal privacy approaches, such as spatial cloaking or aggregated location via differential privacy, and policy-based privacy preferences. The challenges that they present revolve around trust and credibility, reward-based tasking, utility, efficiency, enforcing privacy-enhancing technologies, and raise privacy awareness.
2015	[S27]	Privacy enhancing and challenges in data markets for the IoT	The study introduces privacy enhancing for data markets for IoT devices, focusing on sensing-as-a-service (data analysis of user-aggregated data). It identifies three challenges: Developing IoT middleware for data analysis and autonomous privacy enhancing, autonomous end-user consent acquisition and negotiation, and the autonomous modeling and negotiation of privacy risk and economic reward. The most prevalent privacy-enhancing technologies and strategies they introduce are personal information hubs, onion routing, and data aggregation via differential privacy or k-anonymity.
2014	[S38]	Privacy threats and challenges in the IoT	Classifies the threats and challenges that come along with privacy in IoT applications for individuals into seven categories: Re-identification of individuals through persistent pseudo-identifiers, localization and tracking, profiling for social engineering and price discrimination, information disclosure in life cycle transitions, information linkage of previously separated systems, inventory attacks, and the disclosure of private information to an uninvited audience.

Appendix C. Mappings of privacy- and authenticity-enhancing technologies

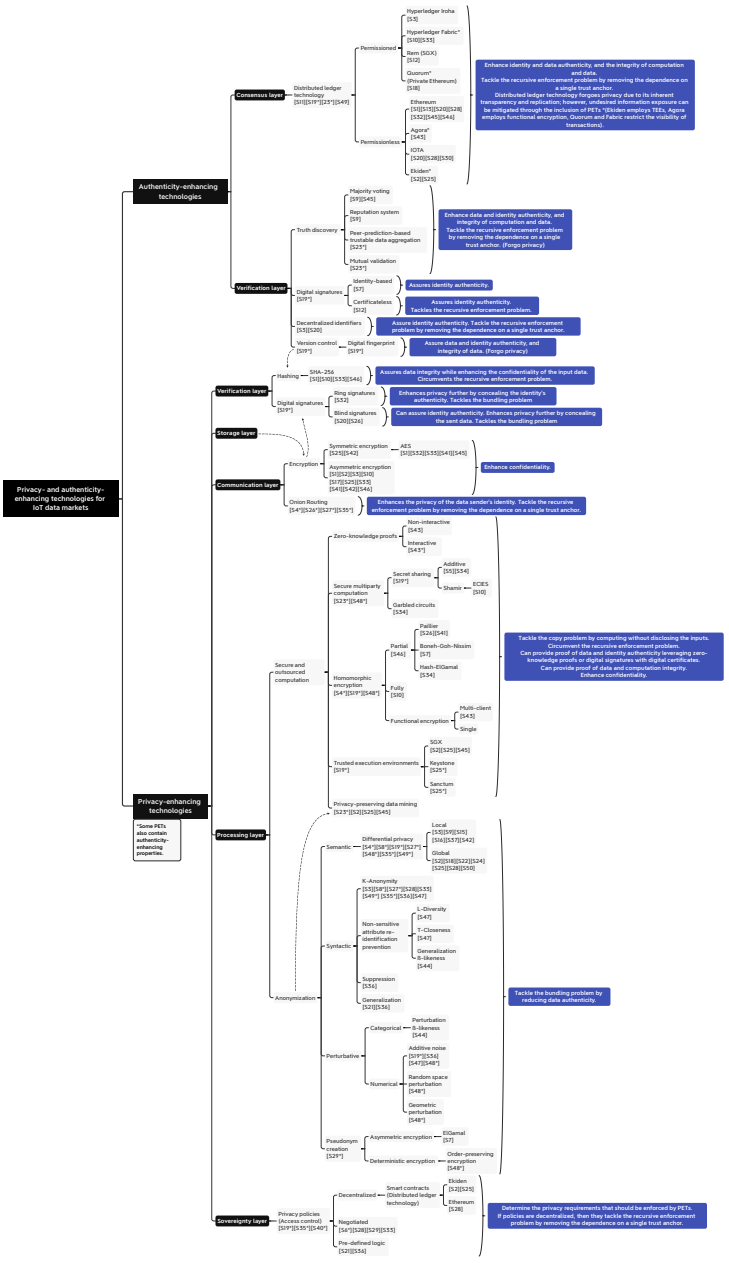

Figure C.11: Classification of the identified privacy- and authenticity-enhancing technologies in this SLR, together with the challenges they tackle. Any other privacy approach encountered in the SLR without a succinct inclusion of the underlying technology was either not included in a leaf node but in a parent node or completely dismissed if too vague. *The publication reviews the technology without delving into it in-depth or using it as a building block of the architecture concept, e.g., the technology is only mentioned in the opportunities for future work.

Appendix D. Summaries of the selected primary studies

Table D.5: Set of examples from the selected studies describing different data marketplace architectures and auctions based on PETs and AETs with a focus on secure computation technologies.

Year	Study	Privacy-enhancing approaches	Description
2020	[S41]	Partially homomorphic encryption, symmetric encryption, and digital signatures	Develops a privacy-enhancing auction for big data trading using Paillier's cryptosystem [18] and a one-time pad. They consider four entities: sellers, buyers, an auctioneer, and an intermediary platform. A data auction is carried out without any entity seeing the data (except the auction winner) or the bid values, which are ordered obliviously by the auctioneer thanks to the homomorphic properties of the ciphertext. Furthermore, to efficiently encrypt the data, the authors use symmetric encryption (AES). Lastly, digital signatures are created with the same homomorphic cryptosystem, which the authors use to encrypt the symmetric keys.
2018	[S7]	Partially homomorphic encryption, digital signatures, and pseudonym creation	Implements a platform for data markets that facilitates data processing and outcome verification while enhancing the privacy of identities and their data. The authors consider four entities: data contributor, service provider, data consumer, and a registration center; in a two-layer system model: data acquisition and trading. Furthermore, the platform synchronizes data processing and signature verification into the same homomorphic ciphertext space (encrypt and then sign). Additionally, they tightly integrate data processing with outcome verification via a set of homomorphic properties. To achieve a trade-off between functionality and performance, they selected a partially homomorphic scheme called Boneh-Goh-Nissim cryptosystem [84].
2018	[S10]	Fully homomorphic encryption, secure multiparty computation, distributed ledger technology, and hashing	Provides a distributed outsourcing computation architecture, whereby data owners may request fully homomorphic computations with a schema called fully homomorphic non-interactive verifiable secret sharing [86]. Moreover, the proposed architecture allows transactions to be verified by the participants of the permissioned blockchain thanks to the immutability properties of the blockchain; Hyperledger Fabric was the selected blockchain architecture. Moreover, the hash value of the shared data is stored in the blockchain, for data recipients to verify the truthfulness of the received data. Furthermore, for secure multiparty computation, the authors implement Shamir's secret sharing [72] with ECIES by leveraging the distributed nature of the blockchain. In this manner, the data owner may share verifiable pieces of information with a set of servers. Then, the servers execute the necessary computations, and when several verified responses are received by the agreed data consumer, the true result is recovered.
2016	[S34]	Partially homomorphic encryption, and secure multiparty computation	Develops an auction cloud-based framework that cryptographically hides the bids from all auction participants until a winner is determined. It achieves this by combining PHE based on the hashed scheme [85] of ElGamal [202], and secure two-party computation through garble circuits and additive secret sharing.
2015	[S26]	Partially homomorphic encryption, digital signatures, and onion routing	Proposes a combinatorial auction [203] mechanism that ensures the privacy of the bidders. The bidders bids are blindly signed through the third party [61] so that the third party does not learn the contents. Later, these signatures are used by the bidder to prove the authenticity of the bids. The winner is determined by the third party through a partially homomorphically encrypted computation using the Paillier cryptosystem [18]. Lastly, the identities of the bidders are enhanced by using onion routing [204].
2012	[S5]	Secure multiparty computation	Implements a privacy-enhancing data mining service market, whereby data donors distribute confidential data among a set of participants employing additive secret sharing. The miners collectively perform secure multiparty computation based on the author's algorithm [205]. Finally, the results are in turn sent to the previously agreed analyst, who combines them to obtain the intelligible output.

Table D.6: Set of examples from the selected studies describing different data marketplace architectures and trading mechanisms based on PETs and AETs with a focus on anonymization technologies.

Year	Study	Privacy-enhancing approaches	Description
2020	[S24]	Differential privacy	Designs a privacy-enhancing crowd-sensed data trading mechanism. First, the data broker orchestrates an auction whereby data consumers bid in a differentially private manner for a data asset. Secondly, to form a data asset, the data broker creates a set of data generation tasks, some of which are fake to protect the privacy of the auction winner. Lastly, the data broker selects data owner outputs in a differentially private manner. More specifically, both the auction-based data pricing and the data collection are based on the differentially private exponential mechanism.
2019	[S37]	Differential privacy	Proposes a differentially private data market auction framework with a fair negotiation method to set the price and noise; this study is extended in [S42]. The entities involved are a data provider, a consumer, and a trusted market manager that matches providers with consumers and enforces Rubinstein bargaining. Firstly, the data provider and consumer enter a negotiation phase that involves the data query, the ε values, and the unit price for ε. Once the negotiation is over, the data provider answers the query with the agreed ε with local differential privacy.
2019	[S42]	Differential privacy and digital signatures	Proposes a differentially private data market framework. This study extends [S37] by specifying the type of differential privacy algorithm, and the digital signature schemata followed to deploy the framework in practice. The authors use local differential privacy for numeric [206] and for categorical [207] data types.
2019	[S15]	Differential privacy	Designs contracts for a data marketplace whereby a data broker matches the required accuracy from a data consumer with the degree of privacy that data owners desire. Furthermore, by handpicking the data sources, the differentially private algorithm incurs a bias that makes the output more accurate while maintaining the desired privacy from the data owners. Lastly, the authors derive an optimal data contract to minimize payment while satisfying accuracy and privacy.
2019	[S50]	Differential privacy	Proposes a framework for counting trading range query results, and designs a pricing approach for the traded results. Firstly, the framework calculates the range counts approximately, and secondly, it protects the results further by using differential privacy, while satisfying the accuracy demands of data consumers. The authors also design the pricing scheme in a way that prevents arbitrage.
2019	[S22]	Differential privacy	Designs an online auction with two stages, whereby a trusted auctioneer aggregates data from data owners and applies differential privacy before selling the data to consumers. In the first stage, the auctioneer selects data owners based on their privacy requirements to maximize profit. In the second stage, the auctioneer applies differential privacy to the aggregated data and subsequently sells the data to a consumer in an auction.
2018	[S2]	Differential privacy, distributed ledger technology, smart contracts, privacy policies, asymmetric encryption, and trusted execution environments	Implements an end-to-end privacy-enhancing decentralized data marketplace for data consumers to train machine learning algorithms. The authors achieve end-to-end privacy by protecting inputs with asymmetric encryption and differential privacy, and the execution with trusted execution environments. More specifically, differential privacy prevents the weights of machine learning algorithms from overfitting to the inputs. Because of the privacy limitations of current distributed ledger technology applications, the authors of this study and of the subsequent publication [S25] created a novel concept unlike any other blockchain-based system. For example, in principle, the smart contracts of their architecture may contain machine learning algorithms which may be executed in trusted execution environments, hold privacy policies and payment logic, and point to where encrypted data and decryption keys are stored privately. On the other hand, data consumers also deploy smart contracts that may interact with the data owners.
2018	[S16]	Differential privacy	Develops an auction framework for privacy-enhancing data aggregation for mobile crowdsensing. The auctioneer chooses data owners based on their sensing capabilities, and the data owners apply differential privacy to their inputs sampling from a noise distribution tailored by the auctioneer for each data owner based on its qualities. The goal of the platform is to optimize task allocation to a set of data owners while minimizing their payment, taking into account accuracy and privacy constraints.
2018	[S9]	Differential privacy and truth discovery	Designs two locally differentially private mechanisms for truth discovery in crowd sensing, so that the answers from edge devices are protected while being useful in aggregate. The second mechanism provides more utility for an equal degree of privacy, and consists on the users randomly selecting a probability distribution, and in turn, adding noise sampled thereof to their truthful answer.

Table D.6: PETs and AETs in the context of anonymization *(continued)*.

Year	Study	Privacy-enhancing approaches	Description
2018	[S21]	Privacy policies and generalization	Proposes a data market framework that models and enforces privacy policies dynamically for data-intensive applications. More specifically, the authors implement a data-flow-focused system with a policy enforcement algorithm defined by users and a context. In data-flow computing, directed graphs embody the application, where edges represent data streams and nodes represent functional operators and data sources or sinks. The data is anonymized based on policies and enforced by generalization, e.g., substituting Munich with Germany. To formalize a language to model the privacy policies, the authors use metric first-order temporal logic.
2017	[S47]	*K*-anonymity and additive noise	Models a data marketplace in which groups of users may actively monetize their data through a mediator and a set of mobile crowd sensing service providers. The authors use a reverse auction, where users bid for performing sensing tasks. Individual users may set their own privacy preferences, and if they are a coalition of users, they are protected by k-anonymity, t-closeness, l-diversity and local noise addition approaches. The total coalition payoff is divided among the cooperative users based on their marginal contributions to the total data quality at the end of the sensing service.
2016	[S36]	*K*-anonymity, additive noise, and privacy policies	Designs a one-to-one privacy enhancing paradigm for a data market place in which privacy policies and data requirements are defined based on the publication record of the data owner. Because published records of a user aggregate over time and thus accrue privacy risk, the paradigm relies on privacy risk management, which is enforced by evaluating the risk associated with revealing yet another piece of information with regard to the privacy requirements. This evaluation is based on the preferences of the user, or if unfeasible, based on current regulation; furthermore, it is based on an assessment of the background information, achieved by semantically analyzing attributes that if released could be linked to externally available information. Ultimately, to privatise the data, the authors propose syntactic technologies such as k-anonymity, suppression, and generalization, and semantic ones like additive noise.

Table D.7: Set of examples from the selected studies describing different data marketplace architectures based on PETs and AETs with an underlying distributed ledger technology.

Year	Study	Privacy-enhancing approaches	Description
2020	[S3]	Distributed ledger technology, smart contracts, decentralized identifiers, digital signatures, *k*-anonymity, and differential privacy	Implements a framework for mobility data markets with six layers, each with a purpose and a technology to execute. Furthermore, the framework focuses on location-based services. The identity layer uses asymmetric identity keys, i.e. a key issued only to a real person, to verify that an entity is a real individual. The privacy layer leverages *k*-anonymity for Geomasking (low utility and high privacy), and when the service needs an exact location, differential privacy for Geo-Indistinguishability (high utility and low privacy). Moreover, the contract layer is based on smart contracts that enforce fair trade and the resolve disputes automatically. For the private communication layer, the authors use decentralized identifiers (DID) [153] issued by the device of a person itself. When devices communicate, the communication has a unique ID based on both of the DIDs, thus, even though the communication data is persisted in a blockchain, it is nontrivial to track the locations of a user. Consecutively, the incentive layer uses smart contracts and data brokers to promote data exchange for a profit; however, this architecture does not tackle the copy problem. The consensus layer is based on a consortium blockchain for distributed governance among non-anonymous honest entities. The blockchain selected was Hyperledger Iroha, chiefly because of its lightweight quality that couples with deployments in IoT devices.
2020	[S43]	Distributed ledger technology, smart contracts, functional encryption, and zero-knowledge proofs	Proposes a privacy-enhancing decentralized data marketplace employing the Agora blockchain with verification technology that enable data prosumers to monetize their data. The privacy-enhancing aspect is achieved by sending encrypted data to brokers employing a primitive called *multi-client functional encryption* [87][88], which ensures that the receiver may only decrypt the output of a formerly agreed-on function. Moreover, consumers may purchase these outputs, together with a proof of correctness from the broker by using non-interactive zero-knowledge proofs. For the decentralized architecture, the authors employ the Agora blockchain, and atomic payments are performed via smart contracts.
2020	[S45]	Distributed ledger technology, smart contracts, trusted execution environments, truth discovery, digital signatures	Proposes a data processing-as-a-service model based on a blockchain-based data trading ecosystem, whereby neither data brokers nor consumers have access to the raw data, only to the analysis. The use of a blockchain (Ethereum) prevents a single point of failure and allows for immutability and transparency in transactions. Furthermore, to protect the data, the analysis results, and the processing itself, the authors use Intel's SGX trusted execution environment [95], in addition to the symmetric encryption algorithm AES-256 to provide encryption and decryption within and outside the secure environment. The architecture uses the conventional Ethereum Virtual Machine (EVM) for traditional smart contracts, while the data analysis contracts are executed in a SGX-protected EVM where an initial key exchange is needed. Lastly, the nodes in the network form a compute market, i.e. multiple nodes execute the analysis and only the most frequent result is delivered to the data consumer, and the corresponding nodes are rewarded.
2020	[S28]	Distributed ledger technology, privacy policies, differential privacy, *k*-anonymity, and digital signatures	Proposes an architecture for a personal data marketplace in which personal data is stored decentrally in a allegedly GDPR compliant manner. To accomplish this, transactions and pointers to the data are encrypted and stored using a distributed ledger technology, namely IOTA. The data is stored either in an interplanetary file system, or in an IOTA-based storage format. In order to access such data, a data aggregator must request permissions through Ethereum-based smart contracts (whitelists) owned by data consumers. Once the permission has been granted, the trusted data aggregator, whose mutually agreed privacy policies are persisted in another Ethereum-based smart contract, waits until enough data owners exist to fulfill a particular analysis, so that the aggregator may perform *k*-anonymity. The data aggregator sells the anonymized data to consumers and remunerates data owners accordingly. However, the presence of a trusted aggregator defeats to some extent the purpose of a decentralized platform. Furthermore, the link between Ethereum and IOTA is carried out by trusted authentication services, which allow data aggregators to decrypt the data. Lastly, in order for data owners to grant access to their data, the authors recommend dynamic threshold encryption [208] over centralized forms of authentication services.
2020	[S18]	Distributed ledger technology, smart contracts, and differential privacy	Proposes a data trading approach in which privacy loss is publicly auditable and data owners set their privacy requirements on publicly available contracts. To accomplish this, the author uses a private Ethereum blockchain called Quorum that supports a set of built-in privacy measures, such as private transactions, messaging, and contracts; however, this design also restricts the interactions that are possible with smart contracts outside the private subset. Furthermore, the data owner applies differential privacy locally before sharing the data with the consumer.

Table D.7: PETs and AETs in the context of DLT *(continued)*.

Year	Study	Privacy-enhancing approaches	Description
2019	[S25]	Distributed ledger technology, trusted execution environments, smart contracts, privacy policies, digital signatures, and differential privacy	Implements an end-to-end privacy-enhancing decentralized data marketplace for data consumers to train machine learning algorithms, among other Turing-complete tasks. The architecture proposed is a mature version of [S2]. Containing all the features of [S2], R. Cheng, et al. [S25] improve the performance of a newly designed distributed ledger technology to allow for horizontal scaling, i.e. the more nodes are added to the network, the more performant the network is, unlike e.g. Bitcoin or Ethereum; furthermore, the authors tackle the problem of confidentiality by separating consensus from execution, whose computations are performed in a trusted execution environment. Horizontal scaling is achieved by allowing for parallel transaction execution, which is, in turn, accomplished by a set of transaction schedulers, and by creating dedicated committees for computation, storage, merging outputs, key management, and consensus. However, scalability through restricting the degree of redundancy entails a security/integrity tradeoff. Key management committees are necessary for the use of trusted execution environments to enable confidential computations. The architecture uses symmetric keys for state encryption and asymmetric encryption for concealing user inputs. The authors achieve end-to-end privacy by protecting inputs with asymmetric encryption and differential privacy, and the execution with trusted execution environments. More specifically, differential privacy prevents the weights of machine learning algorithms from overfitting to the inputs. Because of the privacy limitations of current distributed ledger technology applications, they created a new concept so that smart contracts allow for privacy-enhancing features; this concept was introduced by [S2].
2019	[S32]	Distributed ledger technology, smart contracts, and digital signatures	Proposes a privacy-enhancing fair data trading protocol. The protocol relies on the Ethereum blockchain to achieve a decentralized nature, however, the authors claim the protocol is blockchain agnostic. Nonetheless, despite using a decentralized network, the market manager holds non-negligible authority, as it may trace the identity of sellers so that they can be punished monetarily in case they misbehave. Furthermore, once the buyers have decided which data asset to purchase, the sellers use symmetric keys to encrypt data in chunks before sending it to the buyers. Upon receiving the data chunks, the buyer (i) challenges a set of data chunks, and upon verification of truthful data, (ii) employs similarity learning, a machine learning technology [209], to decide whether to finally purchase the data. Consequently, once the buyer decides to purchase the data, the seller and buyer interact via a payment smart contract and double-authentication-preventing signatures [210] to ensure payment and data decryption. Lastly, in order to enhance the anonymity of the actors, the protocol uses ring signatures [136][137].
2019	[S20]	Distributed ledger technology, smart contracts, digital signatures, and decentralized identifiers	Implements a decentralized data market architecture with secure data processing for the IoT. To achieve decentralization, the authors rely on the distributed ledger technology IOTA, and Ethereum-based smart contracts for subscribing to data streams. The constellation of actors consists of three entities: a data provider, a consumer, and a broker; the former two entities are included in a registry via decentralized identifiers [153]. The product that the consumers purchase is a key to a data stream for a predetermined period of time, created but not accessible by the data broker. For the consumer to attain data access in a private manner, blind signatures [61] are employed, which enable a data broker to verify stream access keys from the data provider without ever accessing these keys. More specifically, the data provider "blinds" the session key with the broker's public key and sends the blinded key to the broker, consequently, the broker certifies the key with its signature and returns the signature to the provider who removes the blinding factor to access the stream. Lastly, to exchange stream data, an inter-planetary file system is employed.
2019	[S12]	Distributed ledger technology, trusted execution environments, and digital signatures	Proposes a distributed IoT data storage system and a data trading scheme. The authors use the blockchain for its distributed nature, immutability, and requester authentication; moreover, their solution is blockchain agnostic. However, for the consensus algorithm, the authors rely on Intel's Software Guard Extension (SGX) [95] to deploy a trusted execution environment, to perform "Proof of Useful Work". The blockchain only contains pointers (addresses) to a distributed hash table, where the data is stored off-chain by peers of the network. Only certified data consumers, e.g. other IoT devices, would be able to query addresses in the blockchain. Furthermore, the authors employ certificateless cryptography so that the key generation center of conventional identity-based encryption does not need to be trusted [152]. To perform the cryptographic operations, edge devices are deployed. Lastly, to share data with purchasers, the authors propose to use either asymmetric encryption or re-encryption [211].
2019	[S1]	Distributed ledger technology, digital signatures, and hashing	Prototypes a decentralized fair data trading platform. The authors rely on the Ethereum blockchain to avoid third-party data brokers and to leverage the ledger's immutability properties. Moreover, data sellers utilize smart contracts to propose their data offers and to interact with sellers. Sellers include the hash of the data in the ledger so that the buyer may initiate a rebuttal if there is an expectation mismatch. Additionally, to ensure accountability the authors rely on digital signatures to verify that the data was encrypted using a specific key that belongs to the seller, and encrypt the data efficiently using symmetric encryption. The asset traded are decryption keys, and buyers may retrieve data as ciphertexts from untrusted storage.

Table D.7: PETs and AETs in the context of DLT *(continued)*.

Year	Study	Privacy-enhancing approaches	Description
2018	[S46]	Distributed ledger technology, smart contracts, partially homomorphic encryption, hashing, and digital signatures	Proposes two secure, and fair data trading decentralized schemata built on the Ethereum blockchain. One scheme enables entities to trade raw data, while the other scheme enables them to exchange statistics. The authors chose blockchain in both schemata for its immutability, smart contracts, P2P payment, and disintermediation. Furthermore, for the second scheme, the authors use partially homomorphic encryption to perform confidential statistics. For the data structure to compute the statistics, the authors chose a Merkle Accumulative Tree, where the leaf nodes hold the encrypted data and the non-leaf nodes contain the hash values and a cumulative sum of homomorphic ciphertexts. The data exchanged is verifiable through digital signatures based on asymmetric encryption.
2017	[S33]	Distributed ledger technology, privacy policies, digital signatures, hashing, and k-anonymity	Prototypes a decentralized data market platform for anonymized data. The underlying distributed ledger technology is Hyperledger Fabric, whose peers act as data brokers. The data brokers may only handle datasets based on a set of privacy policies in the interest of the data owner and dictated by a data domain-specific privacy policy manager. The blockchain acts as an auditable ledger for transactions between data brokers and consumers, while the exchange of data is handled off-chain. Furthermore, the anonymization of data is suggested to be performed employing k-anonymity by the broker upon dataset reception from a secure channel, however, the solution remains anonymization-agnostic. For every actor to verify that the correct anonymized dataset has been shared, the broker sends its hash value using SHA-256 to the blockchain before sending it to the consumer. Upon reception, both the policy manager and data receiver may verify the dataset. Lastly, cryptography technologies are employed to encrypt the dataset symmetrically (128-bit AES) before sharing the dataset to the consumer, and the actors use ECDSA to sign confirmations and transactions.

Appendix E. Distribution of selected studies by publication channels

Table E.8: Publication channels for the studies from our SLR.

#	Publication source	Type	No.	%
1	ACM International Conference Proceeding Series	Conference	2	4
2	IEEE International Conference on Internet of Things	Conference	2	4
3	VLDB Endowment	Journal	2	4
4	ACM Conference on Computer and Communications Security	Conference	1	2
5	ACM International Workshop on Mobile Commerce	Workshop	1	2
6	ACM SIGKDD International Conference on Knowledge Discovery and Data Mining	Conference	1	2
7	ACM SIGMOD Record	Journal	1	2
8	CEUR Workshop	Workshop	1	2
9	Computer Law and Security Review	Journal	1	2
10	Computer Networks	Journal	1	2
11	Concurrency Computation: Practice and Experience	Journal	1	2
12	Conference on Information and Knowledge Management	Conference	1	2
13	Electronic Markets	Journal	1	2
14	IACR Cryptology ePrint Archive	Journal	1	2
15	IEEE Access	Journal	1	2
16	IEEE Cloud Computing	Journal	1	2
17	IEEE Communications Magazine	Journal	1	2
18	IEEE Computer	Journal	1	2
19	IEEE Eurasia Conference on IOT, Communication and Engineering	Conference	1	2
20	IEEE European Symposium on Security and Privacy	Conference	1	2
21	IEEE International Conference on Big Data	Conference	1	2
22	IEEE International Conference on Blockchain and Cryptocurrency	Conference	1	2
23	IEEE International Conference on Collaboration and Internet Computing	Conference	1	2
24	IEEE International Conference on Pervasive Computing and Communications Workshops	Conference	1	2
25	IEEE International Conference on Software Engineering Research, Management and Applications	Conference	1	2
26	IEEE Internet Computing	Journal	1	2
27	IEEE Internet of Things Journal	Journal	1	2
28	IEEE International Conference on Parallel and Distributed Processing with Applications, Big Data and Cloud Computing, Sustainable Computing and Communications, Social Computing and Networking	Conference	1	2
29	IEEE Symposium on Reliable Distributed Systems	Conference	1	2
30	IEEE Transactions on Information Forensics and Security	Journal	1	2
31	IEEE Transactions on Knowledge and Data Engineering	Journal	1	2
32	IEEE Transactions on Network Science and Engineering	Journal	1	2
33	IEEE Transactions on Services Computing	Journal	1	2
34	IEEE Wireless Communications	Journal	1	2
35	Information Sciences	Journal	1	2
36	International Conference on Distributed Computing Systems	Conference	1	2
37	International Conference on Smart Systems and Technologies	Conference	1	2
38	International Conference on Software Engineering	Conference	1	2
39	International Symposium on Mobile Ad Hoc Networking and Computing	Conference	1	2
40	International Workshop on Security and Privacy in Big Data	Workshop	1	2
41	International Workshop on Social Sensing	Workshop	1	2
42	LNCS 7299 – Intelligence and Security Informatics	Journal	1	2
43	Online Information Review	Journal	1	2
44	Security and Communication Networks	Journal	1	2
45	Sensors	Journal	1	2
46	Transportation Research Part C: Emerging Technologies	Journal	1	2
47	Workshop on the Economics of Networks, Systems and Computation	Workshop	1	2
	Total		**50**	**100**

Appendix F. Electronic data sources and inclusion and exclusion criteria

Table F.9: Electronic data sources (SDS) used in automated search.

ID	Name (Acronym)	Website
EDS1	IEEE Xplore (IEEE)	https://ieeexplore.ieee.org/
EDS2	ACM Digital Library (ACM)	https://dl.acm.org/
EDS3	ISI Web of Science (WoS)	https://www.webofknowledge.com
EDS4	ScienceDirect (SD)	https://www.sciencedirect.com/
EDS5	SpringerLink (SL)	https://link.springer.com/
EDS6	Wiley InterScience (WIS)	https://onlinelibrary.wiley.com/
EDS7	SCOPUS (SCOPUS)	https://www.scopus.com/

Table F.10: Selection criteria used to identify relevant papers. Fulfilling only one exclusion criterion discards the publication from being included.

ID	Facet	Inclusion criterion	Exclusion criterion
F1	Coarse focus	The privacy and data market topic must be within the field of computer science and technology	Any other privacy and data market subfield
F2	Narrow focus	The paper must explicitly focus on privacy within data marketplaces within the defined applications	The paper does not explicitly address this research direction
F3	Publication channel type	Conference publication OR journal publication (full text) OR workshop publication	The paper is any other type of publication
F4	Language	English	Non-English
F5	Duplicates	Publications are new to the filtering process	Publication has already been processed
F6	Peer-review	The publication has been peer-reviewed	The publication is a grey publication
F7	Full-text access	TUM-Access granted	TUM-Access not granted

Appendix G. Figures of the metadata analysis

Figure G.12: Map of most active countries in the field of privacy-enhancing data markets for the IoT research.

Table G.11: Classification scheme of research types as described by [212].

Research type	**Description**
Evaluation research	The authors implement existing techniques, and the solutions are evaluated in practice.
Philosophical papers	These studies present a new perspective on existing research by organizing the domain into a taxonomy or a conceptual framework.
Solution proposal	The authors propose a solution to a problem. The solution can be either novel or a significantly enhanced version of an existent technique. A small example or argumentation demonstrates the benefit and applicability of the solution.

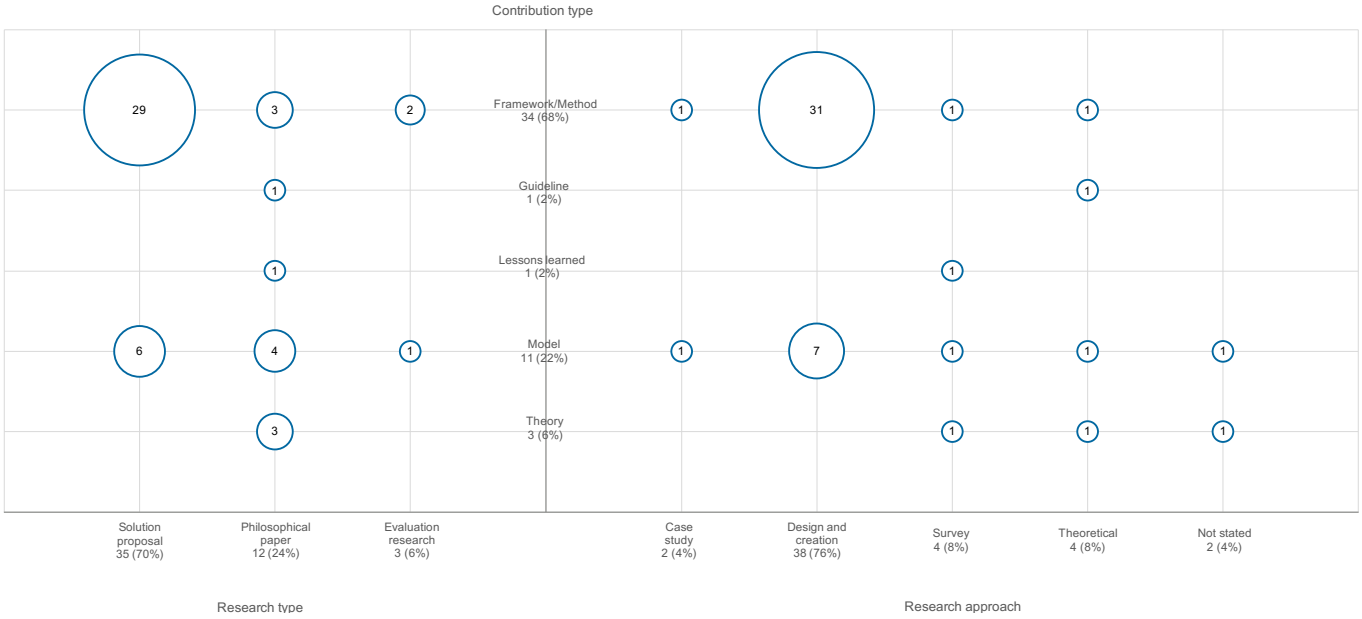

Figure G.13: Mapping of contribution types against research types and approaches.

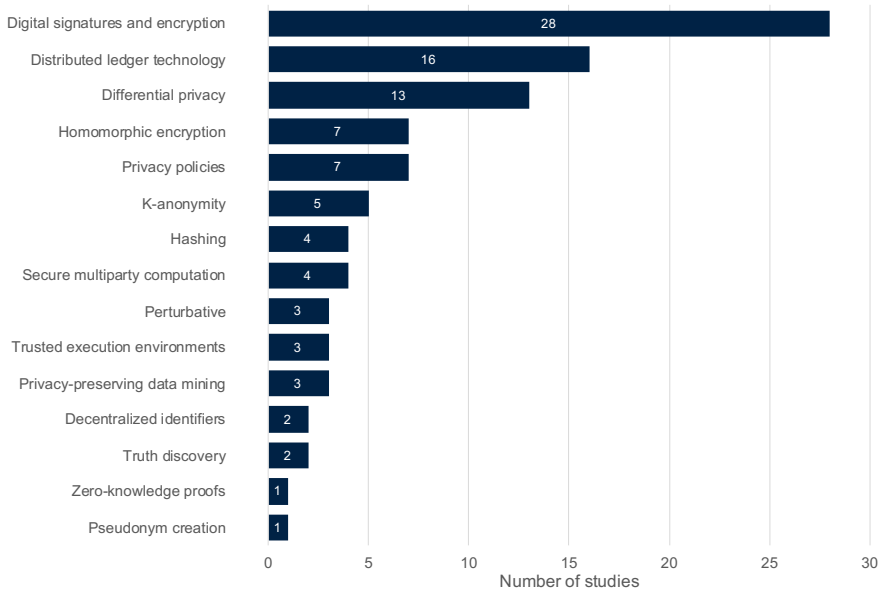

Figure G.14: Distribution of PETs and AETs explicitly employed in the corpus of selected studies.

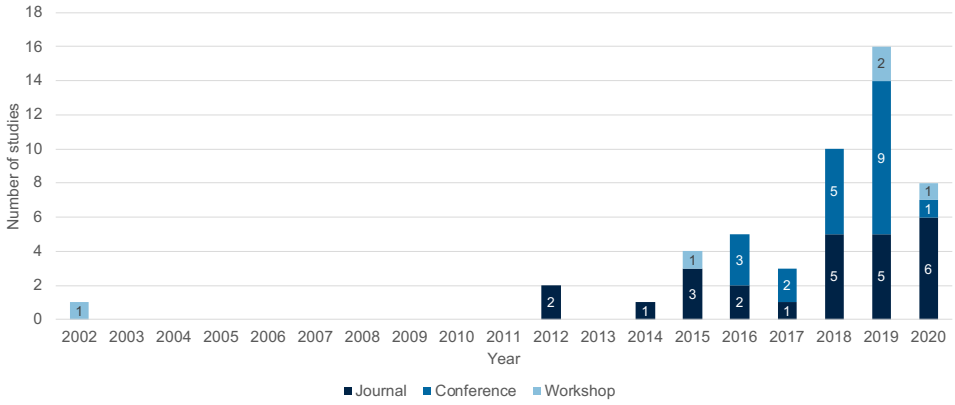

Figure G.15: Distribution of studies over publication domains.

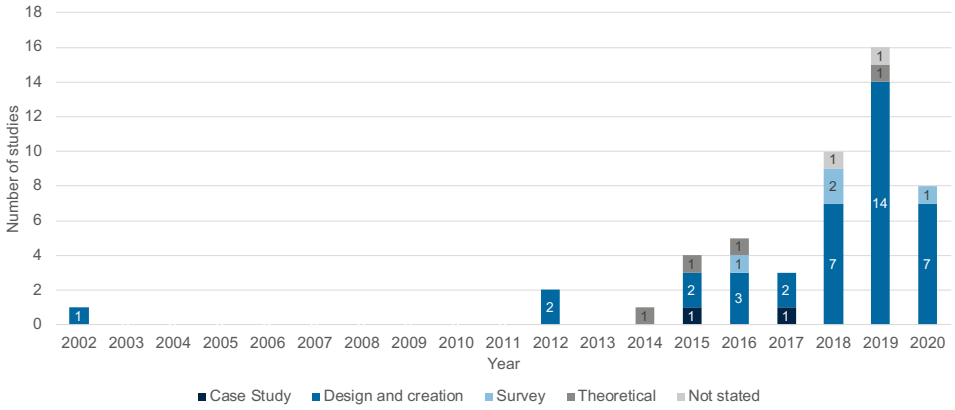

Figure G.16: Number of studies per research approach over time.

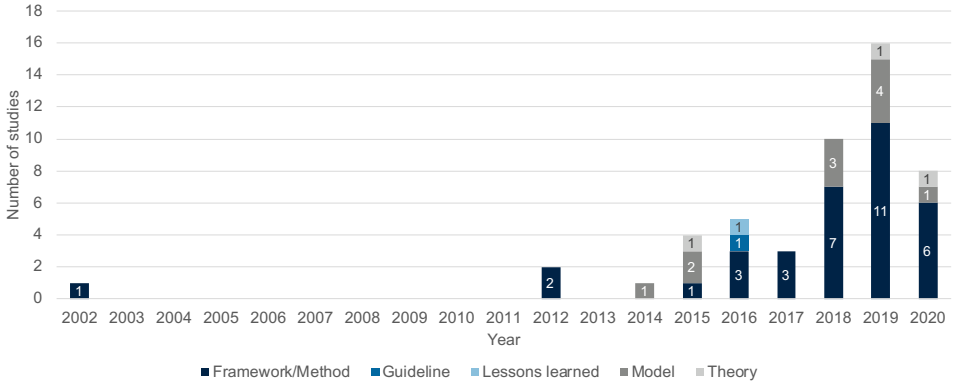

Figure G.17: Distribution of research outcomes over time.

Towards Verifiable Differentially-Private Polling

Gonzalo Munilla Garrido*
gonzalo.munilla-garrido@tum.de
Technical University of Munich
Garching, Germany

Matthias Babel
matthias.babel@fim-rc.de
FIM Research Center,
University of Bayreuth
Bayreuth, Germany

Johannes Sedlmeir
johannes.sedlmeir@fim-rc.de
Fraunhofer FIT, Branch Business &
Information Systems Engineering
Bayreuth, Germany

ABSTRACT

Analyses that fulfill differential privacy provide plausible deniability to individuals while allowing analysts to extract insights from data. However, beyond an often acceptable accuracy tradeoff, these statistical disclosure techniques generally inhibit the verifiability of the provided information, as one cannot check the correctness of the participants' truthful information, the differentially private mechanism, or the unbiased random number generation. While related work has already discussed this opportunity, an efficient implementation with a precise bound on errors and corresponding proofs of the differential privacy property is so far missing. In this paper, we follow an approach based on zero-knowledge proofs (ZKPs), in specific succinct non-interactive arguments of knowledge, as a verifiable computation technique to prove the correctness of a differentially private query output. In particular, we ensure the guarantees of differential privacy hold despite the limitations of ZKPs that operate on finite fields and have limited branching capabilities. We demonstrate that our approach has practical performance and discuss how practitioners could employ our primitives to verifiably query individuals' age from their digitally signed ID card in a differentially private manner.

CCS CONCEPTS

• **Information systems** → *Electronic data interchange*; • **Security and privacy** → **Cryptography**; *Human and societal aspects of security and privacy*; **Privacy-preserving protocols**.

KEYWORDS

Digital wallet, exponential noise, privacy, randomized response, SNARK, survey, zero-knowledge proof

ACM Reference Format:
Gonzalo Munilla Garrido, Matthias Babel, and Johannes Sedlmeir. 2022. Towards Verifiable Differentially-Private Polling. In *The 17th International Conference on Availability, Reliability and Security (ARES 2022), August 23–26, 2022, Vienna, Austria.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3538969.3538992

*Corresponding author.

ARES 2022, August 23–26, 2022, Vienna, Austria

ACM ISBN 978-1-4503-9670-7/22/08…$15.00
https://doi.org/10.1145/3538969.3538992

1 INTRODUCTION

Gathering information through polls to produce statistics regarding, e.g., the health, financial status, or demographics of a population bears the risk of exposing individuals' sensitive data during and after the survey. One approach is to anonymize the gathered data centrally, which implies high costs for implementing security measures and still carries ethical risks. Moreover, since interviewees cannot control that their data is adequately anonymized and protected in this paradigm and their level of trust in the surveyor is limited, their response may be subject to bias, specifically with highly sensitive or embarrassing questions.

A simple means to enhance privacy by design and reduce the risk of bias in such polls is through randomized response [57], its variations [6, 25, 30, 41, 58], or more involved forms of local differential privacy (DP) [3, 31–33, 37, 38], which provide plausible deniability by adding noise to interviewees' answers. The noise distribution of these techniques is typically centered around 0 with finite variance [14], so according to the law of large numbers, the mean of the noisy data converges towards the original mean as the sample size increases, improving accuracy. However, in this local approach, there is a lack of response *verifiability* – the interviewer has no assurance that the interviewee, i.e., the adversary in our system, answered truthfully. This lack of verifiability is arguably severe when rewards encourage malicious participation, e.g., there is a monetary incentive to participate but no willingness to answer truthfully.

Verifiable computation can prove the execution of a particular algorithm from truthful inputs without revealing private information [4]. Accordingly, we suggest combining verifiable computation with local differential privacy (LDP) techniques to prove (i) the interviewees' plausible deniability guarantee derived from randomness and (ii) the truthfulness of their deterministic answer, i.e., the value has been signed by a reputed authority. Similar approaches have been discussed, for instance, in [44, 47]. Our approach hence targets polls where there is cryptographic evidence for the answers, e.g., a digital ID card signed by a government, digital diplomas issued by a certified university, or COVID-19 immunity passports certified by pharmacies or doctors. Such attestations are considered, for instance, in the European digital wallet initiative [15, 46, 48]. In this context, it is particularly helpful that the digital certificates involved in the many implementations of digital wallets are in fact anonymous credentials [7, 50, 52], which allows us to extract a user's attribute values without revealing strongly correlating information. Moreover, we believe that in the private sector, attestations derived, for instance, from cryptographically signed statements of bank accounts or insurance claims and their use in verifiable differentially private surveys could have considerable economic

potential, as data markets require technologies that provide verifiability despite privacy protection [20].

There are two main approaches for verifiable computation: trusted execution environments (TEEs) [45], and non-interactive or interactive zero-knowledge proofs (ZKPs) [4, 23, 53]. Given the numerous known vulnerabilities and attacks on TEEs [1, 40, 54] and Intel's SGX SDK deprecation [35], we decided to focus on ZKP-based approaches. Moreover, non-interactive ZKPs do not require to engage into sequential messaging, so – unlike with interactive ZKP – the prover can convince multiple parties of a claim with a single message [53]. Thus, we opted to use non-interactive ZKPs to enable the verifiability of the computational integrity in the selected DP mechanism.

This paper's scope covers both binary answers, e.g., "*Are you older than* 18*?*", and numerical answers, e.g., "*How old are you?*". We provide plausible deniability for interviewees with differential privacy (DP) mechanisms in the local model, specifically, employing randomized response [57] and exponentially distributed noise [12]. Lastly, we adapt these mechanism such that we can verify their correct execution with ZKPs by employing succinct non-interactive arguments of knowledge (SNARKs) [5, 27], resulting in the primitives represented in Algorithms 1 and 2. We implement the corresponding circuits and evaluate their performance characteristics to assess our approach's practicality.[1]

As randomized response and exponential noise are building blocks for other more complex mechanisms, our scheme could also be extended to prove their verifiability, such as in two-stage randomized response models [30, 41], unrelated question models [25], forced response models [6], LDP models [31–33, 37, 38], private weighted histogram aggregation in crowdsourcing by leveraging multivariate randomized response [58], building histograms [3], or using exponential noise distributions in the central model of DP. Such verifiable forms of DP are also relevant in multilateral protocols that provide economic incentives for participation based on the participants' contribution. In such settings, one should compute fair rewards from the original data without noise, requiring that the computation of both their deterministic contribution and the shared noisy value is verifiable. An example for such a scenario is fair blockchain-based federated learning, studied by Rückel et al. [47].

This paper is structured as follows. We provide preliminaries in Section 2, discuss the SNARK-based approach and its implementation in Section 3, and evaluate it in Section 4. Lastly, we comment on related work in Section 5, discuss our approach in Section 6, and conclude the paper in Section 7.

2 PRELIMINARIES

2.1 Differential Privacy

We consider a collection of records from a population (dataset) D to belong to the universe of possible datasets $\mathcal{D}$. We let $D' \sim D$ denote neighboring datasets, i.e., D and D' differ by only one record. Differential privacy, introduced by Dwork et al. in 2006 [13], formalizes a mathematical definition of privacy whereby an analysis' output distribution is nearly the same across all neighboring datasets. The indistinguishability between datasets is parameterized by $\varepsilon > 0$. The higher ε, the easier it is to identify datasets.

Definition 1. *((ε, δ)-Differential Privacy [14]). A randomized mechanism $\mathcal{M}$ is (ε, δ)-differentially private iff for any neighboring dataset $D' \sim D$, and any set of possible outputs $\mathcal{S} \subseteq Range(\mathcal{M})$,*

$$Pr[\mathcal{M}(D) \in \mathcal{S}] \leq e^{\varepsilon} \cdot Pr[\mathcal{M}(D') \in \mathcal{S}] + \delta.$$

Having a non-zero δ relaxes the strict ε bound for possible but unlikely events; this type of guarantee is called *approximate* DP, whereas with $\delta = 0$, we obtain *pure* DP. A randomized mechanism $\mathcal{M}$ typically ensures DP by adding carefully calibrated random noise to the output of a deterministic function $f(\cdot)$, for example, by adding exponentially distributed noise [12] to a count, average, or median. Furthermore, another factor beyond ε that calibrates noise is the sensitivity of $f(\cdot)$, which measures the maximum variation of the output as the input dataset D changes (denoted as Δ).

Lastly, it is important to note that a DP mechanism $\mathcal{M}$ follows *sequential composition* [14], i.e., if $\mathcal{M}$ is computed n times over a dataset $\mathcal{D}$ with ε_i, in effect, the total ε is given by $\sum \varepsilon_i$. Thus, the results become less private with every query. Yet, a system can effectively impede an attacker from averaging out the noise through a sequence of DP results by blocking subsequent queries or deterministically generating the randomness based on the query parameters.

2.2 Local Differential Privacy

Definition 1 corresponds to the *central model*, in which a trusted curator collects data points and adds noise from a distribution whose variance is tuned by the function's sensitivity (Δ) and the required degree of plausible deniability (ε). In this paper, we focus on the *local model*, whereby the data subject obfuscates the data points directly before sharing. While the local model typically provides less accuracy than the central model, the data subject does not need to trust the curator from a privacy perspective.

For the local setting, we adopt similar notation to [22, 36]. We let $\mathcal{X}$ contain all the possible records of a population, where $x \in \mathcal{X}$ holds a particular individual's data. We let $f : \mathcal{X} \rightarrow [l, u]$ be a function that maps each element of $x \in \mathcal{X}$ to $f(x) \in [l, u]$, where $[l, u]$ is the set of integers between the lower bound $l \in \mathbb{N}$ and the upper bound $u \in \mathbb{N}$. In practice, $f(\cdot)$ provides information about an individual, e.g., their age, income, height, or blood pressure. Let $\mathcal{M} : f(x) \rightarrow [0, n]$ denote a randomized mechanism that maps each deterministic query output $f(x) = i \in [l, u]$ to each possible value $j \in [l, u]$ following a probability distribution that depends on the value i. In this setting, a randomized mechanism $\mathcal{M}$ provides local (ε, δ)-DP iff for every pair of inputs $x, x' \in \mathcal{X}$, and for every possible output $j \in [l, u]$:

$$\Pr[\mathcal{M}(x, f) = j] \leq e^{\varepsilon} \cdot \Pr[\mathcal{M}(x', f) = j] + \delta.$$

We will cover two mechanisms that satisfy local DP: randomized response [57] for binary data and exponentially distributed noise [12] for numerical data.

Randomized response. Warner [57] introduced randomized response in 1965 to provide plausible deniability to interviewees, which encouraged them to answer truthfully, thus reducing bias in surveys. The interviewees would answer queries $f(\cdot)$ of the

[1] The source code can be found at https://github.com/applied-crypto/DPfeatZKP.

form "*Are you a* [...]?". The following algorithm is well-known to be (ln 3, 0)-differentially private [14]:

(1) Flip a coin.
(2) If heads, answer truthfully.
(3) Else, flip the coin again and answer "Yes" if heads and "No" otherwise.

Exponentially distributed noise. Early work by Dwork et al. [12] shows that noise distributed as $\Pr[x] \propto \exp(-\frac{\varepsilon|x|}{\Delta})$ fulfills $(\varepsilon, 0)$-DP. We leverage exponentially distributed noise for locally obfuscating numerical data to answer queries $f(\cdot)$ of the form "*How many* [...]?", i.e., count queries, by adding noise to the deterministic output of $f(\cdot)$ in this manner: $\mathcal{M}(x, f) = f(x) + \textbf{noise}$. As the responses are local, we must ensure indistinguishability between any $f(x)$ and $f(x')$; thus, we must set Δ according to the output range of $f(\cdot)$. In practice, to ensure that an attacker cannot easily distinguish individuals in the extreme-case scenario, e.g., a newborn from a 128-year old person, with the query "*How old are you*?," we set $\Delta = |\max_x f(x) - \min_x f(x)| = |128 - 0|$.

As presented above, both mechanisms fulfill pure DP; however, generating exponentially distributed noise is subject to approximation errors in practical implementations and, thus, we only achieve approximate DP instead [21]. We shed more light on this issue in Section 3.

2.3 Proof Systems and SNARKs

There have been several key milestones in the work towards cryptographically verifiable computations. Babai [2] studied interactive proofs between a *prover* and a *verifier* and analyzed which problems can be checked by a polynomially bounded verifier when adding randomness and interaction. Fiat and Shamir [17] then introduced a heuristic to replace the verifier with a random oracle, which one can implement with a secure hash function. Nonetheless, one uses the word "argument" instead of "proof" in this case because the existence of a secure hash function has not been proven mathematically so far and is rather a working hypothesis, also bound to today's compute and (differential) cryptanalysis capabilities. Goldwasser et al. [23] proved that one could verify a large class of problems probabilistically in this way, where the verifier additionally does not even need to learn anything beyond the statement's correctness. While practical applications of the accordingly termed ZKPs were rare after these early developments, a period of rapid improvements started in the mid-2000s and led to the computationally efficient (quasi-linear complexity) generation of succinct arguments of logarithmic size and verification time, called SNARKs [5, 27].

The research community has since developed other flavors such as scalable transparent arguments of knowledge (STARKs), which differ in the setup procedure, proof size, and cryptographic assumptions, but have similar functional aspects. These new approaches allow succinct or scalable zero-knowledge proofs for the correctness of arbitrary statements and, thus, practical verifiable computation, i.e., proofs for the correct execution of a program without displaying all inputs, outputs, or intermediate steps. Additionally, one can reveal Merkle proofs or other cryptographic relations like the public keys corresponding to a digital signatures by a reputed institution on the inputs to force the prover to use unknown but fixed variables. Arguably, the core area of application of ZKPs today is in distributed ledgers, where because of redundant execution, cheap (succinct) verification without revealing sensitive data is important to solve scalability issues and mitigate excessive data visibility [4, 51].

Table 1: Notation

$\mathcal{X}$	Universe of records
$x \in \mathcal{X}$	Individual record
$f : \mathcal{X} \to [l, u]$	Query function
l	Lower bound, $\min_x f(x)$
u	Upper bound, $\max_x f(x)$
$\mathcal{M} : [l, u] \to [l, u]$	Randomized mechanism
Δ	Sensitivity of $f(\cdot)$, $\lvert u - l \rvert$
ℓ	Noise added to $f(\cdot)$
`nBits`	Number of bits representing ℓ
p_k	Bias of bit k, $\Pr[\ell_k = 1]$
d	Precision

3 IMPLEMENTATION

In this section, we first describe how to adapt standard implementations of uniform randomness generation, randomized response, and exponentially distributed noise (see Section 2) such that ZKPs can verify their use. For our implementation, we employ Circom, a well-known, open-source technology stack for implementing ZKPs [34]. Circom is a domain specific programming language and compiler that translates JavaScript-like arithmetic circuits in a rank one constraint system (R1CS), on behalf of which further libraries (e.g., SnarkJS) can generate SNARKs. A circom-specific variable type to explicitly define constraints is called *Signal*. The programming of these signals is restricted by the underlying quadratic arithmetic program (QAP), which the R1CS encodes, to use only quadratic constraints inside one Signal. Therefore, a Signal can only be assigned once and is immutable. For this reason, calculations often have to be split into multiple sub-calculations. Moreover, branchings and loops can only be used in a restricted way, for instance, the maximum number of iterations must be specified by a constant instead of a variable or Signal.

In what follows, we will use the roles prover for the survey participant and verifier for the surveyor, and the notation specified in Table 1. We also assume that both of them have a dedicated key-pair that they can use for end-to-end encrypted, authenticated communication and for recognizing each other, or another means to bind them to a specific secret key, for instance, through an anonymous credential with private holder binding [7]. Such anonymous credentials can be implemented with specific-purpose ZKPs [7] and with SNARKs [11, 49]. Key-pairs are a common way to facilitate the generation of verifiable randomness, for instance, in Algorand's consensus mechanism [10].

3.1 Verifiable Uniform Randomness

To achieve uniform randomness that cannot be spoiled unilaterally by either the prover or the verifier, we employ two inputs and a hash function as a random oracle [8]. More specifically, we sign a challenge that the verifier specifies with the prover's private

Algorithm 1: Verifiable randomized response for binary data and uniform randomness ("unbiased coins").

Data: v: binary truthful value ("Yes" or "No"); a: prover contribution to randomness (secret key); b: verifier contribution to randomness (challenge).
Result: Differentially private answer.

1 **Function** VerifiableUnifRand(a, b):
2 s = sign(a, b) // sign challenge with secret key
3 r = hash(s) // r is an array of bits
4 **return** r
5 **Function** VerifiableRandomizedResponse(v, a, b):
6 r = VerifiableUnifRand(a, b)
7 **if** $r[0] = 0$ **then**
8 **return** v
9 **else if** $r[1] = 0$ **then**
10 **return** No
11 **else**
12 **return** Yes

key and hash the result. As the private key is determined by the fixed prover's public key, neither of the two parties can bias the resulting randomness without collusion. We use Poseidon[2] – a relatively new hashing algorithm that was specifically developed for use in ZKPs and that is already being used in many blockchain-based applications on Ethereum and, therefore, to some extent battle-tested [24]. We represent this building block as a function in Algorithm 1 between lines 1 and 4, using existing components in Circom for EdDSA signature verification, Poseidon, and conversion of (large) integers to binary representation. Assuming that the Poseidon hash function is a random oracle and the keypair was created without anticipating the survey and the verifier's challenge, this gives us an array of 254 unbiased random bits.

3.2 Verifiable Randomized Response

Randomized response is simple to verify with ZKPs by utilizing the verifiable uniform randomness function (see Algorithm 1). In practice, without loss of generality, we only consider the least two significant bits of the random number generated. For the randomized response algorithm presented in Section 2 and presented formally in Algorithm 1, we need to sample at least once (last bit) and at most twice (second-last bit), depending on the first coin flip. The source code from Fig. 1 implements this in Circom. As if-statements are not natively possible in R1CS and, therefore, only available with restrictions in Circom, we arithmetize the corresponding statements in lines 7 to 12 from Algorithm 1 into the lines 39 to 40 from Fig. 1.

3.3 Verifiable Exponentially Distributed Noise

The exponentially distributed noise adaptation to ZKPs is not as straightforward as with randomized response because it typically involves floating point operations and rounding. After trying different implementations of exponentially distributed noise generation – we briefly cover the journey in Section 6 –we successfully adapted the method proposed by Dwork et al. [12] to ZKP, which we present in Algorithm 2: In their method, Dwork et al. approximated exponentially distributed noise of $\Pr[x] \propto \exp(-\frac{\varepsilon|x|}{\Delta})$ with the Poisson

[2]Using other hashing mechanisms is possible, yet the performance can become considerably worse – for instance, in the case of SHA256, around 30x.

```
1  pragma circom 2.0.0;
2
3  include "./poseidon.circom";        // Poseidon hashing
4  include "./bitify.circom";          // Bit array conversion
5  include "./eddsaposeidon.circom";   // Signature checking
6
7  template Main() {
8     signal input value;         // v
9     signal input challenge;
10    signal input R8[2];         // elliptic curve element of
          signature
11    signal input S;             // field element of signature
12    signal input pk[2];         // public key
13
14    // check signature on challenge against public key
15    component eddsaVerifier = EdDSAPoseidonVerifier();
16    eddsaVerifier.Ax <== pk[0];
17    eddsaVerifier.Ay <== pk[1];
18    eddsaVerifier.S <== S;
19    eddsaVerifier.R8x <== R8[0];
20    eddsaVerifier.R8y <== R8[1];
21    eddsaVerifier.M <== challenge;
22    eddsaVerifier.enabled <== 1;  // checks signature implicitly
23
24    // hash signature and convert this randomness to bit array
25    component hash = Poseidon(3);
26    component bitify = Num2Bits_strict();
27    hash.inputs[0] <== R8[0];
28    hash.inputs[1] <== R8[1];
29    hash.inputs[2] <== S;
30    bitify.in <== hash.out;
31    signal randSeq[254];
32    for(var i = 0; i < 254; i++) {
33       randSeq[i] <== bitify.out[i];
34    }
35
36    // determine result from randomness
37    signal rand;
38    signal output out;
39    rand <== randSeq[0] * randSeq[1];
40    out <== (1 - randSeq[0]) * value + rand;
41
42 }
43
44 component main {public [challenge, pk]} = Main();
```

Figure 1: Circom code for a component that implements verifiable randomized response.

distribution, fulfilling (ε, δ)-DP. Their method samples noise by producing a sequence of biased bits equal in number to the number of bits in the binary expansion of the noise ℓ. The algorithm flips an extra bit to add a sign ($\pm\ell$). The bias of each bit $k \in \{0, ..., \texttt{nBits}\}$ representing ℓ in binary is given in Section 4.1 of [12] by

$$\Pr[\ell_k] := \Pr[\ell_k = 1] = \left(1 + \exp\left(\frac{\varepsilon \cdot 2^k}{\Delta}\right)\right)^{-1}.$$

To generate biased bits from unbiased bits, we include in Algorithm 2 a well-known technique: first, we expand in binary the bias p_k of a bit k. Afterward, the algorithm sequentially examines random unbiased bits until one differs from the corresponding bit in the binary expansion of p_k and, subsequently, outputs the complement of the random unbiased bit [12]. Essentially, this approach allows to simulate biased coins up to a pre-defined precision with unbiased coins. However, the method employed has three limitations.

The *first* limitation entails several issues that relate to representing with a limited precision d the bias of the bits composing p_k,

i.e., d is the number of bits available for representation. Nonetheless, the probability of the inner loop not terminating for $j < d$ and, therefore, raising an error decays with 2^{-d}. Thus, we can easily choose d such that the likelihood of this event is negligible (line 17 of Algorithm 2). Furthermore, we show that we can provide enough precision in our circuit: The randomness generated from a single Poseidon hash could provide a precision of around $2^{252} \approx 10^{75}$, i.e., $d \approx 75$. By using multiple rounds of hashing and signing, we could also generate more random bits and account for higher precision needs. Additionally, we restrict noise values ℓ to the interval $[l, u]$, where u and l are the deterministic function's output upper and lower bounds, respectively. For our experiment on polling individuals' age, we employed the algorithm with $d = 20$ and $\Delta = |u - l| = 128$. These example values require the algorithm to represent ℓ with `nBits`=7 bits and, in turn, generate one instance of noise with $d \cdot 7 + 1 = 141 < 256$ bits ("times d" because each of the 7 bits' bias will be expanded to d bits and one more for the sign). In other words, a single round of hashing and signature verification is sufficient (and would still be sufficient for $d = 35$, which corresponds to an error bound of $2^{-35} \approx 10^{-10}$ when approximating probabilities [12]), and a negligible probability to raise an error (upper bound $7 \cdot 2^{-20}$). We could also aim for the typical machine accuracy of 10^{-16} by using $d > 16 \cdot \log_2(10) \approx 53.1$, i.e., $d \geq 54$, which would involve the creation of two independent random bit arrays.

These design decisions allow us to approximate the Poisson distribution with an error bound that we can determine and control ex-ante, when designing the survey. Thus, we achieve (ε, δ)-DP with a statistical difference of $\delta = \texttt{nBits} \cdot 2^{-d} = 7 \cdot 2^{-20}$ [12]. Consequently, for improving the DP guarantee on ε, we only need to increase d. Moreover, the probability mass outside the considered interval $[-2^d, 2^d]$ is redistributed inside the interval, leading to an additional statistical difference of $2\exp(-(\varepsilon \cdot 2^d)/\Delta)$ that we let the term `nBits` absorb [12].

The *second* limitation is the zero probability assigned to noise values of a binary expansions with more bits than `nBits` (i.e., noise outside of $[l, u]$). Dwork et al. proposed to constrain the algorithm's output, i.e., deterministic answer + **noise**, to `nBits` and return the deterministic answer in case there is an overflow. According to Dwork et al. [12], as the distribution in the range $[l, u]$ is exponential, we maintain the same privacy guarantee by increasing the probability of not adding noise by a *trivial* amount (i.e., δ increases). We execute a modulo operation to remap any output value outside $[l, u]$ back in that range to reduce such an increase (lines 21 and 23 of Algorithm 2). Intuitively, a modulo operation on the output preserves DP as it is a post-processing step and, also, will re-distribute the outputs in the range instead of on one value. Formally, the proof may be found in Lemma 3 of Wang et al. [56].

The *third* limitation comes with flipping an unbiased bit to assign the sign of the noise, which converts a Poisson distribution into a two-sided distribution with double the probability on its center, i.e., of noise = 0. While Dwork et al. did not address this issue in [12], we could follow the approach of Champion et al. [9] of rejecting -0 and executing the algorithm again (section 3.3 of [9]). DP is maintained as the number of failures is independent of the noise. However, instead, to remain computationally performant, we output a uniformly sampled value within $[l, u]$ if -0 (lines 19 to 23 of Algorithm 2), effectively removing the excess probability at 0. Intuitively, we preserve DP by adding more noise to the output distribution. Formally, we provide this justification: Let P_{old} be some DP distribution on N discrete values and P_{new} such that $P_{\text{new}}(x) = \alpha P_{\text{old}}(x) + (1 - \alpha)/N$. Obviously, this is a probability distribution. In our case, $1 - \alpha$ is the probability of obtaining noise -0. For any D and D' that differ in at most one record,

$$
\begin{aligned}
&\Pr[\mathcal{M}_{\text{new}}(D) \in \mathcal{S}] \\
&= \alpha \cdot \Pr[\mathcal{M}_{\text{old}}(D) \in \mathcal{S}] + (1 - \alpha) \cdot \frac{|D|}{N} \\
&\leq \alpha e^{\varepsilon} \cdot \Pr[\mathcal{M}_{\text{old}}(D') \in \mathcal{S}] + \alpha\delta + (1 - \alpha) \cdot \frac{|D|}{N} \\
&= \alpha e^{\varepsilon} \cdot \left(\frac{\Pr[\mathcal{M}_{\text{new}}(D') \in \mathcal{S}]}{\alpha} - \frac{|D'| \cdot (1 - \alpha)}{N\alpha} \right) + \alpha\delta + (1 - \alpha) \cdot \frac{|D|}{N} \\
&= e^{\varepsilon} \cdot \Pr[\mathcal{M}_{\text{new}}(D') \in \mathcal{S}] + \alpha\delta + (1 - \alpha) \cdot \frac{|D| - e^{\varepsilon}|D'|}{N} \\
&\leq e^{\varepsilon} \cdot \Pr[\mathcal{M}_{\text{new}}(D') \in \mathcal{S}] + \alpha\delta + (1 - \alpha) \cdot \frac{1}{N}.
\end{aligned}
$$

Given that $\varepsilon \geq 0$, then $e^{\varepsilon} \geq 1$. Moreover, $|D| - |D'| \leq 1$ because they are neighboring. Thus, since we choose to re-distribute the excess weight for noise -0 ($\alpha = 1 - \frac{1}{2}\text{Prob}(0)$), δ may grow to at most $\frac{1}{2}\text{Prob}(0) \cdot \frac{1}{N}$. The above formulation is a universal upper bound: Essentially, it proves that the convex combination of an (ε, δ_1) mechanism and a uniform distribution with pointwise weight $\delta_2 = \frac{1}{N}$ (which is obviously $(0, \delta_2)$-DP) is (ε, δ)-DP, where δ is the convex combination of δ_1 and δ_2. Future work could focus on obtaining a tighter δ bound specifically for the Poisson distribution or employ the method from Champion et al. [9], which would maintain a δ upper bound of $\texttt{nBits} \cdot 2^{-d}$. Altogether, this approximation allows us to sample exponentially distributed noise preserving (ε, δ)-DP in a way that we can successfully verify with ZKPs. We depict an example of the resulting output distribution in Fig. 2.

As Circom does not allow for branching, i.e., implementing conditional checks and breaking or continuing loops, besides the workaround for if-statements, we had to introduce some additional Signals (see Fig. 3). These Signals allow us to determine the correct return value although all iterations from the loop are simulated in the circuit. We did this by introducing Signal arrays $\texttt{hit}_{1\ldots\text{nBits}}$ with length d that help to identify the firsts unequal pair of bits in one (j) loop, where the loop would break in Algorithm 2. This approach ensures that only the first occurrence of unequal bits (index i) is taken into account for the calculation of the biased randomness. Moreover, we multiply the inverted binary result from `isEqual`, which compares the corresponding bits of the random sequence r_j and the probability B_{j,p_k}, with the bit of the probability and the hit bit-value, which is 1 as long as the result of `isEqual` of the last iteration of the loop was 1, essentially $[1_0, \ldots, 1_{i-1}, 1_i, 0_{i+1}, \ldots, 0_{d-1}]$. Therefore, only at the first inequality of those two bits the probability bit is not multiplied with zero, and, thus, can be taken into account for the noise. In other words, $eval3[k][j]$ is the "running return value" after the $j + 1$st iteration of the loop, and is set to 1 only if the jth bit of the probability is one, the jth bit of the random sequence is 0, and the first time that the probability and random bit array are different occurs at position j as well.

Algorithm 2: Verifiable exponentially distributed noise generation for numerical data. By $x \bmod (l, u)$ we denote $l + (x \bmod (u - l))$.

Data: v: integer-valued truthful value; u: upper bound; l, lower bound; $\Delta = |u - l| \geq 0$: sensitivity of query function; $\varepsilon \geq 0$: privacy parameter; $d \geq 0$: precision of binary expansion; a: prover contribution to randomness (secret key); b: verifier contribution to randomness (challenge).
Result: $v + \text{noise} \sim \text{Pois}(v|\frac{\varepsilon}{\Delta})$.

Function VerifiableExponentialNoise($v, \Delta, \varepsilon, d$):
- $B_K = \text{BinaryExpansion}(\Delta)$
- $B_v = \text{BinaryExpansion}(v)$
- $B_r = [\,]$ // B_r stacks biased bits
- **for** $k \leftarrow 0$ **to** $\text{NumBits}(B_K)$ **do**
  - $p_k = \frac{1}{1+\exp(2^k \frac{\varepsilon}{\Delta})}$
  - $B_{p_k} = \text{BinaryExpansion}(p_k)$
  - // r has at least d bits
  - $r = \text{VerifiableUnifRand}(a, b)$
  - **for** $j \leftarrow 0$ **to** d **do**
    - // Where d is the least significant bit
    - $r_j = \text{r[j]}$ // $r_j \in \{0, 1\}$
    - **if** $r_j = B_{j,p_k}$ **then**
      - **continue**
    - **else**
      - $B_r.\text{push}(B_{j,p_k})$
      - **break**
    - **if** $j = d$ **then**
      - **return** RaiseError
- $\text{noise} = \text{DecimalExpansion}(B_r)$
- $\text{sign} = \text{VerifiableUnifRand}(a, b)[0]$
- **if** (noise = 0 and sign = 0) **then**
  - **return** $\text{DecimalExpansion}(\text{VerifiableUnifRand}(a, b)) \bmod (l, u)$
- **else**
  - **return** $[(v + (2 \cdot \text{sign} - 1) \cdot \text{noise}] \bmod (l, u)$

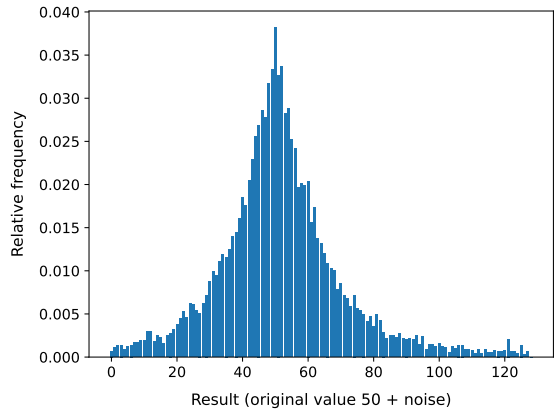

Figure 2: Example histogram for l = 0, u = 128, d = 20, $\varepsilon = 10$, and true value v = 50 with a sample size of 10 000.

```
1
2  // include statements as before, plus modulo component
3  template Main(nBits, d) {
4     signal input challenge, value
5     signal input prob[nBits][d];  // binary expansions of p_k
6     signal input R8[2], S, pk[2]; // signature and public key
7     // check the EdDSA signature of the challenge against pk and
            put it in the hash component to create randSeq, as in
            Figure 1 (lines 10 to 34).
8     ...
9
10    component isEqual[nBits][d];
11    signal noiseBits[nBits];
12    signal eval1[nBits][d];
13    signal eval2[nBits][d];
14    signal eval3[nBits][d + 1];
15    signal hit[nBits][d + 1];
16
17    // run the algorithm to create biased coins
18    for (var i = 0; i < nBits; i++) {
19       for (var j = 0; j < d; j++) {
20          isEqual[i][j] = IsEqual();
21       }
22    }
23    for (var k = 0;  k < nBits; k++) {
24       hit[k][0] <== 1;
25       eval3[k][0] <== 0;
26       for (var j = 0;  j < d; j++) {
27          isEqual[k][j].in[0] <== prob[k][j];
28          isEqual[k][j].in[1] <== randSeq[k * d + j];
29          hit[k][j + 1] <==
30          hit[k][j] * isEqual[k][j].out;
31          eval1[k][j] <== hit[k][j] * (1 - isEqual[k][j].out);
32          eval2[k][j] <== eval1[k][j] * prob[k][j];
33          eval3[k][j + 1] <== eval3[k][j] + eval2[k][j];
34       }
35       noiseBits[k] <== eval3[k][d];
36    }
37
38    component numify[2];
39    // compute exponential noise from its binary representation
40    numify[0] = Bits2Num(nBits);
41    for (var i = 0; i < nBits; i++) {
42       numify[0].in[i] <== noiseBits[i];
43    }
44    signal absNoise <== numify[0].out;
45    signal positiveNoise <== randSeq[nBits * (d + 3)] * (value +
            absNoise);
46    signal noisedResult <== (1 - randSeq[nBits * (d + 3)]) * (value
             - absNoise) + positiveNoise;
47    // generate uniformly distributed noise
48    numify[1] = Bits2Num(nBits);
49    for (var i = 0; i < nBits; i++) {
50       numify[1].in[i] <== randSeq[((d + 2) * nBits) + i];
51    }
52
53    component isZero = IsZero();  // check if noise == -0
54    isZero.in <== absNoise;
55    signal isUnif <== isZero.out * (1 - randSeq[nBits * (d + 3)]);
56    signal unif <== isUnif * numify[1].out;
57    signal result <== (1 - isUnif) * noisedResult + unif;
58
59    component modulo = Modulo();
60    modulo.in <== result;
61    modulo.mod <== 128;
62    signal output out <== modulo.out;
63 }
64
65 component main {public [challenge, pk]} = Main(7, 22);
```

Figure 3: Circom code for generating verifiable LDP noise with $l = 0$ and $u = 128$. Import statements, signal definitions, and EdDSA verification omitted (see also Figure 1).

3.4 Application: Verifiable Differentially-Private Polling with Anonymous Credentials

With the primitives presented in Algorithms 1 and 2, and an implementation of anonymous credentials with Circom, we can now implement verifiable, differentially private polling. Note that a Circom-based implementation of anonymous credentials allows to selectively disclose attributes from a digital certificate that corresponds to a Merkle tree with a signed root, and incorporate authenticity checks, private holder binding, expiration, revocation, and predicate proofs such as range proofs. We first sketch an example of a hypothetical setting. Digitally signed attestations of a person's attributes are stored in their "digital wallet" – a mobile application – in the form of an anonymous credential, which could contain personal information such as the holder's name, age, and gender, as well as a digital signature from an institution that the surveyor trusts regarding the authenticity of the information, e.g., a government or hospital. The digital wallet can respond to so-called proof requests [50] that include requirements from the verifier's side what the survey participant should prove. In our case, this could include the following requirements:

- Prove knowledge of (i) an authentic anonymous credential, issued by some institution, and (ii) knowledge of the secret key associated with the public key for the private holder, which is a binding included as one of the attributes in the anonymous credential.
- Prove that the anonymous credential is (i) not expired (range proof on expiration attribute) and (ii) not revoked (proof about set-inclusion or exclusion, referring to some public accumulator value as specified by the verifier).
- Reveal the result of our implementation of verifiable randomized response or exponentially distributed noise, applied to one of the (boolean or integer-valued) attributes in the credential. The attribute is represented by the issuer's signature on the anonymous credential, for instance, the attribute could be a leaf in a Merkle tree whose root is signed by the issuer.

In the case of a SNARK, the wallet (or the proof request) would also need to contain the structured reference string (proving key) generated in a setup procedure. It is important that while generally this proving key must be generated in a multi-party computation, in this case, it can be generated by the surveyor alone: Any party that knows how the structured reference string was created can fake proofs, but the privacy guarantees are not harmed in this case [18].

When the surveyor does not leak the "toxic waste" used for creating the structured reference string, the ZKP's soundness guarantee provides a chain of trust for the attribute, which is not directly revealed but modified through verifiable noise. Lastly, the verifier (surveyor) can cryptographically check that the attribute and the survey participant's secret key for private holder binding are used as private inputs for the LDP mechanism, and that the challenge as specified by the surveyor is used. We illustrate the survey process with anonymous credentials and verifiable differential privacy in Fig. 4.

4 EVALUATION

In this section, we discuss the performance and practicality of our approach for verifiable LDP. We restrict the discussion to verifiable exponentially distributed noise because it includes strictly more complex operations, so by demonstrating its reasonable performance, we can conclude that verifiable randomized response also is practical. The implementation process of our verifiable LDP approach with exponentially distributed noise was two-tiered. First, we implemented Algorithm 2 in Javascript and verified that it indeed yields an exponential distribution, where values that – after adding the LDP noise – exceed the range of the output are instead displayed without added noise. Secondly, we employed the Poseidon hash function to create a random Oracle that generates verifiable randomness jointly from the prover's and the verifier's input. Next, we implemented the corresponding circuits in Circom.

Our choice of Algorithm 2 and our choice of implementation as displayed in Fig. 3 yielded a highly efficient implementation for creating verifiable, Poisson-distributed LDP: Using the Poseidon hash function, our circuit has 5997 R1CS constraints. On an Ubuntu 20.04 virtual machine with 4 virtual cores that runs on a commercial standard Laptop (Dell Latitude 7400 with an Intel i7 8665U CPU), proof creation – which is typically the bottleneck for using ZKPs – takes around 2.2 s when using the Groth16 proof system [26] on the Barreto-Naehrig curve over a 254 bit prime field (bn128), Web assembly for witness generation, and Javascript for proof generation. The size of the proving and verification key are around 3.4 MB and 3.5 kB, respectively. These sizes suggest that proof generation would also be practical on a web-based mobile application, although future research should validate this assumption.

With an optimized tool, performance is even better: Using an optimized C++-based witness generation and a proof generation based on x86 Assembly for Intel processors ("Rapidsnark", see [28]), proof generation is reduced to only 140 ms. Because to date there is no available optimized tool for proof verification, this operation still takes around 0.8 s in Javascript. Proof verification in a complex, combined survey may even reduce complexity on the surveyor's side because the complexity of SNARK verification is not dependent on the complexity of the original computation for which the survey participant proves integrity. Moreover, we tested the deployment of a smart contract verifier on Ethereum, which could be used for blockchain-based, incentivized differentially private surveys and, therefore, general data protection regulation (GDPR)-compliant applications on personal data. We measured the smart contract's deployment cost at around $1,150,000$ gas and its invocation at around $300,000$ gas.

5 RELATED WORK

While there is current and extensive research in the local model of DP [3, 31–33, 37, 38] and in ZKPs [4, 5, 27], there are only few publications that bridge both technologies.

We identified four studies that are close to ours. Rückel et al. [47] propose an architecture to share weights from federated learning models in a verifiable DP manner and add verifiable noise to the private weights. However, their approach does not acknowledge that their discretization of the Laplace distribution only fulfills *approximate* DP, and does not propose a bound for δ or an approach that

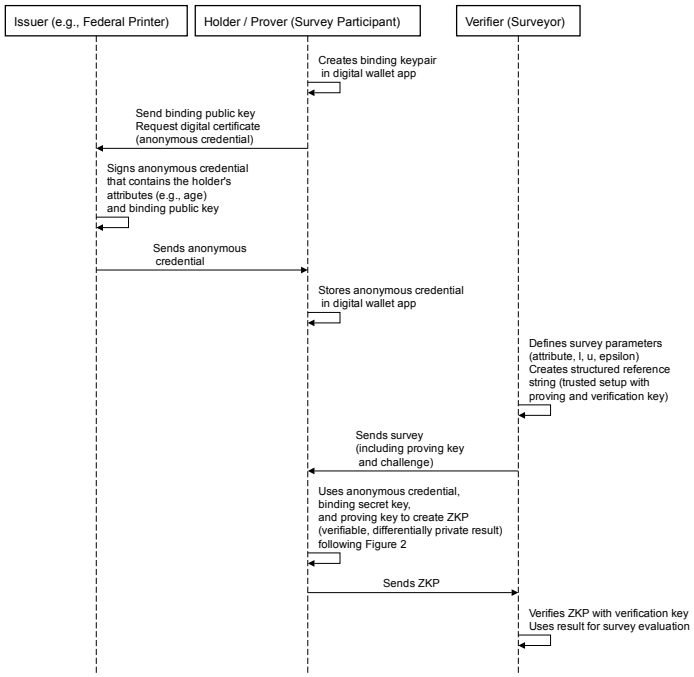

Figure 4: Process of participating in a survey with verifiable differential privacy.

works for high precision requirements, as they use the approximate inverse cumulative distribution function (CDF) as input for the circuit. Moreover, Tsaloli et al. [55] only provide a high-level motivation for using ZKPs for verifiable differential privacy, without implementation details. Furthermore, while Kato et al. [39] provide details on how to create fair randomness with a related technology (secure-multiparty computation), they do not attempt to make the result of the algorithm verifiable, i.e., they cannot provide a cryptographic check of the truthful value from, e.g., an anonymous credential, before adding LDP noise.

Lastly, Narayan et al. [44] discuss the opportunities of verifiable differential privacy, yet they provide no details of their ZKP-based implementation. For instance, they do not elaborate how they consider rounding and how they achieve guarantees on the accuracy of their verifiable *pure* DP proposal. Additionally, their approach focuses on the *central* model of DP instead and shows impractical performance, as it requires 2 hours of proof generation for 32 servers. Nonetheless, when implemented with more recent ZKP libraries, their performance may be closer to ours as advances in proving times over the last years have been dramatic.

6 DISCUSSION

While adapting randomized response to ZKPs is straightforward, we investigated several approaches before successfully implementing exponentially distributed noise with ZKPs. This section discusses the process we followed to arrive to the implementation described in Section 3.

Adapting a DP mechanism that leverages exponential noise to ZKP has two significant challenges. Conceptually, since ZKPs can verify an arbitrary program, sampling from an exponential distribution may seem straightforward. Unfortunately, in practice, repeated operations with floats that involve rounding in classical software are challenging to implement because the range of numbers in the nominator and denominator is bounded by a large prime, and repeated rounding is costly since the complexity of the ZKP always needs to account for the worst possible case. Furthermore, the propagation of the corresponding errors becomes challenging to control. Thus, the generality of computations that ZKPs can cover well is initially limited to arithmetic operations on prime fields and their corresponding primitives such as hash functions and signatures.

The second major challenge is the inability of finite computers to fulfill the definition of DP on the real line. Mironov [43] was the first to demonstrate that implementing a DP mechanism with the floating-point arithmetic of finite computers does not guarantee DP. Mironov proposed to solve this issue while maintaining ε-DP with the Snapping mechanism [43], and recently, Naoise et al. proposed secure random sampling [29]. However, while their output noise is discrete, we must still handle floats that ZKPs cannot process efficiently.

Therefore, we first thought of discretizing the support of the Laplace distribution (well-known to be ε-DP [14]) by sampling from its inverse CDF with a finite input range $\{1, ..., d\}$, where d denotes the precision Circom [34] can handle – similarly to Rückel et al. [47]. However, we were not able to determine provable guarantees on δ for the approximate DP mechanism. Thus, we turned to the *Stone-Weierstrass theorem* to approximate a polynomial so close to the Laplacian probability density function (PDF) that the approximation error would be negligible. Furthermore, because the approximated

PDF would be a polynomial itself, we thought to elegantly prove its use with ZKP. We employed Bernstein polynomials [16] to approximate the PDF in a closed interval and, subsequently, performed rejection sampling. However, we encountered two problems: (i) our approximation was limited to a closed interval, whereas the Laplacian PDF has unbounded support, and (ii) the Bernstein coefficients are in general real numbers, which ZKP cannot process efficiently, and the propagation of errors when rounding with fixed precision is again complex to handle. Specifically, the complexity stems from the very high degree of the polynomials and the lack of homogeneity, i.e., there are many different degrees of monomials that all scale differently for a specific accuracy when multiplying inputs with a large power of 10 and rounding afterward.

Subsequently, we turned to the truncated geometric mechanism (TGM) [22], which coincidentally has the advantage to provide better accuracy for count queries [19], the focus of this study. Additionally, the truncation solved the problem of working with a closed interval (also a limitation of finite computers) by condensing the probability mass outside the interval in its lower and upper bounds. Moreover, the support of the geometric mechanism are integers, which ZKPs can process efficiently. Overall, TGM adapts to finite computers while still providing pure ε-DP. However, the probabilities assigned to each integer still fall on the real line. While we ensured these probabilities became rational numbers by carefully choosing ε [3], which a conventional computer can handle, the integers necessary to represent them were too large for the limited precision available in Circom [34] and other libraries for implementing ZKPs, and we were unable to write a theoretical bound for δ if we approximated the real numbers with finite precision.

To cope with precision limitations and the difficulty to bound δ, we then looked for simple sampling methods that provide bounds on δ, which finally led us to Dwork et al. [12] (see Section 3). This concluded our search, as their method for sampling exponentially distributed noise consists on repeatedly flipping unbiased coins (which is easily implemented in Circom with hashing and conversion to bit arrays), and provides a bound for δ based on the precision we can afford with Circom.

In our implementation, we used verifiable randomness co-created by the verifier (surveyor) and the prover (survey participant). As we noted in Section 2, when a surveyor repeatedly conducts the query in our implementation with different challenges, they could get additional information because by the law of large numbers, the truthful query value without noise can be determined with increasing accuracy. Consequently, in many scenarios, it may be appropriate to use a challenge that is hard coded, derived from the surveyor institution's public key, or even derived solely from the attribute (e.g., the index of the age in the anonymous credential), such that repeated queries, even from different but colluding institutions, would not decrease the degree of plausible deniability ε. Furthermore, note that our verifiable randomized response implementation could be easily extended to flip biased coins by, e.g., generating a verifiable hash and checking whether its normalized value is lower than the desired bias.

7 CONCLUSION

We introduce primitives for implementing verifiable differentially private polls in the local setting. To achieve verifiability, we carefully selected DP mechanisms for binary and numerical data and adapted their implementations to SNARKs. Thanks to these primitives, we can achieve cryptographically verifiable survey responses while providing plausible deniability for survey participants and, in turn, not only reduce but entirely prevent bias in survey participants' answers while giving them the needed privacy guarantees. Furthermore, note that our primitive for verifiable exponentially distributed noise allows for different aggregation queries beyond the count, as it can ingest arbitrary sensitivity – we limited our narrative to count queries for the simplicity of the explanations. Finally, thanks to the evaluations we performed, we conclude that practitioners can deploy our primitives with acceptable performance

We encourage practitioners to develop further primitives that can adapt to other DP mechanisms, e.g., the exponential mechanism for categorical data [42], and other randomized-response [6, 25, 30, 41, 58] and LDP [3, 31–33, 37, 38] approaches. Furthermore, conducting studies about how interviewees would perceive the built-in trust would allow the research community to understand how to frame polls and reassure candidates of their privacy. Lastly, improving the precision limitations of ZKP circuit compilers such as Circom and more literature on frameworks for bounding δ in *approximate* LDP would open the range of practical LDP mechanisms.

ACKNOWLEDGMENTS

We thank the Bavarian Ministry of Economic Affairs, Regional Development and Energy for their funding of the project "Fraunhofer Blockchain Center (20-3066-2-6-14)", and the Ethereum foundation's grant that made this paper possible.

REFERENCES

[1] Fritz Alder, Jo Van Bulck, Jesse Spielman, David Oswald, and Frank Piessens. 2022. Faulty Point Unit: ABI Poisoning Attacks on Trusted Execution Environments. *Digital Threats: Research and Practice* 3, 2, Article 13 (2022), 26 pages. https://doi.org/10.1145/3491264

[2] László Babai. 1985. Trading Group Theory for Randomness. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*. 421–429. https://doi.org/10.1145/22145.22192

[3] Victor Balcer and Salil Vadhan. 2019. Differential Privacy on Finite Computers. *Journal of Privacy and Confidentiality* 9, 2 (2019). https://doi.org/10.29012/jpc.679

[4] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. 2019. Scalable Zero Knowledge with No Trusted Setup. In *Annual International Cryptology Conference*. Springer, 701–732. https://doi.org/10.1007/978-3-030-26954-8_23

[5] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. 2013. SNARKs for C: Verifying Program Executions Succinctly and in Zero Knowledge. In *Annual Cryptology Conference*. Springer, 90–108. https://doi.org/10.1007/978-3-642-40084-1_6

[6] Robert F. Boruch. 1971. Assuring Confidentiality of Responses in Social Research: A Note on Strategies. *The American Sociologist* 6, 4 (1971), 308–311. http://www.jstor.org/stable/27701807

[7] Jan Camenisch and Anna Lysyanskaya. 2001. An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 93–118. https://doi.org/10.1007/3-540-44987-6_7

[8] Ran Canetti, Oded Goldreich, and Shai Halevi. 2004. The Random Oracle Methodology, Revisited. *J. ACM* 51, 4 (2004), 557–594. https://doi.org/10.1145/1008731.10087340

[9] Jeffrey Champion, Abhi Shelat, and Jonathan Ullman. 2019. Securely Sampling Biased Coins with Applications to Differential Privacy. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*. ACM, 603–614. https://doi.org/10.1145/3319535.3354256

[10] Jing Chen and Silvio Micali. 2019. Algorand: A Secure and Efficient Distributed Ledger. *Theoretical Computer Science* 777 (2019), 155–183.

[11] Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, and Bryan Parno. 2016. Cinderella: Turning Shabby X.509 Certificates into Elegant Anonymous Credentials with the Magic of Verifiable Computation. In *Symposium on Security and Privacy*. IEEE, 235–254. https://doi.org/10.1109/SP.2016.22

[12] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. 2006. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *Advances in Cryptology* (2006), Serge Vaudenay (Ed.), Vol. 4004. Springer, 486–503. https://doi.org/10.1007/11761679_29

[13] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer, 265–284. https://doi.org/10.1007/11681878_14

[14] Cynthia Dwork and Aaron Roth. 2013. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3-4 (2013), 211–407. https://doi.org/10.1561/0400000042

[15] European Commission. 2021. Commission Proposes a Trusted and Secure Digital Identity for all Europeans. https://ec.europa.eu/commission/presscorner/detail/en/ip_21_2663

[16] Rida T. Farouki. 2012. The Bernstein Polynomial Basis: A Centennial Retrospective. 29, 6 (2012), 379–419. https://doi.org/10.1016/j.cagd.2012.03.001

[17] Amos Fiat and Adi Shamir. 1986. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In *Conference on the Theory and Application of Cryptographic Techniques*. Springer, 186–194. https://doi.org/10.1007/3-540-47721-7_12

[18] Georg Fuchsbauer. 2018. Subversion-Zero-Knowledge SNARKs. In *IACR International Workshop on Public Key Cryptography*. Springer, 315–347.

[19] Gonzalo Munilla Garrido, Joseph Near, Aitsam Muhammad, Warren He, Roman Matzutt, and Florian Matthes. 2021. Do I Get the Privacy I Need? Benchmarking Utility in Differential Privacy Libraries. http://arxiv.org/abs/2109.10789

[20] Gonzalo Munilla Garrido, Johannes Sedlmeir, Ömer Uludağ, Ilias Soto Alaoui, Andre Luckow, and Florian Matthes. 2021. Revealing the Landscape of Privacy-Enhancing Technologies in the Context of Data Markets for the IoT: A Systematic Literature Review. https://arxiv.org/abs/2107.11905

[21] Ivan Gazeau, Dale Miller, and Catuscia Palamidessi. 2016. Preserving Differential Privacy under Finite-precision Semantics. *Theoretical Computer Science* 655 (2016), 92–108. https://doi.org/10.1016/j.tcs.2016.01.0150

[22] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. 2012. Universally Utility-maximizing Privacy Mechanisms. *SIAM J. Comput.* 41, 6 (2012), 1673–1693. https://doi.org/10.1137/09076828X

[23] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. 1989. The Knowledge Complexity of Interactive Proof Systems. *SIAM J. Comput.* 18, 1 (1989), 186–208. https://doi.org/10.1137/0218012

[24] Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. 2021. Poseidon: A New Hash Function for Zero-Knowledge Proof Systems. In *30th USENIX Security Symposium*. 519–535. https://www.usenix.org/system/files/sec21-grassi.pdf

[25] Bernard G. Greenberg, Abdel-Latif A. Abul-Ela, Walt R. Simmons, and Daniel G. Horvitz. 1969. The Unrelated Question Randomized Response Model: Theoretical Framework. *J. Amer. Statist. Assoc.* 64, 326 (1969), 520–539. http://www.jstor.org/stable/2283636

[26] Jens Groth. 2016. On the Size of Pairing-based Non-interactive Arguments. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 305–326. https://doi.org/10.1007/978-3-662-49896-5_11

[27] Jens Groth, Rafail Ostrovsky, and Amit Sahai. 2006. Perfect Non-Interactive Zero Knowledge for NP. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 339–358. https://doi.org/10.1007/11761679_21

[28] Hermez Network. 2021. Open Sourcing an Ultra-fast zk Prover: Rapidsnark. https://blog.hermez.io/open-sourcing-ultra-fast-zk-prover-rapidsnark/

[29] Naoise Holohan and Stefano Braghin. 2021. Secure Random Sampling in Differential Privacy. In *Computer Security*, Elisa Bertino, Haya Shulman, and Michael Waidner (Eds.). Vol. 12973. Springer, 523–542. https://doi.org/10.1007/978-3-030-88428-4_26

[30] Naoise Holohan, Douglas J. Leith, and Oliver Mason. 2015. Differential privacy in metric spaces: Numerical, Categorical and Functional Data under the One Roof. *Information Sciences* 305 (2015), 256–268. https://doi.org/10.1016/j.ins.2015.01.021

[31] Naoise Holohan, Douglas J. Leith, and Oliver Mason. 2017. Extreme Points of the Local Differential Privacy Polytope. *Linear Algebra Appl.* 534 (2017), 78–96. https://doi.org/10.1016/j.laa.2017.08.011

[32] Naoise Holohan, Douglas J. Leith, and Oliver Mason. 2017. Optimal Differentially Private Mechanisms for Randomised Response. *IEEE Transactions on Information Forensics and Security* 12, 11 (2017), 2726–2735. https://doi.org/10.1109/TIFS.2017.2718487

[33] Justin Hsu, Sanjeev Khanna, and Aaron Roth. 2012. Distributed Private Heavy Hitters. In *Proceedings of the 39th International Colloquium Conference on Automata, Languages, and Programming – Volume Part I*. Springer, 461–472. https://doi.org/10.1007/978-3-642-31594-7_39

[34] iden3. 2018. Circom. https://docs.circom.io/

[35] Intel. 2022. 12th Generation Intel Core Processors. https://cdrdv2.intel.com/v1/dl/getContent/655258

[36] Lefki Kacem and Catuscia Palamidessi. 2018. Geometric Noise for Locally Private Counting Queries. In *Proceedings of the 13th Workshop on Programming Languages and Analysis for Security*. ACM, 13–16. https://doi.org/10.1145/3264820.3264827

[37] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. 2014. Extremal Mechanisms for Local Differential Privacy. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Eds.), Vol. 27. https://proceedings.neurips.cc/paper/2014/file/86df7dcfd896fcaf2674f757a2463eba-Paper.pdf

[38] Vishesh Karwa, Aleksandra B Slavković, and Pavel Krivitsky. 2014. Differentially Private Exponential Random Graphs. In *International Conference on Privacy in Statistical Databases*. Springer, 143–155. https://doi.org/10.1007/978-3-319-11257-2_12

[39] Fumiyuki Kato, Yang Cao, and Masatoshi Yoshikawa. 2021. Preventing Manipulation Attack in Local Differential Privacy Using Verifiable Randomization Mechanism. https://arxiv.org/abs/2104.06569

[40] Fatima Khalid and Ammar Masood. 2022. Vulnerability Analysis of Qualcomm Secure Execution Environment. *Computers & Security* 116 (2022), 102628. https://doi.org/10.1016/j.cose.2022.102628

[41] N. S. Mangat and Ravindra Singh. 1990. An Alternative Randomized Response Procedure. *Biometrika* 77, 2 (1990), 439–442. http://www.jstor.org/stable/2336829

[42] Frank McSherry and Kunal Talwar. 2007. Mechanism Design via Differential Privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 94–103. https://doi.org/10.1109/FOCS.2007.66

[43] Ilya Mironov. 2012. On Significance of the Least Significant Bits for Differential Privacy. In *Proceedings of the Conference on Computer and Communications Security*. ACM. https://doi.org/10.1145/2382196.2382264

[44] Arjun Narayan, Ariel Feldman, Antonis Papadimitriou, and Andreas Haeberlen. 2015. Verifiable Differential Privacy. In *Proceedings of the 10th European Conference on Computer Systems*. Article 28. https://doi.org/10.1145/2741948.2741978

[45] OMTP. 2009. Advanced Trusted Environment: OMTP TR1. , 204 pages. http://www.gsma.com/newsroom/wp-content/uploads/2012/03/omtpadvancedtrustedenvironmentomtptr1v11.pdf

[46] Alexander Rieger, Tamara Roth, Johannes Sedlmeir, Linda Weigl, and Gilbert Fridgen. 2022. Not Yet Another Digital Identity. *Nature Human Behaviour* 6, 1 (2022), 3–3. https://doi.org/10.1038/s41562-021-01243-0

[47] Timon Rückel, Johannes Sedlmeir, and Peter Hofmann. 2022. Fairness, Integrity, and Privacy in a Scalable Blockchain-Based Federated Learning System. *Computer Networks* 202 (2022), 108621.

[48] Sebastian Sartor, Johannes Sedlmeir, Alexander Rieger, and Tamara Roth. 2022. Love at First Sight? A User Experience Stury of Self-Sovereign Identity Wallets. In *Proceedings of the 30th European Conference on Information Systems*. AIS.

[49] Martin Schanzenbach, Thomas Kilian, Julian Schütte, and Christian Banse. 2019. ZKlaims: Privacy-preserving Attribute-Based Credentials using Non-Interactive Zero-Knowledge Techniques. https://arxiv.org/abs/1907.09579

[50] Vincent Schlatt, Johannes Sedlmeir, Simon Feulner, and Nils Urbach. 2021. Designing a Framework for Digital KYC Processes Built on Blockchain-Based Self-Sovereign Identity. *Information & Management* (2021), 103553. https://doi.org/10.1016/j.im.2021.103553

[51] Johannes Sedlmeir, Jonathan Lautenschlager, Gilbert Fridgen, and Nils Urbach. 2022. The Transparency Challenge of Blockchain in Organizations. *Electronic Markets* (2022). https://doi.org/10.1007/s12525-022-00536-0

[52] Johannes Sedlmeir, Reilly Smethurst, Alexander Rieger, and Gilbert Fridgen. 2021. Digital Identities and Verifiable Credentials. *Business & Information Systems Engineering* 63, 5 (2021), 603–613.

[53] Gerardo I Simari. 2002. A Primer on Zero Knowledge Protocols. (2002), 12. http://cs.uns.edu.ar/~gis/publications/zkp-simari2002.pdf

[54] Dimitrios Skarlatos, Mengjia Yan, Bhargava Gopireddy, Read Sprabery, Josep Torrellas, and Christopher W. Fletcher. 2019. MicroScope: Enabling Microarchitectural Replay Attacks. In *Proceedings of the 46th International Symposium on Computer Architecture*. ACM, 318–331. https://doi.org/10.1145/3307650.3322228

[55] Georgia Tsaloli and Aikaterini Mitrokotsa. 2019. Differential Privacy Meets Verifiable Computation: Achieving Strong Privacy and Integrity Guarantees. In *6th International Joint Conference on e-Business and Telecommunications*. 425–430. https://doi.org/10.5220/0007919404250430

[56] Lun Wang, Ruoxi Jia, and Dawn Song. 2022. D2P-Fed: Differentially Private Federated Learning With Efficient Communication. (2022). https://arxiv.org/abs/2006.13039

[57] Stanley L Warner. 1965. Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias. *J. Amer. Statist. Assoc.* 60, 309 (1965), 63–69. https://doi.org/10.1080/01621459.1965.10480775

[58] Qing Yang, Wei Yu, and Yacine Challal (Eds.). 2016. *Wireless Algorithms, Systems, and Applications*. Lecture Notes in Computer Science, Vol. 9798. Springer. https://doi.org/10.1007/978-3-319-42836-9

Lessons Learned: Surveying the Practicality of Differential Privacy in the Industry

Gonzalo M. Garrido
TUM, BMW Group
Germany
gonzalo.munilla-garrido@tum.de

Xiaoyuan Liu
UC Berkeley
USA
xiaoyuanliu@berkeley.edu

Florian Matthes
TUM
Germany
matthes@tum.de

Dawn Song
UC Berkeley
USA
dawnsong@berkeley.edu

ABSTRACT

Since its introduction in 2006, differential privacy has emerged as a predominant statistical tool for quantifying data privacy in academic works. Yet despite the plethora of research and open-source utilities that have accompanied its rise, with limited exceptions, differential privacy has failed to achieve widespread adoption in the enterprise domain. Our study aims to shed light on the fundamental causes underlying this academic-industrial utilization gap through detailed interviews of 24 privacy practitioners across 9 major companies. We analyze the results of our survey to provide key findings and suggestions for companies striving to improve privacy protection in their data workflows and highlight the necessary and missing requirements of existing differential privacy tools, with the goal of guiding researchers working towards the broader adoption of differential privacy. Our findings indicate that analysts suffer from lengthy bureaucratic processes for requesting access to sensitive data, yet once granted, only scarcely-enforced privacy policies stand between rogue practitioners and misuse of private information. We thus argue that differential privacy can significantly improve the processes of requesting and conducting data exploration across silos, and conclude that with a few of the improvements suggested herein, the practical use of differential privacy across the enterprise is within striking distance.

KEYWORDS

Privacy-enhancing technology, data sharing, data analytics, platform, SQL, machine learning, case study, interviews, survey

1 INTRODUCTION

Several factors have spurred the development of more advanced privacy-enhancing technologies (PETs) in the past years. On the one hand, from an adversarial perspective, (i) multiple white-hat *attacks* have shown that "traditional" anonymization techniques such as suppressing names are vulnerable to re-identification across industries [8, 29, 33, 67, 79, 94]. Additionally, between 2020 and 2021, (ii) the total cost of *data breaches* have increased by 10% on average [87]. Moreover, (iii) governments have promulgated *data protection laws* in the past years, such as the European General Data Protection Regulation (GDPR) [30] or the California Consumer Privacy Act [80]. In particular, the GDPR has issued fines as high as $887M [54] and $120M [20]. Furthermore, (iv) beyond the *ethical and moral obligations* of companies to protect people's privacy, providing the best privacy protection available could (v) *differentiate and appreciate their brands* [72], (vi) provide *fairer products and services* that avoid price discrimination [35], and (vii) *increase data collection* as PETs help to surmount regulatory barriers fairly [61]. Aiming to materialize these benefits while mitigating the privacy risks, researchers have turned to differential privacy (DP), which, since its inception in 2006 by Dwork et al. [28], has become the golden privacy standard in academia due to its unique privacy guarantees.

However, despite numerous open-source utilities, only a few tech companies [6, 7, 24] and the US Census Bureau [57] have adopted DP. Accordingly, our work addresses the research gap in bringing DP into organizations' workflows and reaching broader adoption. Dwork et al. [27] partly covered the gap by interviewing DP experts, while our study closes the remaining gap by bringing non-experts into the spotlight. Thus, we interviewed 24 practitioners (19 analysts and 5 data stewards) across 9 major companies that have not yet deployed DP. Overall, our main contributions are:

(i) **Survey Results.** We formulated 5 research questions and derived 24 interview questions thereof. The results of the interviews provide an overview of the current state of data access models (§ 5.1), privacy practices (§ 5.2), motivation behind privacy protection (§ 5.3), and analysis workflows (§ 5.4) in the industry.
(ii) **Key Findings**. From the survey results, we extract 11 key findings, suggest improvements, and answer the 5 research questions about the practicality of DP in the industry (§ 6).
(iii) **Functional Requirements**. Based on the key findings, we propose 10 key desiderata to guide organizations in building privacy-enhancing analytics systems that tackle the privacy-related pain points in their workflows (§ 7.1).
(iv) **Missing Building Blocks**. Given the identified key desiderata, we outline 7 gaps in state-of-the-art DP tooling (§ 7.2).

Privacy officers and *legal practitioners* will find (i) and (ii) helpful in understanding the landscape of privacy and analysis workflows in the industry. *Software engineers* and *developers* will also appreciate (iii) and (iv) as these contributions focus on tooling, and, additionally, will find our early-stage privacy-enhancing analytics system design presented in Appendix H helpful. Overall, notable findings reveal that cumbersome data request processes block analysts for significant periods for every new project. Additionally, we note that SQL was more important than machine learning, and data stewards are more concerned with security than privacy. We conclude that DP could shorten data access processes,

Proceedings on Privacy Enhancing Technologies 2023(2), 151–170

https://doi.org/10.56553/popets-2023-0045

enable data exploration across silos, and is applicable to specific use cases. Moreover, DP tool designers can learn from one another as no tool outperforms the rest in every aspect, and, most importantly, bridging the gap between theory and practice is primarily an engineering problem within striking distance.

2 DIFFERENTIAL PRIVACY

Unlike traditional privacy techniques, which are vulnerable to auxiliary information attacks [8, 29, 33, 67, 79, 94], differential privacy [28] mathematically formalizes a privacy guarantee agnostic to background information. A function guarantees differential privacy (e.g., an analytics query or a machine learning (ML) model) if it bounds the information gain that an attacker can expect from its outputs. Aligned with this adversarial perspective, for the context of this study, we define *privacy* as the prevention of an individual's re-identification [108].

In practice, the outputs of a differentially private function are similarly likely, regardless of an individual's contribution to the input data. This similarity is bounded by the parameter ε, which is inversely proportional to the strength of the privacy protection. A randomized function $\mathcal{M}(\cdot)$ satisfies differential privacy by adding calibrated random noise, typically to a deterministic function's output. Formally, differential privacy is defined as [26]:

Definition 1. *(ε-Differential Privacy). A randomized function $\mathcal{M}(\cdot)$ is ε-differentially private iff for any two datasets D and D' differing on at most one element, and any set of possible outputs $\mathcal{S} \subseteq Range(\mathcal{M})$:*

$$\Pr[\mathcal{M}(D) \in \mathcal{S}] \leq e^{\varepsilon} \times \Pr[\mathcal{M}(D') \in \mathcal{S}].$$

We introduce other concepts useful in the context of this paper:
Sensitivity. Beyond ε, the other parameter that affects the scale of the noise is the sensitivity of the deterministic function, which determines the maximum difference of the function's outputs over all possible neighboring datasets D and D'.
Central/Local Model. An application can add differentially private noise in the *central* model after aggregating data points from different clients or in the *local* model by adding noise to each data point individually. While the local model requires less trust assumptions with the aggregator, it is usually noisier than the central model.
Sequential Composition. Differential privacy algorithms follow *sequential composition* [26], i.e., if one executes a sequence of (possibly different) DP mechanisms n times over D with ε_i, the consumed *privacy budget* $\varepsilon = \sum \varepsilon_i$.
Privacy Budget Tracker. Because the added noise is centered around 0, an attacker could reverse engineer the n outputs by averaging out the noise. Thus, systems should implement privacy budget trackers to prevent this attack.
Floating-Point Vulnerability. Proofs of differential privacy mechanisms work on continuous distributions, which leads to privacy bugs in practice as the implementations rely on floating-point arithmetic [76]. There are a few solutions to this problem. In short, Mironov's Snapping mechanism [76] discards the least-significant bit in a post-processing step, Naoise et. al [51] combine four random samples, and Haney et. al [46] designed a variant of the Laplace mechanism that avoids a precision-based attack.

3 RELATED WORK

Some organizations have developed and deployed differential privacy tooling and have documented their purpose. Specifically, Apple [6, 7], Google [6], and Microsoft [24] employ algorithms based on the local model of differential privacy to collect information from users. The local model is not as predominant in the industry as the global model (our focus), which has seen more deployments in the past years: Google's Plume [4] enables simple statistics (count, mean, sum, variancer, and quantile) over large-scale datasets. Moreover, LinkedIn [63, 88, 89] proposed an API to analyse user data, and the U.S. Census Bureau in 2020 [57] released microdata; however, these two approaches only considered count queries. Additionally, there exist open-source differential privacy *libraries*, *frameworks*, and *systems* from Google [39–42, 107], Harvard [31, 47], IBM [52], Meta [74], OpenMined [83] (experimental product), Tumult Labs [99], and the University of Pennsylvania [78] and Texas [90]. Note that OpenDP encapsulates SmartNoise core [82]. Additionally, researchers have also developed open-source systems focused on *user interfaces* for differentially private analytics: Bittner et. al [12], DPcomp [49], DPP [56], Overlook [97], PSI (Ψ) [32], and ViP [77]. However, only a few libraries have been discussed in a utility benchmark [36]. Moreover, Johnson et al.'s work on differentially private SQL [60] at Uber [59] focused on a quantitative evaluation of the queries without discussing its practicality with practitioners. Unlike the previous literature above, we aim to qualitatively understand the practicality and adaptability of differential privacy in the central model to existing data analysis pipelines within an organization beyond count queries.

Among top searches of surveys related to differential privacy in digital libraries such as IEEE [53], ACM [3], ScienceDirect [92], or ArXiv [19], one may notably find surveys of applications or analysis models for differential privacy in the context of social networks [55], cyber physical systems such as IoT [48], statistical learning [93], location-based services [65], a user survey about privacy in data sharing [15], and lessons learned from employing differential privacy in the US Census [34]. Notably, Kifer et al. [64] distills a set of best practices and implementation details from their experience designing differential privacy systems at Meta, which we consider in our key system desiderata proposal (see section 7.1). However, our work instead explores systems from companies unfamiliar with differential privacy and focuses on answering whether differential privacy could help data analysts in the broader industry. Lastly, the closest work to ours is from Dwork et al. [27]. They interviewed differential privacy experts regarding their implementation specifications. We differentiate from Dwork et al. [27] in that the hereby interviewed practitioners and the organizations as a whole had no significant technical expertise on differential privacy, which are the vast majority in any industry, and, specifically, we sought to understand whether differential privacy could lift the privacy-related roadblocks in their data analysis workflow.

4 RESEARCH METHOD

While a few organizations have successfully deployed differential privacy for data analysis [6, 7, 24, 57, 63], the large majority have not. To understand whether differential privacy in the central model is practical in their analysis workflow, following a method inspired

by Dwork et al. [27], we performed an empirical study of a set of institutions that have not deployed differential privacy yet for their internal analysis workflows in production. Since the focus is learning whether institutions could benefit from differential privacy, the unit of analysis is the institutions themselves.

Our study captures the answers of 24 practitioners from 9 organizations (19 data analysts/engineers and 5 data stewards). These organizations belong to different industries and are of different sizes (see details in Table 2 of Appendix A). The jurisdictions under which the companies operate contextualize our key findings to the EU (5 companies) and the USA (4 companies). In some organizations, we interviewed multiple practitioners to produce a holistic picture of their data analysis ecosystem. Most interviewees held the title of *data analyst*, while a few were data engineers or team leaders. Irrespective of their title, all practitioners had at least two years and at most 10 years of experience in the field (around 5 years on average) and a comprehensive knowledge of their organization's tools and workflows for data analysis.

Interview Format and Research Questions. We interviewed each of the 19 data analysts for approximately one hour through a video conference, except for three via email correspondence, between November 2021 and August 2022. We produced the research questions (RQs) and the questionnaire prior to the interviews and based on the authors' knowledge of differential privacy and feedback from practitioners other than the ones interviewed. The research questions aimed to understand whether differential privacy could enhance their corresponding institutions' analysis workflow by identifying missing opportunities, assessing the impact of differential privacy in their workflow, and identifying roadblocks.

We carefully formulated the questions broadly to enable interviewees to express their views freely, recount their experiences fully, and reduce response bias and priming. Because the organizations have not deployed differential privacy, most interviewees were not familiar with differential privacy; only two had some non-technical familiarity. We tackled this challenge by explaining differential privacy at a high level before starting the questionnaire. We produced the questionnaire for data analysts in Appendix C, whose results are collected in section 5. Only 4 of the 24 questions contained the word "differential privacy", which the interviewees could nonetheless answer without a deeper technical understanding (see Appendix C).

Furthermore, we performed a deep dive in one corporation by interviewing 10 analysts. Additionally, to understand the process and motivation behind this corporation's privacy protection, we interviewed five *data stewards* via video conference or email correspondence with a second questionnaire (see Appendix B). Data stewards control access to and minimize the risk of data interactions, e.g., auditing analysts' purposes before granting them access. Altogether, we distill key findings and answer these 5 RQs:

RQ1: *What is the context of privacy protection in the targeted organization?* The data stewards provided a perspective of their data protection practices, shedding light on their motivation, concerns, and possible improvements of their methods in privacy protection.

RQ2: *Could differential privacy tackle the privacy-related pain points of an analysis workflow in an organization?* The answer draws a picture of the workflow and the improvements analysts would welcome. This holistic picture helps us identify opportunities for differential privacy in organizations' analytics workflows.

RQ3: *When does differential privacy impede an analysis?* Differential privacy is not a silver bullet; thus, we aim to explore the limitations of differential privacy in an organization. Moreover, as the mechanisms to make SQL-like queries fulfill differential privacy are well-understood [40, 59, 60], we investigate whether this type of query is common in analysts' workflows and bring significant benefits in exchange for moderate effort.

RQ4: *How would differential privacy affect the workflow of an analyst?* Analysts are not accustomed to the noisy outputs of differentially private mechanisms. With this RQ, we aim to understand the impact of noise in their analysis and explore their views on different uses of differential privacy.

RQ5: *Can differential privacy be applied to the frequent SQL-like queries analysts execute?* To exclude the impossibility of using differential privacy, we must assess whether analysts can use it in their queries.

5 RESULTS OF THE PRIVACY STUDY

To frame the research questions in the appropriate context, we first depict how the interviewed organizations usually access data and present the state-of-the-art anonymization techniques in the industry. Subsequently, to provide a perspective on the motivation behind privacy protection, we summarize the results of the interviews with the data stewards from the deep-dive organization (RQ1). Finally, we delve into the data analysts' answers to assess whether deploying differential privacy is useful and possible (RQ2-6).

5.1 Data Access Models

This study focuses on practitioners performing data analysis internally, i.e., without publicly releasing the results. The interviewed organizations used one of two models for accessing data internally: *segregated* and *federated* [13]. Fig. 1 provides an informal diagram for a quick intuition of the models. These models used distinct roles: *data owners* in charge of collecting data, *data engineers* building pipelines, *data stewards* assigned to overseeing the data access request processes, and *data analysts* fulfilling analytics use cases.

Analytics teams in the segregated model engaged directly with data owners, whose data are stored in different data centers and regions running different systems. The data owners would provide the data and also act as stewards. Without an established system for automated data exchange and preparation, the analytics teams had data engineers to prepare data for every use case. An improvement over the segregated model is its federation. After collection from multiple sources and pre-processing and anonymization, in the federated model, data from all domains (e.g., demographics, financial, health, etc.) are stored and easily accessible from a single application interface. The data engineers build such data pipeline and are not usually part of an analytics team. Data stewards guard multiple data sources, interact with analysts, and are detached from the data owner role, which is dedicated exclusively to data collection.

In both models, *data protection officers* from the legal department could interact in the dataset request process, namely when analysts requested data for the first time or data were highly sensitive.

While the initial monetary investment to build a federated system could be larger than for the segregated model, the federated

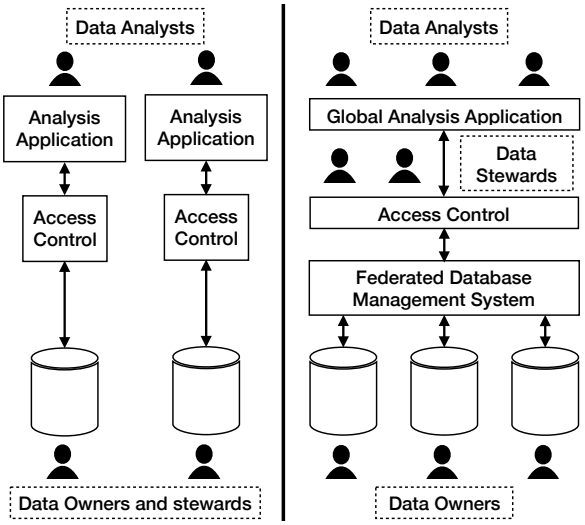

Figure 1: Informal diagram depicting a segregated (left) vs. a federated (right) models for accessing data.

model holds some advantages: it (i) curtails overhead by eliminating the repetition of some tasks in the dataset access request process (e.g., user identification or analysis' purpose specification) and (ii) reduces time-intensive and cumbersome dataset exploration across different systems. Moreover, it (iii) streamlines building data pipelines and defining access request processes by following the same standards across domains and sources. A federated model (iv) simplifies providing precise access control across sources and enforcing policies. Furthermore, it allows to (v) assign non-overlapping roles to practitioners, and (vi) establish re-usable channels between data and analysts. Finally, it could (vii) log the different analyses that other analytics teams have already performed such that others may use them (preventing work duplication). Nonetheless, while the federated model holds such advantages over the segregated model, we observed a similar analyst workflow (see Q7) and an adversarial position for the dataset request process.

5.2 Current Anonymization in the Industry

In this section, we discuss the status quo of the anonymization that companies used to remain compliant without differential privacy, providing a baseline in the context of this work.

Companies can use data collected with user consent exclusively for the agreed *primary* purpose. If companies choose to use data for purposes other than the one agreed (*secondary* purpose), data must be anonymized. All companies had not deployed differential privacy in production or other advanced privacy-enhancing technologies, and employed traditional means of anonymization: *suppressing* direct identifiers such as names, emails, or social security numbers, *truncation* of, e.g., GPS locations and traces, *generalization* (e.g., transforming 28 into [20, 30]), and *dropping* unnecessary attributes and outliers. We consider these techniques *syntactic* [21] because an algorithm transforms the data's syntax following a predetermined model (e.g., GPS locations must only have three decimals). Additionally, data were always encrypted at rest.

Beyond anonymization, to avoid merging multiple sources that could re-identify individuals, some companies did not allow analysts to access multiple datasets at once. In one company, depending on the purpose, stewards granted access solely to a subset of the dataset or a mock dataset for experimenting purposes. Furthermore, for critically sensitive datasets (e.g., illnesses), one company provided access only to an anointed small set of analysts, limited access times, applied anonymization, and restricted analyses to cloud environments. These environments produced logs for later auditing (if needed) and blocked analysts from downloading data. On the other hand, based on user consent for primary use, analysts from one company could access detailed client profiles (names, house prices, mortgages, income, among others). Despite having user consent, we recommend decoupling direct identifiers from the rest of the data (e.g., hashing the direct identifiers) to minimize the consequences of malicious analysts' actions, and encourage the integration of an automated process (or another practitioner) that can only access the analysis output and the direct identifiers to serve the customer (e.g., linked by a hash table only known to the additional process/practitioner).

Altogether, companies applied the principles of *factual anonymity* (i.e., the effort of re-identification is disproportionate to the upside potential of an attacker learning about the individual), *proportionality* (i.e., collection restricted to data necessary to fulfill the primary purpose) [13], *audit logging*, data sharing on a *need-to-know basis*, *data retention* and *purging* [13], *access controls*, and traditional *anonymization*. However, the companies could not measure the privacy achieved by their systems and could only rely on their experience of what is compliant with regulation [30, 80].

5.3 Motivating Privacy

RQ1: *What is the context of privacy protection in the targeted organization?*

(**Q1**) *What is the institution's motivation for privacy protection?* The five stewards agreed on two main motivations: (i) organizations have a legal and moral duty to abide by data-protection laws, (ii) privacy protection is an asset whose "*quality has to be equal to the premium product offered.*"

(**Q2**) *What are your privacy concerns when an analyst has full dataset access?* When proceeding with data protection risk assessments of dataset requests, stewards are predominantly concerned with misappropriation (i.e., unauthorized use of data) and data leakage. While stewards do not expect analysts to be malicious, they are apprehensive of a potential lack of privacy skills, privacy-oriented mindset, and dataset understanding or pure negligence. Specifically, stewards strive to prevent attacks such as unsolicited customer profiling, disclosing data to, or colluding with third parties to take advantage of the customer, combining datasets for re-identification, or using the data for purposes other than the one consented.

(**Q3**) *At what level of data granularity are you protecting and measuring privacy?* The granularity of privacy protection is at the attribute level, and stewards measure privacy based on the fulfillment of data protection regulation. For example, attackers could use the attribute *location* to re-identify individuals; thus, according to GDPR [30], the attribute must be obfuscated so that their home, work, and other points of interest cannot be linked to the individual. Furthermore, the corporation must guarantee the "*security, transparency,*

and legitimacy of the [data] processing." Overall, the anonymization approach strives to achieve the *factual anonymity* principle.

(**Q4**) *What could be improved in the dataset request process?* Data stewards suggested to (i) perform an audit to verify that the executed analysis aligns with the previous commitment, (ii) increase the quality of the datasets' metadata so that analysts can better define a purpose, (iii) increase the privacy training of analysts, (iv) produce privacy-enhanced dataset reports so that after the permission expires analysts can still retain some information, and (v) increment efforts in request process automation.

(**Q5**) *What are your typical questions for the current interview-based full dataset access authorization?* To help other practitioners in the development of their risk assessment process, we gathered the most frequently asked questions from data stewards to data analysts during the dataset request process (see Appendix E). Notably, without a clearly defined data usage purpose, the data stewards would not grant access to analysts.

(**Q6**) *Instead of the interview process, would you be capable to run a program provided by the analyst such that the analysis is carried out without the analyst ever "seeing" the dataset?* While most considered this an efficient, plausible, and necessary step in the future, the five data stewards did not yet have the required technical training, and their system did not enable the functionality. "*At the moment, it is not possible, but it will be a necessary step in the future, if not already today.*" One steward remarked the importance of this functionality, as in some cases, e.g., requesting data from a branch of the company in another country, is extremely challenging.

5.4 The Practicality of Differential Privacy

RQ2: *Could differential privacy tackle the privacy-related pain points of an analysis workflow in an organization?*

(**Q7**) *What is your workflow to analyze data?* Despite the use of either a segregated or a federated model, the workflow was similar across organizations and employed common practices and tools; the main differences were in *dataset exploration*.

(1) **Business Use Case Demand**. A business unit asked an analytics unit to conduct a study for supporting a business need, or analysts continuously studied data from a specific (customer's) domain.

(2) **Dataset Exploration**. Only the companies using the federated model for accessing data could explore datasets' metadata through a *data portal* without requesting access first (unless the dataset was tagged as critically sensitive), making the identification of the suitable dataset for the business need easier. Analysts would find datasets using keywords in a search bar, and datasets provided descriptions, depicted their schema, and had data owners' contact information (analysts sometimes interviewed them to further understand the suitability of the dataset).

In the deep-dive organization, analysts could additionally perform any SQL aggregation query on the anonymized dataset prior to access (e.g., counts, averages, etc.), which they used for data understanding and quality checking (e.g., number of nulls and duplicates or measuring skewness). However, for privacy reasons, analysts could only retrieve a few rows when executing `SELECT *` query types and aggregations could time out (preventing excessive execution costs). Analysts used this *preview* functionality frequently "[...] *to get a feeling for the data*" and found it useful for exploration "*The preview query is the best feature.*" Companies without a federated model could not explore datasets, required data engineers for each use case, and analysts relied either on leveraging their contact network or on an experienced team lead to find promising datasets within the company.

(3) **Dataset Access Request**. Once the analysts identified a promising dataset, they formally requested access, which involved filling standard forms about the details and purpose of the analysis so that data stewards could assess the privacy risks. Except for three small companies, the request entailed interviewing with stewards, where they asked questions such as the ones in Appendix E.

(4) **Visual Inspection and Preparation**. With full dataset access, analysts would sometimes visually inspect the data values, types and schema. Analysts deemed these checks necessary because of the flaws sometimes found in the pipelines and dataset descriptions of the federated data portal or the data provided by the data owners in the segregated model. Moreover, as datasets consisted of many tables, analysts often checked which joins were possible and which attributes were most suitable for primary and foreign keys. With this information, they performed retrieval SQL queries with GROUP BY, WHERE, and `JOIN` clauses to build a sub-dataset fine-tuned for their analysis. Many analysts also performed quality (double) checks and data wrangling using the Python's Pandas library [84] instead of SQL.

(5) **Data Analysis**. Once analysts had checked the quality and wrangled the data, they primarily performed their analysis or ML model training in Python Jupyter Notebooks [5], and if the analyst dealt with *big data*, they employed PySpark clusters [85].

(6) **Output Interpretation and Model Deployment**. If the use case required building a model for online prediction, the analysts would sometimes load the model into a more performant language like Scala before deployment. However, analysts frequently only needed to report statistics and visualizations, from which the business units drew actionable information.

Most of the platforms and workflows employed AWS analytics tools [9] namely S3 buckets (storage), Glue (data preparation), Athena (SQL querying), Sage Maker (data analyses), and analysts also used Python for visualization (one used R) and two of them complemented their results with Tableau [95]. Additionally, two analysts used Knime [66] for drag-and-drop analysis and visualization, and another two employed SAP data management software tailored to their department's needs.

The small interviewed companies had a few major differences, namely, they used a hybrid between the central (all datasets stored in a single data warehouse) and the segregated model. Because of their small customer pool (managed centrally), they collected data from their customers or purchased user-data products from other companies to analyse or train ML models with more data, which required interaction with a segregated set of external data owners. Furthermore, because of the small size of some companies, they had no need for formal dataset request processes as most employees were aware of the activities of the rest; their overhead was at the time of signing the initial contract with customers, which included data access policies and non-disclosure agreements. They also employed traditional anonymization techniques and only retrieved with SQL data stored in, e.g., Google Cloud [38], if strictly needed

(less data for building the model and testing, and more data for the final training or analysis).

(**Q8**) *Why do you need full dataset access?* The main reasons given for accessing all the records of a dataset instead of, e.g., through solely a query interface, were:

(i) *Obtaining a Holistic Understanding of Data.* All analysts worked uncomfortably if they could not make preliminary statistics or visualizations that encompassed all records "*I need to see the entire dataset to understand the data,*" "*I am not necessarily sure of what I need to look at until I look at it. It is an improvisation, you start with a broad question and then you delve into it.*"
(ii) *Less Effort.* A few analysts could fulfill their analysis with only SQL aggregation queries (e.g., counts and averages) and produced visualizations afterward; however, some found using other tools easier: "*Having access to the entire dataset allows me to use Pandas.*"
(iii) *Cleaning Data.* Given that there could be flaws in previous data preparation steps, analysts tended to (double) check all data for quality.
(iv) *Wrangling Data.* In the federated model, data engineers often built datasets without precisely knowing the purpose of a data analyst; thus, analysts sometimes took an engineering role, creating features for ML models or further tailor the dataset for their analysis by grouping or executing queries with `JOIN` clauses.
(v) *Debugging ML Models.* Analysts frequently needed to debug their ML models when testing and training, as there might be corrupted data points.
(vi) *Visually Inspecting Values.* Some use cases, such as root-cause analysis, required analysts to check specific IDs and attribute values, and at times analysts needed to check whether an output table is feasible or map (truncated) GPS traces to street names for the analysis to be interpreted.

Other analysts, however, did not always require access to all records because their ML model already converged, did not overfit, and provided enough accuracy: "*Since I am normally only doing exploratory work, I usually do not need access to the full dataset to prove that the given problem can be solved.*"

(**Q9**) *How often do you request full dataset access? How long does it usually take?* Among the large companies, the request frequency varied widely between 4 times a month to once every 6 months, with an average between once and twice a month. Likewise, regarding waiting times, the minimum hovered around one to three days, while the maximum was two months, with an average between one to two weeks. If another country hosted the data, the first request could take 9 months. Overall, analysts from the interviewed large organizations were blocked for at least one week for every new requested dataset, which they solicited on average once a month. Specifically, in the deep-dive organization, analysts requested 5073 datasets altogether in 2021 (around 14 requests per day, which increased to 18 as of 2022). Out of all the requests in 2021, stewards rejected around 5.6%, amounting to fruitless weeks of revisions[1]. Moreover, the number of requests was more than double the number of available datasets in the deep-dive organization in 2021 (a sign of significant duplication of work, accruing more costs). On the other hand, three of the small organizations did not have such a formal request process, making them agile.

(**Q10**) *What do you think about the process to request full dataset access in your organization?* While analysts at small and US-based organizations were satisfied with the request process, there was an overall consensus at the EU-based large organizations on the following statement: "*The process to get customer data is slow. It might take from three days to weeks, to months*" and for some, even "*Two to three days is too slow.*" In the worst-case scenario, an analyst could wait weeks for a rejection.

Some analysts thought the interviews with stewards were primarily for building trust, and once built "*I always receive access. I do not see the point of waiting and interviewing every time.*" Furthermore, frequently there were too many practitioners involved, leading to lengthy discussions about which dataset to use and often suffered a dilemma because responsibility entailed accountability in one organization "*If there is more than one data steward responsible, then it seems no one takes full responsibility for the acceptance or rejection of the request.*" On the other hand, there were bottlenecks in the vacation season when only one steward was responsible.

Analysts agreed that accessing data has become better since they moved from a segregated model to a federated model; however, the process was still cumbersome, so much so that some teams incurred into the malpractice of entrusting a single analyst to manage the process. One analyst summarized the inefficiency of the segregated model: "*There is a lot of bureaucracy and everyone is extremely reluctant to grant access to a full dataset. Even for internal problems and non-sensitive data. It is cumbersome to request full dataset access because there is no central point where the dataset access can be requested and no central entity which manages access control and usage control for all datasets. For every instance, the process is a bit different depending on the responsible department, underlying workflow and data pipeline.*" In the segregated setting, the process was lengthier, and an analyst could not explore what others had analyzed or requested, sometimes leading to redundant work.

(**Q11**) *What features do you think are missing in your organization's data analysis workflow?* The most notable proposed improvements were: (i) including rich information regarding dataset metadata (preferably with visualizations) and their access request process, (ii) improve real-time analytics performance, (iii) enabling full analytics in data portals such that an analyst does not need to transfer data to other tools, (iv) limiting access times to improve security, and, from a data engineering perspective, (v) automating sensitive data detection and (vi) improving quality and automated checks to minimize visual inspections.

RQ3: *When does differential privacy impede an analysis?*

(**Q12**) *In which analytics use cases have you been involved?* Most analysts worked on *descriptive* use cases. Some of these use cases focused on reporting conclusions from the past by performing root cause (error), cost down, and warranty costs analysis. Other analysts strove to increase the situational awareness of the company by analyzing location-based time series of users (identify

[1]The daily rejection rate went from 0.8 in 2021 to 0.9 in 2022, potentially indicating updated stricter policies.

points-of-interest or common traces), their behavior when using a product or a service (frequently used features, A/B tests, purchases or component performance), and demographics (user-base analysis or advertisement). Additionally, some analysts focused on alerting internal stakeholders of quality defects in real-time, and another analyst performed correlation analysis to better understand the interplay of different variables in products and services. Most of these use cases required performing aggregate statistics (including visualizations to report to management), namely for situational awareness, while others demanded visually inspecting exact values (namely for error detection or financial data), and one analyst used classification ML models for quality checks.

The minority of interviewed analysts were involved in *predictive* use cases: forecasting product lifetime, labelling spam and inappropriate images, user behavior, the company's profit and loss, claim costs, and predictive maintenance and creating automated underwriting models. While these use cases relied on basic statistics, some used vanilla ML such as linear regression (for underwriting models). Nonetheless, the interviewed analysts agreed that using ML was rare; thus, most analysts relied on aggregation and visualizations, as the business units demanded *quick* and easily *interpretable* results.

(**Q13**) *Is SQL-meaningful for your work? How many SQL-like queries do you make weekly?* Most of the interviewees employed SQL, chiefly during exploration, and they deemed SQL an important part of their workflow "*SQL is amazing, everyone who tells you SQL is going away is wrong*," because they could quickly look into rows and performed preliminary statistics, and, with `JOIN` clauses, prepare a dataset for their use case. The least adept analyst executed 5 weekly queries, while the most assiduous SQL user performed 250, being the average around 50 queries per week.

(**Q14**) *How often do you need machine learning to fulfill your analysis in contrast to using SQL?* Two interviewees always needed ML to fulfill their analysis, while another 4 used ML for some of their use cases. The analysts who were allowed to explore datasets used SQL for exploration, and three used SQL to generate statistics and completely fulfill their analysis (complemented with visualizations), while the rest preferred Python or other tools for analysis. Furthermore, analysts often visualized data to accompany their results with other tools (see Q7) and employed retrieval SQL queries for visual data inspections (e.g., for error analysis) or building tailored datasets for their analysis.

(**Q15**) *What are your most used machine learning models?* The 6 analysts employing ML most often resourced to decision trees and linear regression because they are easy to debug, interpret and visualize the results. These analysts also mentioned the use of random forests or XGboost (preferred), Bayesian approaches, support-vector machines, and, for time series, they used outlier detection techniques for error analysis and autoregressive integrated moving average for forecasting. Analysts avoided neural networks because they are hard to interpret; nonetheless, one practitioner indicated they were working on deploying neural networks in the future for underwriting models. In particular, one analyst employed PyCaret [86] for automated ML workflows, as in the corresponding department "*It is more important to be quick and give a good-enough overview than having well trained precise models,*" "*Complex machine learning is often never required.*" Other analysts voiced that such is often the case.

(**Q16**) *If you were to use differential privacy to fulfill your analysis, when and how much accuracy would you be willing to forgo?* The willingness to forgo accuracy depended on the use case, with a spectrum ranging from the need for absolute accuracy for quality, error, or financial analyses, to indifference for accuracy in exploratory use cases (only enough accuracy to prove a solution works). For the rest of the use cases, while the interviewees would need to estimate the minimum accuracy formally, they informally reported on average that an accuracy of around 98% would be sufficient, and none reported below 95%. Some financial analyses could also allow errors in the magnitude of cents of a monetary unit, and one analyst reported the need for at least 99% accuracy for finding suitable primary keys for joins. Additionally, comments such as "*I am scared of introducing noise into the analysis. [...] From all the analyses I do every year, there will be some that will be wrong. [...] How well you are compensated depends on how well you do. [...] Because you are paid to have an opinion, you are not allowed to be wrong," suggest that organizations' incentive systems for data scientists, e.g., bonuses, should change to account for errors due to differential privacy.*

RQ4: *How would differential privacy affect the workflow of an analyst?*

(**Q17**) *How much would the noise affect your analysis?* Depending on how much the noise could affect an analysis, we observed three categories for use cases: (i) suffer adverse effects, (ii) reach a tradeoff, and (iii) robust to noise. The first one relates to analyses reporting error, quality, or financial results, where noise could have catastrophic consequences, e.g., a defective component is installed in a product, or yearly budgets are inflated. Moreover, analysts sometimes dealt with low data quality (notably from sensors) that noise could worsen, e.g., GPS locations may already have a 10m error, making a points-of-interest analysis noisy in itself. Adding noise to the aggregation might produce unusable results.

The second type concerns aggregation and visualization reports, where, given enough data, the noise would not affect the interviewees' analysis workflow (e.g., demographics or product usage studies); however, analysts would prefer working with error bounds to report confidently to management. The third type of noise relates to analysts testing solutions "*Since my work is exploratory and we mostly try to prove that the problem can potentially be solved, noise would not have any negative effects for my analysis.*"

(**Q18**) *Would you find it helpful to execute differentially private SQL queries to explore and fully analyse datasets without the standard permissions?* We theorized that given the plausible deniability guarantees of a differentially private analysis, which can be argued to comply with the identifiability notion in GDPR [30, 50], some uses cases that heavily rely on aggregation might abate or not need the standard dataset request processes. From this perspective, most analysts found differentially private SQL queries helpful, in summary, because "*If having differentially private SQL queries for data exploration implies reduced bureaucracy and easier access, then this would save a lot of time and discussions.*"

Notably, one interviewee saw the potential of differentially private queries for data exploration: companies could expose data

externally through an API, allowing others to understand their data products by conducting preliminary analyses. Another analyst favored integrating differential privacy into, e.g., AWS Athena [9]. On the other hand, few analysts did not see the value of differential privacy because their use cases required, e.g., visual inspections for error detection, or their organizations were already agile in accessing data. Lastly, two analysts voiced a general concern "*[Differential privacy] is a double edge sword. You could get quick [data] access, but then [results are] noisy,*" "*I think I would find it annoying, since it adds an additional step and obfuscates the results*," and it could lead to confusion as analysts usually work with accurate data.

(**Q19**) *Only based on the information extracted from a dataset exploration with differential privacy, could you write a script to fulfill your analysis goal?* A couple of interviewees shared their inability to program their script as they needed to see the data (e.g., error analysis), and the others shared their skepticism by highlighting the problem of low data quality. Even if an analyst developed an intuition for the data through differentially private aggregation queries, programming other statistics, visualizations, or ML models would likely require debugging, which may lead to visual inspections.

(**Q20**) *What are the minimum properties for you as an analyst such that you are confident to write an analysis script without full dataset access?* Assuming enough data quality and a use case that does not require visually inspecting data, the interviewees indicated that for tentatively writing code without dataset access, they needed: good metadata from the dataset, such as attribute descriptions, knowledge about the events that trigger data collection, primary keys, data types (IDs, dates, timestamps, floats, strings), dataset size (number of rows and columns), and attribute distributions to learn about sparsity in the form of histograms or box plots.

(**Q21**) *Would you find it helpful to use a dynamic dashboard that visualizes dataset information with differential privacy?* Since data platforms may not expose sensitive data on a dashboard for exploration, we conceptualized enabling this functionality with differential privacy. All but one interviewee considered such a dashboard helpful for finding a suitable dataset faster and with a better user experience than their available utilities (static and scant summaries or using SQL). Specifically, an interviewee commented that, in general, one should be able to visualize the data and get basic statistics before requesting access, and another analyst would have liked to preview similar information as the "`describe`" method of a Pandas dataframe [84] (count, mean, standard deviation, minimum, quartiles, maximum). Nonetheless, one analyst noted that a dashboard is a nice-to-have because it is only more convenient than SQL. Lastly, another interviewee underlined a problem that may arise when an analyst does not trust the data provided by the visualization, e.g., when the plot seems implausible. The interviewee suggested that a dashboard should enable the analyst to drill down or provide contact information from a data owner to verify correctness.

RQ5: *Can differential privacy enhance the privacy of the frequent SQL-like queries analysts execute?*

(**Q22**) *What are your top SQL-like queries before you have full dataset access?* If analysts could explore datasets, most would usually conduct a metadata analysis with SQL to assess data quality: finding the number of duplicates, outliers, nulls, and not-a-number values and measuring the skewness. Analysts would also explore the dataset for data understanding using `COUNT`, `DISTINCT`, `MAX`, `MIN`, `AVG`, and `VARIANCE` functions with `WHERE` and `GROUP BY` clauses. Analysts were typically interested in frequent values within a column (see details in Appendix D). Furthermore, the deep-dive organization allowed to use `SELECT * LIMIT(X)` for a few `X` rows so that analysts could have a "*feeling*" for the data. On the other hand, fewer analysts performed retrieval queries (limited in output rows) to verify whether an ID was present or two tables could be joined.

(**Q23**) *What are your top SQL-like queries after you have full dataset access?* Analysts who could not explore the dataset prior to having dataset access would execute queries such as those in Q22 first (see Appendix D for details). Afterward, if they did not already retrieve the necessary information from the exploration, they resorted to Python and other visualization tools to fulfill the use case. Some analysts performed additional retrieval SQL queries with `JOIN` and `SELECT *` clauses with different filters to visually inspect data points (e.g., IDs or potential errors), identify cut-offs (e.g., where an attribute data type changes), or fetch the specific data they needed.

(**Q24**) *What is the ratio between aggregation queries and queries to retrieve items?* While the interviewees would need to calculate the percentage formally, they reported informally, on average, that around 30% of their queries were for aggregation, being the lowest 0% and the highest 90%. Another three analysts used SQL for retrieval and Python for aggregation or vice versa.

6 DISCUSSION

In this section, we present selected key findings (KF) distilled from the data stewards' and analysts' answers to the 24 interview questions of section 5, most accompanied by succinct recommendations. Lastly, we answer the research questions proposed in section 4.

6.1 Key Findings

(KF1) *Data stewards seem to be more concerned about security than privacy.*[2] Data misappropriation and leakage retain the most attention (Q2), which is reflected in the established cumbersome dataset request processes that dictate access controls and accountability in the name of building trust with analysts. However, we highlight that privacy lacks such attention, even when some companies still allow their analysts to "*see*" or download sensitive data. Attacks on privacy (Q2) could enable malicious analysts to misuse data for spying on or leaking secret information of celebrities, acquaintances, "friends", or relatives [100], blackmailing, or discriminating individuals in social or commercial transactions online [35]. Despite the risks, companies predominantly use traditional and potentially vulnerable anonymization techniques (e.g., pseudo-anonymization or k-anonymity), as demonstrated by the research community [8, 29, 33, 67, 79, 94]. Thus, we suggest companies increase efforts to research and deploy more advanced PETs.
(KF2) *Running analysts' scripts without "seeing" the data is a distant reality for the interviewed companies.* We explored multiple ways for analysts to run scripts without direct dataset access. In

[2] In the context of this work, we refer to *security* as the measures for blocking *unauthorized* data access, while *privacy* focuses on limiting harm by *authorized* analysts [13].

Q6, stewards declared their technical inability to execute scripts that analysts could share and, thus, avoid granting them access (saving time). With the proper tooling, a non-technical steward could potentially run the script; however, current systems do not offer such abstracted functionality and this option would relay the responsibility to the stewards instead. An alternative to transferring the trust to stewards consists on executing analytics in trusted execution environments[3]. Additionally, analysts altogether gave six reasons why they needed full dataset access (Q8) and reported skepticism when asked about writing a script (beyond aggregation) based on a differentially private exploration (Q19 and Q20).

The main impediment reported was data quality, which often led to visual inspections of dataset values. As encouraged by point Q11-(v), we suggest companies prioritize increasing data quality as it will indirectly improve privacy and increase the technical training of data stewards and owners, enabling more data security options.

(KF3) *Given the analysis workflow, differential privacy could have a significant impact on dataset exploration* (see Q7). As long as exploration does not require visually inspecting a particular ID or an exact attribute value, differential privacy can provide noisy statistics for the analyst to familiarize with the data (e.g., number of rows, averages, quantiles, etc.), which is often enough to assess the dataset's suitability. Furthermore, while analysts were not allowed to explore critically sensitive datasets with SQL, employing differential privacy could arguably enable their exploration by adding an extra layer of protection. Additionally, platforms could provide privacy-enhanced dataset previews (e.g., only revealing a few rows or producing dummy or synthetic data with or without differential privacy). Overall, differential privacy could facilitate exploration that otherwise might not be possible or timely.

(KF4) *Analysts could employ differentially private mechanisms to fulfill certain use cases* (see Q7). If the analysis requires summary statistics and visualizations, a differentially private analysis could fulfill the privacy-utility tradeoff given enough data. Consequently, analysts could fulfill use cases without exact outputs, avoiding potential privacy leaks. Regarding ML, while its differentially private implementations are at an early stage, researchers and practitioners could explore systems to assess whether a model shows signs of converging with enough accuracy after training on a sample of the target dataset. Such a system could help analysts to determine the validity of the model or the dataset. Lastly, we suggest exploring whether differential privacy can enable more accurate analyses than the current organizations' anonymization processes.

(KF5) *After fulfilling the use cases, the interviewed companies do not have a human-supported privacy auditing step.* The last reported step of the workflow in Q7 was "*output interpretation and model deployment*". Aligned with a steward in Q4: "*perform an audit to verify* [*alignment with analysis commitment*]," we suggest privacy officers in companies add a randomly-sampled auditing step with a human in the loop after the conclusion of the use case. We also suggest audit logs, which one of the interviewed companies produced for every execution on sensitive data in secured machines, where analysts could not download data or install new software.

[3]Hardware and software designed to run applications securely against unsolicited retrieval of sensitive information or key material [81].

(KF6) *Given the six reasons analysts shared for fully accessing datasets, differentially private mechanisms could help in (i) "obtaining a holistic understanding of data" by providing dataset summary statistics* (see Q8). Additionally, we suggest substituting tedious SQL analyses with dashboards for visualization, so that tasks require (ii) "*less effort*". We also suggest engineers develop and integrate tools that enable analysts to (iii) "*clean*" and (iv) "*wrangle data*" without visually inspecting the values (i.e., no complete data access required). With such tools, filtering values, imputing, removing duplicates and outliers, fixing wrong formattings, handling missing data, or creating new attributes would also help with (v) "*debugging ML*." Moreover, aligning data engineers with analysts could improve data quality, e.g., by involving engineers in the conversations between stewards and analysts. Lastly, researchers could investigate how differentially private set union mechanisms [22, 45] could help analysts to (vi) "*visually inspect values*." Meanwhile, we suggest increased security measures for such cases.

(KF7) *Analysts are frequently blocked for significant periods every time they request access to datasets* (see Q9). There are a few consequences of such delays. Data stewards and privacy officers must also invest their time in reviewing the requests. From our conversations with the interviewees, we also learned that long waiting times could hamper analysts' bursts of creativity and productivity, which indirectly negatively affect the quality of work. Additionally, an interviewee recounted the malpractice of deferring all the dataset request process responsibility to a single analyst in the team (see Q10). Such practice overburdens an individual with the responsibilities of the entire team for, e.g., a data leakage, creating an unhealthy imbalance in accountability. This practice further increases the company's privacy risk by potentially having the other analysts handle data without privacy training. We suspect this malpractice is a sign of over-complicated dataset request processes and long waiting times; thus, we suggest privacy officers streamline their processes and prompt teams to refrain from overburdening a single analyst.

Differential privacy's stronger guarantee could reduce the complexity of the interactions between practitioners by offloading their data protection demands and, thus, reduce the costs accrued by these human-intensive processes. Lastly, given that there were multiple requests for the same datasets from different teams, we encourage companies to build interfaces depicting privacy-enhanced summaries of past fulfilled use cases per dataset. An example is the repository designed by Johnson et al. [58] in the health industry.

(KF8) *Differential privacy could arguably reduce the time to access data.* As differential privacy brings a higher and formal guarantee of privacy, it could relax the inquisitiveness of data stewards, eliminate (steps of) the request process, and enable exploration that was otherwise not possible. By enabling exploration, analysts reduce the likelihood of investing time in request processes that could even result in accessing a non-suitable dataset. With exploration and higher privacy guarantees, differential privacy could also speed up requesting data from other countries, which seemed the most significant bottleneck (see Q9). Additionally, differential privacy could potentially prolong access times (if these are limited) and shorten development cycles with an earlier data access by testing algorithms and applications with noisy data or outputs. Regarding

applications specifically, once finished, the customers can confirm whether the product works appropriately with real data.

We have also observed that, once analysts have access, much of the data protection and accountability lies on their shoulders, which differential privacy could lift to a degree by protecting beyond trust and policy. However, the analyst somewhat familiar with differential privacy pointed out that, unless data quality is improved (as also suggested in Q11), "*There is a still a ways to go to deploy differential privacy,*" because the need to debug by visually inspecting data will prevail. To increase privacy protection in those cases, we suggest using differential privacy with limited visual inspection.

(KF9) *Most analysts employed aggregations and visualizations to fulfill their use case in a timely manner, while machine learning was not as predominant* (see Q12). We found that analysts could employ differential privacy to explore datasets suitable for all the identified use cases. However, for the analysis itself, the interviewees voiced that the noise would invalidate the use cases related to quality, error, and (some) financial analyses because mistakes in safety decisions and financial planning are company critical. Nonetheless, for the use cases that required aggregation and visualization, with enough data, we suggest analysts fulfill these use cases with differentially private queries such as counts, averages, and percentiles, among others (e.g., user behavior, demographics, and some location-based analyses). However, the available tools for differentially private ML are not mature for widespread adoption. Thus, we encourage researchers and practitioners to improve and build systems around existing proposals in future work, e.g., location-date analysis [105], heavy hitter identification [71], mining frequent itemsets [114], deep and supervised learning, random forests, and linear regression, among others [1, 52, 70, 111, 113].

(KF10) *For the interviewed companies, SQL was more important than machine learning and was considered a meaningful tool frequently employed in their workflow* (see Q13 and Q14). Additionally, on average, 30% of the top SQL queries executed before and after full dataset access were for aggregation (see Q22, Q23, and Q24), which researchers have already adapted to fulfill differential privacy [40, 59, 60]. Therefore, there is still a gap between what researchers have enabled and what practitioners need for enhancing the privacy of their frequently used SQL queries—a gap we intend to partly cover in section 7 by proposing 10 key system desiderata that an integrable privacy-enhancing analysis system should fulfill. Beyond SQL, differential privacy and its available tools are also suitable even when analysts preferred using Python for aggregation and ML use cases that allowed for lower precision. In particular, we encourage using Python libraries such as IBM's diffprivlib [36, 52] that provide many off-the-shelf differentially private ML models that could provide enough precision for the intended purpose, such as for the linear regression model one company used for underwriting (see Q15). However, practitioners will require further engineering to limit Python to strictly privacy-enhancing libraries and amenable standard functionalities (e.g. by using policy enforcement paradigms such as Wang et al.'s Data Capsule [102]).

(KF11) *Analysts confirm that differential privacy would be helpful for dataset exploration, fulfill certain use cases, and for enabling privacy-enhancing dashboards for dataset visualization* (see Q18 and Q21). For aggregation-based use cases where noise has no detrimental effects, analysts informally reported, on average, a required accuracy of 98% (see Q16 and Q17). While such a figure might seem high, given the large amount of data handled, analysts could potentially find enough for aggregations that fulfill their privacy/utility tradeoff. For example, as of early 2022, the deep-dive organization had roughly 2260 datasets in its federated system amounting to 3.4PB (1.5TB per dataset on average) with an average daily query execution of over 900TB. However, size might not be enough for some use cases, as the analysis could be sensitive to outliers or corrupted data. Lastly, we observe that it is critical for analysts to know whether the accuracy is above their required level, which would consume privacy budget and be hard to estimate, e.g., when the analysis needs post-processing (clamping or truncation).

6.2 Answers to the Research Questions

RQ1: *What is the context of privacy protection in the targeted organization?* The deep-dive organization invests more resources to security than privacy-enhancing analysis—pattern also present in the other organizations. Moreover, stewards consider privacy an asset and strive to provide the best standard for their customers. However, organizations still employ traditional anonymization techniques. Furthermore, companies today are unable to tangibly measure the privacy of their process (see Q3), and while there are specific privacy and security measures such as access controls, they are hard to quantify formally. Lastly, we observed that the interviewed companies are far from having "*data at their fingertips*", one of the reasons being the onerous dataset request processes, which confuse access hardship with access protection.

RQ2: *Could differential privacy tackle the privacy-related pain points of an analysis workflow in an organization?* Yes, to a large extent—In essence, the main problems are (i) lengthy and cumbersome dataset request processes. Moreover, given that analysts can sometimes "*see*", download, and share the data once they are granted access, and even collude with other co-workers with access to linkable datasets, (ii) *only* policy protects data once stewards grant access. Based on our work, we argue that differential privacy can reduce time-to-data by enabling exploration of critically sensitive data or across third-party data sources, relax the current data access restrictions thanks to its formal privacy guarantee, is applicable to some aggregation-based use cases, and, for some use cases, engineers should consider building solutions that block analysts from "*seeing*" the data.

RQ3: *When does differential privacy impede an analysis?* The answer to this RQ heavily depends on the use case and whether the analysts are willing to forgo accuracy. On the one hand, noise addition-based differential privacy is useful in aggregations performed by the interviewees (e.g., querying demographics or frequently used product features). Moreover, on average, interviewees were comfortable with 98% accuracy. However, differential privacy is not a silver bullet, as some of the interviewees' use cases cannot rely on it (e.g., error analyses or critical financial estimations). Therefore, we suggest building systems that enable differential privacy while maintaining the flexibility of allowing non-differentially private queries when the use case strictly needs them.

RQ4: *How would differential privacy affect the workflow of an analyst?* If differential privacy enabled previously unavailable exploration and provided data for privacy-enhanced dashboards, analysts

would have a better user experience in their workflow and lower time spent on processes and exploration, but would also need to accustom to working with noisy data.

RQ5: *Can differential privacy be applied to the frequent SQL-like queries analysts execute?* Yes—While not as frequent as retrievals, around a third were aggregations amenable to differential privacy.

7 TOWARDS PRACTICAL DIFFERENTIAL PRIVACY

This section provides a set of critical system desiderata a differential privacy (DP) analytics system should satisfy for practical deployments. Subsequently, we identify requirements fulfilled by state-of-the-art tools (see Table 1) and highlight the gaps in practice.

7.1 Key System Desiderata

In secondary use cases, an alternative to *syntactic* anonymization (see section 5.2) for sharing data is an inherently private analysis, i.e., the analysis satisfies a *semantic* privacy definition such as DP [21], which uniquely provides a measure of privacy (ε). With DP, organizations do not necessarily need to use potentially vulnerable syntactic techniques (e.g., rounding or truncation) because the analysis itself already enhances individuals' privacy. Based on the (i) interviewees' description of their analytics workflows and systems, (ii) the authors' knowledge in the domain of privacy, and (iii) the feedback provided by additional privacy practitioners and researchers who work closely with/in our lab, we propose 10 key desiderata. The desiderata correspond to a system that enables differentially-private analyses in the central model and focuses on dataset exploration and fulfilling use cases requiring aggregations (see use cases in Q12). These use cases often rely on SQL-like queries such as counts, averages, etc. Additionally, we inspired some of the characteristics of the key desiderata related to (III) *Security* and (V) *Visualization* from Kifer. et. al [64] and Nanayakkara. et. al [77].

(I) *Differentially Private Analytics.* The system bestows DP to a learning function (e.g., a query or an ML algorithm) by adding calibrated noise to the deterministic outputs (or by other means). The system supports the (i) aggregation queries: `COUNT`, `MAX`, `MIN`, `AVG`, `VAR`, and `SUM`, (ii) provides a complementary ML feature, and stores executed queries for future retrieval. The queries (iii) allow for `WHERE`, `GROUP BY`, and `JOIN` clauses.

(II) *Usability.* The system provides logic to preserve (i) the semantic consistency of queries (e.g., variance > 0) and across overlapping domains (e.g., the sum of noisy element counts is not larger than the noisy total). Moreover, the system presents the option to (ii) estimate the sensitivity of a query without user input, and (iii) recommends or sets privacy parameters automatically depending on the dataset and query.

(III) *Security.* The system (i) provides a stochastic tester or other functions to automatically verify whether the algorithm fulfills DP, (ii) employs cryptographically secure pseudo-random number generation with careful seed management, (iii) generates noise impervious to floating-point vulnerabilities [46, 76]. Furthermore, the system (iv) blocks the user from "*seeing*" the data, i.e., while analysts can execute queries, they cannot download or visually inspect the dataset, (v) does not allow to execute arbitrary code, (vi) executes heuristic optimizers only at post-processing, and (vii) protects against timing attacks [64]. A libraries' and frameworks' scope limits to fulfilling (i), (ii), and (iii).

(IV) *Synthetic Data Generation.* When the goal is to develop an application or explore whether an ML model is suitable for a task, the system produces synthetic data. After testing, the analyst can proceed with the real data (without "*seeing*" it). Synthetic data generation could rely on simple techniques (e.g., sampling from a normal distribution with the same mean and standard deviation as the target attribute), ML [18, 96, 112], or combining DP with either. If the analyst is only interested in the data schema, the system produces dummy data, preserving only the schema and data types.

(V) *Visualization.* The system presents a dashboard depicting interactive plots (e.g., histograms) relying on DP queries for quick and intuitive (i) dataset exploration. Additionally, the dashboard visualizes an analysis' expected (ii) accuracy and (iii) disclosure risk, (iv) uncertainty (i.e., a measure of how the same mechanism can produce different outputs with the same input arguments), (v) statistical inference (i.e., privacy parameter estimation with confidence intervals), and (vi) budget splitting (i.e., help in splitting the privacy budget across queries) [77].

(VI) *Privacy Budget.* The system (i) tracks the budget spent (ε "odometer"), (ii) blocks further queries if analysts exhaust their budget, and (iii) accommodates the budget for growing datasets. (iv) It should enable data stewards to specify budgets for teams, individual analysts, and use cases depending on the data's sensitivity.

(VII) *Accuracy Adjustment.* The system allows the user to propose a desired accuracy level. Alternatively, after the query execution, the system provides either information about the noise scales (without additional budget) or a confidence interval (spending budget) [101].

(VIII) *Query Sensitivity.* The system enables a practitioner, e.g., a data analyst or steward, to input the attributes' bounds (maximum and minimum values) as function parameters or in the dataset schema so that the system calculates the query's sensitivity.

(IX) *Privacy-Sensitive Data Annotation.* The system enables data stewards to allowlist attributes based on teams, roles, and use cases. The system automatically obfuscates attributes outside the allowlist.

(X) *Authentication and Access Controls.* The system easily integrates with existing authentication and access control services and enables data stewards to define their access policies.

7.2 Gaps in Differential Privacy Practice

Despite available open-source tooling, one company found it hard to find external partners that could bring DP into practice in their internal analysis workflow. Furthermore, another company stated after exploring the use of DP that, while it seemed helpful, "*[Deploying differential privacy] was more expensive than doing nothing.*" Instead, the department decided to upload syntactically anonymized data to a highly secured system, with limitations on access time, downloads, number of analysts, and audit logs. We kindly argue that their over-statement was due to the intangible costs of the dataset request processes and the lack of integrability of current DP tooling, which makes deployment a complex endeavor.

Overall, our findings indicate a gap between the theory and practice of DP. Working towards bridging the gap, we qualitatively mapped in Table 1 our key system desiderata with DP tools to highlight areas of future work for the privacy community. We selected

Table 1: Mapping between open-source tools and user interfaces and the key system desiderata. Legend: ✓= functionality fully available; ✗= limited functionality or not available; N/A = not applicable; P. = Privacy; DP = Differential P.; TF = TensorFlow; I.i = Enables aggregation queries; I.ii = Enables machine learning; I.iii = Enables query clauses (e.g., `JOIN`); II.i = Query semantic consistency; II.ii = DP sensitivity calculation; II.iii = Privacy parameter search; III.i = DP correctness verification; III.ii = Cryptographically secure pseudo-random number generation; III.iii = Protection against floating-point vulnerability; III.iv = Block data visibility; III.v = Block arbitrary code; VI.i = Budget accountant; VI.ii = Query blocker.

Table 1A: Libraries, frameworks, and systems for differential privacy analytics.

Tool/Desiderata	(I) DP Analytics	(II) Usability	(III) Security	(IV) Synthetic Data	(V) Visuals	(VI) Privacy Budget	(VII) Accuracy Adjustment	(VIII) Query Sensitivity	(IX) Data Annotation	(X) Access Controls
Libraries[†]										
diffprivlib [52]	I.i, ii ✓	II.i ✓	III.ii, iii ✓	✗	N/A	VI.i ✓	✗	✓	N/A	N/A
Google DP [41]	I.i ✓	II.ii ✓	✓	✗	N/A	VI.i ✓	✗	✓	N/A	N/A
Opacus [74]	I.ii ✓	✗	III.ii ✓	✗	N/A	VI.i ✓	✗	✓	N/A	N/A
OpenDP [47]	I.i ✓	II.iii ✓	III.ii, iii ✓	✗	N/A	VI.i, ii ✓	✓	✓	N/A	N/A
TF Privacy [39]	I.ii ✓	✗	✗	✗	N/A	VI.i ✓	✗	✓	N/A	N/A
Frameworks[†]										
Chorus [60]	I.i, iii ✓	✗	✗	✗	N/A	VI.i ✓	✗	✓	✓	N/A
PipelineDP [83]	I.i ✓	✗	III.ii, iii ✓	✗	N/A	VI.i ✓	✗	✓	✗	N/A
P. on Beam [42]	I.i ✓	II.ii ✓	✓	✗	N/A	VI.i ✓	✗	✓	✗	N/A
Tumult Analy.[99]	I.i, iii ✓	✗	✓	✗	N/A	VI.i, ii ✓	✗	✓	N/A	N/A
ZetaSQL [40]	I.i, iii ✓	II.ii ✓	✓	✗	N/A	✗	✗	✓	✗	N/A
Systems										
Airavat [90]	I.i, ii ✓	✗	III.iv ✓	✗	✗	VI.i, ii ✓	✗	✓	✗	✓
DJoin [78]	I.i, iii ✓	✗	III.ii, iv, v ✓	✗	✗	VI.i, ii ✓	✓	✓	✗	✗

[†]Libraries' and frameworks' (III) Security scope is limited to three sub-desiderata (i), (ii), and (iii).

Table 1B: User interfaces for differential privacy analytics (cf. adapted [77]).

User Interface/Desiderata	(V.i) Dataset Exploration	(V.ii) Accuracy Visualization	(V.iii) Risk Visualization	(V.iv) Uncertainty Visualization	(V.v) Statistical Inference	(V.vi) Budget Splitting
Bittner et. al [12]	✗	✓	✗	✗	✗	✗
DPcomp [49]	✓	✓	✗	✗	✗	✗
DPP [56]	✗	✓	✓	✗	✗	✗
Overlook [97]	✓	✓	✗	✓	✗	✗
PSI (Ψ) [32]	✓	✗	✗	✗	✗	✓
ViP [77]	✓	✓	✓	✓	✓	✓

the tools from the related work in section 3 that offer open-source implementations for the central model of DP (see tool descriptions in Appendix F). We must highlight that some of these tools are *libraries* (provide specific functions) or *frameworks* (abstractions used to build specific applications) and, thus, lack functionalities that a *system* (end-to-end application) like Airavat [90] could provide, such as (III.iv) *Blocking the visibility of data* or (X) *Authentication and access controls*. Note that libraries and frameworks assume analysts have data access. Additionally, we regard *user interfaces* (systems focused on visualizations and providing analytics metadata) as a set of tools that should fulfill key desiderata specific to (V) *Visualization*. Accordingly, we assign each open-source software to its category in Tables 1A and B for an appropriate comparison.

We must highlight that the mapping of Table 1 provides high-level guidance, as there are (out-of-scope) nuances Table 1 does not capture. For example, user interfaces such as Bittner et. al [12] and DPP [56] in Table 1B provide exploratory results for using DP ML and for disclosure risk, respectively; however, they do not help understanding the dataset, which is a critical requirement for data analysts. Regarding the tools in Table 1A, diffprivlib [52] offers multiple ML models (PCA, Naive Bayes, liner and logistic regression, k-means) while others focus on deep learning (Opacus [74] and TensorFlow (TF) Privacy [39]) or MapReduce functionality (Airavat). Additionally, the frameworks are designed for large-scale datasets. We note that Google DP [41] provides the building blocks for ZetaSQL [40] and Privacy on Beam [42] (and PipelineDP [83]), which add more functionality for considering datasets with multiple individual's contributions. Lastly, most tools provide only an "odometer" for privacy budgeting, while a few block new queries if the budget is spent (e.g., OpenDP [47] and Tumult Analytics [99]), and Google DP offers functionality to distribute budget across different DP mechanisms [14, 23]. None, however, account for growing datasets, which is a challenge recently tackled in [68]. One may find more of these nuances in [36].

Based on the non-availability or limited implementations of some desiderata in Table 1, we conclude that *differential privacy tool designers can learn from one another, no tool outperforms the rest in every aspect*, and, most importantly, that *bridging the gap is primarily an engineering problem*. Subsequently, we identify the major gaps in differential privacy practice:

Gap 1: (II) *Usability.* While semantic consistency is sometimes desirable for analysts, it can also introduce more error/bias in some scenarios. Only diffprivlib implements mechanisms to fulfill DP and consistency for specific queries (e.g., variance > 0), whereas Google DP or Tumult Analytics only truncate values in post-processing. Furthermore, only Google DP can calculate the query sensitivity in a privacy-enhancing manner without any user input, which is necessary when an analyst lacks domain knowledge of the application (i.e., input bounds). Thus, none of the tools in Table 1 completely fulfill the usability desiderata. *Guidance*: [2, 41, 52, 88, 89, 103, 104]
Gap 2: (III) *Security.* The tools do not provide many security features individually. E.g., most lack stochastic testers to verify that an analysis fulfills DP, and none implement protections against time-attacks [64]. Wrt to secure random number generation: TF Privacy inherits TF's insecure RNG [43, 44] and Airavat employs the insecure utility `java.util.Random` [91] in contrast to DJoin, which relies on FairplayMP [10, 11]. Moreover, while TF Privacy developers are aware [73], we encourage them to include floating-point protections in their deep learning models or rely on discrete noise distributions [16, 37]. Moreover, most tools should tackle their precision-based attack vulnerability [46]. Lastly, we highlight some of the good practices Kifer et. al [64] proposed: open-sourcing systems (the community can check for vulnerabilities) and performing code audits and unit tests to ensure correctness in DP, privacy accounting, and noise sampling. *Guidance*: [4, 41, 64]
Gap 3: (IV) *Synthetic Data Generation* (SDG). Similarly to tools offering DP ML [39, 52, 74, 90], we suggest developers package and include DP SDG logic. *Guidance*: [17, 18, 96, 98, 106, 109, 110, 112].
Gap 4: (V) *Visualization.* While there is enough research on user interfaces, the most popular frameworks and libraries do not adopt them. We suggest packaging available DP user interfaces for patching analytics tools. *Guidance*: [77, 97].
Gap 5: (VI) *Privacy Budget.* A surprisingly high number of tools implement privacy "odometers" without a logic to block queries after exceeding the budget. *Guidance*: [42, 78, 90].
Gap 6: (VII) *Accuracy Adjustment.* While most user interfaces provide some form of accuracy calculation and visualization, many other tools overlook such feature. *Guidance*: [77, 101].
Gap 7: (IX & X) *Functionality for Data Stewards.* Only a few tools enable data stewards and owners to (IX) annotate sensitive data and (X) define and enforce access controls. Developers do not need to reinvent the wheel, as they adopt current best practices from popular cloud platforms [9, 38, 75]. *Guidance*: [56, 60, 90].

Given that most functional requirements are fulfilled in components across tools, we conclude that *engineering efforts are within striking distance.* To complement these building blocks, we offer an early stage, high-level system design blueprint in Appendix H. The blueprint aims to spark interest in practitioners to develop holistic analytics tooling that follows the identified key system desiderata.

8 FURTHER CHALLENGES

Beyond the engineering and organizational challenges discussed in the previous sections, there exist other critical technical challenges in DP. In combination: Managing privacy budgets on large-scale user data streams [68] with unknown domains and user contributions on multiple records [4] across different systems while adapting the noise level as the budget diminishes. Furthermore, fitting a mathematical model to such a system's semantics and verifying DP fulfillment with, e.g., unit tests, poses additional difficulties [64]. Additionally, DP might not be *fair* [62] in some use cases where a DP calculation determines a critical outcome, e.g., a user's financial support in an underwriting model (see challenges in Appendix G).

Our work highlights the challenges blocking the broader adoption of DP in organizations' workflows. Dwork et al. [27] partly studied these challenges by interviewing DP experts, while our study brings non-experts into the discussion. Dwork et. al distilled four main challenges from their interviews (section 3.6 [27]), which overlap with a few of our findings: (i) *Part of the challenges deploying DP were design based.* In section 7, we highlight that current DP tools still require engineering effort to be easily deployable in organizations. (ii) *DP deployment complexity is also institutionally based.* A common theme of the interviewed companies was their intricate networks of stakeholders and processes, which hamper goal alignment and technology deployments. (iii) *There was no consistency in DP approaches across institutions, indicating a need for shared learning.* One of our conclusions in section 7 signals that tool designers can learn from one another. (iv) *Transparency and testable privacy statements can benefit companies in the regulatory landscape.* Similarly, section 7 advocates for transparency in system designs and moving towards DP-centered systems and away from syntactic privacy definitions that only guarantee *factual anonymity.*
Future work. We suggest privacy practitioners fill the gaps highlighted in section 7 and tackle the challenges of Appendix G. Moreover, specifically for privacy researchers, we encourage (i) improving guidance on selecting ε [27, 56] and (ii) studying and communicating to non-experts how mechanism designs affect utility. For example, studying how output consistency can imbue bias [2, 103, 104] or floating-point protection may provide less utility. As new DP deployments increasingly resort to more complicated algorithms [69], we suggest (iii) studying the unpredictable artifacts these algorithms may introduce (e.g., in the 2020 US Census [34]). Lastly, we encourage improving current proposals of differentially private (iv) ML and (v) synthetic data generation.

9 CONCLUSION

We conclude that DP can improve the work of data scientists across industries by enabling sensitive data exploration across silos, potentially shortening data access times by relaxing the adversity of data request processes, and can fulfill some types of use cases. Furthermore, analysts meaningfully and frequently employed analyses amenable to DP and, on average, would feel comfortable with a 98% of accuracy. Therefore, we suggest companies focus on privacy-enhancing analysis to harvest these benefits, not mainly on security. Moreover, we regard enabling analysts to work without "*seeing*" data and providing analysis accuracy expectation as critical, multifaceted challenges for the research community to solve. We also highlight that current open-source tools do not facilitate easy deployments, a problem requiring engineering effort within striking distance. Consequently, we encourage the community of privacy practitioners to tackle this engineering problem and ease deployments by enabling interactive dashboards, accuracy expectation measurements, improving usability and security, and integration of data annotation and access control capabilities, for ultimately bridging the identified gap between theory and practice.

ACKNOWLEDGMENTS

We sincerely thank the BMW Group for generously funding this project.

REFERENCES

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep Learning with Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (Vienna, Austria) *(CCS '16)*. Association for Computing Machinery, New York, NY, USA, 11 pages. https://doi.org/10.1145/2976749.2978318

[2] John Abowd, Robert Ashmead, and Ryan Cumings-Menon. 2021. An Uncertainty Principle is a Price of Privacy-Preserving Microdata. *NeurIPS* (2021), 13. https://arxiv.org/abs/2110.13239

[3] ACM. 1947. ACM Digital Library. https://dl.acm.org/. Online; accessed 20 May 2022.

[4] Kareem Amin, Jennifer Gillenwater, Matthew Joseph, Alex Kulesza, and Sergei Vassilvitskii. 2022. Plume: Differential Privacy at Scale. https://doi.org/10.48550/ARXIV.2201.11603

[5] Apache Spark. 2014. PySpark Documentation. https://spark.apache.org/docs/latest/api/python/. Online; accessed 6 Mar 2022.

[6] Apple and Google. 2021. Exposure Notification Privacy-preserving Analytics (ENPA). https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ENPA_White_Paper.pdf. Online; accessed 15 August 2022.

[7] Apple Differential Privacy Team. 2017. Learning with privacy at scale. http://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html. Online; accessed 18 February 2022.

[8] Maryam Archie, Sophie Gershon, Abigail Katcoff, and Aileen Zeng. 2018. Who's Watching? De-anonymization of Netflix Reviews using Amazon Reviews.

[9] AWS. 2006. Analytics on AWS. https://aws.amazon.com/big-data/datalakes-and-analytics/. Online; accessed 7 Mar 2022.

[10] Assaf Ben-David, Noam Nisan, and Benny Pinkas. 2008. FairplayMP: A System for Secure Multi-Party Computation. In *Proceedings of the 15th ACM Conference on Computer and Communications Security* (Alexandria, Virginia, USA) *(CCS '08)*. Association for Computing Machinery, New York, NY, USA, 257âĂŞ266. https://doi.org/10.1145/1455770.1455804

[11] Ben-David, Assaf and Nisan, Noam and Pinkas, Benny. 2015. FairplayMP repository. https://github.com/FaiplayMP/FairplayMP/blob/master/runtime/src/utils/PRG.java#L11. Online; accessed 29 July 2021.

[12] Daniel M Bittner, Alejandro E Brito, Mohsen Ghassemi, Shantanu Rane, Anand D Sarwate, and Rebecca N Wright. 2021. Understanding Privacy-Utility Trade-offs in Differentially Private Online Active Learning. *Journal of Privacy and Confidentiality* 10, 2 (2021). https://doi.org/10.29012/jpc.720

[13] Courtney Bowman, Ari Gesher, John K. Grant, and Daniel Slate. 2015. *The Architecture of Privacy*. O'Reilly.

[14] Mark Bun and Thomas Steinke. 2016. Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In *Theory of Cryptography*, Martin Hirt and Adam Smith (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 635–658.

[15] Andre Calero Valdez and Martina Ziefle. 2019. The usersâĂŹ perspective on the privacy-utility trade-offs in health recommender systems. *International Journal of Human-Computer Studies* 121 (2019), 108–121. https://doi.org/10.1016/j.ijhcs.2018.04.003 Advances in Computer-Human Interaction for Recommender Systems.

[16] ClÃĬment L. Canonne, Gautam Kamath, and Thomas Steinke. 2021. The Discrete Gaussian for Differential Privacy. *arXiv:2004.00010 [cs, stat]* (Jan. 2021). http://arxiv.org/abs/2004.00010 arXiv: 2004.00010.

[17] Dingfan Chen, Tribhuvanesh Orekondy, and Mario Fritz. 2020. GS-WGAN: A Gradient-Sanitized Approach for Learning Differentially Private Generators. *Neural Information Processing Systems (NeurIPS)* (2020).

[18] Rui Chen, Qian Xiao, Yu Zhang, and Jianliang Xu. 2015. Differentially Private High-Dimensional Data Publication via Sampling-Based Inference. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Sydney, NSW, Australia) *(KDD '15)*. Association for Computing Machinery, New York, NY, USA, 10 pages. https://doi.org/10.1145/2783258.2783379

[19] Cornell University. 1991. ArXiv. https://arxiv.org/. Online; accessed 28 Jul 2022.

[20] Tech Crunch. 2020. France fines Google and Amazon. https://uk.news.yahoo.com/france-fines-google-120m-amazon-085553384.html. Online; accessed 8 April 2022.

[21] Sabrina De Capitani Di Vimercati, Sara Foresti, Giovanni Livraga, and Pierangela Samarati. 2012. Data privacy: Definitions and techniques. *International Journal of Uncertainty, Fuzziness and Knowlege-Based Systems* 20, 6 (2012), 793–817. https://doi.org/10.1142/S0218488512400247

[22] Damien Desfontaines, James Voss, Bryant Gipson, and Chinmoy Mandayam. 2020. Differentially private partition selection. https://doi.org/10.48550/ARXIV.2006.03684

[23] Differential Privacy Team, Google. 2022. Privacy Loss Distribution. https://eprint.iacr.org/2018/820.pdf. Online; accessed 23 August 2022.

[24] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. 2017. Collecting Telemetry Data Privately. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA) *(NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 3574-3583.

[25] Josep Domingo-Ferrer, David SÃąnchez, and Alberto Blanco-Justicia. 2021. The limits of differential privacy (and its misuse in data release and machine learning). *Commun. ACM* 64, 7 (2021), 33–35. https://doi.org/10.1145/3433638

[26] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. 2006. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *Advances in Cryptology - EUROCRYPT 2006*, Serge Vaudenay (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 486–503.

[27] Cynthia Dwork, Nitin Kohli, and Deirdre Mulligan. 2019. Differential Privacy in Practice: Expose your Epsilons! *Journal of Privacy and Confidentiality* 9, 2 (2019). https://doi.org/10.29012/jpc.689

[28] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284. https://link.springer.com/chapter/10.1007/11681878-14 Online; accessed 30 December 2021.

[29] Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. 2017. Exposed! A Survey of Attacks on Private Data. *Annual Review of Statistics and Its Application* 4, 1 (2017), 61–84. https://doi.org/10.1146/annurev-statistics-060116-054123 Publisher: Annual Reviews.

[30] European Parliament and Council of the European Union. 4 May 2016. REGULATION (EU) 2016/679 Directive 95/46/EC (General Data Protection Regulation): General Data Protection Regulation. , 88 pages. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679

[31] Marco Gaboardi, Michael Hay, and Salil Vadhan. 2020. A Programming Framework for OpenDPâĹŰ. (2020), 31.

[32] Marco Gaboardi, James Honaker, Gary King, Kobbi Nissim, Jonathan Ullman, Salil Vadhan, and Jack Murtagh. 2016. PSI (ĺĺ): a Private data Sharing Interface. In *Theory and Practice of Differential Privacy*. New York, NY. https://arxiv.org/abs/1609.04340

[33] Xianyi Gao, Bernhard Firner, Shridatt Sugrim, Victor Kaiser-Pendergrast, Yulong Yang, and Janne Lindqvist. 2014. Elastic pathing: your speed is enough to track you. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '14 Adjunct*. ACM Press, Seattle, Washington, 975–986. https://doi.org/10.1145/2632048.2632077

[34] Simson L. Garfinkel, John M. Abowd, and Sarah Powazek. 2018. Issues Encountered Deploying Differential Privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society* (Toronto, Canada) *(WPES'18)*. Association for Computing Machinery, New York, NY, USA, 133-137. https://doi.org/10.1145/3267323.3268949

[35] Rod Garratt and Maarten R.C. van Oordt. 2018. Privacy as a Public Good: A Case for Electronic Cash. *Journal of Political Economy* (2018). https://doi.org/10.1086/714133

[36] Gonzalo Munilla Garrido, Joseph Near, Aitsam Muhammad, Warren He, Roman Matzutt, and Florian Matthes. 2021. Do I Get the Privacy I Need? Benchmarking Utility in Differential Privacy Libraries. (2021). arXiv:2109.10789 http://arxiv.org/abs/2109.10789

[37] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. 2012. Universally Utility-maximizing Privacy Mechanisms. *SIAM J. Comput.* 41, 6 (2012), 1673–1693. https://doi.org/10.1137/09076828X

[38] Google. 2008. Website. https://cloud.google.com/. Online; accessed 8 Mar 2022.

[39] Google. 2019. TensorFlow Privacy repository. https://github.com/tensorflow/privacy. Online; accessed 8 May 2022.

[40] Google. 2020. ZetaSQL repository. https://github.com/google/zetasql. Online; accessed 22 August 2022.

[41] Google. 2021. Google DP repository. https://github.com/google/differential-privacy. Online; accessed 22 August 2022.

[42] Google. 2021. Privacy on Beam repository. https://github.com/google/differential-privacy/tree/main/privacy-on-beam. Online; accessed 13 June 2022.

[43] Google Brain. 2015. TensorFlow Random Number Generator. https://www.tensorflow.org/api-docs/python/tf/random/Generator. Online; accessed 23 August 2022.

[44] Google Brain. 2019. TensorFlow Privacy Gaussian Query. https://github.com/tensorflow/privacy/blob/e826ec717a8dd93542aa7038868c8a75213836a9/tensorflow-privacy/privacy/dp-query/gaussian-query.py#L20. Online; accessed 23 August 2022.

[45] Sivakanth Gopi, Pankaj Gulhane, Janardhan Kulkarni, Judy Hanwen Shen, Milad Shokouhi, and Sergey Yekhanin. 2020. Differentially Private Set Union. https:

//doi.org/10.48550/ARXIV.2002.09745
[46] Samuel Haney, Daniel Desfontaines, Luke Hartman, Ruchit Shrestha, and Michael Hay. 2022. Precision-based attacks and interval refining: how to break, then fix, differential privacy on finite computers. https://tpdp.journalprivacyconfidentiality.org/2022/papers/HaneyDHSH22.pdf. *Theory and Practice of Differential Privacy, ICML 2022* (2022). Online; accessed 22 August 2022.
[47] Harvard. 2021. OpenDP repository. https://github.com/opendp/opendp. Online; accessed 22 August 2022.
[48] Muneeb Ul Hassan, Mubashir Husain Rehmani, and Jinjun Chen. 2020. Differential Privacy Techniques for Cyber Physical Systems: A Survey. *IEEE Communications Surveys Tutorials* 22, 1 (2020), 746–789. https://doi.org/10.1109/COMST.2019.2944748
[49] Michael Hay, Ashwin Machanavajjhala, Gerome Miklau, Yan Chen, Dan Zhang, and George Bissias. 2016. Exploring Privacy-Accuracy Tradeoffs Using DPComp. In *Proceedings of the 2016 International Conference on Management of Data* (San Francisco, California, USA) *(SIGMOD '16)*. Association for Computing Machinery, New York, NY, USA, 2101âĂŞ2104. https://doi.org/10.1145/2882903.2899387
[50] J. Hoelzel. 2019. Differential Privacy and the GDPR. 5, 2 (2019), 184–196. https://doi.org/10.21552/edpl/2019/2/8
[51] Naoise Holohan and Stefano Braghin. 2021. Secure Random Sampling in Differential Privacy. In *Computer Security âĂŞ ESORICS 2021*, Elisa Bertino, Haya Shulman, and Michael Waidner (Eds.). Vol. 12973. Springer International Publishing, Cham, 523–542. https://doi.org/10.1007/978-3-030-88428-4-26 Series Title: Lecture Notes in Computer Science.
[52] IBM. 2020. diffprivlib repository. https://github.com/IBM/differential-privacy-library. Online; accessed 22 August 2022.
[53] IEEE. 1963. IEEE Xplore. https://ieeexplore.ieee.org/Xplore/home.jsp. Online; accessed 20 May 2022.
[54] Insider. 2021. Amazon has been fined a record $887 million. https://www.businessinsider.com/amazon-eu-fine-data-privacy-gdpr-luxembourg-european-union-2021-7. Online; accessed 8 April 2022.
[55] Honglu Jiang, Jian Pei, Dongxiao Yu, Jiguo Yu, Bei Gong, and Xiuzhen Cheng. 2021. Applications of Differential Privacy in Social Network Analysis: A Survey. *IEEE Transactions on Knowledge and Data Engineering* (2021), 1–1. https://doi.org/10.1109/TKDE.2021.3073062
[56] Mark F. St. John, Grit Denker, Peeter Laud, Karsten Martiny, Alisa Pankova, and Dusko Pavlovic. 2021. Decision Support for Sharing Data using Differential Privacy. In *2021 IEEE Symposium on Visualization for Cyber Security (VizSec)*. 26–35. https://doi.org/10.1109/VizSec53666.2021.00008
[57] John M. Abowd, Gary L. Benedetto, Simson L. Garfinkel, Scot A. Dahl, Aref N. Dajani, Matthew Graham, Michael B. Hawes, et al. 2021. The modernization of statistical disclosure limitation at the U.S. Census bureau. https://www.census.gov/library/working-papers/2020/adrm/modernization-statistical-disclosure-limitation.html. Online; accessed 18 February 2022.
[58] Alistair E W Johnson, David J Stone, Leo A Celi, and Tom J Pollard. 2017. The MIMIC Code Repository: enabling reproducibility in critical care research. *Journal of the American Medical Informatics Association* 25, 1 (09 2017), 32–39. https://doi.org/10.1093/jamia/ocx084 arXiv:https://academic.oup.com/jamia/article-pdf/25/1/32/34149701/ocx084.pdf
[59] Noah Johnson, Joseph P. Near, and Dawn Song. 2018. Towards Practical Differential Privacy for SQL Queries. *Proc. VLDB Endow.* 11, 5 (January 2018), 526-539. https://doi.org/10.1145/3177732.3177733
[60] Joseph P. Near. 2020. Chorus repository. https://github.com/uvm-plaid/chorus. Online; accessed 22 August 2022.
[61] Nesrine Kaaniche and Maryline Laurent. 2016. Attribute-Based Signatures for Supporting Anonymous Certification. In *Computer Security – ESORICS 2016*, Ioannis Askoxylakis, Sotiris Ioannidis, Sokratis Katsikas, and Catherine Meadows (Eds.). Springer International Publishing, Cham, 279–300. https://www.semanticscholar.org/paper/Attribute-Based-Signatures-for-Supporting-Anonymous-Kaaniche-Laurent-Maknavicius/3b0624ff32b9258ca2351c894d320d83a546fcd6
[62] Christopher T. Kenny, Shiro Kuriwaki, Cory McCartan, Evan T. R. Rosenman, Tyler Simko, and Kosuke Imai. 2021. The use of differential privacy for census data and its impact on redistricting: The case of the 2020 U.S. Census. 7, 41 (2021), eabk3283. https://doi.org/10.1126/sciadv.abk3283
[63] Krishnaram Kenthapadi and Thanh T. L. Tran. 2018. PriPeARL: A Framework for Privacy-Preserving Analytics and Reporting at LinkedIn. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (Torino, Italy) *(CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 2183-2191. https://doi.org/10.1145/3269206.3272031
[64] Daniel Kifer, Solomon Messing, Aaron Roth, Abhradeep Thakurta, and Danfeng Zhang. 2020. Guidelines for Implementing and Auditing Differentially Private Systems. https://doi.org/10.48550/ARXIV.2002.04049
[65] Jong Wook Kim, Kennedy Edemacu, Jong Seon Kim, Yon Dohn Chung, and Beakcheol Jang. 2021. A Survey Of differential privacy-based techniques and their applicability to location-Based services. 111 (2021), 102464. https://doi.org/10.1016/j.cose.2021.102464
[66] Knime. 2006. Website. https://www.knime.com/. Online; accessed 7 Mar 2022.
[67] Daniel Kondor, Behrooz Hashemian, Yves-Alexandre de Montjoye, and Carlo Ratti. 2020. Towards Matching User Mobility Traces in Large-Scale Datasets. *IEEE Transactions on Big Data* 6, 4 (Dec. 2020), 714–726. https://doi.org/10.1109/TBDATA.2018.2871693
[68] Mathias Lécuyer, Riley Spahn, Kiran Vodrahalli, Roxana Geambasu, and Daniel Hsu. 2019. Privacy Accounting and Quality Control in the Sage Differentially Private ML Platform *(SOSP '19)*. Association for Computing Machinery, New York, NY, USA, 181âĂŞ195. https://doi.org/10.1145/3341301.3359639
[69] Chao Li, Michael Hay, Gerome Miklau, and Yue Wang. 2014. A Data- and Workload-Aware Algorithm for Range Queries under Differential Privacy. *Proc. VLDB Endow.* 7, 5 (Jan. 2014), 341âĂŞ352. https://doi.org/10.14778/2732269.2732271
[70] Zekun Li and Shuyu Li. 2017. Random forest algorithm under differential privacy. In *2017 IEEE 17th International Conference on Communication Technology (ICCT)*. 1901–1905. https://doi.org/10.1109/ICCT.2017.8359960
[71] Min Lyu, Dong Su, and Ninghui Li. 2017. Understanding the Sparse Vector Technique for Differential Privacy. *Proc. VLDB Endow.* 10, 6 (feb 2017), 637–648. https://doi.org/10.14778/3055330.3055331
[72] McKinsey & Company. 2020. Four ways to accelerate the creation of data ecosystems. https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/four-ways-to-accelerate-the-creation-of-data-ecosystems. Online; accessed 18 April 2022.
[73] H. Brendan McMahan, Galen Andrew, Ulfar Erlingsson, Steve Chien, Ilya Mironov, Nicolas Papernot, and Peter Kairouz. 2018. A General Approach to Adding Differential Privacy to Iterative Training Procedures. https://doi.org/10.48550/ARXIV.1812.06210
[74] Meta. 2021. Opacus repository. https://github.com/pytorch/opacus Online; accessed 8 May 2022.
[75] Microsoft. 2010. Azure. https://azure.microsoft.com/en-us/. Online; accessed 28 July 2022.
[76] Ilya Mironov. 2012. On significance of the least significant bits for differential privacy. In *Proceedings of the 2012 ACM conference on Computer and communications security - CCS '12*. ACM Press, Raleigh, North Carolina, USA, 650. https://doi.org/10.1145/2382196.2382264
[77] Priyanka Nanayakkara, Johes Bater, Xi He, Jessica Hullman, and Jennie Rogers. 2022. Visualizing Privacy-Utility Trade-Offs in Differentially Private Data Releases. 2022, 2 (2022), 601–618. https://doi.org/10.2478/popets-2022-0058
[78] Arjun Narayan and Andreas Haeberlen. 2012. DJoin: Differentially Private Join Queries over Distributed Databases. (2012), 14.
[79] Arvind Narayanan and Vitaly Shmatikov. 2008. Robust De-anonymization of Large Sparse Datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*. IEEE, Oakland, CA, USA, 111–125. https://doi.org/10.1109/SP.2008.33 ISSN: 1081-6011.
[80] Office of the Attorney General. 2018. California Consumer Privacy Act (CCPA) | State of California - Department of Justice - Office of the Attorney General. https://oag.ca.gov/privacy/ccpa.
[81] OMTP. 2009. Advanced Trusted Environment: OMTP TR1. , 204 pages. http://www.gsma.com/newsroom/wp-content/uploads/2012/03/omtpadvancedtrustedenvironmentomtptr1v11.pdf
[82] OpenDP. 2021. SmartNoise Core. https://github.com/opendp/smartnoise-core. Online; accessed 22 August 2022.
[83] OpenMined. 2022. PipelineDP repository. https://github.com/OpenMined/PipelineDP. Online; accessed 27 July 2022.
[84] Pandas developer community. 2008. Pandas Documentation. https://pandas.pydata.org/docs/. Online; accessed 6 Mar 2022.
[85] Project Jupyter. 2015. Jupyter Documentation. https://jupyter.org/. Online; accessed 6 Mar 2022.
[86] PyCaret community. 2020. PyCaret open-source repository. https://github.com/pycaret/pycaret. Online; accessed 22 August 2022.
[87] IBM report. 2022. Cost of a Data Breach Report. https://www.ibm.com/downloads/cas/OJDVQGRY. Online; accessed 11 April 2022.
[88] Ryan Rogers. 2020. A Differentially Private Data Analytics API at Scale. In *2020 USENIX Conference on Privacy Engineering Practice and Respect (PEPR 20)*. USENIX Association. https://www.usenix.org/conference/pepr20/presentation/rogers
[89] Ryan Rogers, Subbu Subramaniam, Sean Peng, David Durfee, Seunghyun Lee, Santosh Kumar Kancha, Shraddha Sahay, and Parvez Ahammad. 2021. LinkedIn's Audience Engagements API: A Privacy Preserving Data Analytics System at Scale. 11, 3 (2021). https://doi.org/10.29012/jpc.782
[90] Indrajit Roy, Srinath T V Setty, Ann Kilzer, Vitaly Shmatikov, and Emmett Witchel. 2010. Airavat: Security and Privacy for MapReduce. (2010), 16.

[91] Roy, Indrajit and Setty, Srinath T V and Kilzer, Ann and Shmatikov, Vitaly and Witchel, Emmett. 2013. Airavat repository. https://github.com/bochunz/Airavat/blob/master/AiravatJvmTest.java#L11. Online; accessed 29 July 2021.
[92] ScienceDirect. 1997. ScienceDirect Digital Library. https://www.sciencedirect.com/. Online; accessed 20 May 2022.
[93] Lingling Shen, Xiaotong Wu, Datong Wu, Xiaolong Xu, and Lianyong Qi. 2020. A Survey on Randomized Mechanisms for Statistical Learning under Local Differential Privacy. In *2020 IEEE 22nd International Conference on High Performance Computing and Communications; IEEE 18th International Conference on Smart City; IEEE 6th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*. 1195–1202. https://doi.org/10.1109/HPCC-SmartCity-DSS50907.2020.00155
[94] Latanya Sweeney, Akua Abu, and Julia Winn. 2013. Identifying Participants in the Personal Genome Project by Name. *SSRN Electronic Journal* (2013). https://doi.org/10.2139/ssrn.2257732
[95] Tableau Software, Inc. 2003. Tableau website. https://www.tableau.com/. Online; accessed 6 Mar 2022.
[96] Yuchao Tao, Ryan McKenna, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau. 2021. Benchmarking Differentially Private Synthetic Data Generation Algorithms. https://doi.org/10.48550/ARXIV.2112.09238
[97] Pratiksha Thaker, Mihai Budiu, Parikshit Gopalan, Udi Wieder, and Matei Zaharia. 2020. Overlook: Differentially Private Exploratory Visualization for Big Data. In *Workshop on Theory and Practice of Differential Privacy*. https://arxiv.org/abs/2006.12018
[98] Reihaneh Torkzadehmahani, Peter Kairouz, and Benedict Paten. 2019. Dp-cgan: Differentially private synthetic data and label generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*.
[99] Tumult Labs. 2021. Tumult Analytics repository. https://gitlab.com/tumult-labs/analytics. Online; accessed 15 August 2022.
[100] United States Court of Appeals, Ninth Circuit. 2012. United States v. Huping Zhou. https://caselaw.findlaw.com/us-9th-circuit/1600563.html. Online; accessed 28 July 2022.
[101] Elisabet Lobo Vesga, Alejandro Russo, and Marco Gaboardi. 2020. A Programming Framework for Differential Privacy with Accuracy Concentration Bounds. *2020 IEEE Symposium on Security and Privacy (SP)* (2020), 411–428.
[102] Lun Wang, Joseph P. Near, Neel Somani, Peng Gao, Andrew Low, David Dao, and Dawn Song. 2019. Data Capsule: A New Paradigm for Automatic Compliance with Data Privacy Regulations. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare*, Vijay Gadepally, Timothy Mattson, Michael Stonebraker, Fusheng Wang, Gang Luo, Yanhui Laing, and Alevtina Dubovitskaya (Eds.). Springer International Publishing, Cham, 3–23.
[103] Tianhao Wang, Zitao Li, Ninghui Li, Milan Lopuhaä-Zwakenberg, and Boris Skoric. 2019. Consistent and Accurate Frequency Oracles under Local Differential Privacy. *CoRR* abs/1905.08320 (2019). arXiv:1905.08320 http://arxiv.org/abs/1905.08320
[104] Tianhao Wang, Milan Lopuhaa-Zwakenberg, Zitao Li, Boris Skoric, and Ninghui Li. 2020. Locally Differentially Private Frequency Estimation with Consistency. In *Proceedings 2020 Network and Distributed System Security Symposium* (San Diego, CA, 2020). Internet Society. https://doi.org/10.14722/ndss.2020.24157
[105] Yufeng Wang, Minjie Huang, Qun Jin, and Jianhua Ma. 2018. DP3: A Differential Privacy-Based Privacy-Preserving Indoor Localization Mechanism. *IEEE Communications Letters* 22, 12 (2018), 2547–2550. https://doi.org/10.1109/LCOMM.2018.2876449
[106] Matthew Wilchek and Yingjie Wang. 2021. Synthetic Differential Privacy Data Generation for Revealing Bias Modelling Risks. In *2021 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking (ISPA/BDCloud/SocialCom/SustainCom)*. 1574–1580. https://doi.org/10.1109/ISPA-BDCloud-SocialCom-SustainCom52081.2021.00211
[107] Royce J Wilson, Celia Yuxin Zhang, William Lam, Damien Desfontaines, Daniel Simmons-Marengo, and Bryant Gipson. 2019. Differentially Private SQL with Bounded User Contribution. https://doi.org/10.48550/ARXIV.1909.01917
[108] Felix T Wu. 2012. Defining Privacy and Utility in Data Sets. *84 University of Colorado Law Review 1117 (2013); 2012 TRPC* (2012), 1117–1177. https://doi.org/10.2139/ssrn.2031808
[109] Bangzhou Xin, Wei Yang, Yangyang Geng, Sheng Chen, Shaowei Wang, and Liusheng Huang. 2020. Private FL-GAN: Differential Privacy Synthetic Data Generation Based on Federated Learning. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2927–2931. https://doi.org/10.1109/ICASSP40776.2020.9054559
[110] Chugui Xu, Ju Ren, Deyu Zhang, Yaoxue Zhang, Zhan Qin, and Kui Ren. 2019. GANobfuscator: Mitigating Information Leakage Under GAN via Differential Privacy. *IEEE Transactions on Information Forensics and Security* 14, 9 (2019), 2358–2371. https://doi.org/10.1109/TIFS.2019.2897874
[111] Depeng Xu, Shuhan Yuan, and Xintao Wu. 2017. Differential Privacy Preserving Causal Graph Discovery. In *2017 IEEE Symposium on Privacy-Aware Computing (PAC)*. 60–71. https://doi.org/10.1109/PAC.2017.24
[112] Jun Zhang, Graham Cormode, Cecilia M. Procopiuc, Divesh Srivastava, and Xiaokui Xiao. 2017. PrivBayes: Private Data Release via Bayesian Networks. *ACM Trans. Database Syst.* 42, 4, Article 25 (oct 2017), 41 pages. https://doi.org/10.1145/3134428
[113] Jun Zhang, Xiaokui Xiao, Y. Yang, Zhenjie Zhang, and Marianne Winslett. 2013. PrivGene: differentially private model fitting using genetic algorithms. In *SIGMOD '13*.
[114] Ding Zhe, Chunwang Wu, Zhao Jun, and Binyong Li. 2020. Frequent Itemsets Mining Algorithm based On Differential Privacy and FP-Tree. In *2020 17th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP)*. 271–274. https://doi.org/10.1109/ICCWAMTIP51612.2020.9317373

A INTERVIEWED COMPANIES OVERVIEW

Table 2 presents a summary of the characteristics of the interviewed companies. The companies belong to a diverse set of industries, predominantly SW development, and 4 of the 9 companies are significantly large, with over 100, 000 employees. 5 companies are under the jurisdiction of the EU with regulations such as the General Data Protection Regulation [30], and 4 companies operate under US law, e.g., the California Consumer Privacy Act [80].

Table 2: Overview of the interviewed companies. Legend: SW dev. = Software development

Industry (focus)	Size (employees)	Team's Location
Team operates internationally		
Automotive (car manufacturer)	$> 100,000$	Germany
Insurance (health)	$> 100,000$	Germany
SW dev. (data processing)	$< 2,000$	Germany
SW dev. (subscription newsletters)	$< 2,000$	USA
Team operates nationally		
Consultancy (banking and big pharma)	$> 100,000$	Spain
Entertainment (finance)	$< 2,000$	USA
SW dev. (business operations)	$> 100,000$	Germany
SW dev. (data processing)	$< 2,000$	USA
SW dev. (smart sound system)	$< 2,000$	USA

B INTERVIEW QUESTIONNAIRE FOR DATA STEWARDS

RQ1: *What is the context of privacy protection in the targeted organization?*

Q1: *What is the institution's motivation for privacy protection?*

Q2: *What are your privacy concerns when an analyst has full dataset access?*

Q3: *At what level of data granularity are you protecting and measuring privacy?*

Q4: *What could be improved in the dataset request process?*

Q5: *What are your typical questions for the current interview-based full dataset access authorization?*

Q6: *Instead of the interview process, would you be capable to run a program provided by the analyst s.t. the analysis is carried out without the analyst ever "seeing" the dataset?*

C INTERVIEW QUESTIONNAIRE FOR DATA ANALYSTS

As we interviewed non-experts in differential privacy, we minimized the number of questions that contained the words or required knowledge of "*differential privacy*". We kept a few because we aimed to assess whether systems offering differential privacy functionality could be valuable to analysts. First, we briefly explained differential privacy in a simplified manner: "*Differential privacy is a technique that adds noise to analytics results so that one cannot reverse engineer the outputs to a specific person.*" Additionally, if we perceived the interviewees were disoriented with Q18 or Q19, we explained that the hypothetical system would be the same as the one they used every day, the only difference being that the results would slightly differ from the deterministic outputs. Picturing the system they used daily was very helpful for imagining one where the outputs are noisy. Furthermore, we carefully parsed their answers to assess whether they understood the concept or its integration into their system. If they did not, we kindly repeated the procedure above.

RQ2: *Could differential privacy tackle the privacy-related pain points of an analysis workflow in an organization?*
- **Q7:** *What is your workflow to analyze data?*
- **Q8:** *Why do you need full dataset access?*
- **Q9:** *How often do you request full dataset access? How long does it usually take?*
- **Q10:** *What do you think about the process to request full dataset access in your organization?*
- **Q11:** *What features do you think are missing in your organization's data analysis workflow?*

RQ3: *When does differential privacy impede an analysis?*
- **Q12:** *In which analytics use cases have you been involved?*
- **Q13:** *Is SQL-meaningful for your work? How many SQL-like queries do you make weekly?*
- **Q14:** *How often do you need machine learning to fulfill your analysis in contrast to using SQL?*
- **Q15:** *What are your most used machine learning models?*
- **Q16:** *If you were to use differential privacy to fulfill your analysis, when and how much accuracy would you be willing to forgo?*

RQ4: *How would differential privacy affect the workflow of an analyst?*
- **Q17:** *How much would the noise affect your analysis?*
- **Q18:** *Would you find it helpful to execute differentially private SQL queries to explore and fully analyse datasets without the standard permissions?*
- **Q19:** *Only based on the information extracted from a dataset exploration with differential privacy, could you write a script to fulfill your analysis goal?*
- **Q20:** *What are the minimum properties for you as an analyst such that you are confident to write an analysis script without full dataset access?*
- **Q21:** *Would you find it helpful to use a dynamic dashboard that visualizes dataset information with differential privacy?*

RQ5: *Can differential privacy be applied to the frequent SQL-like queries analysts execute?*
- **Q22:** *What are your top SQL-like queries before you have full dataset access?*
- **Q23:** *What are your top SQL-like queries after you have full dataset access?*
- **Q24:** *What is the ratio between aggregation queries and queries to retrieve items?*

D FREQUENT QUERIES

In Table 3, we include the most frequent SQL queries recorded during the interviews with the data analysts before and after accessing a dataset. Note that not all analysts were allowed to explore datasets and a few did not employ SQL for data preparation or analytics; instead, they resorted to Python scripts for statistical analysis, ML, and visualization or tools such as Tableau [95], Knime [66], or proprietary SAP data management software. For exploring the dataset prior to access, 14 analysts resourced to `SELECT *` to get a "*feeling*" for the data. Furthermore, `COUNT` and `DISTINCT`, and `WHERE` and `GROUP BY` were the most frequently used functions and clauses, respectively.

Table 3: Most frequent queries before (data exploration) and after (data preparation/analysis) data access. Legend: Freq. = Frequency (i.e., number of analysts who used such query).

Query	Freq. before access	Freq. after access
Function		
`COUNT`	7	7
`DISTINCT`	6	4
`MAX`	4	5
`MIN`	4	5
`AVG`	4	4
`VAR`	2	3
Statement		
`SELECT * LIMIT`	14	12
Clause		
`WHERE`	13	10
`GROUP BY`	12	9
`JOIN`	2	8

E FREQUENTLY ASKED QUESTIONS FROM DATA STEWARDS TO ANALYSTS

We compiled the most frequently asked questions data stewards make to data analysts during the data access request process.

- Could you describe in detail the analytics use case?
- Is the use case approved by the corresponding internal stakeholders?
- Why is the dataset needed?
- Is the dataset adequate regarding quality, volume, and use case?
- Could you reach the goal without dataset access?
- Is the entire dataset needed or only a set of attributes?
- Is the dataset already available, or must a data engineer create a new dataset?

- Is the dataset classified as very sensitive? If affirmative, additional access control measures and monitoring must be defined in detail.

F OPEN-SOURCE TOOLS DESCRIPTIONS

We provide a quick description of each of the selected open-source tools mapped to the key system desiderata in section 7 appearing in Table 1.

Libraries

diffprivlib: IBM researchers developed a general-purpose Python library to execute differentially private aggregation queries and machine learning in the context of data science (namely Notebooks) [52].

Google DP: Google researchers developed a library in multiple languages (C++, Go, and Java) that an expert may use to build new applications supporting differential privacy [41].

Opacus: Meta researchers developed a library dedicated to training machine learning models offered by PyTorch in a differentially private manner [74].

OpenDP: Harvard implemented a flexible architecture for differentially private analysis, consisting of a (pluggable) runtime in Rust wrapped around a Python API, in addition to a "validator" that calculates parameters such as the sensitivity of a query. [47].

TensorFlow Privacy: Google researchers developed a library that includes TensorFlow differentially-private optimizers for training machine learning models [39].

Frameworks

Chorus: Johnson et al.'s [59, 60] wrote a framework in Scala that works in cooperation with existing infrastructure (a SQL database) to explore the use of differentially private SQL queries at scale.

PipelineDP (experimental): OpenMined, in collaboration with Google, propose a framework to execute differentially private aggregations in large-scale datasets using batch processing systems (Apache Spark and Apache Beam) [83].

Privacy on Beam: Similarly to PipelineDP, Privacy on Beam [42] proposes a solution based on Apache Beam and Google DP [41] for executing differentially private analytics at scale.

Tumult Analytics: Tumult Labs provides a Python library built atop a framework similar to OpenDP for computing aggregate statistics over tabular data at scale [99].

ZetaSQL: Google researchers wrote a framework for SQL that defines a language, a parser, and an analyzer meant to work with an existing database engine [40].

Systems:

Airavat: Roy et al. [90] designed a MapReduce-base system written in Java for distributed computations on sensitive data that integrates differential privacy and access control with policies defined by data owners/stewards.

DJoin: Narayan et al. [78] built a system capable of processing a wide range of differentially private SQL queries across datasets from different organizations and leverages homomorphic primitives to hide inputs.

User Interfaces:

Bittner et. al [12]: With a focus on ML, Bittner et. al aim to help researchers decide which algorithm to use by offering an interface that quantifies the disclosure risk of different algorithms.

DPcomp: A web-based system enabling researchers to assess the utility of differentially private algorithms and understand their respective incurred error [49].

DPP: This user interface specifically helps data owners to set the noise level per the disclosure risk of an attribute. The underlying mechanism relies on a novel parameter selection procedure for differential privacy [56].

Overlook: Thaker et al. [97] designed a system for differentially private data exploration that supports counts with an interactive browser-interfacing dashboard (namely visualizing histograms).

PSI (Ψ): Harvard's Privacy Tools Project works on a data sharing interfaces for researchers to explore datasets with differential privacy [32].

ViP: Visualizing Privacy is an interface that provides information about the relationships between utility, ε, and disclosure risk (among others), allowing users to adjust the privacy parameters of their analysis based on visualizations of expected risk and accuracy [77].

G DIFFERENTIAL PRIVACY CHALLENGES

This section enumerates other critical challenges we encourage researchers and system designers to investigate.

(1) While DP is highly adaptable to use cases (e.g., using the local or central model) and algorithms (e.g., queries or ML), the adaptations are non-trivial and have often led to erroneous implementations [25]. Thus, practitioners should exercise extreme care to ensure the correctness of their DP implementation with the same sentiment as "*do not write your own crypto*."

(2) *Fairness* could be another obstacle to DP adoption, which Harvard researchers also highlighted when referring to the US Census of 2020 [62]. Specifically, one analyst underlined the topic of fairness when asked about how noise would affect their analysis (Q17). If analysts add differentially private noise during training underwriting linear regression models, users might be over- or under-funded. While the company would not incur a loss as the predictions would be "right" on average, the effect noise has on their users could impact their brand perception.

(3) Managing *user-level* privacy budgets in user data streams [68].

(4) Tracking the privacy budget across systems and adapting the noise level based on the remaining budget.

(5) There is a significant difference between the local and central model noise levels.

(6) Choosing ε and other privacy parameters [27].

(7) Building systems that fulfill DP for (i) large-scale datasets (ii) when users make contributions to multiple records (iii) with unknown domains [4].

(8) Verifying DP compliance of a complex system by proving and fitting a mathematical model to the system's semantics [64] and developing unit tests to ensure the system conforms to such model.

H SYSTEM DESIGN

Although there are many potential ways to construct privacy-enhancing analytics systems, to show the feasibility of covering all system desiderata presented in section 7, this section discusses one

design to guide practitioners in their development. The design is in an early stage, and, thus, we cannot discuss the components in detail. Instead, we sketch the system's primary components, aiming to spark interest in further system development and research in the community.

We consider two roles interacting with the system: (i) *data stewards* have the authority to access the original data and the legal background for data management. Stewards can authorize data access inside an organization and ensure compliance. (ii) *Data analysts* analyze data to fulfill use cases. Analysts often need to access data by employing SQL aggregation or retrieval and python scripts. In the current system design, we mainly consider SQL aggregation and retrieval. Lastly, we assume that analysts cannot share query results with unauthorized recipients. In such a setting, we present our high-level system design blueprint in Fig. 2.

The components are the following:

Database Schema: The system requires one dedicated component to manage the database schema and ensure its consistency at all places to make sure all components have a consistent view of the processed data format.

Policy Panel: Data stewards create and update the configuration stored in the policy panel to authorize data access from data analysts to satisfy desiderata **(X)**, annotate data sensitivity to satisfy desiderata **(VIII)** and **(IX)**, and ensure compliance. Other system components rely on the configuration in the policy panel to decide whether to proceed with particular requests or queries.

User Registration Service: The user registration service component maintains a user system to standardize the onboarding procedure of data analysts to satisfy desiderata **(X)**; thus, the system can distinguish between different data analysts with different data access requirements and permissions.

Statistic Dashboard: The statistic dashboard is a privacy-enhanced visualization for database statistics, which will help authorized data analysts explore datasets, thus satisfying desiderata **(V)**.

Query Gateway: The query gateway reads annotated data schema from the policy panel and uses it to analyze the query structure, parsing the query for later stages. The system can thus run a preliminary policy check on the incoming query and route it to the corresponding database proxy.

Original Database: The database service stores the original sensitive data securely to satisfy desiderata **(III)**, ideally with encrypted storage and restricted access for the data steward and other necessary system security components to fully comply with **(III)**.

Budget Manager: With the information about both the database schema and the sensitivity annotation from the policy panel, the budget manager models the differential privacy budget and keeps track of the budget consumption in various queries.

Differentially Private Database Proxy: Before executing the query from the data analyst on the tables in the original database, the proxy analyzes how to apply differential privacy by transforming the query and also calculates the budget consumption to satisfy desiderata **(I), (II)** and **(VI)**. Before returning the query result, it also outputs the query's accuracy estimation to satisfy **(VII)**.

Synthesized Database: The synthesized database maintains the dummy or differential-privately synthesized versions of tables in the original database to satisfy desiderata **(IV)**.

Synthetic Database Proxy: Upon receiving queries to the dummy or differential-privately synthesized data, the proxy checks whether the required version of the tables has already been generated in the synthesized database. If the required version is missing, the proxy orchestrates the generation procedure from the original database on-demand.

Lastly, we describe the communication between system components to explain the workflow to access different privacy-enhancing analytic functionalities.

(1) Data Analyst User Registration. Once the system is correctly set up, data analysts should begin to create their user accounts in the system with the user registration service.

(2) Submission of the Query Request. The request should include both the SQL query and a piece of metadata that specify the privacy details like accuracy requirements or whether to use the dummy or synthesized data. The query gateway checks the query request to see if it is compliant with the system policy and routes it to the corresponding database proxy.

(3) Exploring Data Statistics on the Dashboard. Data analysts can use the dashboard to explore dataset statistics that are periodically gathered from the original database.

(4) Data Steward Adjustment for Data Access Policy. In addition to allowing the data steward to configure the data access policy manually, ideally, the policy panel should also make suggestions on potentially useful policy changes. Such suggestions can be based on the data access application or frequently rejected requests to other system components during user registration.

(5) User Registration Following Policies in the Policy Panel. If the data steward decides to include specific steps during the registration procedure (e.g., signing acknowledgments, reading materials, finishing tutorials), the registration procedure would reflect such requirements.

(6) User Access to Query and Dashboard Service Determined by Policy. Considering both queries and the statistic dashboard exploration reveal information about the original data, the policy panel should control the access of data analysts to both services.

(7) Query Execution With One Proxy. Depending on the privacy-related metadata, one of the proxies executes the query with its transformation and returns the query result with privacy details like result accuracy or budget consumption. If the result consumes the privacy budget, the proxy also notifies the budget manager to track the change.

(8) On-Demand Data Synthesis. If the synthetic database proxy cannot find the required version of the dummy of synthetic tables from the synthesized database, it triggers the generation of that required version.

(9) Unified Schema Synchronization Between System Components. The database schema component enforces consistency by tracking the changes in the original database. It locks the whole system for schema changes until the updates are applied to all system components.

As of September 2022, our GitHub hosts an early-stage open-source effort to benchmark libraries and frameworks suitable for some of the system components in Fig. 2:

https://github.com/camelop/dp-lab

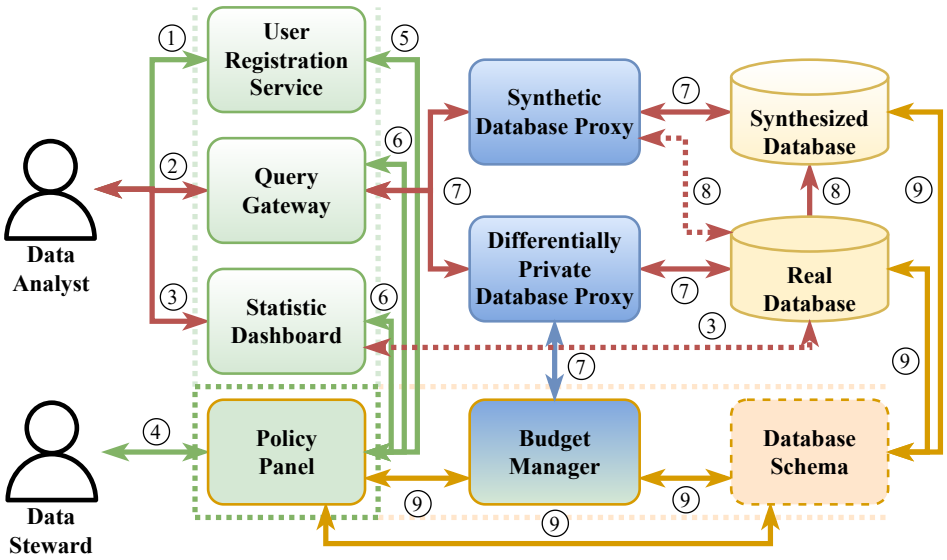

Figure 2: High-level system design blueprint of a privacy-enhancing analytics tool. The components and communication links are described in Appendix H. Solid lines represent communications between components triggered by all relevant query events, while dashed lines represent communications that happen periodically or only under certain circumstances. We specify those circumstances in the workflow description.