

Guiding the Practical Adoption of Federated Machine Learning in Organizations

Tobias Müller

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. Jens Großklags

Prüfende der Dissertation:

1. Prof. Dr. Florian Matthes
2. Prof. Dr.-Ing. Stefan Tai

Die Dissertation wurde am 08.11.2023 bei der Technischen Universität München eingereicht
und durch die TUM School of Computation, Information and Technology am 07.05.2024
angenommen.

Abstract

Problem Statement: The ever-increasing wealth of digitized data powers the disruptive potential of Machine Learning (ML) and its immense economic impact. Still, a significant amount of the world’s data is scattered and locked in data silos. Accessing the siloed data becomes more difficult with privacy concerns and legal regulations, which leaves the full potential and therefore economic opportunity of the stored data largely untapped. Federated Machine Learning (FedML) is a novel model-to-data approach that enables the training of ML models on decentralized, potentially siloed data without the need for direct data sharing. Despite its potential, most FedML projects remain as prototypes or simulations and do not lead to a productive use of the technology. This dissertation investigates this challenge and offers a qualitative exploration of the practical adoption of FedML in organizations. In particular, we address research themes that have not been investigated in previous literature: the influencing factors for a successful practical adoption, the socio-technical aspects of FedML collaborations, and its software development life cycle.

Research Design: We follow an inductive, qualitative research strategy. First, we built on an expert interview study with a reflexive thematic data analysis strategy to explore the influencing factors for the practical adoption of FedML. The resulting insights demonstrated the need to investigate research gaps regarding collaborations in the ML domain and the complex FedML implementation process. Consequently, we investigated the socio-technical challenges of collaborative FedML through a systematic literature review, focus group discussions, and semi-structured expert interviews. Finally, we followed the design science research strategy to develop a high-level process model and detailed activity models for the FedML software development life cycle. Herein, we conducted literature reviews, focus groups, and expert interviews to build our knowledge base, gather requirements, and collect feedback.

Results: First, we reveal the reasons practitioners leverage FedML and offer an exhaustive overview of the risks, challenges, and factors that influence the successful practical adoption of FedML in organizations. Second, this dissertation provides a comprehensive overview of the socio-technical challenges that arise due to the decentralized and collaborative characteristics of FedML. Thereby, we compiled a catalog of guiding questions that need to be considered in the collaboration creation within FedML projects. Third, we developed a high-level process model to provide a comprehensible understanding of the project flow and interrelations within FedML projects. Fourth, we designed detailed activity models to structure and guide practitioners through the corresponding process.

Contribution: This thesis offers qualitative insights into the organizational adoption of FedML and contributes to various literature streams on FedML, ML adoption, and ML software development life cycles. First, we identified the influencing factors for the organizational adoption of FedML. Thereby, we uncovered novel factors specific to FedML and demonstrated how its decentralized paradigm affects prevalent ML adoption factors. Practitioners can leverage these insights to counteract potential challenges preemptively and actively drive the successful adoption of FedML through better-informed decisions. The results also showcase potential research avenues for future work. Second, we compiled a comprehensive overview of the socio-technical

challenges related to collaborations within FedML projects. Organizations can use these results as a reference book for the relevant aspects in the creation of collaboration agreements within FedML projects. Third, we constructed a representation of the FedML software development life cycle and complemented extant ML development life cycles with a perspective on FedML. Practitioners can utilize these artifacts as guidance to enhance their understanding of the project structure and interdependencies. Furthermore, the artifacts can aid practitioners in planning their projects, tracking progress, and identifying the next steps.

Limitations: Our research has several limitations. The conducted literature reviews may be subject to shortcomings since the search process might not cover all relevant papers and the coding might be biased. We mitigated these risks, by establishing rigorous search procedures, performing comprehensive forward as well as backward searches, and letting multiple researchers code the articles independently. As for the focus groups and expert interviews, the researchers might be susceptible to bias and we could only identify a limited sample size. Although we involved a diverse variety of voices and only closed studies after reaching theoretical saturation, the sample sizes limit generalizability. We minimized subjective bias by ensuring observer and data triangulation. Additionally, the evaluation of our artifacts comes with limitations since we could not assess the performance of the artifacts in varying scenarios. Lastly, the scope of our models is tailored to the most prevalent FedML architectural pattern.

Future Research: As our work is exploratory, we encourage future researchers to conduct follow-up studies to quantitatively validate, refine, and extend our results. For example, future studies could determine the significance and correlation of the influencing factors or focus on selected challenges. On this note, we argue that advances on the challenges of data interoperability or clarifying the legal situation of FedML would make a significant contribution to the widespread adoption of FedML. Also, while we identified a multitude of socio-technical challenges, practitioners still lack guidance on how to address these challenges. Hence, performing observational studies and deriving best practices of how current first or early adopters handle these challenges would be a promising avenue for future work. Given the limited scope, our models focus on server-client architectures with a central orchestrating server. As this only represents a fraction of possible architectural patterns, researchers could further investigate how other patterns influence the development process and designing corresponding models.

Zusammenfassung

Motivation: Die ständig wachsende Menge an digitalisierten Daten bildet die Grundlage für den disruptiven Erfolg von Machine Learning (ML) und dessen enormen wirtschaftlichen Auswirkungen. Dennoch ist ein großer Teil der weltweiten Daten verstreut und in Datensilos isoliert. Die Nutzung dieser Daten wird durch Datenschutzbedenken und gesetzliche Vorschriften erschwert, wodurch das volle Potenzial und damit die wirtschaftlichen Möglichkeiten der gespeicherten Daten weitgehend ungenutzt bleiben. Federated Machine Learning (FedML) ist ein neuartiger Ansatz, der es ermöglicht, ML-Modelle auf dezentralen, möglicherweise isolierten Daten zu trainieren, ohne die Daten selbst zu teilen. Trotz dieses Potenzials verbleiben die meisten FedML-Projekte als Prototypen oder Simulationen und führen nicht zu einer produktiven Nutzung der Technologie. Diese Dissertation versucht diese Forschungslücke zu schließen und bietet eine qualitative Untersuchung der praktischen Anwendung von FedML in Organisationen. Angesichts dessen untersuchen wir Forschungsthemen, die in der bisherigen Literatur wenig berücksichtigt wurden: die Einflussfaktoren für eine erfolgreiche praktische Umsetzung, die sozio-technischen Aspekte von FedML Kollaborationen und den Softwareentwicklungslebenszyklus von FedML Systemen.

Forschungsdesign: In dieser Arbeit verfolgen wir eine induktive, qualitative Forschungsstrategie. Mit Hilfe von Experteninterviewstudie mit einer reflexiven thematischen Datenanalysestrategie haben wir die Einflussfaktoren für die praktische Anwendung von FedML untersucht. Unsere Ergebnisse zeigen, dass die meisten Projekte aufgrund von Problemen scheitern, welche durch die Kollaboration und den komplexen FedML Implementierungsprozess entstehen. Daraufhin haben wir die sozio-technischen Herausforderungen von FedML Kollaborationen durch eine systematische Literaturrecherche, Fokusgruppen und Experteninterviews analysiert. Weiterhin nutzten wir Design Science Research für die Entwicklung eines Prozessmodells sowie detaillierte Aktivitätsmodelle für die Softwareentwicklung. Hierfür führten wir Literaturrecherchen, Fokusgruppen und Experteninterviews durch, um unsere Wissensbasis aufzubauen, Anforderungen zu sammeln und Feedback einzuholen.

Ergebnisse: Zunächst beleuchten wir die Gründe, warum Organisationen FedML einsetzen, und bieten einen umfassenden Überblick über die Risiken, Herausforderungen und Faktoren, die den erfolgreichen Einsatz von FedML in Unternehmen beeinflussen. Außerdem geben wir einen umfangreichen Einblick in die sozio-technischen Herausforderungen, die sich aus den dezentralen und kollaborativen Aspekten von FedML ergeben. Dabei haben wir eine Liste von Leitfragen zusammengestellt, die bei der Gestaltung von Kooperationsvereinbarungen in FedML-Projekten berücksichtigt werden sollten. Zusätzlich entwickelten wir ein Prozessmodell, das den Projektablauf und die Zusammenhänge in FedML-Projekten transparent macht. Abschließen erstellten wir detaillierte Aktivitätsmodelle, um Organisationen bei der Umsetzung zu unterstützen.

Beitrag: Diese Arbeit liefert qualitative Einblicke zur organisatorischen Umsetzung von FedML und trägt zur Forschung im Bereich von FedML, ML-Einführung, ML-Softwareentwicklungszyklen bei. Zunächst haben wir die Einflussfaktoren für die organisatorische Umsetzung von FedML identifiziert, wobei wir neue, FedML-spezifische Faktoren

hervorgehoben haben und aufzeigten, wie die Dezentralität von FedML herkömmliche ML-Einführungsfaktoren beeinflusst. Diese Erkenntnisse dienen dazu, potenziellen Herausforderungen proaktiv entgegenzuwirken und die erfolgreiche Umsetzung durch informierte Entscheidungen zu fördern. Außerdem bieten wir einen umfassenden Überblick über die sozio-technischen Herausforderungen in FedML Kollaborationen. Organisationen können diese Ergebnisse bei der Gestaltung von Kooperationsvereinbarungen nutzen, um alle relevanten Aspekte zu berücksichtigen. Schließlich erstellten wir Darstellungen des FedML-Softwareentwicklungslebenszyklus und ergänzten bestehende ML-Entwicklungslebenszyklen um FedML-spezifische Aspekte. Diese Artefakte helfen Organisationen, um Projekte besser zu strukturieren und um Abhängigkeiten zu identifizieren. Darüber hinaus helfen die Artefakte bei der Projektplanung, der Überwachung des Fortschritts und der Planung nächster Schritte.

Limitationen: Unsere Forschung unterliegt mehreren Einschränkungen. Die durchgeführten Literaturrecherchen könnten möglicherweise nicht alle relevanten Artikel erfasst haben und die Kodierung könnte subjektiv sein. Um dem entgegenzuwirken haben wir Suchverfahren vordefiniert, umfassende Vorwärts- und Rückwärtssuchen durchgeführt und mehrere Forscher unabhängig voneinander kodieren lassen. Hinsichtlich Fokusgruppen und Experteninterviews, könnten die Forscher subjektiv gehandelt haben und unsere Stichprobengröße war begrenzt. Obwohl wir eine hohe Diversität von Stimmen einbezogen und die Studien erst nach Erreichen theoretischer Sättigung abgeschlossen wurden, ist die Verallgemeinerbarkeit aufgrund der begrenzten Stichprobengröße limitiert. Durch die Triangulation von Beobachtern und Daten, probierten wir mögliche subjektive Verzerrungen zu minimieren. Darüber hinaus ist die Evaluation unserer Artefakte eingeschränkt, da wir die Artefakte nicht in unterschiedlichen Anwendungsszenarien testen konnten. Außerdem sind unsere Modelle auf das am häufigsten verwendete FedML-Architekturmuster ausgerichtet.

Ausblick: Da unsere Arbeit explorativ ist, könnten zukünftige Forscher Folgestudien durchführen, um unsere Ergebnisse quantitativ zu validieren, zu verfeinern und zu erweitern. Beispielsweise könnten zukünftige Studien die Signifikanz einzelner Einflussfaktoren, sowie deren Korrelation bestimmen oder sich auf ausgewählte Herausforderungen konzentrieren. Vor allem Fortschritte bezüglich Dateninteroperabilität oder Transparenz der rechtlichen Situation von FedML könnten einen wesentlichen Beitrag zur Einführung von FedML leisten würden. Des Weiteren haben wir eine Vielzahl soziotechnischer Herausforderungen identifiziert, aber fehlt es noch an klaren Anleitungen, wie man diesen Herausforderungen begegnen kann. Daher könnten künftige Arbeiten Beobachtungsstudien durchführen und daraus ableiten, wie Anwender mit diesen Herausforderungen umgehen. Angesichts des begrenzten Umfangs konzentrieren sich unsere Modelle auf Server-Client-Architekturen mit einem zentralen Orchestrierungsserver. Da dies nur einen Bruchteil der möglichen Architekturmuster darstellt, könnte weitere Forschung analysieren, wie andere Architekturmuster den Entwicklungsprozess beeinflussen und entsprechende Modelle entwerfen.

Acknowledgment

This dissertation emerged from countless coffee talks, stimulating discussions, and the support of many brilliant people who helped me grow personally and professionally.

First and foremost, I want to take the opportunity to express my gratitude to my supervisor Prof. Dr. Florian Matthes for providing the best possible conditions, extraordinary support, and most importantly for your trust. Your mentorship and support have been invaluable. Furthermore, I want to thank Prof. Dr.-Ing. Stefan Tai for being the second supervisor of my dissertation and Prof. Jens Großklags, Ph.D. for being the examination chair.

This dissertation profited immensely from the excellent environment and camaraderie at the sebis chair. I would like to thank all my colleagues from sebis for their great support, guidance, and invaluable feedback! Looking back, I can not imagine a better environment to grow as a researcher and a person.

I am profoundly thankful to my colleagues at SAP SE, especially to my industrial supervisor Dr.-Ing. Nemrude Verzano, for your unwavering guidance, mentorship, and support throughout this research endeavor. Your expertise, open mindset, and dedication have been instrumental in shaping this work. Beyond that, I would like to extend my sincere appreciation to the entire team at SAP SE for providing me with the necessary resources, fruitful discussions, and an inspiring environment needed for this thesis. Special thanks go to Alexander Kläger, Stefan Wagner, Rüdiger Eichin, Oliver Frendo, Jonas Böhler, and Nadine Gärtner!

I would like to express my deepest gratitude to Milena Zahn, whose unwavering dedication, enthusiasm, and invaluable contributions were instrumental in the completion of this dissertation. Without her exceptional insights, contributions, and tireless effort this work would not have been possible. I am immensely grateful for her commitment to this project and I am fortunate to have had the opportunity to work alongside her.

Finally, I want to thank my parents, Ferdinand and Monika, and my brother Stefan. Thank you for supporting me throughout my entire life - without you, I would not be where I am today and I can not thank you enough for this!

This dissertation would not have been possible without the collective contributions and support of all the involved individuals and institutions. Thank you for being a part of this significant chapter in my academic and professional life.

Garching b. München, 03.11.2023
Tobias Müller

Table of Contents

Part A	1
1 Introduction	2
1.1 Motivation	2
1.2 Research Questions	4
1.3 Structure of the Dissertation	6
2 Theoretical Background	13
2.1 Federated Machine Learning	13
2.1.1 Training Process	14
2.1.2 Variations of Federated Learning	14
2.2 Technology Adoption Models	16
2.2.1 Diffusion of Innovation	16
2.2.1.1 Innovation-Decision Process	17
2.2.1.2 Rate of Adoption and Organizational Innovativeness	18
2.2.2 Technology-Organization-Environment Model	19
2.3 Machine Learning Systems Development	21
2.3.1 Software Development Life Cycles	21
2.3.2 Cross-Industry Standard Process for Data Mining	23
2.3.2.1 Reference Model	23
2.3.2.2 User Guide	25
2.3.2.3 Extensions and Variations	25
3 Research Design	27
3.1 Research Paradigm	27
3.2 Research Methods	28
3.2.1 Literature Reviews	28
3.2.2 Focus Groups	30
3.2.3 Expert Interviews	31
3.2.4 Design Science Research Methodology	32
Part B	35
4 Core Contributions	36
4.1 (P1) Revealing the Impacting Factors for the Adoption of Federated Machine Learning in Organizations	38

Table of Contents

4.2	(P2) Unlocking the Potential of Collaborative AI - On the Socio-Technical Challenges of Federated Machine Learning	39
4.3	(P3) A Process Model for the Practical Adoption of Federated Machine Learning	40
4.4	(P4) A Pathway for the Practical Adoption of Federated Machine Learning Projects	41
Part C		43
5	Summary of Results and Discussion	44
5.1	RQ1: Influencing Factors for the Organizational Adoption	44
5.2	RQ2: Socio-Technical Challenges of Federated Machine Learning	47
5.3	RQ3: Software Development Life Cycle of Federated Machine Learning Projects .	49
6	Contributions and Implications	55
6.1	Theoretical Implications	55
6.2	Practical Implications	56
7	Limitations	59
8	Conclusion and Future Work	61
8.1	Conclusion	61
8.2	Future Work	63
Bibliography		67
Publications		85
Abbreviations		87
A	Embedded Publications in Original Format	89
B	Guiding Questions for Federated Machine Learning Collaboration Agreements	140

List of Figures

1.1	Structure of the dissertation	7
2.1	FedML training process	14
2.2	FedML data partitioning variants (based on [RAA ⁺ 21])	15
2.3	FedML architectural patterns (based on [BPS ⁺ 23])	16
2.4	Five stages in the Innovation-Decision-Process (based on [Rog83c, Sah06])	17
2.5	Variables of diffusion of innovation in organizations (based on [Rog83a, Rog83d])	18
2.6	Technology-Organization-Environment framework (based on [TF90, OM11])	20
2.7	Major software development life cycle stages (based on [AOO ⁺ 20])	22
2.8	Phases of the CRISP-DM process model (based on [WH00]).	24
3.1	Interrelation and influence of interview guidelines (based on [MN07])	32
3.2	Design science research process in Information Systems (based on [PTRC07])	33
5.1	Overview of insights for the organizational adoption of FedML (based on P1).	46
5.2	Overview of socio-technical challenges of collaborative ML projects (based on P2).	48
5.3	Process model of an end-to-end FedML project.	52

List of Tables

1.1	Overview of core and additional publications	11
3.1	Overview of research methods applied in the embedded publications	28
5.4	Fact sheet on publication P1	38
5.1	Fact sheet on publication P2	39
5.2	Fact sheet on publication P3	40
5.3	Fact sheet on publication P4	41
5.1	Overview of FedML software development life cycle stages with their goals	50
5.2	Overview of key results	53
8.1	Potential avenues for future research	65

Part A

"Pressure on businesses to keep delivering is immense — and only collaborating and building ecosystems will ensure that growth and innovation persist." [For23]

This quote from the World Economic Forum [For23] illustrates that collaboration is key to fostering innovativeness and ensuring the growth of organizations. In this dissertation, we explore collaborations of organizations within the Machine Learning (ML) domain. More specifically, we take a step beyond traditional centralized ML systems and empirically analyze the organizational adoption of the decentralized paradigm *Federated Machine Learning (FedML)*. First, we investigate the crucial factors for the successful adoption of FedML in organizations. Motivated by the insights, we acquire an understanding of the socio-technical challenges that result from collaboration in the ML domain. Moreover, we analyze how the decentralized paradigm affects the Software Development Life Cycle (SDLC) and develop a corresponding process model.

Following, we start by motivating and introducing the topic of this dissertation (see Section 1.1). Subsequently, we outline the research gaps and present the accompanying Research Questions (RQs) that guide this thesis (see Section 1.2) and outline the structure of this dissertation including a summary of the embedded publication (see Section 1.3).

1.1. Motivation

In recent years, technological advances sparked a widespread interest in the field of ML. Unmistakably, ML systems had an immense effect on the economy [FS19] and played a pivotal role in business model innovation [LSRB19]. The integration of ML in organizations has proven to unlock unprecedented value by providing decision support, automating processes, driving customer and employee engagement, as well as enabling new products and services [BLS⁺21]. Considered

a true General Purpose Technology (GPT) [Ras20, GTT23], ML has applications in a variety of use cases across all industries [RKGR17], but as a GPT, the complete spectrum of potential applications remains currently unknown. Some studies anticipate that ML will attain full efficacy within the next decades and substantially impact nearly all aspects of our lives similar to the previous industrial revolutions [Mak17, Lou18]. Regardless of its actual influence, ML is expected to have one of the largest technological impacts in the next years [RFGG21].

The primary catalyst behind this disruptive technology is the rise of big data, serving as the fundamental basis for ML systems. However, the lack of available and suitable training data also remains a persisting barrier to the implementation of ML systems [PFW⁺21, BvDK20, HAAY⁺23]. Even though the volume of data is constantly growing, around 80% of industrial data is never used and therefore its economic potential never realized [Uni22]. Often referred to as the data silo issue [KHC⁺16], a significant portion of the unused data usually remains isolated and is hardly accessible. This problem is strengthened by various challenges such as privacy concerns [VPPE⁺14], data security issues [SLZ20] and legal regulations [KHC⁺16] like the General Data Protection Regulation or the California Consumer Privacy Act. These regulations seek to safeguard individuals' privacy, however, further restrict the exchange of data between different organizations [LZZ⁺22] and thereby strengthen data silos.

From a technical perspective, the emerging ML paradigm FedML has the capability to train ML systems on distributed data silos from multiple sources without the need for data sharing [LDCH22]. Thereby, it has the technical potential to overcome the data silo issue. FedML was first introduced by McMahan et al. [MMR⁺17] with the goal of utilizing local data sets on mobile phones for training a ML model while keeping the data private. Hereby, the authors proposed a novel decentralized ML paradigm that adheres to the following procedure: First, a global ML model is distributed across all participating clients and each mobile device receives the current model version. Then, each client individually improves the model by locally training it on the phone's individual dataset. The changes to the model are summarized and only this update is sent to the cloud, where it is averaged with updates from other users to improve the global model. Through this communication protocol and on-device learning, multiple distributed data sets can be leveraged while all training data remains at its source. With this model-to-data approach, FedML promises to enhance privacy by design, achieve higher utility, enhance communication efficiency, and reduce computational overhead [YML⁺22].

Since the first emergence of FedML, researchers have investigated FedML from various perspectives: from FedML algorithms [NSU⁺18], incentive mechanisms [ZZH⁺21] or domain-specific applications [AAdCK⁺22, KSH⁺21] to research on its privacy and security [MPP⁺21]. Despite increasing interest and its potential to enable novel business applications [BAA⁺22], FedML has seen a relatively low adoption rate in organizations, and most projects do not evolve beyond the prototype or simulation stage [LLW⁺21]. The missing operationalization and low organizational adoption of FedML remain largely unexplored and the current ML and FedML literature corpus shows corresponding research gaps. First, extensive work has been done on the factors that influence the adoption of ML [ACM19, PTH19, CRDB21, HAAY⁺23, KKG21, BKKP22]. However, extant literature does not offer an understanding on the barriers and success factors for the organizational adoption of FedML. Second, despite significant work on value drivers and business model innovation of traditional ML applications [LSRB19, RÅE20, MT21, ÅRP22],

current literature does not offer information on the dynamics of joint value creation and collaboration management within FedML. Third, compared to the growing body of knowledge in the standards, processes, and methodologies underpinning the development of ML applications [WH00, ABB⁺19, SBD⁺21, LBM⁺22, KKH23], there is no such perspective on FedML. Within the scope of this dissertation, we aim to address these research gaps.

1.2. Research Questions

The following outlines the research gaps and motivates the RQs that guide this dissertation.

Research Gap 1: *Currently, most FedML projects do not evolve beyond the prototype stage. While research on ML adoption grows, there is a lack of understanding of the barriers and success factors for the organizational adoption of FedML.*

The adoption of novel technologies needs to be approached from a socio-technical and multidisciplinary perspective since it is dependent on the needs of a diverse range of stakeholders and is subject to a large variety of influences [BvHS05, KVSO13, MNH21]. Thus, organizations need to take a multidisciplinary view beyond the technical dimension to successfully utilize emerging technologies. In the case of ML, organizations face a new realm of challenges that is different from the adoption of previous technologies due to their non-deterministic behavior [CHTB20]. Furthermore, exacerbating the situation, FedML incorporates the stochastic behavior of ML and additionally introduces another layer of complexity due to its decentralized paradigm. While previous literature provides insights into the drivers and barriers of successful ML adoption [ACM19, PTH19, CRDB21, HAAY⁺23], it remains unclear which novel aspects result from the decentralized characteristics of FedML and how it affects currently known influences. Therefore, it is required to revisit the adoption factors in the context of FedML to ensure that organizations have an overview of what factors influence a successful adoption. Consequently, we articulate the first RQ:

Research Question 1 (RQ1)

What are the major factors influencing the successful adoption of Federated Machine Learning systems in organizations?

We fill this research gap by leveraging established *Technology Adoption Frameworks* and conducting an expert interview study to draw relevant practitioners in the field. Thereby, we aim to identify the main challenges, risks, and influencing factors regarding the organizational adoption of FedML. We structure the resulting insights according to their technological, environmental, and organizational context. This approach allows us to provide a systemized overview of the factors that influence a successful adoption and to uncover novel aspects that are unique to FedML. It shows that the success of a FedML project is mainly dependent on factors regarding the collaboration of multiple organizations, its complex technical implementation, and regulatory uncertainties. Given the permanence of legal and regulatory frameworks, we opted to focus on research gaps regarding collaboration in the ML domain (see RQ2) and the complex implementation process (see RQ3).

Research Gap 2: *Current literature offers little knowledge on the dynamics of joint value creation and collaboration management in the ML domain.*

The increasing interest in ML-based systems also led to a growing literature corpus on how organizations create value and leverage ML as a catalyst for business model innovation [LSRB19, RÅE20, MT21, ÅRP22]. Hereby, FedML constitutes a distinctive case which requires further considerations beyond traditional ML systems. With its privacy-enhancing features, FedML has the potential to overcome data silos and to drive unprecedented use cases [BKB20]. However, overcoming data silos requires the collaboration of multiple parties, and consequently the formulation of a collaboration agreement including a business model for the joint endeavor. Hence, to create value from FedML, organizations need to mitigate challenges of both domains: ML and inter-organizational collaborations. While extant literature offers information on the socio-technical factors either regarding ML-enabled business cases [BPLW21, DVPHE20, JYK⁺23] or collaborative business models [PLAK21, CCM15, RBBC⁺14, DC18], none studied the intersection of both domains. Thus, current literature does not offer an overview of the socio-technical aspects regarding joint value creation in the ML context. We aim to fill this research gap and thereby follow the call of Enholm et al. [EPMK22] to investigate the dynamics and potential conflicts of interest in collaborations around ML. Therefore, we raise the following RQ:

Research Question 2 (RQ2)

What are the distinct socio-technical challenges of collaborative Federated Machine Learning projects?

To investigate this research gap, we employ a three-step approach. First, we systemize knowledge on the socio-technical challenges of prevailing collaborative business models. Second, we conduct a focus group discussion to initially explore the challenges regarding the collaborative characteristics of FedML. Third, we complement previous insights through expert interviews. Thereby, we provide a holistic overview of general and FedML-specific socio-technical aspects that need to be considered in the collaboration agreement within FedML projects.

Research Gap 3: *The existing studies do not report on the standards and processes of FedML systems development.*

Although many organizations are eager to adopt ML systems to improve their performance through novel unprecedented business opportunities [BLS⁺21], some experience difficulties due to the lack of standards and guidance of the development process [RKW20, ADD21, SBD⁺21]. Motivated by this, researchers have investigated best practices regarding ML development and proposed a set of ML life cycles models [WH00, ABB⁺19, SBD⁺21, LBM⁺22, KKH23] or domain-specific activity models [TWN19, KBM⁺20, VHC⁺20]. However, in contrast to traditional ML pipelines, FedML trains a joint model across multiple clients through a model-to-data approach in an iterative and decentralized process. Accordingly, FedML also requires software system designs that differ from traditional ML system architectures [LLZ⁺22] which postulates adaptations to the development process. While current literature on FedML offers information on architectural patterns and system designs of the training process [BEG⁺19, AAdCK⁺22, LLZ⁺22], it does not report on the corresponding development process and holistic project life cycle. Yet, considering all life cycle stages is a crucial and existing problem for the implementation of

industrial-level FedML systems [ZBO20]. Accordingly, it is necessary to investigate the SDLC of FedML projects which yields the following RQ:

Research Question 3 (RQ3)

What is the software development life cycle of Federated Machine Learning projects?

We address this RQ by following the Design Science Research (DSR) [HMMP04, PTRC07] methodology to conceptualize a process model and detailed activity models of an end-to-end FedML project life cycle. The models incorporate best practices from literature on SDLC, ML life cycles, and FedML as well as the expertise and feedback of project teams and experts.

1.3. Structure of the Dissertation

This thesis consists of three parts and comprises four publications that answer the three RQs. Figure 1.1 illustrates the logical flow and structure of this publication-based dissertation.

Part A consists of three chapters. Chapter 1 begins with the overall motivation of the research problem (see Section 1.1), summarizes the research gap by presenting the three RQs (see Section 1.2) and explains the outline of this thesis with an overview of the embedded publications (see Section 1.3). Chapter 2 introduces the theoretical background of the concepts that were applied in the embedded publications. Hereby, we explain the fundamentals related to FedML (see Section 2.1), Technology Adoption Models (see Section 2.2), and Machine Learning Systems Development (see Section 2.3). Chapter 3 concludes Part A with a description of the underlying research paradigm (see Section 3.1) and applied research methods (see Section 3.2).

Part B provides an overview and fact sheets of the four publications included in this thesis. The published papers can be found in their original format in Appendix A. The first publication (P1) offers results for RQ1 and analyzes the reasons, challenges, risks, and influencing factors for the practical adoption of FedML in organizations (see Section 4.1). The second publication (P2) aims to answer RQ2 by identifying and systemizing the socio-technical challenges regarding the collaborative characteristics of FedML (see Section 4.2). The two subsequent publications (P3, P4) address RQ3 and propose a process model of an end-to-end FedML SDLC (see Section 4.3) and detailed activity models (see Section 4.4).

Part C comprises four chapters and concludes this dissertation. Chapter 5 summarizes the results of each RQ and discusses the findings in relation to extant literature. In Chapter 6, we present the contributions and implications for research (see Section 6.1) and practice (see Section 6.2). Subsequently, Chapter 7 delineates the limitations of the embedded publications and therefore this dissertation regarding the research approach and scope. Lastly, Chapter 8 closes this thesis by reflecting the formulated RQs with corresponding results (see Section 8.1) and outlining potential directions for future research activities (see Section 8.2).

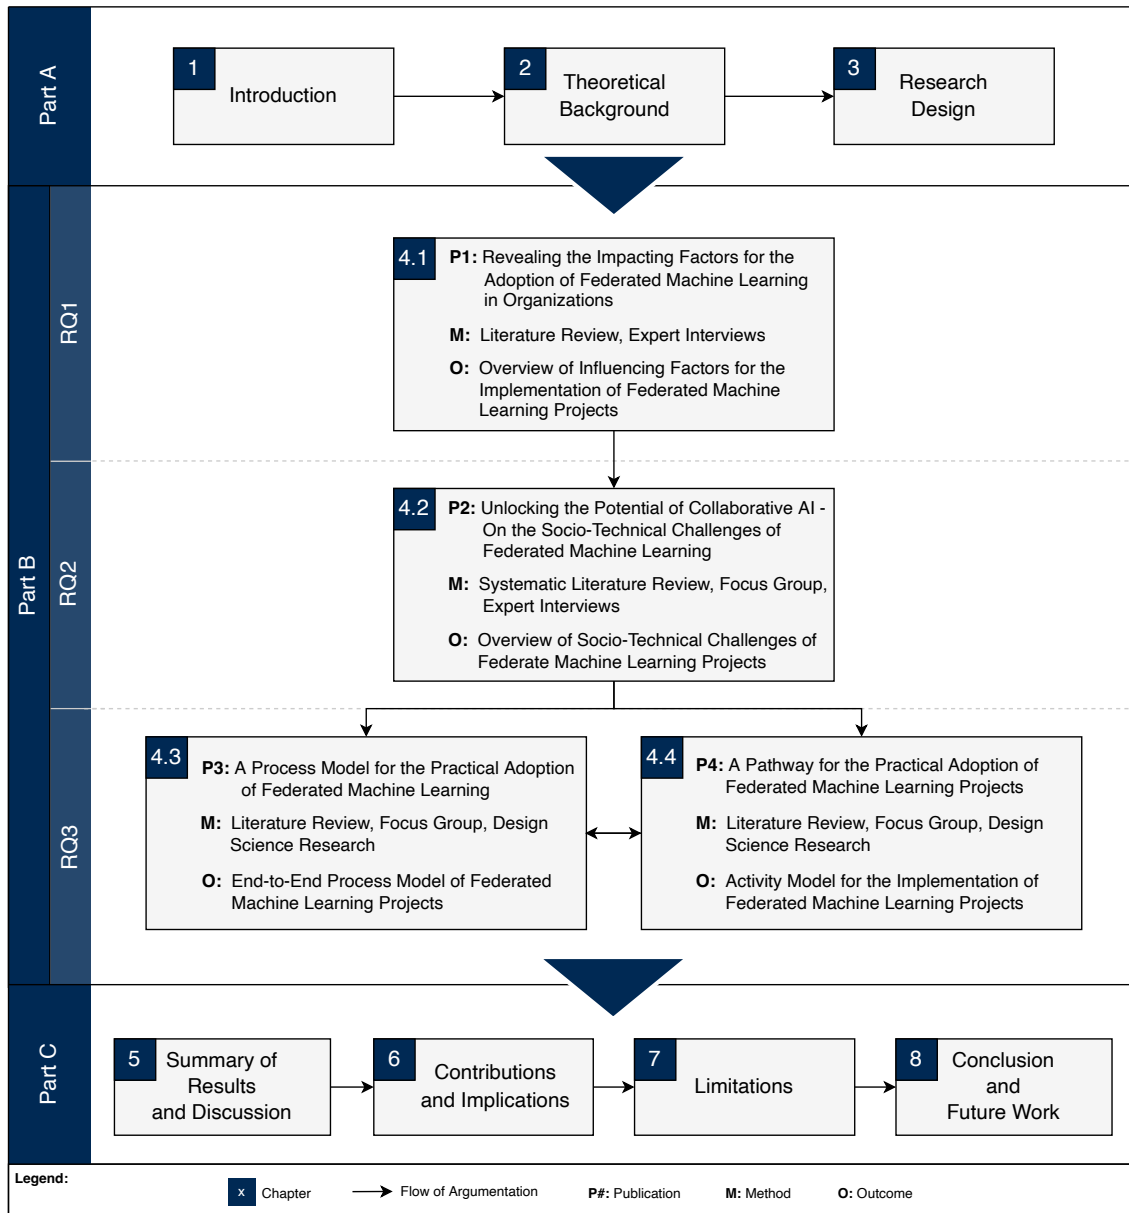


Figure 1.1.: Structure of the dissertation.

In the following, we present an overview of the published papers that are embedded in Part B. For each publication (P), we outline the underlying research problem, research method, and main contribution. Furthermore, we elaborate on additional publications that have been conducted within the course of this thesis. Table 1.1 gives an overview of the embedded and additional studies in relation to the RQs.

P1: Revealing the Impacting Factors for the Adoption of Federated Machine Learning in Organizations. While the body of knowledge on ML adoption grows, it remains unclear how the decentralized characteristics of FedML affect the adoption in organizations. This publication [MZM24] addresses this knowledge gap. Through an expert interview study, this article compiles a comprehensive overview of 19 influencing factors for the organizational adoption of FedML. The identified factors are structured according to their technological, organizational, or environmental context. The insights coincide with known factors on ML adoption and additionally uncovered novel factors that are specific to FedML projects. Moreover, this article reveals the practitioner’s reasons for utilizing FedML and describes their main encountered challenges and risks. Overall, this publication extends the current literature on ML adoption with a view on FedML and aids practitioners to preemptively counteract pitfalls throughout the project life cycle. The study showed that practitioners are currently mainly challenged by issues regarding collaboration, the complex technical implementation of FedML, and regulatory uncertainties. Given the immutability of legal and regulatory frameworks, these insights provided the motivation to focus on the collaboration of multiple organizations within the ML domain (P2) and to facilitate the complex implementation process (P3, P4).

P2: Unlocking the Potential of Collaborative AI - On the Socio-Technical Challenges of Federated Machine Learning. FedML enables organizations to overcome data silos and jointly train a ML model across company boundaries. However, this requires collaboration between the companies, and consequently the formulation of a collaboration agreement including a respective collaborative business model. Hereby, the participants need to agree on various aspects regarding the scope, financials, and distribution of responsibilities. While extant literature studies collaborative business models and the accompanying socio-technical challenges, there is no guidance regarding collaborations in the ML domain. The first publication [MZM23d] addresses this knowledge gap and investigates the socio-technical challenges of collaborative FedML projects. This paper systemizes current knowledge of the challenges of collaborative business models and subsequently explores challenges related to collaborative FedML through a focus group discussion and semi-structured expert interviews. Thereby, this article provides a systemized overview of the identified challenges and unveils novel aspects related to collaborative FedML projects. Based on these insights, this study presents a comprehensive set of guiding questions to assist practitioners in the creation of a collaboration agreement. Amongst others, this study revealed the difficulty of assigning responsibilities and accountabilities among the participants. This challenge strengthened the motivation for the two subsequent publications (P3, P4).

P3: A Process Model for the Practical Adoption of Federated Machine Learning. The second publication [MZM23c] focuses on aiding practitioners with the division and communication of responsibilities and accountabilities. A dedicated focus group discussion confirmed that this challenge prevents many FedML projects from leaving the prototype stage. Hereby,

the barriers arise in the project initiation and planning phase. This study [MZM23c] focuses on the project initiation stage while P3 addresses the planning stage. The focus group concluded that an end-to-end process model could facilitate the challenge of structuring the process flow and communicating responsibilities during the project initiation phase. By following the DSR approach, this study designs a corresponding model with an abstract overview of the process flow with the interrelations of the required resources, resulting artifacts, roles, and activities. Thereby, this article provides a holistic overview of the FedML SDLC and takes an initial step toward standards for the FedML development process. From a practical view, business stakeholders and solution architects can use the models as a basis to outline their project plan in the project initiation phase and to facilitate communication with potential participants.

P4: A Pathway for the Practical Adoption of Federated Machine Learning Projects.

As already recognized in P3, barriers during the project initiation and project planning stage inhibit FedML projects to evolve beyond prototypes. This study [MZM23b] addresses issues related to the project planning phase. Hereby, practitioners are challenged by the intricate implementation process due to the decentralized and collaborative paradigm of FedML. The focus group stated that a detailed step-by-step model that depicts the sequence of activities could provide guidance for the planning of the implementation process. This study follows a DSR approach to construct corresponding activity models and thereby complements the process model of P3. The evaluation showed that the activity models aid product owners and project managers in understanding the required steps of the development process and making better-informed decisions. While P3 and P4 offer representations of the FedML project flow on different dimensions, P1 provides additional information on the critical factors that influence the adoption of FedML, and P2 provides guidelines for the creation of the collaboration agreement.

Additional Publications. In addition to the four embedded publications, this research endeavor produced four additional studies. Although the additional studies provide complementary insights to the discussed topics and RQs, we did not select these publications as the core contributions of this dissertation. The following will give a short overview of the additional publications.

The first position paper (P5) [MGVM22] was created during the first stages of a large industrial lighthouse project and aimed to discuss the challenges of FedML adoption that arose within the project. In the starting phase of this project, we noted that practitioners were challenged with the cumbersome implementation and customizability of Privacy-Enhancing Technologies (PETs) in the FedML context which resulted in the termination of various FedML endeavors. However, the discussed concerns of this article were rendered obsolete by subsequent studies (P2, P3, P4, P8). Within a follow-up position paper (P8) [MZM23a] we discuss the empirically grounded observations that current project teams are missing clarity over the dependencies and interrelations within FedML projects. Therefore, we make a case for the need for a process model. This position paper marked the starting point of P3 and P4. In retrospect, our embedded publications validated the observations of P8.

Related to RQ1, we conducted a Systematization of Knowledge (SoK) of the literature corpus on applied FedML (P6) [MSG⁺23]. Thereby, we analyzed 74 papers and systemized information on the stated motivational drivers, application domains, and encountered challenges. Additionally, we derived trends and characteristics regarding the distribution of the published papers, publi-

cation channels, affiliated countries, and research types. Our report suggests that the interest in applied FedML has risen significantly in recent years, with China and the United States as the most active countries and conferences as the most prevalent publication channel. Hereby, the majority of work proposes technological solutions and utilizes various neural network architectures as preferred predictive models. The literature predominantly states that privacy is the main motivational driver and that FedML is mostly used within the medicine sector. The stated challenges were mostly of a technological nature and revolved around system heterogeneity, scalability, privacy, hardware restrictions, and the sophisticated communication protocol.

Lastly, we published a study on the design of a technology selection tool (P7) [ZMM24]. The tool aims to support managerial decision-making and to avoid an inappropriate fit between the technology and the use case. This tool was developed following the DSR methodology and is divided into two steps. The first step is based on a decision tree that evaluates the reasonability and necessity of utilizing FedML within the use case. The second step consists of a survey with a scoring system that helps practitioners assess the complexity of implementing FedML within their given use case and framework.

For the sake of completeness, it should be mentioned that we have also contributed to various publications that are not related to this thesis [MSS⁺21, MRSA21, ASS⁺23].

RQ	No.	Authors	Title	Outlet (Type)	Core Rank
Core Contributions					
RQ1	1	Tobias Müller Milena Zahn Florian Matthes	Revealing the Impacting Factors for the Adoption of Federated Machine Learning in Organizations	HICSS 2024 (Conf)	B
RQ2	2	Tobias Müller, Milena Zahn Florian Matthes	Unlocking the Potential of Collaborative AI - On the Socio-Technical Challenges of Federated Machine Learning	ECIS 2023 (Conf)	A
RQ3	3	Tobias Müller Milena Zahn Florian Matthes	A Process Model for the Practical Adoption of Federated Machine Learning	AMCIS 2023 (Conf)	A
	4	Tobias Müller Milena Zahn Florian Matthes	A Pathway for the Practical Adoption of Federated Machine Learning Projects	PACIS 2023 (Conf)	A
Additional Publications					
RQ1	5	Tobias Müller Nadine Gärtner Nemrude Verzano Florian Matthes	Barriers to the Practical Adoption of Federated Machine Learning in Cross-company Collaborations*	ICAART 2022 (Conf)	B
	6	Tobias Müller Maximilian Stäbler Hugo Gascon Frank Köster Florian Matthes	SoK: Assessing the State of Applied Federated Machine Learning	ArXiv (-)	-
	7	Milena Zahn Tobias Müller Florian Matthes	Supporting Managerial Decision-Making for Federated Machine Learning: Design of a Technology Selection Tool	HICSS 2024 (Conf)	B
RQ3	8	Tobias Müller Milena Zahn Florian Matthes	On the Adoption of Federated Machine Learning: Roles, Activities and Process Life Cycle*	ICEIS 2023 (Conf)	C
Outlet				Type	
ECIS	European Conference on Information Systems			Conf	Conference
PACIS	Pacific-Asia Conference on Information Systems			*	Position Paper
AMCIS	Americas Conference on Information Systems				
HICSS	Hawaii International Conference on System Sciences				
ICEIS	International Conference on Enterprise Information Systems				
ICAART	International Conference on Agents and Artificial Intelligence				

Table 1.1.: Overview of core and additional publications.

Theoretical Background

This chapter entails three sections and describes the theoretical background that forms the fundamental basis of this dissertation and all embedded publications. First, we introduce the concept of *Federated Machine Learning* (see Section 2.1). Thereafter, we present *Technology Adoption Frameworks* (see Section 2.2) as the foundation of P1. Lastly, we describe the relevant background on *Machine Learning Systems Development* (see Section 2.3), which builds the groundwork mainly for P3 and P4.

2.1. Federated Machine Learning

FedML is a novel ML paradigm that was introduced in 2016 by McMahan et al. [MMR⁺17] and varies from conventional ML approaches. In traditional ML systems, training data is usually accumulated in a central location, where the ML model is subsequently trained. For this centralization of individual datasets into a single repository, data owners are required to share their data with a central server. By that, data owners sacrifice their data sovereignty, which leads to an increased dependency on third parties and a potential loss of Intellectual Property (IP).

In contrast to traditional ML systems, FedML allows the training of ML models on decentralized data without the need for direct data sharing. Originally introduced in 2016, McMahan et al. [MMR⁺17] followed a model-to-data approach to train a language model on a loose federation of mobile devices. By eliminating the requirement for sharing training data, FedML enables the privacy-enhancing usage of siloed data and thereby has also the potential to enable collaborations across different organizations [TSP22]. Since its emergence, FedML has sparked large research interest which resulted in a high variety of methods, architectural patterns, and theoretical concepts [WMOT21, LLZ⁺22].

2. Theoretical Background

The following Section 2.1.1 first introduces the fundamental methods behind FedML. Subsequently, Section 2.1.2 presents an overview of possible FedML system variants.

2.1.1. Training Process

Formally, FedML is concerned with the problem of learning a joint predictive model by multiple parties who own separate sets of data. The original algorithm follows a model-to-data approach to enable the collaborative learning of a global ML model while keeping their data locally stored [MMR⁺17]. Generally, the basic FedML training process can be summarized into four distinct steps:

1. The server specifies the structure of the global model, which will be used across all clients. The design of this initial global model should be suitable for the use case and underlying data structure. Additionally, this model can be pre-trained on an initial dataset.
2. The server distributes the global model amongst all clients. This model will be the starting point for the local model training.
3. Each client trains the received model on their own locally stored dataset. This local training step is similar to conventional ML training, where data is used to update the model's parameters through an optimization algorithm. Each client stores the resulting parameter updates (gradients).
4. The clients send their individually computed gradients back to the server. These gradients are then aggregated based on a pre-defined aggregation scheme to create an updated global model.

Steps 2-4 can be repeated over several iterations until a certain accuracy level is reached or until the accuracy converges. This process is illustrated in Figure 2.1.

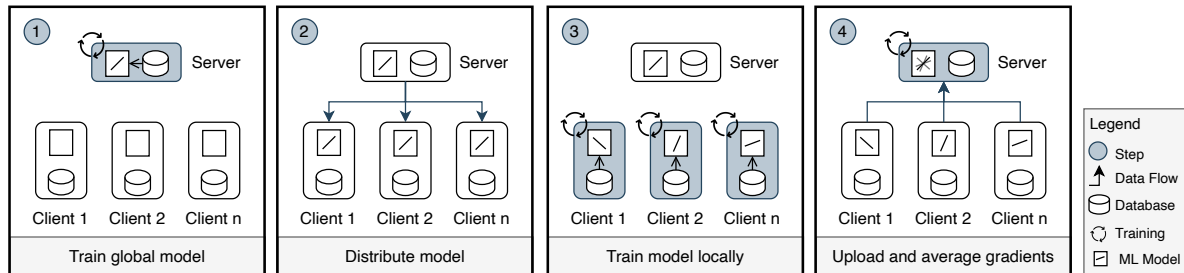


Figure 2.1.: Simplified illustration of the FedML training process (own work).

2.1.2. Variations of Federated Learning

Depending on the modality and system characteristics, FedML systems can be divided into different categories. Current literature proposes a multitude of different taxonomies, each with a varying focus. However, based on [LLW⁺21], the variations of FedML systems are most com-

monly differentiated regarding the *Scale of Federation*, *Data Partitioning*, and *Communication Architecture*.

Scale of Federation. FedML systems can be implemented in two settings: *cross-device* or *cross-silo* [LWW⁺23]. *Cross-device* FedML refers to settings with a large number of participating clients, in which each party has a relatively small amount of data as well as computational power. In contrast, *cross-silo* FedML describes settings with a relatively small number of participants. Here, each party has a comparably large amount of data and computational power.

Data Partitioning. Based on the data partitioning regarding features and sample space across the participants, FedML systems can be divided into three categories: *horizontal*, *vertical* and *transfer* [RAA⁺21]. In *horizontal* FedML, datasets *A* and *B* share similar features but differ in data samples. If instead, the same samples are present in all datasets but feature spaces are heterogeneous, the setup is known as *vertical* FedML. Highly heterogeneous settings with different sample and feature spaces across participants are referred to as *transfer* FedML. Transfer FedML is inspired by transfer learning, where already learned knowledge is re-used for a related task via feature extraction and/or fine-tuning [TSK⁺18]. The difference between the data partitioning variants can be seen in Figure 2.2. Most research considers horizontal FedML systems [LWW⁺23].




Figure 2.2.: Illustration of data partitioning variants (own work based on [RAA⁺21]).

Communication Architecture. While conventional FedML operates on a client-server architecture, alternatives have been proposed that do not rely on a central orchestrating server [LLZ⁺22]. Classified according to their network topology, these FedML systems can be grouped into *client-server*, *partially-connected* or *peer-to-peer* architectures [RAA⁺21, BPS⁺23]. Figure 2.3 illustrates the different architectures. The *client-server* architecture consists of several clients, which perform local training, and a central orchestrating server for aggregation. Each client is only directly connected to the central server. Opposed to *partially-connected* architectures, in which a subset of clients are capable of communicating through direct links. Finally, *peer-to-peer* architectures eliminate an orchestrating server by establishing a peer-to-peer network amongst all clients. Currently, most implementations use client-server architectures [RAA⁺21].

Further Considerations. Apart from these categorizations, further characteristics have been used to group FedML systems. Some studies additionally classified FedML systems according to the underlying ML algorithms [LWW⁺23, JSD22, KSH⁺21], privacy mechanism [LWW⁺23, JSD22, BPS⁺23, KSH⁺21], or business-related aspects such as incentivization mechanism [KSH⁺21]. However, these additional categories vary depending on the emphasis of the present study and are not considered in the majority of the literature corpus on FedML taxonomies. Consequently, we also only focus on the characteristics presented above.




Figure 2.3.: Illustration of the FedML architectural patterns (own work based on [BPS⁺23]).

In this dissertation, we aim to provide guidance for the practical adoption of FedML in organizations. Therefore, we mainly target the most widely used system configuration. According to current literature, a large majority of the work solely considers *client-server* architectures [RAA⁺21] with *horizontal* data partitioning [LWW⁺23]. Hence, if not stated otherwise, we refer to *horizontal, client-server* FedML systems in the following. Furthermore, we consider intra-organizational FedML systems as *cross-device*, and inter-organizational FedML as *cross-silo*.

2.2. Technology Adoption Models

The process of adopting innovative technologies has been a widely studied area within Information Systems (IS). Research in this domain yielded a variety of different models and theories with the goal of identifying, predicting, and describing the variables that affect adoption behavior [DEZ20]. These models can be grouped according to their objective, depending on whether they study the adoption behavior of groups, individuals, or organizations [LMJ08].

Most prominent models, such as the *Technology Acceptance Model* [Dav85] or *Theory of Planned Behavior* [Ajz85], were originally developed to target the individual level [LMJ08]. Even though some adaptations and novel theories were proposed to analyze group behavior [KTB11, SVS05], literature on group models is relatively sparse [LMJ08]. From an organizational perspective, the most widely used models include the *Diffusion of Innovation (DOI)* theory [Rog83b] and *Technology-Organization-Environment (TOE)* Framework [TF90] [MJ17]. Since this dissertation investigates the adoption of FedML in organizations, we will only focus on organizational-based frameworks. The following introduces DOI theory and the TOE framework.

2.2.1. Diffusion of Innovation

The *DOI* by Rogers [Rog83b] is a theory that investigates how, why, and at what rates innovations and new technologies are adopted at an individual and firm level. More specifically, Rogers [Rog83b] defines *diffusion* as "[...] the process by which an innovation is communicated through certain channels over time among the members of a social system". Hence, per definition, DOI

theory focuses on *innovation*, *communication channels*, *time*, and *social systems* as the four key components that influence the innovation-decision process.

2.2.1.1. Innovation-Decision Process

As illustrated in Figure 2.4, the innovation-decision process follows the five steps from *knowledge*, *persuasion*, *decision*, *implementation* to *confirmation* in a sequential manner [Rog83c].

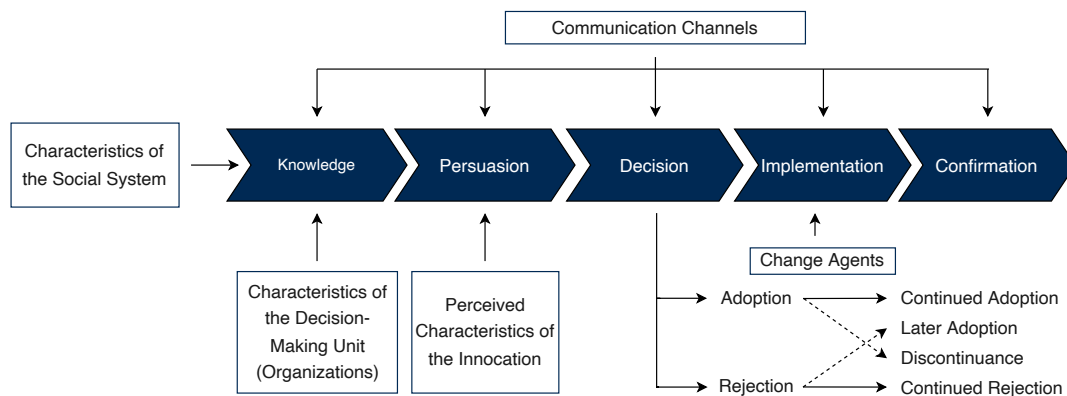


Figure 2.4.: Five stages in the Innovation-Decision-Process (own work based on [Rog83c, Sah06]).

(1) Knowledge Stage. In the first stage, individuals and organizations learn about the existence of an innovation and aim to determine [...] *what the innovation is and how and why it works*" [Rog83c]. This step is influenced by the nature of the social systems and therefore by their norms, previous practices, subjective needs, problems, and innovativeness. Additionally, the knowledge stage is impacted by the characteristics of the decision-making unit. In an organizational context, this can be translated to characteristics of organizational innovativeness. Section 2.2.1.2 will go into more detail about the factors that influence the innovativeness of an organization.

(2) Persuasion Stage. The persuasion stage forms the attitude toward an innovative technology and shapes the decision for the adoption or rejection of an innovation. However, Rogers [Rog83c] emphasizes, that a favorable or unfavorable attitude does not directly or indirectly lead to a decision regarding the adoption. The persuasion stage is impacted by the perceived characteristics of the innovation. Section 2.2.1.2 elaborates on these factors.

(3) Decision Stage. After learning about the innovation and forming an attitude, the organization (or individual) makes a decision about the adoption or rejection of the innovation. Usually, an innovation is adopted more quickly if it has a partial trial bias [Sah06]. It is important to note, that rejection is possible in every stage of the process, not only in the decision stage. The decision might be confirmed or revised in the fifth stage.

(4) Implementation Stage. In the fourth stage, the innovation is implemented and put into practice. Due to the novelty of innovation, there is a degree of uncertainty about the outcomes, which might force an organization to modify the planned way of implementation.

2. Theoretical Background

(5) **Confirmation Stage.** Lastly, the organization (or individual) looks for support and reinforcement of the made decision. Depending on the received feedback, the decision can be reversed. Thereby, a later adoption of the innovation can be triggered or an already adopted technology can be discontinued.

2.2.1.2. Rate of Adoption and Organizational Innovativeness

There is a multitude of varying factors, which influence the five stages of the innovation-decision process and accordingly the rate of adoption. Overall, these factors can be grouped into aspects regarding the *perceived attributes of the innovation*, *type of innovation-decision*, *communication channel*, *change agents*, and *characteristics of the decision-making unit*. In an organizational context, the characteristics of the decision-making unit can be translated to the factors regarding *organizational innovativeness*. An overview of the variables, which determine the diffusion of innovation in organizations can be seen in Figure 2.5. Based on [Rog83a, Rog83d, Sah06], the following elaborates on each category.

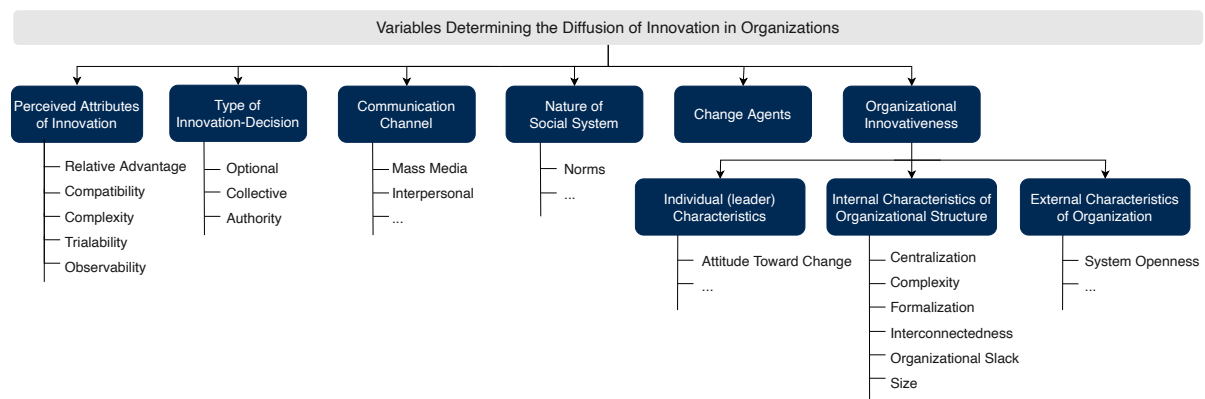


Figure 2.5.: Factors influencing the diffusion of innovation in organizations (own work based on [Rog83a, Rog83d]).

Perceived Attributes of Innovation. According to Rogers [Rog83a], there are five perceived attributes of innovations: *relative advantage*, *compatibility*, *complexity*, *trialability*, and *observability*. It is important to note that the *perception* of the organization (or individual) about the characteristics influences the rate of adoption.

Regarding *relative advantage*, cost and social status are the main drivers for quick adoption. Innovators are usually more status-motivated, while status is less significant for the late majority. Similarly, financial incentives which can be achieved directly or indirectly through the innovation can increase its rate of adoption. The adoption can be further accelerated if the innovative technology is *compatible* with current infrastructure, practices, and social systems. Additionally, low *complexity* as well as the possibility to *try* and *observe* the implications of the technology reinforces a positive choice in favor of the innovation.

Innovation-Decision. DOI differentiates between *optional*, *collective*, and *authoritative* decisions. The fastest rate of adoption is usually achieved through authoritarian decisions. Optional

(or individual) decisions can be made faster than collective decisions with multiple individuals. However, authoritative decisions tend to be circumvented during implementation. In an organizational setting, the type of decision is usually either collective or authoritative.

Communication Channel. The type of communication channels determines how participants create and share information with each other to reach a mutual understanding and influences each stage of the innovation-decision process. Here, two prominent channels are *mass media* and *interpersonal communication*. The influence of each type differs on the characteristics of the situation. While *mass media* is suitable to eliminate a deficit of awareness, *interpersonal communication* results in a bigger change of prevailing attitudes [Orr03].

Nature of Social System. The social system is defined as "[...] a set of interrelated units engaged in joint problem solving to accomplish a common goal" [Rog83b]. Hence, the social system has current practices and aims to solve a common problem with specific needs. Additionally, social systems are influenced by their norms and degree of innovativeness. Naturally, the rate of adoption is quicker if the innovative technology ties into their norms, current practices and solves the problem of the social systems.

Change Agents. During implementation, *change agents* can help to reduce the degree of uncertainty about the results and consequences of the innovation decision. Thereby, the involvement of change agents such as external consultants increases the rate of adoption.

Organizational Innovativeness. The organizational innovativeness heavily influences whether an innovation is adopted and is determined by independent variables regarding the *individual (leader) characteristics*, *internal structures*, and *external characteristics* of the organization [Rog83d, OM11]. A positive attitude toward change from an *individual* (or leader) strengthens the decision in favor of adoption. Regarding the *internal organizational structure*, a great extent of centralization and formalization are usually correlated with less innovativeness. In contrast, a higher degree of organizational complexity, interconnectedness, organizational slack, and larger size are found in more innovative organizations. In terms of *external characteristics*, more system openness corresponds positively with organizational innovativeness.

2.2.2. Technology-Organization-Environment Model

The *TOE* framework by Tornatzki and Fleischer [TF90] aims to investigate how the context of the organization influences the adoption of innovations [Bak12]. *TOE* thereby focuses on higher-level attributes instead of detailed behaviors and explains adoption decisions through three different elements of an organization's context: *technological context*, *organizational context*, and the *environmental context* [TF90, Bak12]. As illustrated in Figure 2.6, each context is influenced by the others and the sum of all three contexts set "[...] constraints and opportunities for technological innovation" [TF90]. The following introduces the three aspects.

Technological Context. The technological context describes all relevant technologies to the firm. This includes internal technologies, which are already in use, as well as external technologies which are not in use but available in the marketplace [TF90, Hag80, Sta76]. The internal technologies set a broad limit on the scope, showcase the technological context, and correspondingly set the pace of technological change within the organization [CHH88]. More-

2. Theoretical Background

over, the externally available technologies demonstrate the innovation possibilities and how the organization could evolve through innovations [Bak12].

Organizational Context. The organizational context incorporates all descriptive measures about an organization [TF90], which also corresponds to the characteristics of organizational innovativeness as described by DOI [OM11]. Hence, this context refers to factors such as intra-firm communication processes, size, organizational slack, or the degree of centralization [Bak12].

Environmental Context. The environmental context refers to all external forces such as the structure of the industry, competitors, governmental influence, or legal regulations [TF90, OM11, Bak12].

Overall, the TOE framework is consistent with DOI theory on how, why, and at what rate innovations are diffused within organizations. Both frameworks include individual, internal, and external characteristics of an organization as the influencing factors for its technology adoption and innovativeness [OM11]. TOE additionally considers environmental factors as an important aspect of the innovativeness of an organization. For example, competition usually increases the incentive to invest in research, thereby fostering innovation [Gil06], whereas legal regulations could potentially hinder the adoption of technologies. Therefore, TOE is considered to make DOI theory more complete to explain organizational innovation adoption [HKD06]. Summarized, TOE describes a holistic and frequently used analytical tool in IS research to understand the process of innovation adoption in organizations through these three dimensions [OM11].

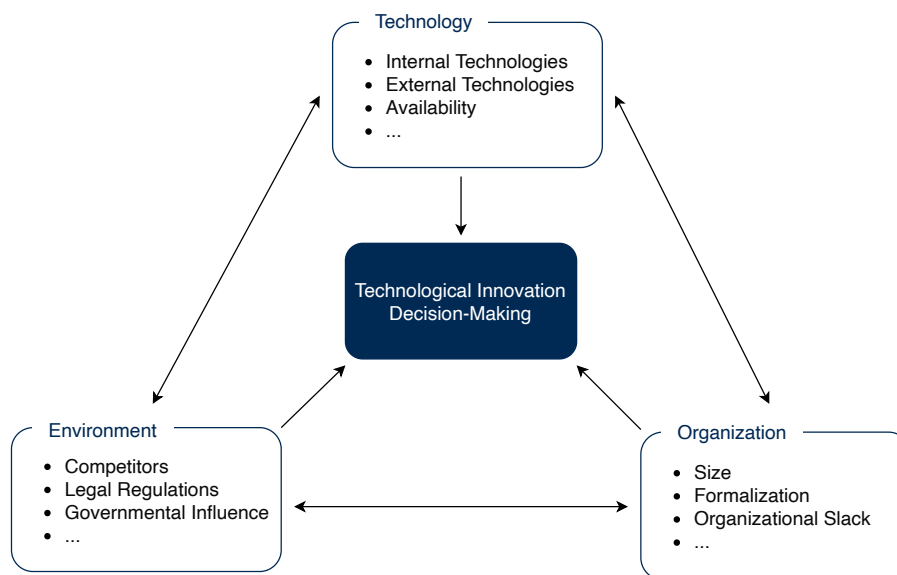


Figure 2.6.: Technology-Organization-Environment framework (own work based on [TF90, OM11]).

2.3. Machine Learning Systems Development

The data-driven and stochastic characteristics of ML systems with their uncertainty in preliminary experiments and non-deterministic outcomes require a shift of mindset in software development practices [WXML21]. Since ML capabilities are usually served as part of large software-intensive systems, practitioners are faced with the problem of implementing non-deterministic ML systems into traditional, deterministic software products [Gir21]. Necessarily, organizations are forced to evolve their development process and interweave ML-specific procedures, such as the process of feature engineering, into their current practices [ABB⁺19]. Similarly, the production of ML systems requires the field of ML to progress beyond mere ML model development to entire life cycle activities [LRC⁺20]. To alleviate these challenges, many organizations rely on *Cross-Industry Standard Process for Data Mining (CRISP-DM)* [WH00], which is the de-facto standard in the industry [SKMG21]. Accordingly, our endeavor to provide guidance for the adoption of FedML systems in organizations builds heavily on CRISP-DM and the corresponding software engineering best practices.

The following lays the theoretical groundwork for traditional software development life cycles in Section 2.3.1. Subsequently, we elaborate on CRISP-DM in Section 2.3.2.

2.3.1. Software Development Life Cycles

A well-articulated SDLC model details the phases from requirements analysis to the final finished product and aids in improving the development process [AOO⁺20]. By breaking down increasingly complex tasks into smaller subtasks, SDLC models help to plan and monitor work as well as support cooperation and communication between team members to ensure efficient development of high-quality software [Kne17]. Each life cycle model defines the development process differently and has varying goals, advantages, and drawbacks. The selection and usage of a suitable SDLC is a challenging task, but highly significant to deliver high-quality software, that fulfills customer requirements within a given schedule and budget [MNB12, AOO⁺20]. According to Benediktsson et al. [BDT06], typical SDLCs can be categorized into *sequential*, *incremental*, *evolutionary*, and *agile* approaches. The following will give a short description of the categories with exemplary SDLC models.

Sequential. Sequential approaches, such as the *Waterfall Model* [Roy87] or *V-Model* [Boe79], follow a sequential path of executing software development processes [BDT06]. The system development process is separated into distinct phases, whereas each phase should be completed before proceeding with the subsequent phase [MOJ05]. The progress is linear and reaching pre-defined milestones determines when a phase is considered accomplished and the next phase should be started. Thereby, sequential approaches are highly structured with a maximum amount of control over the process. Strictly following sequential approaches makes the software development process inflexible and resistant to change [MOJ05, BDT06].

Incremental. Incremental approaches, such as the correspondent *Incremental Model*, follow a phased development and develop elements in stages by splitting a plan into smaller segments. These segments (increments) are seen as isolated, meaningful subsets of the project, that can be developed, tested, and implemented independently from the remaining segments. Hence,

2. Theoretical Background

each increment is implemented comparably to a "mini-waterfall" process [MG10]. These distinct increments are developed in multiple, parallel mini-cycles [BDT06]. Each mini-cycle adds further functionalities and capabilities, which slowly builds the system incrementally [LB03]. In contrast to evolutionary or agile methods, the increments have pre-defined specifications followed by time-boxed development [LB03, AB15].

Evolutionary. In contrast to incremental approaches, the requirements and specifications for the increments are not fully pre-determined before implementation but might evolve during the development process [BDT06]. Hence, evolutionary approaches, such as the Evolutionary Prototyping Model [Jal12], start with an initial estimate and preliminary specification but allow for details to be throughout during the process. These approaches are feedback-driven and the finite goals evolve based on user needs and changes along the development, which allows for fast adaptations in change-intensive projects [MZ96, PSDB20].

Agile. Agile development approaches are a group of methods based on incremental and iterative development [LLTT12, BM12]. Through a certain set of principles and values, agile approaches aim to reduce software development overhead and allow to adapt to quick changes [ASSA20]. This set of principles and values originated through an official alliance of 17 software engineering consultants, which published the *Agile Software Development Manifesto* [BBVB⁺01] in 2001. The manifesto prioritizes the response to change in the development life cycle rather than strictly following a defined plan [ASSA20], where the requirements evolve through user feedback and the collective endeavor of self-organizing cross-functional teams [GG17]. Summarized, in agile approaches, the team should self-organize, adjust to requirement changes at any stage, actively involve customers, and work at a pace best for their creativity and productivity [DNBM12]. Currently, the most widely adopted agile methods are Scrum and Extreme Programming [HA13, GG17].

The variety of existing SDLC approaches breaks down complex tasks into smaller subtasks and determines how teams develop software. Even though each model describes the development process differently, they usually all include phases regarding the requirements analysis, design, implementation, testing, maintenance, and deployment [AOO⁺20]. One iteration of the major phases of SDLC models is illustrated in Figure 2.7. In which of the described modalities these stages are approached, if the development process incorporates further stages, or if there is a specific focus on certain stages, depends on the chosen SDLC.




Figure 2.7.: Major software development life cycle stages (own work based on [AOO⁺20]).

In addition to the development process, each process model usually defines an abstract representation of participating roles, activities, and resulting artifacts [CKO92]. The role descriptions help to understand the involvement of participating entities and provide information about their distinct set of behaviors and responsibilities. Similarly, the activity aims to summarize the units

of work that are performed by a role during a certain stage. The artifact description provides information on the tangible by-products that are produced, modified, or used by a process [Anw14]. Again, the set of participating roles, required activities, and resulting artifacts are dependent on the chosen SDLC. For example, Scrum [SS11] requires specific roles such as the Scrum master which are not required in other SDLC models.

2.3.2. Cross-Industry Standard Process for Data Mining

Developing ML-enabled applications require the integration of non-deterministic ML systems into algorithms that were "hard-coded" by humans [Gir21]. However, the unique characteristics of ML systems demand specific background knowledge, skills, and particular approaches, which collide with traditional software development practices [Gir21, LRB⁺19]. The self-learning and stochastic nature of ML adds uncertainty to the systems, which not only require additional knowledge of statistics and information theory from practitioners but also change several SDLC stages [WXMLM21]. For example, collecting requirements involves comparably more preliminary experiments, data needs to be thoroughly preprocessed, and testing becomes non-deterministic [ASM⁺19, Gir21, LRB⁺19, WXMLM21].

Against this backdrop, Wirth and Hipp [WH00] argued that a common process model increases the chance of the project's success by helping to link people with diverse skills together to form an efficient and effective project. Moreover, a process model helps to structure the project, provides advice for each task, and gives guidance to practitioners, especially novices. Thereupon, the authors proposed a **Cross-Industry Standard Process for Data Mining (CRISP-DM)**, an industry- and technology-agnostic standard process for data mining projects. Since then, a multitude of varying ML process models have been proposed based on CRISP-DM. However, CRISP-DM is still the current de-facto industry standard for applying data mining and therefore ML projects [SKMG21]. CRISP-DM can be separated into a *Reference Model* and a *User Guide* [WH00].

2.3.2.1. Reference Model

The CRISP-DM *Reference Model* describes *what to do* in data mining projects and provides an overview of the project's life cycle. The overall life cycle is broken down into six distinct phases, whereas each phase comprises a set of tasks and their corresponding outputs. The authors emphasized, that the sequence of the phases is not strict and that the outcome of each phase determines which phase has to be performed next [WH00]. However, the most frequent dependencies between phases are illustrated in Figure 2.8. Overall, CRISP-DM follows a cyclic approach and provides possibilities for the iterative refinement of each phase and respectively the whole project cycle. The following will give a short description of the CRISP-DM phases as described by the originally proposed reference model [WH00] and CRISP-DM consortium [CCK⁺00].

Business Understanding. In the first phase, practitioners focus on determining the business objective including the project requirements, and success criteria. This step involves assessing the current situation by analyzing the resource inventory, project requirements, assumptions,

2. Theoretical Background




Figure 2.8.: Phases of the CRISP-DM process model (own work based on [WH00]).

constraints, risks, costs, and benefits. Thereupon, business objectives can be translated into data mining problems and a data mining goal. Finally, a project plan can be derived.

Data Understanding. After collecting an initial dataset, practitioners first aim to understand the coarse properties of the available data, such as the data format, quantity, and attributes. Subsequently, the available data is explored more thoroughly through querying, visualization, and reporting. These analyses should address directly the data mining goal and provide insights regarding data quality, potential subsets, hidden information, or distribution of key attributes.

The first and second phases are closely linked. It is required to understand the available data to a certain degree in order to formulate the business and data mining goals in the previous phase. Hence, as illustrated in Figure 2.8, there are frequent dependencies between *Business Understanding* and *Data Understanding*.

Data Preparation. This phase constructs the final dataset, which will be used for model training. Data preparation includes steps such as data selection, data cleaning, reduction to relevant attributes, data transformation, and feature engineering. These steps are usually performed over multiple iterations without any pre-defined order.

Modeling. The modeling phase generates the data mining tool (e.g., the ML model). Usually, various models are selected and applied to the prepared dataset to test the model’s quality and validity. The varying models are compared through a pre-defined test procedure with quality criteria. The design of this test procedure varies depending on the use case, dataset, and modeling technique. Accordingly, the parameter settings of the models are calibrated to optimal values based on the test design as well as the domain knowledge of the expert.

As illustrated in Figure 2.8, the *Data Preparation* and *Modeling* phases are frequently linked. Data problems may only occur during modeling or the modeling phase provides inspiration for new features. Therefore, these two phases are often performed iteratively.

Evaluation. Before deployment, it is required to evaluate the results of the final model and to review the approved model as well as the process that was needed to construct the model. This phase assesses whether the model properly achieves the business objectives or if an important business goal can not be sufficiently addressed. Finally, a decision on the deployment is being made.

Deployment. Finally, the model is being deployed. This can take various forms such as a productized ML model, a repeatable data mining process, or simply a generated report of the results. A deployment plan as well as model monitoring and maintenance help to make

appropriate use of the model. Lastly, the project is reviewed, documented, and organized in a way that the customer can use it.

2.3.2.2. User Guide

The CRISP-DM *User Guide* describes *how to do* data mining projects with detailed tips for each phase and task [WH00, CCK⁺00]. The user guide goes beyond the descriptions of the reference model and provides comprehensive guidance on how to perform the different tasks. For each phase, the *User Guide* is structured into multiple task descriptions. These descriptions are accompanied by the corresponding rationale, desired output, required activities, hints at pitfalls and problem areas, as well as tips for the process. This advice may also take the form of checklists, questionnaires, tools, techniques, sequences of steps, and decision points. Additionally, the CRISP-DM *User Guide* provides descriptions of the desired outputs and contents for the most important reports. The *User Guide* is available at a generic level in [CCK⁺00].

2.3.2.3. Extensions and Variations

In recent years, many extensions and variations of CRISP-DM have been proposed that are tailored to specific focus topics. Some studies aimed to address shortcomings of CRISP-DM and conceptualized models that specifically concentrate on integrating ML development with the overall software engineering process [HSM⁺19], emphasizing requirements collection [ABB⁺19], data infrastructure and monitoring [BCN⁺17], or focusing on quality assurance [SBD⁺21]. Additionally, some companies introduced their own processes based on CRISP-DM, for example, IBM uses *Lightweight IBM Cloud Garage Method for Data Science*¹ or Microsoft utilizes their own *Team Data Science Process*² to provide a standardized life cycle to structure their data science projects. So far, none of these models have been widely adopted and CRISP-DM is still seen as the de-facto industry standard [SKMG21].

¹<https://developer.ibm.com/articles/the-lightweight-ibm-cloud-garage-method-for-data-science/>

²<https://learn.microsoft.com/en-us/azure/architecture/data-science-process/overview>

The research design represents an overall plan for the activities and procedures that need to be performed to fulfill the research objectives adequately [Bha12, Cre09]. This plan involves multiple decisions, from the broad assumptions of the research inquiry to the specific methods that are used to collect and analyze data [Cre09]. Accordingly, the following sections elaborate on the broad paradigm of our research (see Section 3.1) and correspondingly on the selected research methods (see Section 3.2).

3.1. Research Paradigm

Depending on the research goal, the research endeavor can either be of *inductive* or *deductive* nature [Bha12]. *Inductive* research aims to infer theoretical concepts and patterns from observed data, whereas *deductive* research tests these already known theoretical concepts and patterns through new empirical data [Bha12]. Following the described RQs, we pursue an *inductive* research strategy. Hereby, data can be collected through *qualitative*, *quantitative*, and *mixed* strategies [Cre09]:

- *Qualitative* strategies, such as expert interviews or focus groups [GSTC08], aim to understand and interpret complex social and organizational phenomena [SC14]. Qualitative research does not aim at generalizing results since the complexity of the phenomena usually delimits the available units of analysis [SC14].
- *Quantitative* strategies, such as surveys or experiments [Cre09], intend to explain phenomena by collecting numerical data [Suk07]. Usually, the numerical data is then analyzed through mathematically based methods such as statistics [Cre09].

3. Research Design

- *Mixed* strategies combine quantitative and qualitative methods with the goal of developing richer insights into phenomena, that cannot be understood through qualitative or quantitative methods alone [VBB13].

In this work, we chose a qualitative strategy since we aim to understand the practical adoption of FedML, a complex, emerging technology with limited units of analysis. Hence, we follow an *inductive, qualitative* research approach.

3.2. Research Methods

In our studies, we mainly gathered data through expert interviews (P1, P2) and focus groups (P1-P4). Additionally, we used literature reviews in all studies (P1-P4) to assimilate knowledge and ground our research in current literature. We followed the Design Science Research (DSR) methodology by Peffers et al. [PTRC07] for the design and creation of solution artifacts (P3, P4). Table 3.1 provides an overview of methods that have been applied for each embedded core contribution. The subsequent sections briefly introduce the four used methods: literature reviews (see Section 3.2.1), focus groups (see Section 3.2.2), expert interviews (see Section 3.2.3), and design science research (see Section 3.2.4). The specific application of each research method is described in detail in each separate publication.

3.2.1. Literature Reviews

Literature reviews are an essential method that allows researchers to understand, synthesize, and get an overview of the extant body of knowledge [Sch15, PTJK15]. By conducting literature reviews, researchers can get a sound theoretical foundation for the intended research goal [Bak00]

No.	Title	Literature Review	Focus Group	Expert Interviews	Design Science
P1	Revealing the Impacting Factors for the Adoption of Federated Machine Learning in Organizations	○		●	
P2	Unlocking the Potential of Collaborative AI - On the Socio-Technical Challenges of Federated Machine Learning	●	●	●	
P3	A Process Model for the Practical Adoption of Federated Machine Learning	○	○		●
P4	A Pathway for the Practical Adoption of Federated Machine Learning Projects	○	○		●
Legend: ● Primary Method Used in the Publication ○ Secondary Method Used in the Publication					

Table 3.1.: Overview of research methods applied in the embedded publications.

and identify research gaps [vBSN⁺09, Row12]. Thereby, researchers can justify the proposed study, and incrementally build on and extend the research of others [LE06, PTJK15]. Hence, a literature review poses a relevant first step and essential feature when conducting a research project [Bak00, WW02].

Depending on the research objective, a literature review can either seek to summarize prior knowledge, aggregate/integrate data, build explanations, or critically assess extant literature [PTJK15]. For the research goal of this dissertation, we leveraged literature reviews to summarize prior knowledge, thereby building our knowledge base and theoretical foundation of our research endeavor. Depending on the goal and search strategy, these literature reviews can either take the form of a *descriptive*, *narrative*, or *scoping* review [PTJK15]. In the embedded core publications, we conducted a structured *descriptive* review (P2) and multiple *narrative* reviews (P1, P3, P4). Therefore, the following sections will elaborate on these types of literature reviews.

Descriptive Reviews. A descriptive review usually employs a structured search method [PTJK15] and aims to reveal interpretable patterns from existing literature [GJK87]. Each individual study is treated as one data point and used to identify patterns and trends amongst the surveyed literature corpus [KH05, PTJK15]. Accordingly, descriptive reviews produce quantification, such as frequency counts on research outcomes to determine the extent to which a literature corpus reveals trends [GJK87, KH05]. The result of a descriptive review is usually claimed to represent the state of the art in a research domain [KH05].

In P2, we used a *descriptive* literature review to identify, aggregate, and synthesize critical attributes of business models for inter-organizational collaborations. We leveraged the resulting insights as a fundamental basis to investigate the socio-technical challenges of FedML projects. We employed the structured search strategy as proposed by Zhang et al. [ZBT11]. This research strategy follows a five-step approach. First, the search process starts by identifying relevant publication venues and search engines. Thereupon, a subsequent manual search screens the resulting papers and gathers a collection of relevant studies, thereby establishing a *quasi-gold standard*. Thirdly, search strings for the automated search are elicited from the *quasi-gold standard* or defined subjectively. Consequently, the selected digital libraries are searched automatically by using the search strings. Finally, the search performance is evaluated in relation to the *quasi-gold standard* [ZBT11].

Narrative Reviews. A narrative review usually does not follow a systematic search process [Dav00] and aims to present a broad perspective on a given subject [GJA06, PTJK15]. Against this backdrop, it is important to note, that a narrative review may be vulnerable to subjectivity due to its unstructured nature [GJA06]. Often, narrative reviews do not seek to generalize knowledge from the reviewed literature corpus [Dav00]. Thereby, narrative reviews are usually opportunistic and intend to identify the diversity of the available literature, including selected grey literature [Dav00]. Since there might be a long lag time between the submission and publication of evidence, the inclusion of grey literature helps to ensure the most current picture of a research subject [Pae17]. This makes narrative reviews suitable to provide up-to-date knowledge about a specific topic [Rot07].

In P1, we used the findings of a *narrative* literature review on ML adoption studies to assess and contextualize the results of an expert interview study. In P3 and P4, we built our knowledge

base by conducting *narrative* literature reviews on SDLC models and ML life cycles. We assessed which of the identified practices, procedures, and information is relevant and applicable to the development of the FedML process model and the activity model.

3.2.2. Focus Groups

Focus group interviews are an effective qualitative data collection method to gather a broad range and variety of views, attitudes, and ideas on a specific topic [MWM⁺05, AA05]. A focus group consists of a small, relatively homogeneous group of six to nine people with relevant expertise including at least one moderator [AA05, BWG20]. Led by the moderator, the focus groups usually last 90-120 minutes and are carefully planned discussions that enable participants to build on the responses of others [Den17]. Thereby focus groups leverage the group dynamic to increase the richness of obtained information [KLB04, Den17]. A focus group usually incorporates the following main steps [Mor88, KLB04, NWDM18]:

1. *Define Research Problem:* The first step defines the objectives of the study and identifies participants. Additionally, it is required to select a suitable location [NWDM18].
2. *Data Collection:* Prior to the discussion, the moderator should get familiar with the script, group dynamics, and the environment [NWDM18]. Then the discussion should comprise a preliminary introduction, the discussion itself, and a conclusion [NWDM18]. Hereby, the moderator should create a comfortable atmosphere, which fosters the discussion and group dynamic [Den17]. The discussion should be recorded [BWG20] and the recording can take the form of handwritten notes, transcripts, audio recordings, or video recordings [KLB04, BWG20]
3. *Data Analysis:* After data collection, the recorded discussion and corresponding data will be analyzed [NWDM18]. The method of analysis varies depending on the underlying form of transcript [NWDM18].
4. *Results and Reporting:* Finally, the researcher needs to decide on the suitable target audience and if the results should be reported to academics, policymakers, practitioners, and/or the participants of the study [NWDM18].

The focus group method is suitable for types of research objectives that require exploring a new topic, generating new ideas, and gathering information or feedback from practitioners [Edm99, Bel12]. Therefore, focus groups can be a valuable tool for the initial evaluation of potential solutions or to collect characterizing information about current practices or problems [KLB04].

In P2, we used the focus group method to initially explore the socio-technical challenges of inter-organizational FedML projects due to the novelty of the topic. In P3 and P4, we leveraged focus groups to explore the problem space, iteratively evaluate the current state of our solution, and gather additional information on current practices of ML and FedML project life cycles.

3.2.3. Expert Interviews

The qualitative interview is one of the most commonly used data collection techniques in qualitative research [Sea99, MN07]. Interviews are often leveraged to collect opinions, experiences, interpretative perspectives, and historical data from the memories of the interviewees [Sea99, MWM⁺05]. An interview is usually an exchange between a researcher and an expert on the investigated subject [KB18]. However, there are various types of qualitative interviews, from *group interviews* with multiple interviewees, *structured interviews* with a complete interview script to *unstructured* and *semi-structured interviews* with an incomplete script [FF00]. Hereof, *semi-structured interview* are the most used type of interviews in qualitative research in IS [MN07]. The following will only focus on *semi-structure interviews* since this is the only type of interview we conducted in our research.

The *semi-structured interview* is characterized by its incomplete interview guide with pre-defined topics and open-ended questions [MN07, BWG20]. The interview study usually involves the following steps [Ada15, Sea99]:

1. *Selecting Respondents and Arranging Interviews:* Depending on the research objective, the size of the target group should represent a manageable random group of the stratified sample [Ada15]. The identification and request of knowledgeable experts should be done such that the overall interview group is a randomly chosen sample of experts that eliminates the biasing effect of convenience samples [Ada15]. After identification and first contact, an individual appointment should be scheduled [Ada15].
2. *Drafting the Interview Guide:* In semi-structured interviews, the interview guide should be a more-or-less partially developed script with questions that was prepared beforehand [MN07]. The interview guide should involve at least an introduction of the researcher, the purpose of the interview, key questions regarding the research objective, and possible follow-ups or ending questions [MN07, Ada15]. Since potentially all data is useful [Sea99], the questionnaire should allow open-ended questions [BWG20].
3. *Starting and Conducting the Interview:* The interview should start with an introduction of the researcher, a short explanation of the research objective, and the matter of confidentiality should be discussed [Sea99, Ada15]. The call or meeting could be recorded or the interviewer might take notes during the interview [Sea99, Ada15]. The interview and interviewer must be flexible enough to allow unforeseen information and to judge if an answer needs to be cut off in case the conversation has wandered too far [Sea99]. After the interview, the data is coded and analyzed [Sea99].

In our studies, we followed the guidelines of Myers and Newman [MN07], who argue that the qualitative interview is a social interaction that should be seen as a drama. Through the dramaturgical analogy, Myers and Newman [MN07] derive seven guidelines for conducting semi-structured interviews: (1) The researcher should situate themselves as both, interviewer and interviewee before the interview takes place. (2) Social dissonance should be minimized. More specifically, the researcher should minimize anything that might make the interviewee feel uncomfortable. Thereby, the quality of disclosure can be increased. (3) The researcher should interview a variety of voices and bias should be avoided. (4) Also, each participant is a creative

3. Research Design

interpreter of their world. Therefore, an interview produces the creation of multiple texts (initial transcript and its interpretation). (5) By mirroring parts and phrases of the interviewees' comments, the researcher could construct subsequent questions. Ideally, these questions are open rather than closed and the questions should move from general to specific. (6) Due to the incomplete script, the interviewer is required to be flexible, and open, and to be able to improvise in relation to the interviewees' attitudes. (7) Lastly, the collected data should be confidential as agreed at the beginning of the interview. Figure 3.1 displays the interrelation and influence of the described guidelines on the interviewer and interviewee in relation to the interview context [MN07].

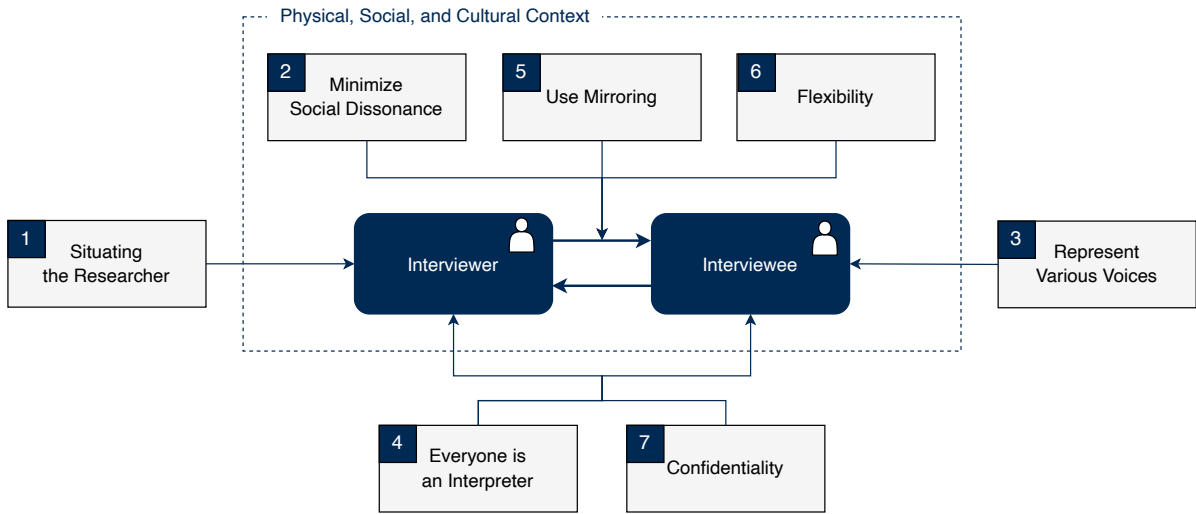


Figure 3.1.: Interrelation and influence of interview guidelines (based on [MN07]).

In P1, we used semi-structured interviews to collect data on the influencing factors for the adoption of FedML in organizations. In our studies, we followed the guidelines of Myers and Newman [MN07]. In P2, we conducted semi-structured expert interviews to draw from the experiences of experts on the socio-technical challenges of FedML projects. Thereby, we complemented the results of a focus group discussion, gathered more in-depth insights, and enriched the captured variety of voices.

3.2.4. Design Science Research Methodology

DSR provides a methodical, rigorous approach for creating and evaluating innovative, purposeful artifacts for a relevant organizational problem [HMPR04]. This research framework builds on the business needs of an appropriate application domain (environment) to ensure relevance and the applicable scientific foundation (knowledge base) to establish rigor [HMPR04]. The interplay of the design activities with the environment is commonly referred to as the *Relevance Cycle*, whereas the connection of the knowledge base with the design activities denotes the *Rigor Cycle*. Additionally, iterating between the core activities of designing the artifact is referred to as the *Design Cycle* [Hev07].

Hevner et al. [HMPR04] proposed seven DSR guidelines to ensure scientific rigor for effective DSR in IS: (1) DSR must produce a viable artifact in the form of a construct, model, method, or instantiation. (2) This artifact should be a technology-based solution, which addresses a relevant business problem and (3) it should be rigorously demonstrated through well-executed evaluation methods. (4) Additionally, the design artifact, design foundations, or design methodologies must provide a clear and verifiable research contribution. (5) The construction and evaluation of the artifact must rely on the application of rigorous research methods and (6) should be appropriate to the context of the environment. (7) The research must be effectively presented to technology-oriented as well as management-oriented audiences.

Peppers et al. [PTRC07] proposed a DSR methodology for IS research that incorporates these guidelines, principles, and procedures. The DSR methodology process is as follows [PTRC07]:

1. *Problem Identification and Motivation*: The first step aims to clarify the problem and highlights the importance of finding a solution.
2. *Objectives of a Solution*: Based on the identified problem, the objectives of a solution should be inferred rationally.
3. *Design and Development*: Following the preparatory activities, the third stage is the core of design science and aims to generate the artifact. This step includes the specification of the desired functionalities, the artifact architecture, and finally its creation.
4. *Demonstration*: Subsequently, the use of the artifact should be demonstrated in one or more instances of the problem to showcase its ability to solve the identified problem.
5. *Evaluation*: More comprehensively, this step evaluates how well the artifact supports a solution to the problem by comparing the objectives of a solution to the actual observed results. After evaluation and interpretation of the results, practitioners could iterate back to step 2 or 3 to refine the process and improve the effectiveness of the artifact.
6. *Communication*: Finally, the practitioners communicate the problem and the designed artifact including its relevance, utility, design process, and effectiveness. This information should be communicated to other researchers and further relevant audiences. Therefore, communication requires knowledge of the research domain and its disciplinary culture.

As displayed in Figure 3.2, the DSR process is not necessarily nominally sequential but can comprise iterations and a researcher could start at almost any step [PTRC07].




Figure 3.2.: Design science research process in Information Systems (based on [PTRC07]).

In P3, we followed the DSR methodology to develop a process model of an end-to-end FedML

3. Research Design

project life cycle. In P4, we used DSR to provide comprehensive activity models that outline the required tasks in the development of FedML systems. Both publications adhere to the seven guidelines of effective DSR as proposed by Hevner et al. [HMPR04].

Part B

Core Contributions

The publication-based dissertation comprises the following publications (P1-P4):

- (P1) Müller, T., Zahn, M., and Matthes, F. (2024). Revealing the Impacting Factors for the Adoption of Federated Machine Learning in Organizations. *Proceedings of the 57th Hawaii International Conference on System Sciences*. 7343-7352.
- (P2) Müller, T., Zahn, M., and Matthes, F. (2023). Unlocking the Potential of Collaborative AI - On the Socio-Technical Challenges of Federated Machine Learning. *Proceedings of the 31st European Conference on Information Systems. Research Papers*, 245.
- (P3) Müller, T., Zahn, M., and Matthes, F. (2023). A Process Model for the Practical Adoption of Federated Machine Learning. *Proceedings of the 29th Americas Conference on Information Systems*, 1.
- (P4) Müller, T., Zahn, M., and Matthes, F. (2023). A Pathway for the Practical Adoption of Federated Machine Learning Projects. *Proceedings of the 27th Pacific-Asia Conference on Information Systems*, 6.

The following sections give an overview of each publication via a fact sheet with key facts about the publication and the abstract of the paper.

Listed for completeness, the author contributed to the following publications during the time of my doctoral studies. Insights from these studies complement the content of this dissertation, however, are not part of the core contributions:

- (P5) Müller, T.; Gärtner, N.; Verzano, N. and Matthes, F. (2022). Barriers to the Practical Adoption of Federated Machine Learning in Cross-company Collaborations. *Proceedings of the 14th International Conference on Agents and Artificial Intelligence*, pages 581-588.
- (P6) Müller, T., Stäbler, M., Gascon, H., Köster, F., and Matthes, F. (2023). SoK: Assessing the State of Applied Federated Machine Learning. *arXiv*.
- (P7) Zahn, M., Müller, T., and Matthes, F. (2024). Supporting Managerial Decision-Making for Federated Machine Learning: Design of a Technology Selection Tool. *Proceedings of the 57th Hawaii International Conference on System Sciences*.
- (P8) Müller, T.; Zahn, M.; and Matthes, F. (2023). On the Adoption of Federated Machine Learning: Roles, Activities and Process Life Cycle. *Proceedings of the 25th International Conference on Enterprise Information Systems*, pages 525-531.

4.1. (P1) Revealing the Impacting Factors for the Adoption of Federated Machine Learning in Organizations

Table 5.4. Fact sheet on publication P1

Authors	Müller, Tobias* Zahn, Milena* Matthes, Florian* *Technical University of Munich, TUM School of Computation, Information and Technology Chair of Software Engineering for Business Information Systems, Boltzmannstraße 3, D-85748 Garching, Germany
Outlet	57th Hawaii International Conference on System Sciences (HICSS-57)
Page Number	10
Status	Published
Contribution of first author	Problem definition, research design, data collection, data analysis, interpretation, writing, reporting
Citation	Müller, T., Zahn, M., and Matthes, F. (2024). Revealing the Impacting Factors for the Adoption of Federated Machine Learning in Organizations. <i>Proceedings of the 57th Hawaii International Conference on System Sciences</i> . 7343-7352.

Abstract. The success of Machine Learning is driven by the ever-increasing wealth of digitized data. Still, a significant amount of the world’s data is scattered and locked in data silos, which leaves its full potential and therefore economic value largely untapped. Federated Machine Learning is a novel model-to-data approach that enables the training of Machine Learning models on decentralized, potentially siloed data. Despite its potential, most Federated Machine Learning projects fail to actualize. The current literature lacks an understanding of the crucial factors for the adoption of Federated Machine Learning in organizations. We conducted an interview study with 13 experts from seven organizations to close this research gap. Specifically, we draw on the Technology-Organization-Environment framework and identified a total of 19 influencing factors. Thereby, we intend to facilitate managerial decision-making, aid practitioners in avoiding pitfalls, and thereby ease the successful implementation

4.2. (P2) Unlocking the Potential of Collaborative AI - On the Socio-Technical Challenges of Federated Machine Learning

Table 5.1. Fact sheet on publication P2

Authors	Müller, Tobias* Zahn, Milena* Matthes, Florian* *Technical University of Munich, TUM School of Computation, Information and Technology Chair of Software Engineering for Business Information Systems, Boltzmannstraße 3, D-85748 Garching, Germany
Outlet	31st European Conference on Information Systems (ECIS 2023)
Page Number	14
Status	Published
Contribution of first author	Problem definition, research design, data collection, data analysis, interpretation, writing, reporting
Citation	Müller, T., Zahn, M., and Matthes, F. (2023). Unlocking the Potential of Collaborative AI - On the Socio-Technical Challenges of Federated Machine Learning. <i>Proceedings of the 31st European Conference on Information Systems. Research Papers</i> , 245.

Abstract. The disruptive potential of AI systems roots in the emergence of big data. Yet, a significant portion is scattered and locked in data silos, leaving its potential untapped. Federated Machine Learning is a novel AI paradigm enabling the creation of AI models from decentralized, potentially siloed data. Hence, Federated Machine Learning could technically open data silos and therefore unlock economic potential. However, this requires collaboration between multiple parties owning data silos. Setting up collaborative business models is complex and often a reason for failure. Current literature lacks guidelines on which aspects must be considered to successfully realize collaborative AI projects. This research investigates the challenges of prevailing collaborative business models and distinct aspects of Federated Machine Learning. Through a systematic literature review, focus group, and expert interviews, we provide a systemized collection of socio-technical challenges and an extended Business Model Canvas for the initial viability assessment of collaborative AI projects.

4.3. (P3) A Process Model for the Practical Adoption of Federated Machine Learning

Table 5.2. Fact sheet on publication P3

Authors	Müller, Tobias* Zahn, Milena* Matthes, Florian* *Technical University of Munich, TUM School of Computation, Information and Technology Chair of Software Engineering for Business Information Systems, Boltzmannstraße 3, D-85748 Garching, Germany
Outlet	29th Americas Conference on Information Systems (AMCIS 2023)
Page Number	10
Status	Published
Contribution of first author	Problem definition, research design, data collection, data analysis, interpretation, artifact design, verification, writing, reporting
Citation	Müller, T., Zahn, M., and Matthes, F. (2023). A Process Model for the Practical Adoption of Federated Machine Learning. <i>Proceedings of the 29th Americas Conference on Information Systems</i> , 1.

Abstract. The wealth of digitized data forms the fundamental basis for the disruptive impact of Machine Learning. Yet a significant amount of data is scattered and locked in data silos, leaving its full potential untouched. Federated Machine Learning is a novel Machine Learning paradigm with the ability to overcome data silos by enabling the training of Machine Learning models on decentralized, potentially siloed data. Despite its advantages, most Federated Machine Learning projects fail in the project initiation phase due to their decentralized structure and incomprehension interrelations. The current literature lacks a comprehensible overview of the complex project structure. Through a Design Science Research approach, we provide a process model of a Federated Machine Learning life cycle including required activities, roles, resources, artifacts, and interrelations. Thereby, we aim to aid practitioners in the project initiation phase by providing transparency and facilitating comprehensibility over the entire project life cycle.

4.4. (P4) A Pathway for the Practical Adoption of Federated Machine Learning Projects

Table 5.3. Fact sheet on publication P4

Authors	Müller, Tobias* Zahn, Milena* Matthes, Florian* *Technical University of Munich, TUM School of Computation, Information and Technology Chair of Software Engineering for Business Information Systems, Boltzmannstraße 3, D-85748 Garching, Germany
Outlet	27th Pacific Asia Conference on Information Systems (PACIS 2023)
Page Number	16
Status	Published
Contribution of first author	Problem definition, research design, data collection, data analysis, interpretation, artifact design, verification, writing, reporting
Citation	Müller, T., Zahn, M., and Matthes, F. (2023). A Pathway for the Practical Adoption of Federated Machine Learning Projects. Proceedings of the 27th Pacific-Asia Conference on Information Systems, 6.

Abstract. Big data forms the fundamental basis for the success of Machine Learning. Yet, a large amount of the world’s digitized data is locked up in data silos, leaving its potential untapped. Federated Machine Learning is a novel Machine Learning paradigm with the potential to overcome data silos by enabling the decentralized training of Machine Learning models through a model-to-data approach. Despite its potential advantages, most Federated Machine Learning projects fail to actualize due to their decentralized structure and incomprehensive interrelations. Current literature lacks clear guidelines on which steps need to be performed to successfully implement Federated Machine Learning projects. This study aims to close this research gap. Through a design science research approach, we provide three distinct activity models which outline required tasks in the development of Federated Machine Learning systems. Thereby, we aim to reduce complexity and ease the implementation process by guiding practitioners through the project life cycle.

Part C

Summary of Results and Discussion

This dissertation entails four publications that address the three RQs. In this chapter, we summarize the key results and discuss how the publications contribute to answering the RQs. We give a brief resume of the research process and then present the key findings. We start by summarizing the insights on RQ1 and the influencing factors for FedML adoption (see Section 5.1). Subsequently, we describe the main results of RQ2 and address the socio-technical challenges of FedML (see Section 5.2). Finally, we present the findings on RQ3 and the FedML SDLC (see Section 5.3). Additionally, every section discusses the findings in relation to the existing literature. All key findings are summarized in Table 5.2.

5.1. RQ1: Influencing Factors for the Organizational Adoption

Research Question 1 (RQ1)

What are the major factors influencing the successful adoption of Federated Machine Learning systems in organizations?

In P1, we conducted an expert interview study to gather the influencing factors for the adoption of FedML in organizations and address RQ1. In the interview study, we gathered the experiences of 13 relevant experts from seven organizations with a broad spectrum of backgrounds, job roles, and experiences. Hereby, we first investigated why the practitioners adopted FedML (1) as well as the overall challenges and risks (2). Subsequently, we drew from the TOE Framework and specifically explored influencing factors regarding their *Technological*, *Organizational*, and *Environmental* context (3). The data was coded according to the guidelines of the *Reflexive*

Thematic Analysis process [BCHT18]. As illustrated in Figure 5.1, we identified a total of three reasons for adoption, six overall challenges, and 19 TOE factors.

(1) *Reasons of Adoption.* Our results showcase three main reasons for the adoption of FedML. Practitioners leverage FedML to overcome data privacy issues or to establish novel fields of application by using currently untapped sensitive data. Also, the communication-efficient nature of FedML is a major motivator for the usage of FedML.

(2) *Challenges and Risks.* Most experts stated that due to the novelty of the technology, most organizations face challenges as first or early adopters. This means that organizations need to tackle particularly complex compliance assessments and uncertainties regarding regulatory or emerging standards. Additionally, first and early adopters need to develop novel business cases and cannot rely on existing ones, making it difficult to allocate budget and get top management support. Besides, organizations are currently mainly challenged by various aspects regarding the collaboration with other organizations and by the complex technical implementation of FedML. We aimed to facilitate the complex technical implementation through the artifacts of P3 and P4.

(3) *Technological, Organizational, and Environmental Factors.* Through our study, we identified 19 influencing factors for the adoption of FedML in organizations and grouped these factors according to the TOE framework. Out of the 19 influencing factors, we allocated nine factors to the technological context, five factors to the organizational context, and five factors to the environmental context. All experts unanimously agreed that collaboration management is an essential and complex factor for the successful development of FedML projects. We specifically addressed this challenge in P2. From an organizational perspective, the degree of top management support and expertise within the organization heavily impacts the success of FedML projects. Regarding technological aspects, it is most important that organizations ensure high data quality and the interoperability of the various data sources. As for the environmental factors, data privacy issues and the missing legal clarity impede the development FedML systems the most. Figure 5.1 provides an overview of all 19 identified influencing factors.

Based on this study, we provide a thorough understanding of the risks and influencing factors in the practical adoption of FedML. By revealing the critical factors, we hope to aid management-oriented and technology-oriented audiences to avoid pitfalls and overcome challenges in the planning and development process.

Discussion. ML is a general-purpose technology with the potential to transform current businesses [IL20, Cra21]. Hereby, the main value drivers of ML and thereby its reasons for adoption are the possibility of unprecedented business opportunities, as well as the use of ML to provide decision-support, drive customer engagement, and automate processes [BLS⁺21]. Our results build on these insights and show that practitioners leverage FedML to establish further application areas through collaboration and by overcoming data privacy issues (P1). Additionally, we recognized that the communication efficiency of FedML is another main motivator for its use (P1). Despite its advantages, ML and FedML systems show a relatively low organizational adoption rate [ZKB⁺21, HN21, LLW⁺21].

Motivated by this, IS literature sparked various studies that investigate the influencing factors for the organizational adoption of ML systems [ACM19, PTH19, CRDB21, HAAY⁺23]. We

5. Summary of Results and Discussion

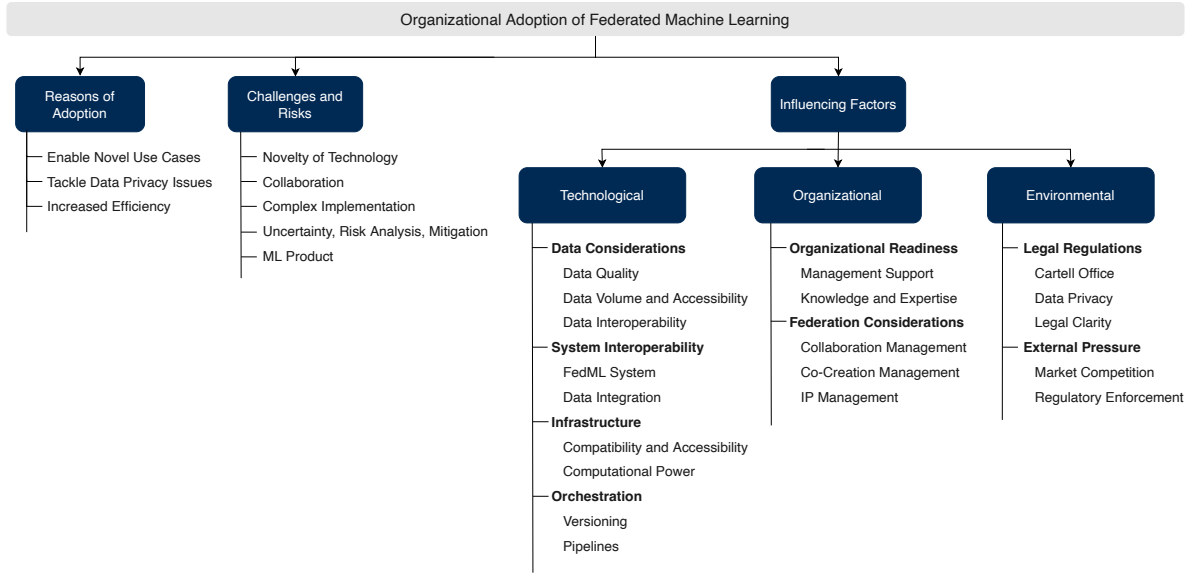


Figure 5.1.: Overview of insights for the organizational adoption of FedML (based on P1).

argue that the decentralized nature of FedML adds distinct influencing factors for its practical adoption that needs to be investigated (P1). In accordance with current literature, our results show that the degree of top management support, lack of expertise, infrastructure, as well as data quality, and quantity are among the most critical influencing factors in the adoption of ML [ACM19, KKG21, BKKP22] and FedML systems (P1). However, the novelty and decentralization of FedML introduce additional technology-specific influencing factors. Organizations that currently leverage FedML are early adopters that can not rely on best practices which yields additional complex compliance assessments due to regulatory uncertainties and ambiguous legal clarity regarding data privacy and antitrust laws (P1). Also, through the lack of best practices practitioners are currently missing guidance for the implementation process which challenges the successful realization of FedML projects (P1). In our studies P3 and P4, we aimed to address this issue. The decentralized paradigm additionally requires the interoperability of data and participating systems. Moreover, there is a need for collaboration and co-creation management (P1), which we further analyze in P2.

Through P1, we confirmed the influencing factors for the adoption of ML systems and recognized that FedML also incorporates the factors. However, we also found that FedML introduces further impacting factors that need to be considered due to its novelty and decentralized setting.

5.2. RQ2: Socio-Technical Challenges of Federated Machine Learning

Research Question 2 (RQ2)

What are the distinct socio-technical challenges of collaborative Federated Machine Learning projects?

In P2, we used a tripartite research approach to identify the socio-technical challenges of collaborative FedML projects and thereby address RQ2. While P1 considers the whole SDLC, this study places a specific focus on collaboration creation. First, we performed a systematic literature review to gain an overview of the challenges of prevailing collaborative business models (1). Subsequently, we explored the socio-technical challenges of collaborative FedML (2) through a focus group discussion with five experts from two project teams that worked on cross-company FedML projects. Finally, we complemented the insights through additional semi-structured expert interviews with five experts from three organizations. We closed the interview study after reaching theoretical saturation.

(1) *Challenges of Prevailing Collaborative Business Models.* Based on a systematic literature review, we compiled an extensive overview of the challenges of prevailing collaborative business models. In the scope of our study, we explicitly focused on network-related challenges. Based on Diirr and Cappelli [DC18], we grouped the identified network-related challenges into *Management Challenges*, *Business Process Challenges*, and *Collaboration Challenges*. Out of the 18 identified challenges, we assigned five management challenges, ten business process challenges, and three collaboration challenges. Based on the number of occurrences, we recognized that the lack of commitment and trust between participating organizations as well as the decision-making and coordination slowness are the most critical challenges. Also, agreeing on the distribution of financials and aligning the collaboration with the organizations' own objectives, culture and ethics are among the most important challenges in the creation of collaborative business models.

(2) *Socio-Technical Challenges of Collaborative FedML Projects.* Building on the previous insights, we used a focus group discussion and expert interviews to specifically investigate the socio-technical challenges of collaborative FedML projects. Through the socio-technical challenges, we provide a detailed understanding of the challenges that need to be addressed in the creation of a business model in collaborative ML use cases. We proposed to structure the identified aspects into four distinct clusters: *Collaboration Management*, *Co-Creation Management*, *Co-Creation Practices*, and *FedML Product*. Our findings reveal that most challenges occur regarding the *Collaboration Management* and *Co-Creation Management*. Thereby, we recognized that the highest priority should be set on jointly clarifying the allocation of model ownership with the accompanying accountabilities and responsibilities. This also goes along with a thorough agreement on the IP management and distribution of financials. Regarding *Collaboration Management*, it is essential to rigorously describe the collaboration structure and how decision-making is handled and coordinated within the collaboration. Finally, we compiled an exhaustive list of guiding questions to offer a list of aspects that need to be considered in the collaboration

5. Summary of Results and Discussion

creation. Thereby, we assist practitioners in the creation of the collaboration agreement and prepare them for challenges related to collaboration. Figure 5.2 provides an overview of all identified socio-technical aspects and the set of FedML-specific guiding questions can be seen in Appendix B.

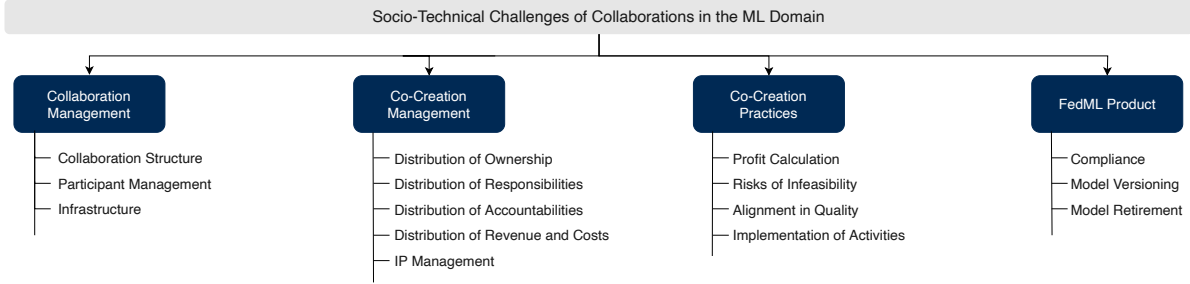


Figure 5.2.: Overview of socio-technical challenges of collaborative ML projects (based on P2).

Discussion. The advent of ML systems transformed the business models of various domains, from the insurance [ZC23] sector to health care [GL19] or education [RH20]. To leverage the opportunities of ML and successfully create value, organizations need to develop a solid business case [ACMA20]. The creation of these ML-enabled business cases is heavily dependent on a variety of socio-technical factors that have been extensively studied in recent literature [BPLW21, DVPHE20, JYK⁺23]. However, as compared to traditional ML systems, FedML involves multiple clients that collaboratively train a ML model [KMA⁺21]. This implies that the creation of FedML business cases also needs to consider challenges of collaborative business models. Our results strengthen this assumption and show the significant intersection between the persisting challenges regarding business models and collaborative FedML (P2). Consistent with our findings on collaborative FedML (P2), prevailing collaborative business models are especially challenged by aspects such as participant management [PLAK21], change management [CCM15, RBBC⁺14], or alignment on the quality of the product [DC18].

However, our study also identified distinct socio-technical challenges of collaborative FedML systems and adaptations to previously known challenges. While the distribution of financials, costs, and revenues is a known issue for collaborative business models [PLAK21, CCM15], these challenges need to be revisited in the context of FedML (P2). For example, collaborators also need to factor in the value of their data and its influence on the final model, when negotiating and agreeing on the allocation of the ML model ownership and incentivization (P2). In this context, the question of how data contributors can be adequately and fairly compensated for their (processed) data sparked a multitude of studies [KXN⁺19, ZLQ⁺20, ZZH⁺21, OMY⁺22] but still remains an open problem [YJLC22]. In sum, we argue that FedML introduces novel socio-technical challenges that need to be further explored and addressed. In P2, we take an initial step and reveal the socio-technical challenges for the collaboration creation of FedML projects.

5.3. RQ3: Software Development Life Cycle of Federated Machine Learning Projects

Research Question 3 (RQ3)

What is the software development life cycle of Federated Machine Learning projects?

In P3 and P4, we investigated the software development life cycle of FedML projects and thereby address RQ3. In both studies, we used DSR methodology to design a process model (1) as well as an activity model (2) for the FedML SDLC. In the rigor cycle, we performed a literature review on SDLCs, ML life cycles, as well as state-of-the-art FedML architectures and algorithms. Additionally, we conducted iterative focus group discussions with varying participants to complement our knowledge base and to ensure that our models are aligned with best practices and meet the requirements of the practitioners. We demonstrated the models in two iterations of demonstrations with subsequent survey-based evaluations. In total, each model was evaluated by 14 experts from eight organizations. Hereby, the evaluation groups consisted of six technical experts and eight target users.

Through an initial focus group discussion, we recognized the need for two distinct models, a process model (P3) and activity models (P4), that address different target groups and serve differing purposes. At this point, it is important to note the differences between the two models. In the focus group discussion, we identified two reasons why FedML projects fail to actualize. First, the project flow is complicated to understand, and the division of tasks is not clear. Therefore, communicating the tasks throughout the team and potential collaborators is cumbersome. Second, practitioners find it challenging to structure the complex implementation process due to its decentralized and collaborative nature. The focus group participants concluded that a process model, which gives a holistic overview of the process, and a detailed activity model with a comprehensive sequence of tasks would alleviate both barriers. Against this backdrop, we decided to conduct two distinct studies.

(1) *Process Model*. The process model (P3) provides a highly abstracted and holistic overview of an end-to-end FedML project including the activities, required resources, role distributions, and resulting artifacts. By providing a comprehensible overview of the sequence and interrelations of these components, we aid business stakeholders and solution architects in the project initiation phase outline the general project, and facilitate communication with potential participants.

Our study shows that the FedML development life cycle is divided into a total of five stages: *Project Initiation*, *Project Validation*, *Project Setup*, *System Design and Development*, and *Deployment and Maintenance*. The *System Design and Development* stage contains the FedML training process and is divided into three more sub-stages: *Global Model Design*, *Local Model Training*, and *Global Model Aggregation*. An overview of the stages with the corresponding goals can be seen in Table 5.1.

Additionally, we identified nine different roles that are required for FedML projects. Roles with strategic responsibilities comprise *Business Stakeholders*, *Project Managers*, *Subject Matter Experts*, *Solution Architects*. Operational roles are largely involved in the development process

5. Summary of Results and Discussion

Stage	Goal
Project Initiation	This initial step lays the projects' foundation by broadly defining the goal of the project and its corresponding high-level requirements.
Project Validation	Validates the feasibility from a business and technical perspective.
Project Setup	Preparing the project by arranging a potential collaboration and setting up the technical foundation for the implementation of the project.
Global Model Design	Yields the data specifications and source code of the initial ML model which will be the technical foundation of the FedML training process
Local Model Training	Details activities regarding the local training on the data contributor's side. Each participant performs the local training and sends the updates to the server.
Global Model Aggregation	Performs the aggregation process and builds the updated global ML model based on the collected model updates.
Deployment and Maintenance	Deploys and maintains the packaged ML model and yields the model service.

Table 5.1. Overview of FedML SDLC stages with their goals.

and encompass *Data Scientists*, *Data Engineers*, *Software Engineers*, *Development-Operations Engineers*. Lastly, a *Legal Representative* needs to be involved. We provide role descriptions in the embedded publication P3 with their mapping to the corresponding activities (see Appendix A). We also showcased the required resources for each activity and the accompanying resulting artifacts. Hereby, one of the critical and most complex FedML-specific activities is the creation of a collaboration agreement. The previous study P2 offers a thorough analysis of its content and provides assistance for practitioners. The final process model with the interrelations between the stages, activities, roles, resources, and artifacts can be seen in Figure 5.3.

(2) *Activity Model*. The activity model (P4) provides a detailed activity sequence through the needed steps of implementing a FedML project. These comprehensive activity descriptions should guide product owners and project managers in the planning phase and the development process. Overall, we proposed four activity models. The first activity model comprises all relevant activities from the *Project Initiation*, *Project Validation*, and *Project Setup* stages. Following, the second model describes the activity sequence of the *Global Model Design* stage, whereas another model incorporates the *Local Model Training* and *Global Model Aggregation* stage. The last model guides through the *Deployment and Maintenance* stage. Each model displays one iteration of the consecutive activities within the stages, however, in practice the stages are usually performed iterative to refine the process. To increase comprehensiveness and for the sake of simplicity, we did not display the iterations between the activity models. These relations are illustrated in the process model (P3). The resulting activity models can be seen in the embedded publication P4 (see Appendix A).

Overall, we investigated the SDLC of a FedML project by developing an abstracted process model and detailed activity models that further specify the process model. The FedML SDLC is divided into a total of seven stages. As seen in the process model, the FedML SDLC is inherently

iterative with multiple iteration loops. However, as shown in the activity models, some activities need to be performed consecutively. With a combination of both artifacts from P3 and P4, we reconstructed the entire FedML SDLC. The suitability, completeness, and usefulness of the tool are shown based on two rounds of demonstrations and evaluations with technical experts and target users. The findings for RQ3 are summarized in Table 5.2.

Discussion. The integration of ML capabilities into larger software-intensive systems requires the alignment of traditional software engineering processes with the unique characteristics of ML [Gir21]. These unique characteristics not only require additional skill sets and roles [KZDB16, LERH20] but also adaptations in the SDLC [IY19, WXML19]. This challenge becomes more difficult with the current lack of standards and development processes that guide practitioners through the ML-specific SDLC [RKW20, ADD21, SBD⁺21]. Similarly, for FedML, we recognized that many FedML projects come to a halt since practitioners struggle to structure the project flow with the corresponding development process and division of tasks (P3, P4). In addition to the characteristics of ML, FedML incorporates another layer of complexity due to its decentralized paradigm which further complicates the multi-faceted process flow. As most FedML projects fail to actualize and never leave the prototype or simulation stage [LLW⁺21], we argue that this can be attributed to a lack of guidance and standards. Therefore, the development of a standard process might alleviate this challenge and facilitate the implementation of production-ready FedML applications (P3, P4).

For traditional ML systems, CRISP-DM [WH00] is used as the current de-facto industry standard process model [SKMG21] to guide practitioners through the ML development process. Since CRISP-DM was developed in 2000 [WH00] and mainly focuses on data mining without explicitly covering ML scenarios, practitioners see the need for revised standard processes and life cycle models with a distinct focus on ML applications [SBD⁺21]. Accordingly, multiple life cycle models for ML projects have been proposed in recent literature [BCN⁺17, ABB⁺19, SBD⁺21, LBM⁺22, KKH23]. Due to the collaborative nature of FedML that alters the development process, we argue that distinct process models for FedML projects are required (P3, P4). Institute of Electrical and Electronics Engineers (IEEE) also recognized this need and published a "*IEEE Guide for Architectural Framework and Application of Federated Machine Learning*" [Gro21]. This guide focuses on the FedML training process but does not consider the whole project life cycle. Our work complements existing ML life cycle models with a perspective on FedML systems as well as the IEEE guide through a holistic view of the whole life cycle (P3, P4).

Consistent with our findings, novel ML process models such as *CRISP-ML(Q)* [SBD⁺21] build on CRISP-DM [WH00] and incorporate emerging practices such as Machine Learning Operations (MLOps) processes as well as ML-specific iteration loops. Kreuzberger et al. [KKH23] additionally compiled a list of roles that are necessary to realize ML projects within the MLOps framework. We identified concurring roles for FedML systems (P3), which implies that FedML does not require additional role profiles beyond regular ML systems. However, we still found that practitioners are challenged by the novel, collaborative paradigm (P3, P4). To help practitioners, we proposed a high-level process model to provide a comprehensible overview of the holistic project structure and interrelations (P3) with additional detailed activity sequence descriptions (P4).

5. Summary of Results and Discussion

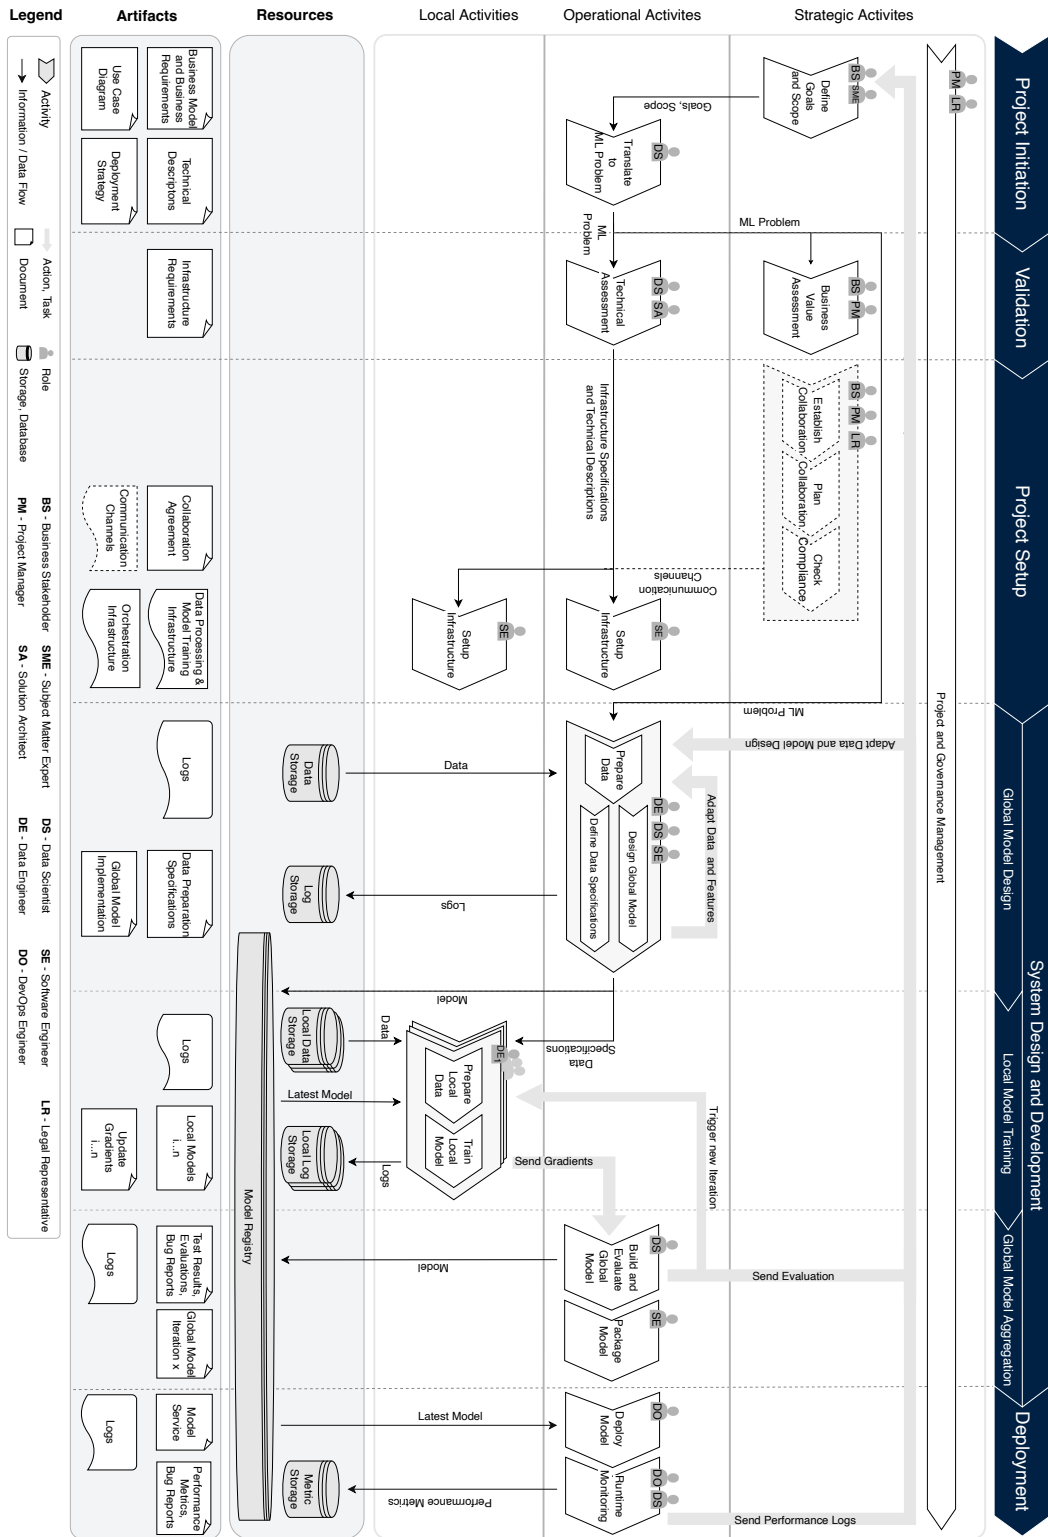


Figure 5.3.: Process model of an end-to-end FedML project.

RQ	No.	Key Findings
RQ1	P1	<ul style="list-style-type: none"> ● Practitioners mainly adopt FedML to overcome data privacy issues, establish novel fields of application, or due to communication efficiency ● Major challenges and risks arise due to its novelty and in collaborations ● Particularly, practitioners can not rely on best practices since currently only a limited number of first or early adopters use FedML systems <ul style="list-style-type: none"> ○ Most mentioned risks are: <ul style="list-style-type: none"> - Regulatory uncertainties and compliance - Collaboration issues - Complex technical implementation ● Identified 19 influencing factors for the organizational adoption of FedML <ul style="list-style-type: none"> ○ Thereof: 9 technological, 5 organizational, 5 environmental factors ○ Most mentioned factors are: <ul style="list-style-type: none"> - Collaboration management - Top management support and expertise - Quality and interoperability of data - Potential data privacy concerns - Missing legal clarity
RQ2	P2	<ul style="list-style-type: none"> ● Identified 18 network-related challenges of collaborative business models <ul style="list-style-type: none"> ○ Thereof: 5 management challenges, 10 business process challenges, and 3 collaboration challenges ○ Most stated challenges are: <ul style="list-style-type: none"> - Lack of commitment and trust - Slow decision-making and coordination - Agreeing on the distribution of financials - Aligning collaboration goals with own objectives, culture, and ethics ● Identified 13 socio-technical challenges of collaborative ML projects <ul style="list-style-type: none"> ○ Thereof: 3 collaboration management, 3 Co-creation management, 4 co-creation practice, 3 FedML product challenges ○ Most stated challenges are: <ul style="list-style-type: none"> - Allocating model ownership with responsibilities and accountabilities - Agreement on IP management and distribution of financials - Clarifying the collaboration structure - Establishing how decision-making is handled and coordinated ● Compiled a set of guiding questions for the creation of FedML collaborations
RQ3	P3	<ul style="list-style-type: none"> ● Practitioners struggle in structuring and communicate the project flow and division of tasks within the team and with collaboration partners <ul style="list-style-type: none"> ○ Creation of a process model to provide a comprehensible overview ● Practitioners have difficulties structuring the FedML development process <ul style="list-style-type: none"> ○ Creation of an activity model to provide a detailed activity sequence
	P4	<ul style="list-style-type: none"> ● FedML SDLC consists of seven distinct stages ● FedML SDLC requires the involvement of nine roles ● The FedML SDLC is inherently iterative but encompasses linear and consecutive activity sequences that are illustrated in the activity models

Table 5.2. Overview of key results.

Contributions and Implications

This dissertation and its embedded publications have several contributions and implications for theory and practice. Overall, we offer the first exploration of the practical adoption of FedML in organizations. Thereby, we contribute to the field of FedML, ML adoption, and ML SDLCs. The following summarizes the most important theoretical (see Section 6.1) and practical implications (see Section 6.2).

6.1. Theoretical Implications

First, our results have several implications for the literature on FedML in general. This thesis initially advances the understanding of FedML by providing an extensive list of risks, challenges, and influencing factors of FedML that have a direct or indirect impact on its organizational adoption (P1). Our findings initially explore the adoption factors of FedML and additionally disclose essential aspects that need to be clarified for its broad practical adoption. Thereby, we emphasize the importance of best practices regarding technical implementation and collaboration management. Also, compliance aspects and legal regulations are currently ambiguous and need to be clarified (P1). Therefore, our insights provide a profound basis for future research on FedML adoption. Subsequently, we investigated the socio-technical challenges that arise due to the collaborative and decentralized characteristics of FedML (P2). As a result, we synthesize empirical knowledge from isolated studies on challenges of collaborative business models, complement these insights with FedML-specific aspects, and elucidate the relationship between these challenges. Our results show that FedML projects need to address the socio-technical challenges that can be found in the literature on collaborative business models [PLAK21, CCM15, RBBC⁺14, DC18] but also introduce further distinct challenges. The newly discovered factors shed light on the challenges regarding collaboration management, co-creation management, co-creation practices,

and additional aspects for the resulting FedML product. Consequently, we provide a comprehensive overview of relevant socio-technical challenges for FedML projects. Thereby, we enhance the understanding of how closely technology and humans are intertwined in the context of FedML and how collaboration affects this interrelation (P2). Moreover, we complement the literature on FedML by providing a comprehensible representation of the FedML SDLC in the form of an abstracted process model (P2) and detailed activity models (P3). Hereby, we structure the implementation of FedML systems and reveal the dependencies and interrelations within the project.

Second, we contribute the literature on ML adoption with a perspective on FedML by providing an itemized list of potential factors influencing the successful adoption of FedML in organizations (P1). For this, the TOE framework [TF90] served as a basis that was populated with empirically derived FedML-specific factors. Thereby, we obtained an inclusive list of technological, organizational, and environmental factors that apply to ML in general with further aspects that are distinctly relevant to FedML. In alignment with studies on ML adoption factors [ACM19, PTH19, CRDB21, HAAY⁺23], our results show that the degree of top management support, lack of expertise as well as data quality and quantity remain the most critical factors for the successful ML and FedML adoption. However, newly discovered factors that are specifically important for FedML include aspects such as regulatory uncertainties, collaboration management, or data interoperability (P1). Additionally, we recognized the importance of standards regarding development processes and best practices (P3, P4) to help practitioners in the complex technical implementation process (P1). These factors show how the decentralized process of FedML adds complexity and changes the impacting factors for the organizational adoption of ML. FedML systems can establish novel fields of application, enhance communication efficiency, and overcome data privacy issues due to their model-to-data approach (P1). Summarized, we are advancing the understanding of the current state, challenges, risks, and open problems of organizationally applied FedML and thereby taking a step in maturing this line of research.

Third, our studies extend the literature on ML SDLCs. We observed that practitioners have difficulties structuring the FedML development process and confirm the need of ML-specific SDLC [RKW20, ADD21, SBD⁺21] (P1, P3, P4). By following the design science research approach, we constructed the FedML SDLC to provide structure to the development process. With these models, we enhance the understanding of the interplay between the different components that are required in the development process and showcase how the activities, interacting roles, required resources, and resulting artifacts are interrelated. We incorporated best practices of prevailing ML SDLCs [BCN⁺17, ABB⁺19, SBD⁺21, LBM⁺22, KKH23] and revealed how FedML differs and changes known SDLCs. In total, we showed that an end-to-end FedML project consists of seven distinct stages and involves a total of nine roles. Hence, we provide an outlook on the development of decentralized ML systems, while extant literature focuses mostly on traditional, centralized architectures.

6.2. Practical Implications

Adding to the theoretical contributions, this dissertation offers several implications for practice and provides practitioners, decision-makers, and organizations with actionable insights into what

is required for the adoption of FedML. Overall, this thesis aims to offer a guide for organizations that plan to develop or are currently in the process of adopting FedML.

First, our findings reveal the influencing factors that need to be considered by practitioners for the successful practical adoption of FedML in organizations. Through this systemized list, we provide an overview of the various factors influencing the adoption of FedML in organizations from a technological, organizational, and environmental perspective (P1). Through this study, we present structured insights into the complex processes and the interrelation of technology and humans in adopting FedML. These identified factors represent potential pitfalls for the integration of FedML systems in organizations and help practitioners become aware of these aspects. Hence, the insights can be used by organizations that are currently planning or developing FedML projects to identify and counteract challenges preemptively. In turn, these insights also elucidate factors that could be enforced to actively drive the successful adoption of FedML through better-informed decisions. More abstractly, organizations could also systematically assess the provided factors to evaluate their maturity for FedML adoption. Thereby, organizations can validate the viability of the project and identify potential areas of improvement that can be addressed by decision-makers. Overall, our study provides practical guidance on how decision-makers and practitioners can influence the adoption of FedML systems by creating awareness of the technological, organizational, and environmental influencing factors.

Second, our results aid practitioners in the creation and management of collaborations with different participants or organizations within FedML projects (P2). As seen in P1, all interviewed experts unanimously stated that collaboration is one of the main challenges that influence the successful adoption of FedML. For this purpose, we compile a systemized list of socio-technical challenges that represent critical aspects and potential pitfalls for collaborations within the FedML context (P2). By this, we not only support organizations to identify potential collaboration-related challenges early but also allow for an initial viability assessment and to counter potentially emerging challenges preemptively. Since FedML works in a decentralized and potentially collaborative setting, participants need to define a *collaboration agreement* before the development process can be initiated (P3). The extensive set of guiding questions offers practical guidance for the creation of such an agreement and provides information on the necessary aspects that need to be considered. This comprehensive list can be used as a reference to help organizations consider the known relevant challenges of collaborative business models as well as novel FedML-specific aspects (P2).

Third, this thesis illustrates how the FedML SDLC is structured (P3, P4). The results can be used by practitioners and organizations as a blueprint for their implementation procedure or as guidance for the project plan. Currently, practitioners are challenged by the complex technical implementation process as well as a lack of standards, best practices, and guidance (P1). We observed that this lack of guidance leads to complications in the project initiation and project planning phase (P3, P4), which creates a barrier to the productization of FedML systems. Against this backdrop, we developed a high-level process model (P3) and detailed activity models (P4) to initially provide a comprehensible and structured overview of the FedML SDLC. The process model gives a holistic overview and can help practitioners gain an understanding of the overall project structure with the interrelations of the different activities, roles, required resources, and resulting artifacts. In P3, we showed that this structured and simplified illustra-

tion, helps practitioners to communicate the distribution of tasks with potential collaboration partners and thereby eases the project initiation phase. Additionally, the activity models offer more in-depth descriptions of the activity sequences for the end-to-end FedML project. In P4, the evaluation of the activity models with 14 experts (E1-E14) showed a multitude of practical implications. The models aid in "[...] *understanding the FedML process*" (E12) and make the FedML process for "[...] *non-technically proficient shareholder easier to comprehend*" (E11). Consequently, the models "[...] *help make better-informed decisions*" (E9). Moreover, technical practitioners can use the models to facilitate communication with stakeholders since "*this is very much what decision-makers would require from practitioners in order to understand what needs to be done [...]*" (E2). Thus, practitioners and non-technical stakeholders could use the models for guidance in the project planning phase, track the development process, and identify the next steps. Overall, the combination of the process model and activity models provides an initial step toward a standard process for the development of FedML systems. Hereby, the process model mainly intends to aid business stakeholders and solution architects in the project initiation phase (P3), whereas the activity models target product owners and project managers in the planning phase and the development process (P4).

The embedded publications and thereby this dissertation are subject to several limitations. While each publication offers a detailed discussion of its limitations, we will summarize and discuss the most important limitations regarding our research methods, data sources, and scope in the following.

Literature Reviews. All publications used literature reviews to build the theoretical foundation for further research endeavors. Literature reviews are limited by their reliance on the search and coding process. In this regard, our search process might not have covered all papers related to the given topic since we might have missed alternative phrases for our search terms or relevant outlets. We tried to mitigate this limitation by conducting comprehensive forward and backward searches to identify further studies [WW02]. Additionally, we searched for articles in journals and conference proceedings as well as grey literature from the computer science, social sciences, and IS literature corpus. Thereby, we aimed to ensure the inclusion of all subject areas including novel insights that have not been published. Moreover, the analysis of the literature review and coding process is prone to bias as coding depends on the researcher's subjective interpretation of the interrelations. This applies especially to P2, where we contextualized and systemized the challenges of collaborative business models from isolated empirical studies. To address this issue, we introduced pre-defined selection criteria and multiple researchers independently coded the relevant articles. All conflicts were mutually resolved through regular discussions.

Focus Groups. In P2, P3, and P4 we performed focus group discussions which were limited by the data collection and analysis process. Here, the researcher is the primary means for data collection and is therefore susceptible to bias. To minimize this problem, we ensured observer triangulation by involving two participating researchers in each focus group [RH09]. Additionally, we were limited by the available number of experts. Since FedML is an emerging technology, only first and early adopters leverage FedML in their organizations. As a result, we

could only identify a limited amount of suitable experts and project teams. The participants were mainly affiliated with larger enterprises, research institutes, or start-ups. Medium-sized enterprises were only sparsely represented. Therefore, the findings could be biased and might have led to results tailored to the participating affiliations that might not be fully generalizable to medium-sized enterprises. We tried to counteract this bias by carefully selecting highly involved experts with different backgrounds and affiliations from each enterprise size to maximize the variety of voices to the best of our possibilities [MN07]. Similar to literature reviews, the coding process also introduces potential bias and we mitigated this risk through data triangulation and involving multiple researchers that coded the collected data.

Expert Interviews. There are also limitations to semi-structured expert interviews, which we conducted in P1 and P2. Expert interviews introduce the same limitations as focus group discussions. On this note, it is important to emphasize that our studies are based on the experiences and expertise of a limited sample size due to the novelty of FedML. Even though we only concluded our studies when theoretical saturation was reached, we cannot rule out that further interviewees might have added more insights which might have altered our results. Hence, more perspectives and consequently more data from a bigger and more diverse set of interviewees might have enriched our results.

Design Science Research. Also, the use of design science research in P3 and P4 yields some limitations. Here, our studies are mainly limited by researcher bias in the design process and by the evaluation. We tried to mitigate potential researcher bias by establishing design principles and iteratively assessing the artifacts with independent focus groups throughout the design process. In IS literature, there is little guidance on how to evaluate the contributions of DSR artifacts thoroughly, and describing the complexities of an artifact is only partly possible [GH13, CW23]. In this regard, we focused on evaluating the artifacts' technical completeness and usability through demonstrations with subsequent survey-based evaluations within a large industrial lighthouse project as part of a project ideation process. While this is a reasonable evaluation method [PCWA14], we could not assess the performance of the artifacts in further use cases. More rigorous case studies as well as additional iterations and user feedback might be required. Therefore, we seek to evaluate the models in more scenarios and gather feedback for improvement.

Scope. Regarding the scope, we mainly focused on FedML systems with a client-server architecture and orchestrating server since this is the most widely used architectural pattern [LLW⁺21]. However, more architectural patterns such as completely decentralized peer-to-peer processes were also suggested in recent literature [LLZ⁺22]. Therefore, the generalizability of the process model (P3) and activity models (P4) is limited to the scope of this thesis. As discussed in P4, the choice of architectural pattern only requires revisiting the models in the *Local Model Training* and *Global Model Aggregation* stage, whereas the remainder of the models can be reused. Additionally, we aimed at generically usable models that are not specific to use cases or domains. To evaluate the generalizability, we gathered a diverse group of evaluation participants from multiple areas, domains, enterprise sizes, and job profiles. However, we conceivably only depict a subset of the potential user demographic and can not discount potentially missing domain-dependent and process-specific characteristics. Nevertheless, the models are a reasonable basis for extension and further development.

Conclusion and Future Work

With this chapter, we conclude this dissertation and outline future work. More specifically, we will first recapitulate the research endeavor with concluding remarks (see Section 8.1) and finally describe potential avenues for future research that emerged from the findings of this thesis (see Section 8.2).

8.1. Conclusion

The lack of available and suitable training data is a persisting barrier to the implementation of ML systems [PFW⁺21, BvDK20, HAAY⁺23]. Organizations could overcome this barrier by voluntary data sharing and collaboration. However, the companies' willingness to share data is low due to privacy concerns and potential loss of IP [SLZ20]. FedML is an emerging ML paradigm with the promise to enable the training of a joint ML model while keeping the contributor's data private [MMR⁺17]. Through its model-to-data approach, the training data remains within company borders and thereby permits cross-company collaboration across multiple data silos without having to share data. However, despite its potential, most FedML projects come to a halt in the simulation or prototype stage and consequently fail to actualize [LLW⁺21]. Although there is an increasing literature corpus on FedML [AAdCK⁺22], there is still no work investigating its organizational adoption. Against this backdrop, the purpose of this dissertation was to close this research gap and contribute to the existing body of knowledge by investigating three emergent research questions regarding the organizational adoption of FedML: the influencing factors for a successful organizational adoption (RQ1), its distinct socio-technical challenges related to the unique characteristics of collaborative FedML (RQ2), and the FedML SDLC (RQ3).

Conclusion related to RQ1. This dissertation initially reveals the influencing factors for the organizational adoption of FedML. We compiled a list of three main reasons for FedML

adoption, six main risks, and 19 influencing factors. Thereby, we extend the literature corpus on ML adoption with a view on FedML systems and enhance the understanding of how the decentralized nature of FedML adds complexity and changes the impacting factors of ML adoption. Researchers can use our findings as a basis for further investigation and to determine new best practices for tackling identified novel challenges. Moreover, practitioners can utilize the insights to become aware and avoid potential obstacles. Additionally, organizations could systematically analyze the proposed list of influencing factors to determine their maturity regarding FedML adoption and identify possible areas of improvement.

Conclusion related to RQ2. The findings of this dissertation provide a systemized overview of the socio-technical challenges that emerge from the unique collaborative characteristics of FedML. Through a systematic literature review, focus group discussions, and subsequent expert interviews, we compiled a list of novel socio-technical challenges and a set of corresponding guiding questions. Thereby, we offer a better understanding of how collaboration affects the interrelation of humans and technology within the ML context. Practitioners can use the obtained results as a reference book for critical aspects and potential pitfalls. Thereby practitioners can preemptively mitigate arising collaboration-related challenges. Additionally, the systemized list can be leveraged as practical guidance for the creation of a *collaboration agreement* that needs to be defined within FedML projects.

Conclusion related to RQ3. Through this dissertation, we illustrated the SDLC of FedML projects through an abstracted process model and detailed activity models. We recognized that practitioners are challenged by the complex implementation process and that standards or guidance can facilitate this challenge. In this light, we designed a set of artifacts that structures, simplifies, and provides guidance for the implementation process through a holistic overview of the overall project flow and in-depth representations of the required activity sequences. Thereby, we complemented the literature corpus on ML SDLCs with a perspective on FedML. Practitioners can utilize our results to gain an understanding of the project structure and interrelations within the FedML projects, which eases the communication of technological-oriented practitioners with non-technically proficient stakeholders. Moreover, practitioners can use our results as guidance to plan the project, track progress, and identify the next steps. Overall, these models offer a first step towards a standardized FedML development process and help make better-informed decisions for project initiation and planning.

Overall, this dissertation takes a step toward an exploration of the adoption of FedML in organizations. We identified the distinct socio-technical challenges of collaborative FedML, proposed FedML SDLCs, and revealed the influencing factors for its organizational adoption. Amongst other things, practitioners can utilize our findings as guidance for their projects. Additionally, organizations can use our results as a reference book for the creation of a *collaboration agreement* and to avoid potential pitfalls during the project. Thereby, we also provide the fundamental basis for future research endeavors within the field of organizational FedML adoption.

8.2. Future Work

In the context of this work, several research gaps emerged that could not be addressed in the scope of this thesis and provide promising avenues for future research. The following will outline potential starting points for new research efforts and is summarized in Table 8.1.

Future research regarding RQ1. To explore the influencing factors for the adoption of FedML, we followed a qualitative research approach across diverse sample demography. The exploratory and qualitative approach of our study yields multiple promising possibilities for future research. First, since we conducted a qualitative study with a limited sample size, it might be interesting to perform follow-up studies to validate, extend, and quantify our results. Second, we only revealed the influencing factors without putting the different factors in relation to each other. Therefore, future research could empirically analyze the significance and correlation of the identified adoption factors and determine correlation coefficients or weightings for the attributes. Third, our work involves various industries and is by nature industry-agnostic. However, as seen in Cubric et al. [Cub20], varying industries may introduce differing value drivers or barriers to ML adoption. Thus, future research can build on our results and focus on distinct industries to contrast the similarities and differences of FedML adoption across domains.

Additionally, the identified adoption factors offer the theoretical basis for a variety of potential research opportunities. For example, one of the main challenges revolves around data quality and how interoperability across participants can be assured. Against this backdrop, we encourage researchers to study data quality, interoperability features and how corresponding assessments can be conducted within the FedML context. These studies could additionally develop tools to aid practitioners in the assessment. Also, practitioners are missing legal clarity (P1) and there is still ambiguity regarding currently emerging legal frameworks such as the *Artificial Intelligence (AI) Act* of the European Union or the *United States Algorithmic Accountability Act* [VZB21, MJWF22] and how antitrust laws respond to FedML collaborations [MLP21]. As an initial step, it might be interesting for future research to assess the current legal situation and corollary open questions of FedML systems regarding AI regulations and antitrust laws. We expect these results to be crucial for the broad adoption of FedML systems.

Future research regarding RQ2. The current research in the field of FedML is dominated by work on challenges and problem statements with a technical focus [NDR20, ZXB⁺21, BAA⁺22]. However, the organizational adoption of novel technologies, such as FedML, is dependent on more than the technical dimension [MSH91, GDR14]. With our work on the socio-technical aspects of collaborative FedML, we took an initial step beyond technological aspects. As this thesis is predominantly exploratory, we encourage researchers to perform follow-up studies to quantitatively validate, refine, and extend our results. While we identified a multitude of challenges, current literature still lacks best practices to address these challenges. Therefore, we motivate researchers to perform observational studies, identify success stories, and derive best practices by investigating how current first and early adopters solve these challenges. In this regard, one of the most cited challenges involved the distribution of model ownership and financials. Correspondingly, investigating how production-level FedML systems deliver value and understanding the accompanying value streams between organizations is a promising direction for future research. Moreover, there are persisting issues regarding the collaboration structure,

decision-making, and IP management. Thus, future work could analyze the underlying governance framework of current FedML systems and rigorously conceptualize possible governance frameworks, methods, and tools.

Future research regarding RQ3. In this thesis, we designed a process model and activity models to structure, simplify, and provide guidance for the implementation of FedML projects. The models proved to be useful in our evaluations, however, we did not validate the models in industrial projects due to the limited scope of this thesis. Therefore, we encourage researchers and practitioners to apply the artifacts in case studies to further validate the model’s usefulness and gather feedback for improvement. In our work, we focused on implementation projects with a server-client architecture and a central orchestrating server. While this is the most prevalent and utilized architectural pattern [LLW⁺21], other architectural patterns have been proposed as well [LLZ⁺22]. Thus, future research could empirically evaluate the generalizability of our proposed models and design models that are tailored to other architectural patterns, such as partially connected architectures or completely decentralized peer-to-peer networks [BPS⁺23]. Since privacy is one main reason for adopting FedML, there is increasing research on the additional implementation of PETs in FedML algorithms [WLD⁺20, ZLX⁺20, MPP⁺21, EOA22]. Similar to the previous point, future work could also contribute to this research stream with a view on the engineering life cycle and investigate how the additional utilization of PETs alters the FedML development process.

Through our models, we also recognized various opportunities to develop helpful tools and aid practitioners through the development process. For example, our set of guiding questions (P2) facilitates the creation of the *collaboration agreement* and in an additional publication, we designed a design support tool for the technology selection process (see P7) [ZMM24]. Future studies could complement this tool set to help practitioners through the development of FedML projects. This research avenue can span from pattern catalogs that focus on stakeholder-related concerns and patterns to technological frameworks or tools that automate incentive mechanisms or establish governance.

RQ	No.	Potential Research Avenues
RQ1	P1	<ul style="list-style-type: none"> • Perform quantitative studies to validate, refine, and extend our results • Empirically study the significance and correlation of the adoption factors • Investigate similarities and differences of adoption factors across domains • Study data quality and interoperability features within the FedML context and how corresponding assessments could be done • Assess current legal situation and corollary open questions of FedML systems regarding AI regulations and antitrust laws
RQ2	P2	<ul style="list-style-type: none"> • Perform quantitative studies to validate, refine, and extend our results • Conduct case studies to identify success stories and derive best practices that address the socio-technical challenges • Investigate how production-level FedML systems deliver value and understand the value streams between participating organizations • Analyze underlying governance mechanisms of production-level FedML systems and conceptualize governance frameworks, methods, and tools
RQ3	P3	<ul style="list-style-type: none"> • Apply models in case studies to further validate the usefulness of the model's gather feedback for improvement • Empirically evaluate the generalizability of our model and investigate how other architectural patterns influence the development life cycle
	P4	<ul style="list-style-type: none"> • Develop process and activity models for other architectural patterns • Analyze how the utilization of PETs alters the development process • Rigorously develop tools that help practitioners through project stages

Table 8.1. Potential avenues for future research.

Bibliography

- [AA05] Garry Anderson and Nancy Arsenault. *Fundamentals of educational research*. Routledge, 2 edition, 2005.
- [AAAdCK⁺22] Rodolfo S. Antunes, Cristiano André da Costa, Arne Küderle, Imrana A. Yari, and Björn Eskofier. Federated learning for healthcare: Systematic review and architecture proposal. *ACM Transactions on Intelligent Systems and Technology*, 13(4):1–23, 2022.
- [AB15] Adel Alshamrani and Abdullah Bahattab. A comparison between three sdlc models waterfall model, spiral model, and incremental/iterative model. *International Journal of Computer Science Issues*, 12(1):106–111, 2015.
- [ABB⁺19] Saleema Amershi, Andrew Begel, Christian Bird, Robert DeLine, Harald Gall, Ece Kamar, Nachiappan Nagappan, Besmira Nushi, and Thomas Zimmermann. Software engineering for machine learning: A case study. In *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice*, pages 291–300, Montreal, QC, Canada, 2019. IEEE/ACM.
- [ACM19] Sulaiman Abdallah Alsheibani, Yen Cheung, and Chris Messom. Factors inhibiting the adoption of artificial intelligence at organizational-level: A preliminary investigation. In *Proceedings of the 25th Americas Conference on Information Systems*, volume 2, pages 1–10, Cancun, Mexico, 2019.
- [ACMA20] Sulaiman Abdallah Alsheibani, Yen Cheung, Chris Messom, and Mazoon Alhosni. Winning AI strategy: Six-steps to create value from artificial intelligence. In *Proceedings of the 26th Americas Conference on Information Systems*, volume 1, pages 1–10, Virtual, 2020.
- [Ada15] William C. Adams. Conducting semi-structured interviews. In Kathryn E. Newcomer, Harry P. Hatry, and Joseph S. Wholey, editors, *Handbook of practical program evaluation*, pages 492–505. Wiley Online Library, 4 edition, 2015.
- [ADD21] Jumana Almahmoud, Robert DeLine, and Steven M. Drucker. How teams communicate about the quality of ML models: A case study at an international technol-

- ogy company. *Proceedings of the ACM on Human-Computer Interaction*, 5:1–24, 2021.
- [Ajz85] Icek Ajzen. From intentions to actions: A theory of planned behavior. In *Action control: From cognition to behavior*, pages 11–39. Springer, Berlin, Heidelberg, 1985.
- [Anw14] Ashraf Anwar. A review of rup (rational unified process). *International Journal of Software Engineering*, 5(2):12–19, 2014.
- [AOO⁺20] Jide E. T. Akinsola, Afolakemi S. Ogunbanwo, Olatunji J. Okesola, Isaac J. Odun-Ayo, Florence D. Ayegbusi, and Ayodele A. Adebisi. Comparative analysis of software development life cycle models (SDLC). In *Intelligent Algorithms in Software Engineering*, pages 310–322, Cham, 2020. Springer International Publishing.
- [ÅRP22] Josef Åström, Wiebke Reim, and Vinit Parida. Value creation and value capture for AI business model innovation: A three-phase process framework. *Review of Managerial Science*, 16(7):2111–2133, 2022.
- [ASM⁺19] Moayad Alshangiti, Hitesh Sapkota, Pradeep K. Murukannaiah, Xumin Liu, and Qi Yu. Why is developing machine learning applications challenging? A study on stack overflow posts. In *Proceedings of the 13th International Symposium on Empirical Software Engineering and Measurement*, pages 1–11, Porto de Galinhas, Recife, Brazil, 2019.
- [ASSA20] Samar Al-Saqqqa, Samer Sawalha, and Hiba AbdelNabi. Agile software development: Methodologies and trends. *International Journal of Interactive Mobile Technologies*, 14(11):246–270, 2020.
- [BAA⁺22] Syreen Banabilah, Moayad Aloqaily, Eitaa Alsayed, Nida Malik, and Yaser Jararweh. Federated learning review: Fundamentals, enabling technologies, and future applications. *Information processing & management*, 59(6):1–24, 2022.
- [Bak00] Michael J Baker. Writing a literature review. *The marketing review*, 1(2):219–247, 2000.
- [Bak12] Jeff Baker. The technology–organization–environment framework. *Information Systems Theory: Explaining and Predicting Our Digital Society*, 1:231–245, 2012.
- [BBVB⁺01] Kent Beck, Mike Beedle, Arie Van Bennekum, Alistair Cockburn, Ward Cunningham, Martin Fowler, James Grenning, Jim Highsmith, Andrew Hunt, Ron Jeffries, Jon Kern, Brian Marick, Robert C. Martin, Steve Mellor, Ken Schwaber, Jeff Sutherland, and Dave Thomas. Manifesto for agile software development, 2001.
- [BCHT18] Virginia Braun, Victoria Clarke, Nikki Hayfield, and Garreth Terry. Thematic analysis. In Pranee Liamputtong, editor, *Handbook of Research Methods in Health Social Sciences*, pages 1–18. Springer Singapore, 2018.

- [BCN⁺17] Eric Breck, Shanqing Cai, Eric Nielsen, Michael Salib, and D. Sculley. The ML test score: A rubric for ML production readiness and technical debt reduction. In *Proceedings of the 2017 IEEE International Conference on Big Data*, pages 1123–1132, Boston, MA, USA, 2017. IEEE.
- [BDT06] Oddur Benediktsson, Darren Dalcher, and Helgi Thorbergsson. Comparison of software development life cycles: A multiproject experiment. *IEE Proceedings-Software*, 153(3):87–101, 2006.
- [BEG⁺19] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, Timon Von Overveldt, David Petrou, Daniel Ramage, and Jason Roselander. Towards federated learning at scale: System design. In *Proceedings of the 1st Machine Learning and Systems*, pages 374–388, Stanford, California, USA, 2019.
- [Bel12] France Belanger. Theorizing in information systems research using focus groups. *Australasian Journal of Information Systems*, 17(2):109–135, 2012.
- [Bha12] Anol Bhattacharjee. *Social science research: Principles, methods, and practices*. Create Space, Scotts Valley, California, USA, 2 edition, 2012.
- [BKB20] Bouziane Brik, Adlen Ksentini, and Maha Bouaziz. Federated learning for UAVs-enabled wireless networks: Use cases, challenges, and open problems. *IEEE Access*, 8:53841–53849, 2020.
- [BKKP22] Kuldeep Bhalerao, Anuj Kumar, Arya Kumar, and Purvi Pujari. A study of barriers and benefits of artificial intelligence adoption in small and medium enterprise. *Academy of Marketing Studies Journal*, 26:1–6, 2022.
- [BLS⁺21] Aline F. Borges, Fernando J. Laurindo, Mauro M. Spínola, Rodrigo F. Gonçalves, and Claudia A. Mattos. The strategic use of artificial intelligence in the digital era: Systematic literature review and future research directions. *International Journal of Information Management*, 57:1–16, 2021.
- [BM12] Sundramoorthy Balaji and Sundararajan M. Murugaiyan. Waterfall vs. V-model vs. agile: A comparative study on sdlc. *International Journal of Information Technology and Business Management*, 2(1):26–30, 2012.
- [Boe79] Barry W. Boehm. Guidelines for verifying and validating software requirements and design specifications. In *Proceedings of the European Conference on Applied Information Technology of the International Federation for Information Processing*, pages 711–719, London, England, 1979.
- [BPLW21] Thommie Burström, Vinit Parida, Tom Lahti, and Joakim Wincent. AI-enabled business-model innovation and transformation in industrial ecosystems: A framework, model and outline for further research. *Journal of Business Research*, 127:85–95, 2021.

- [BPS⁺23] Enrique T. Beltrán, Mario Q. Pérez, Pedro M. Sánchez, Sergio L. Bernal, G r me Bovet, Manuel G. P rez, Gregorio M. P rez, and Alberto H. Celdr n. Decentralized federated learning: Fundamentals, state-of-the-art, frameworks, trends, and challenges. *IEEE Communications Surveys & Challenges*, 1:1–31, 2023.
- [BvDK20] Markus Bauer, Clemens van Dinther, and Daniel Kiefer. Machine learning in SME: an empirical study on enablers and success factors. In *Proceedings of the 26th Americas Conference on Information Systems*, volume 3, pages 1–10, Virtual, 2020.
- [BvHS05] Jenine Beekhuyzen, Liisa von von Hellens, and Mark Siedle. Cultural barriers in the adoption of emerging technologies. In *Proceedings of the 11th International Conference on Human-Computer Interaction*, volume 10, Las Vegas, Nevada, USA, 2005. Citeseer.
- [BWG20] Loraine Busetto, Wolfgang Wick, and Christoph Gumbinger. How to use and assess qualitative research methods. *Neurological Research and Practice*, 2(14):1–10, 2020.
- [CCK⁺00] Pete Chapman, Julian Clinton, Randy Kerber, Thomas Khabaza, Thomas Reinartz, Colin Shearer, R diger Wirth, et al. CRISP-DM 1.0: Step-by-step data mining guide. *SPSS Inc*, 9(13):1–73, 2000.
- [CCM15] Angela Carid , Maria Colurcio, and Monia Melia. Designing a collaborative business model for SMEs. *Sinergie Italian Journal of Management*, 33:233–253, 2015.
- [CHH88] Paul D. Collins, Jerald Hage, and Frank M. Hull. Organizational and technological predictors of change in automaticity. *Academy of Management Journal*, 31(3):512–543, 1988.
- [CHTB20] Crispin Coombs, Donald Hislop, Stanimira K Taneva, and Sarah Barnard. The strategic impacts of intelligent automation for knowledge and service work: An interdisciplinary review. *The Journal of Strategic Information Systems*, 29(4):101600, 2020.
- [CKO92] Bill Curtis, Marc I. Kellner, and Jim Over. Process modeling. *Communications of the ACM*, 35(9):75–90, 1992.
- [Cra21] Nicholas Crafts. Artificial intelligence as a general-purpose technology: An historical perspective. *Oxford Review of Economic Policy*, 37(3):521–536, 2021.
- [CRDB21] Sheshadri Chatterjee, Nripendra P. Rana, Yogesh K. Dwivedi, and Abdullah M. Baabdullah. Understanding AI adoption in manufacturing and production firms using an integrated TAM-TOE model. *Technological Forecasting and Social Change*, 170:120880, 2021.
- [Cre09] John W. Creswell. *Research design: Qualitative, quantitative, and mixed methods approaches*. Sage Publications, Inc., 3 edition, 2009.

-
- [Cub20] Marija Cubric. Drivers, barriers and social considerations for AI adoption in business and management: A tertiary study. *Technology in Society*, 62:101257, 2020.
- [CW23] Marcel Cahenzli and Robert Winter. Writing DSR articles for maximum impact. In *Proceedings of the 31st Conference on Information Systems. Research Papers*, volume 407, pages 1–15, Kristiansand, Norway, 2023.
- [Dav85] Fred D. Davis. *A technology acceptance model for empirically testing new end-user information systems: Theory and results*. PhD thesis, Massachusetts Institute of Technology, 1985.
- [Dav00] Philip Davies. The relevance of systematic reviews to educational policy and practice. *Oxford Review of Education*, 26(3-4):365–378, 2000.
- [DC18] Bruna Diirr and Claudia Cappelli. A systematic literature review to understand cross-organizational relationship management and collaboration. In *Proceedings of the 51st Hawaii International Conference on System Sciences*, pages 145–154, Honolulu, Hawaii, USA, 2018.
- [Den17] Martyn Denscombe. *The good research guide: For small-scale social research projects*. McGraw-Hill Education (UK), 2017.
- [DEZ20] Thando Dube, Rene Van Eck, and Tranos Zuva. Review of technology adoption models and theories to measure readiness and acceptable use of technology in a business organization. *Journal of Information Technology and Digital World*, 2(4):207–212, 2020.
- [DNBM12] Torgeir Dingsøy, Sridhar Nerur, VenuGopal Balijepally, and Nils B. Moe. A decade of agile methodologies: Towards explaining agile software development. *Journal of systems and software*, 85(6):1213–1221, 2012.
- [DVPHE20] Assunta Di Vaio, Rosa Palladino, Rohail Hassan, and Octavio Escobar. Artificial intelligence and business models in the sustainable development goals perspective: A systematic literature review. *Journal of Business Research*, 121:283–314, 2020.
- [Edm99] Holly Edmunds. The focus group research handbook. *The Bottom Line*, 12(3):46–46, 1999.
- [EOA22] Ahmed El Ouadrhiri and Ahmed Abdelhadi. Differential privacy for deep and federated learning: A survey. *IEEE access*, 10:22359–22380, 2022.
- [EPMK22] Ida Merete Enholm, Emmanouil Papagiannidis, Patrick Mikalef, and John Krogstie. Artificial intelligence and business value: A literature review. *Information Systems Frontiers*, 24(5):1709–1734, 2022.
- [FF00] Andrea Fontana and James H. Frey. The interview: From structured questions to negotiated text. In Normal K. Denzin and Yvonna S. Lincoln, editors, *Handbook of Qualitative Research*, pages 645–670. Thousand Oaks, CA: Sage, 01 2000.

- [For23] World Economic Forum. Competition vs collaboration: rethinking how businesses innovate and grow, 2023.
- [FS19] Jason Furman and Robert Seamans. AI and the economy. *Innovation policy and the economy*, 19(1):161–191, 2019.
- [GDR14] Hemlata Gangwar, Hema Date, and A. D. Raoot. Review on IT adoption: Insights from recent technologies. *Journal of enterprise information management*, 27(4):488–502, 2014.
- [GG17] Swadha Gupta and Deepali Gouttam. Towards changing the paradigm of software development in software industries: An emergence of agile software development. In *Proceedings of the 2017 IEEE International Conference on Smart Technologies and Management for Computing, Communication, Controls, Energy and Materials*, pages 18–21, Chennai, India, 2017. IEEE.
- [GH13] Shirley Gregor and Alan R. Hevner. Positioning and presenting design science research for maximum impact. *Management of Information Systems quarterly*, 37(2):337–355, 2013.
- [Gil06] Richard J. Gilbert. Competition and innovation. *Journal of Industrial Organization Education*, 1(1):1–23, 2006.
- [Gir21] Görkem Giray. A software engineering perspective on engineering machine learning systems: State of the art and challenges. *Journal of Systems and Software*, 180:111031, 2021.
- [GJA06] Bart N. Green, Claire D. Johnson, and Alan Adams. Writing narrative literature reviews for peer-reviewed journals: Secrets of the trade. *Journal of chiropractic medicine*, 5(3):101–117, 2006.
- [GJK87] Richard A. Guzzo, Susan E. Jackson, and Raymond A. Katzell. Meta-analysis analysis. *Research in organizational behavior*, 9(1):407–442, 1987.
- [GL19] Massimo Garbuio and Nidhida Lin. Artificial intelligence as a growth engine for health care startups: Emerging business models. *California Management Review*, 61(2):59–83, 2019.
- [Gro21] IEEE Shared Machine Learning Working Group. IEEE guide for architectural framework and application of federated machine learning. *IEEE Std 3652.1-2020*, pages 1–69, 2021.
- [GSTC08] Paul Gill, Kate Stewart, Elizabeth Treasure, and Barbara Chadwick. Methods of data collection in qualitative research: Interviews and focus groups. *British dental journal*, 204(6):291–295, 2008.
- [GTT23] Avi Goldfarb, Bledi Taska, and Florenta Teodoridis. Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings. *Research Policy*, 52(1):104653, 2023.

-
- [HA13] Amani Hamed and Hisham Abushama. Popular agile approaches in software development: Review and analysis. In *Proceedings of the International Conference on Computing, Electrical and Electronic Engineering*, pages 160–166, Khartoum, Sudan, 2013.
- [HAAY⁺23] Omar M. Horani, Ahmad S. Al-Adwan, Husam Yaseen, Hazar Hmoud, Waleed M. Al-Rahmi, and Ali Alkhalifah. The critical determinants impacting artificial intelligence adoption at the organizational level. *Information Development*, pages 1–25, 2023.
- [Hag80] Jerald Hage. Theories of organizations: Form, process, and transformation. *John Wiley & Sons*, 1980.
- [Hev07] Alan R. Hevner. A three cycle view of design science research. *Scandinavian Journal of Information Systems*, 19(2):87–92, 2007.
- [HKD06] Pei-Fang Hsu, Kenneth L. Kraemer, and Debora Dunkle. Determinants of e-business use in US firms. *International Journal of Electronic Commerce*, 10(4):9–45, 2006.
- [HMPR04] Alan R. Hevner, Salvatore T. March, Jinsoo Park, and Sudha Ram. Design science in information systems research. *Management Information Systems Quarterly*, 28(1):75–105, 2004.
- [HN21] Mia Hoffmann and Laura Nurski. What is holding back artificial intelligence adoption in Europe? *Bruegel Policy Contribution*, 24, 2021.
- [HSM⁺19] Marc Hesenius, Nils Schwenzfeier, Ole Meyer, Wilhelm Koop, and Volker Gruhn. Towards a software engineering process for developing data-driven applications. In *Proceedings of the 7th International Workshop on Realizing Artificial Intelligence Synergies in Software Engineering*, pages 35–41, Montreal, QC, Canada, 2019. IEEE/ACM.
- [IL20] Marco Iansiti and Karim R. Lakhani. *Competing in the age of AI: Strategy and leadership when algorithms and networks run the world*. Harvard Business Press, 2020.
- [IY19] Fuyuki Ishikawa and Nobukazu Yoshioka. How do engineers perceive difficulties in engineering of machine-learning systems?-questionnaire survey. In *Proceedings of the 7th International Workshop on Conducting Empirical Studies in Industry (CESI) and 6th International Workshop on Software Engineering Research and Industrial Practice (SER&IP)*, pages 2–9, Montreal, QC, Canada, 2019. IEEE.
- [Jal12] Pankaj Jalote. *An integrated approach to software engineering*. Springer Science & Business Media, 2 edition, 2012.
- [JSD22] Divya Jatain, Vikram Singh, and Naveen Dahiya. A contemplative perspective on federated machine learning: Taxonomy, threats & vulnerability assessment and challenges. *Journal of King Saud University - Computer and Information Sciences*, 34(9):6681–6698, 2022.

- [JYK⁺23] Philip Jorzik, Anil Yigit, Dominik K. Kanbach, Sascha Kraus, and Marina Dabić. Artificial intelligence-enabled business model innovation: Competencies and roles of top management. *IEEE Transactions on Engineering Management*, pages 1–13, 2023.
- [KB18] Steinar Kvale and Svend Brinkmann. *Doing interviews*. Sage Publications Ltd, 2018.
- [KBM⁺20] Diana Koshtura, Myroslava Bublyk, Yurii Matseliukh, Dmytro Dosyn, Liliya Chyrun, Olga Lozynska, Ihor Karpov, Ivan Peleshchak, Mariya Maslak, and Oleg Sachenko. Analysis of the demand for bicycle use in a smart city based on machine learning. In *Proceedings of the 2nd International Workshop on Modern Machine Learning Technologies and Data Science*, pages 172–183, Lviv-Shatsk, Ukraine, 2020.
- [KH05] William King and Jun He. Understanding the role and methods of meta-analysis in IS research. *Communications of the Association for Information Systems*, 16:665–686, 2005.
- [KHC⁺16] Jinkyu Kim, Heonseok Ha, Byung-Gon Chun, Sungroh Yoon, and Sang K. Cha. Collaborative analytics for data silos. In *Proceedings of the 32nd IEEE International Conference on Data Engineering*, pages 743–754, Helsinki, Finland, 2016. IEEE.
- [KKG21] Sudatta Kar, Arpan Kumar Kar, and Manmohan P. Gupta. Modeling drivers and barriers of artificial intelligence adoption: Insights from a strategic management perspective. *Intelligent Systems in Accounting, Finance and Management*, 28(4):217–238, 2021.
- [KKH23] Dominik Kreuzberger, Niklas Kühn, and Sebastian Hirschl. Machine learning operations (MLOps): Overview, definition, and architecture. *IEEE Access*, 11:31866–31879, 2023.
- [KLB04] Jyrki Kontio, Laura Lehtola, and Johanna Bragge. Using the focus group method in software engineering: Obtaining practitioner and user experiences. In *Proceedings of the 2004 International Symposium on Empirical Software Engineering*, pages 271–280, Redondo Beach, CA, USA, 2004. IEEE.
- [KMA⁺21] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2):1–210, 2021.
- [Kne17] Ralf Kneuper. Sixty years of software development life cycle models. *IEEE Annals of the History of Computing*, 39(3):41–54, 2017.
- [KSH⁺21] Latif U. Khan, Walid Saad, Zhu Han, Ekram Hossain, and Choong Seon Hong. Federated learning for internet of things: Recent advances, taxonomy, and open challenges. *IEEE Communications Surveys & Tutorials*, 23(3):1759–1799, 2021.

-
- [KTB11] Stefan Koch, Aysegul Toker, and Philip Brulez. Extending the technology acceptance model with perceived community characteristics. *Information Research*, 16(2):16–2, 2011.
- [KVSO13] Saeed Khanagha, Henk Volberda, Jatinder Sidhu, and Ilan Oshri. Management innovation and adoption of emerging technologies: The case of cloud computing. *European Management Review*, 10(1):51–67, 2013.
- [KXN⁺19] Jiawen Kang, Zehui Xiong, Dusit Niyato, Shengli Xie, and Junshan Zhang. Incentive mechanism for reliable federated learning: A joint optimization approach to combining reputation and contract theory. *IEEE Internet of Things Journal*, 6(6):10700–10714, 2019.
- [KZDB16] Miryung Kim, Thomas Zimmermann, Robert DeLine, and Andrew Begel. The emerging role of data scientists on software development teams. In *Proceedings of the 38th International Conference on Software Engineering*, pages 96–107, Austin, TX, USA, 2016.
- [LB03] Craig Larman and Victor R Basili. Iterative and incremental developments. A brief history. *Computer*, 36(6):47–56, 2003.
- [LBM⁺22] Samuli Laato, Teemu Birkstedt, Matti Mäantymäki, Matti Minkkinen, and Tommi Mikkonen. AI governance in the system development life cycle: Insights on responsible machine learning engineering. In *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*, pages 113–123, Pittsburgh, PA, USA, 2022.
- [LDCH22] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated learning on non-iid data silos: An experimental study. In *Proceedings of the 38th IEEE International Conference on Data Engineering*, pages 965–978, Kuala Lumpur, Malaysia, 2022. IEEE.
- [LE06] Yair Levy and Timothy J Ellis. A systems approach to conduct an effective literature review in support of information systems research. *The International Journal of an Emerging Transdiscipline*, 9:181–212, 2006.
- [LERH20] Hanyan Liu, Samuel Eksmo, Johan Risberg, and Regina Hebig. Emerging and changing tasks in the development process for machine learning systems. In *Proceedings of the 14th International Conference on Software and System Processes*, pages 125–134, Seoul, Republic of Korea, 2020.
- [LLTT12] Yu B. Leau, Wooi K. Loo, Wai Y. Tham, and Soo F. Tan. Software development life cycle: Agile vs. traditional approaches. In *Proceedings of the 2nd International Conference on Information and Network Technology*, volume 37, pages 162–167, Chennai, India, 2012.
- [LLW⁺21] Sin K. Lo, Qinghua Lu, Chen Wang, Hye-Young Paik, and Liming Zhu. A systematic literature review on federated machine learning: From a software engineering perspective. *ACM Computing Surveys*, 54(5):1–39, 2021.

- [LLZ⁺22] Sin K. Lo, Qinghua Lu, Liming Zhu, Hye-Young Paik, Xiwei Xu, and Chen Wang. Architectural patterns for the design of federated learning systems. *Journal of Systems and Software*, 191:111357, 2022.
- [LMJ08] Zhenhua Liu, Qingfei Min, and Shaobo Ji. A comprehensive review of research in IT adoption. In *Proceedings of the 4th International Conference on Wireless Communications, Networking and Mobile Computing*, pages 1–5, Dalian, China, 2008. IEEE.
- [Lou18] António Loureiro. There is a fourth industrial revolution: The digital revolution. *Worldwide Hospitality and Tourism Themes*, 10(6):740–744, 2018.
- [LRB⁺19] Lucy E. Lwakatare, Aiswarya Raj, Jan Bosch, Helena H. Olsson, and Ivica Crnkovic. A taxonomy of software engineering challenges for machine learning systems: An empirical investigation. In *Proceedings of the 20th International Conference on Agile Software Development*, volume 355, pages 227–243, Montreal, QC, Canada, 2019. Springer International Publishing.
- [LRC⁺20] Lucy E. Lwakatare, Aiswarya Raj, Ivica Crnkovic, Jan Bosch, and Helena H. Olsson. Large-scale machine learning systems in real-world industrial settings: A review of challenges and solutions. *Information and Software Technology*, 127:106368, 2020.
- [LSRB19] Jaehun Lee, Taewon Suh, Daniel Roy, and Melissa Baucus. Emerging technology and business model innovation: The case of artificial intelligence. *Journal of Open Innovation: Technology, Market, and Complexity*, 5(3):44, 2019.
- [LWW⁺23] Qinbin Li, Zeyi Wen, Zhaomin Wu, Sixu Hu, Naibo Wang, Yuan Li, Xu Liu, and Bingsheng He. A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering*, 35(4):3347–3366, 2023.
- [LZZ⁺22] Junyang Li, Chaobo Zhang, Yang Zhao, Weikang Qiu, Qi Chen, and Xuejun Zhang. Federated learning-based short-term building energy consumption prediction method for solving the data silos problem. In *Building Simulation*, volume 15, pages 1145–1159. Springer, 2022.
- [Mak17] Spyros Makridakis. The forthcoming artificial intelligence (AI) revolution: Its impact on society and firms. *Futures*, 90:46–60, 2017.
- [MG10] Nabil M. Munassar and Aliseri Govardhan. A comparison between five models of software engineering. *International Journal of Computer Science Issues*, 7(5):94–101, 2010.
- [MJ17] Sebastian Molinillo and Arnold Japutra. Organizational adoption of digital information and technology: a theoretical review. *The Bottom Line*, 30(01):33–46, 2017.

-
- [MJWF22] Jakob Mökander, Prathm Juneja, David S. Watson, and Luciano Floridi. The US Algorithmic Accountability act of 2022 vs. the EU Artificial Intelligence Act: What can they learn from each other? *Minds and Machines*, 32(4):751–758, 2022.
- [MLP21] Robert Mahari, Sandro C. Lera, and Alex Pentland. Time for a new antitrust era: Refocusing antitrust law to invigorate competition in the 21st century. *Stanford Computational Antitrust*, 1, 2021.
- [MMR⁺17] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera Y. Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54, pages 1273–1282, Lauderdale, Florida, USA, 2017.
- [MN07] Michael D. Myers and Michael Newman. The qualitative interview in IS research: Examining the craft. *Information and organization*, 17(1):2–26, 2007.
- [MNB12] Rupa Mahanti, Madhumita S. Neogi, and Vandana Bhattacharjee. Factors affecting the choice of software life cycle models in the software industry - An empirical study. *Journal of Computer Science*, 8(8):1253–1262, 2012.
- [MNH21] Marco Marabelli, Sue Newell, and Valerie Handunge. The lifecycle of algorithmic decision-making systems: Organizational choices and ethical challenges. *The Journal of Strategic Information Systems*, 30(3):101683, 2021.
- [MOJ05] Kjetil Molokken-Ostfold and Magne Jorgensen. A comparison of software project overruns-flexible versus sequential development models. *IEEE Transactions on Software Engineering*, 31(9):754–766, 2005.
- [Mor88] David L. Morgan. *Focus groups as qualitative research*. Sage Publications, Inc., 1988.
- [MPP⁺21] Virraaji Mothukuri, Reza M. Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantanha, and Gautam Srivastava. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115:619–640, 2021.
- [MSH91] Philip H. Mirvis, Amy L. Sales, and Edward J. Hackett. The implementation and adoption of new technology in organizations: The impact on work, people, and culture. *Human Resource Management*, 30(1):113–139, 1991.
- [MT21] Shrutika Mishra and Asha Ram Tripathi. AI business model: An integrative business approach. *Journal of Innovation and Entrepreneurship*, 10(1):1–21, 2021.
- [MWM⁺05] Natasha Mack, Cynthia Woodsong, Kathleen M. MacQueen, Greg Guest, and Emily Namey. *Qualitative research methods: A data collector’s field guide*. Family Health International, 2005.
- [MZ96] Elaine L. May and Barbara A. Zimmer. The evolutionary development model for software. *Hewlett Packard Journal*, 47:39–41, 1996.

- [NDR20] Solmaz Niknam, Harpreet S. Dhillon, and Jeffrey H. Reed. Federated learning for wireless communications: Motivation, opportunities, and challenges. *IEEE Communications Magazine*, 58(6):46–51, 2020.
- [NSU⁺18] Adrian Nilsson, Simon Smith, Gregor Ulm, Emil Gustavsson, and Mats Jirstrand. A performance evaluation of federated learning algorithms. In *Proceedings of the second workshop on distributed infrastructures for deep learning*, pages 1–8, Rennes, France, 2018.
- [NWD18] Tobias O. Nyumba, Kerrie Wilson, Christina J. Derrick, and Nibedita Mukherjee. The use of focus group discussion methodology: Insights from two decades of application in conservation. *Methods in Ecology and Evolution*, 9(1):20–32, 2018.
- [OM11] Tiago Oliveira and Maria F. Martins. Literature review of information technology adoption models at firm level. *Electronic journal of information systems evaluation*, 14(1):110–121, 2011.
- [OMY⁺22] Abdul-Rasheed Ottun, Pramod C Mane, Zhigang Yin, Souvik Paul, Mohan Liyanage, Jason Pridmore, Aaron Yi Ding, Rajesh Sharma, Petteri Nurmi, and Huber Flores. Social-aware federated learning: Challenges and opportunities in collaborative data training. *IEEE Internet Computing*, 27(2):36–44, 2022.
- [Orr03] Gregg Orr. Diffusion of innovations, by everett rogers (1995), 2003.
- [Pae17] Arsenio Paez. Gray literature: An important resource in systematic reviews. *Journal of Evidence-Based Medicine*, 10(3):233–240, 2017.
- [PCWA14] Nicolas Prat, Isabelle Comyn-Wattiau, and Jacky Akoka. Artifact evaluation in information systems design-science research—a holistic view. In *Proceedings of the 19th Pacific-Asian Conference on Information Systems*, Chengdu, China, 2014.
- [PFW⁺21] Luisa Pumplun, Mariska Fecho, Nihal Wahl, Felix Peters, and Peter Buxmann. Adoption of machine learning systems for medical diagnostics in clinics: Qualitative interview study. *Journal of Medical Internet Research*, 23(10):e29301, 2021.
- [PLAK21] Tommi Pauna, Hannele Lampela, Kirsi Aaltonen, and Jaakko Kujala. Challenges for implementing collaborative practices in industrial engineering projects. *Project Leadership and Society*, 2:1–14, 2021.
- [PSDB20] Debasis Pradhan, Sasank Sekhar Dalai, and Mandakini P. Behera. A comparative study on evolutionary model for software development. *International Journal Engineering Research & Technology*, 8(1):1–3, 2020.
- [PTH19] Luisa Pumplun, Christoph Tauchert, , and Margareta Heidt. A new organizational chassis for artificial intelligence - Exploring organizational readiness factors. In *Proceedings of the 27th European Conference on Information Systems*, Stockholm & Uppsala, Sweden, 2019.
- [PTJK15] Guy Paré, Marie-Claude Trudel, Mirou Jaana, and Spyros Kitsiou. Synthesizing information systems knowledge: A typology of literature reviews. *Information & Management*, 52(2):183–199, 2015.

-
- [PTRC07] Ken Peppers, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. A design science research methodology for information systems research. *Journal of management information systems*, 24(3):45–77, 2007.
- [RAA⁺21] Jawadur K. Rahman, Faisal Ahmed, Nazma Akhter, Mohammad Hasan, Ruhul Amin, Kazi E. Aziz, Muzahidul A. Islam, Saddam H. Mukta, and Najmul A. Islam. Challenges, applications and design aspects of federated learning: A survey. *IEEE Access*, 9:124682–124700, 2021.
- [RÅE20] Wiebke Reim, Josef Åström, and Oliver Eriksson. Implementation of artificial intelligence (AI): A roadmap for business model innovation. *AI*, 1(2):180–191, 2020.
- [Ras20] Vladislav E. Rasskazov. Financial and economic consequences of distribution of artificial intelligence as a general-purpose technology. *Finance Theory and Practice*, 24(2):120–132, 2020.
- [RBBC⁺14] Tobias Redlich, Sissy-Ve Basmer, Sonja Buxbaum-Conradi, Pascal Krenz, Jens Wulfsberg, and Franz-L: Bruhns. Openness and trust in value co-creation: Inter-organizational knowledge transfer and new business models. In *Proceedings of the 2014 Portland International Center for Management of Engineering and Technology*, pages 217–225, Portland, OR, USA, 2014. IEEE.
- [RFGG21] Ricardo Francisco Reier Forradellas and Luis Miguel Garay Gallastegui. Digital transformation and artificial intelligence applied to business: Legal regulations, economic impact and perspective. *Laws*, 10(3):70, 2021.
- [RH09] Per Runeson and Martin Höst. Guidelines for conducting and reporting case study research in software engineering. *Empirical software engineering*, 14:131–164, 2009.
- [RH20] André Renz and Romy Hilbig. Prerequisites for artificial intelligence in further education: Identification of drivers, barriers, and business models of educational technology companies. *International Journal of Educational Technology in Higher Education*, 17(1):1–21, 2020.
- [RKGR17] Sam Ransbotham, David Kiron, Philipp Gerbert, and Martin Reeves. Reshaping business with artificial intelligence: Closing the gap between ambition and action. *MIT Sloan Management Review*, 59(1):1–17, 2017.
- [RKW20] Lars Reimann and Günter Kniesel-Wünsche. Achieving guidance in applied machine learning through software engineering techniques. In *Companion Proceedings of the 4th International Conference on Art, Science, and Engineering of Programming*, pages 7–12, Porto, Portugal, 2020.
- [Rog83a] Everett M. Rogers. Attributes of innovations and their rate of adoption. In *Diffusion of Innovations*. The Free Press, 3 edition, 1983.
- [Rog83b] Everett M. Rogers. Elements of diffusion. In *Diffusion of Innovations*. The Free Press, 3 edition, 1983.

- [Rog83c] Everett M. Rogers. The innovation-decision process. In *Diffusion of Innovations*. The Free Press, 3 edition, 1983.
- [Rog83d] Everett M. Rogers. Innovation in organizations. In *Diffusion of Innovations*. The Free Press, 3 edition, 1983.
- [Rot07] Edna T. Rother. Systematic literature review x narrative review. *Acta paulista de enfermagem*, 20:5–7, 2007.
- [Row12] Frantz Rowe. Toward a richer diversity of genres in information systems research: New categorization and guidelines, 2012.
- [Roy87] Winston W. Royce. Managing the development of large software systems: Concepts and techniques. In *Proceedings of the 9th international conference on Software Engineering*, pages 328–338, Monterey, California, USA, 1987.
- [Sah06] Ismail Sahin. Detailed review of rogers’ diffusion of innovations theory and educational technology-related studies based on Rogers’ theory. *Turkish Online Journal of Educational Technology*, 5(2):14–23, 2006.
- [SBD⁺21] Stefan Studer, Thanh B. Bui, Christian Drescher, Alexander Hanuschkin, Ludwig Winkler, Steven Peters, and Klaus-Robert Müller. Towards CRISP-ML (Q): A machine learning process model with quality assurance methodology. *Machine learning and knowledge extraction*, 3(2):392–413, 2021.
- [SC14] Anselm Strauss and Juliet Corbin. *Basics of qualitative research*. Sage publications, 4 edition, 2014.
- [Sch15] Guido Schryen. Writing qualitative is literature reviews—guidelines for synthesis, interpretation, and guidance of research. *Communications of the Association for Information Systems*, 37(1):286–325, 2015.
- [Sea99] Carolyn B. Seaman. Qualitative methods in empirical studies of software engineering. *IEEE Transactions on Software Engineering*, 25(4):557–572, 1999.
- [SKMG21] Christoph Schröder, Felix Kruse, and Jorge Marx Gómez. A systematic literature review on applying CRISP-DM process model. *Procedia Computer Science*, 181:526–534, 2021.
- [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*, 30:649–665, 2020.
- [SS11] Ken Schwaber and Jeff Sutherland. The scrum guide. *Scrum Alliance*, 21(1):1–38, 2011.
- [Sta76] William H. Starbuck. Organizations and their environments. *Handbook of Industrial and Organizational Psychology*, 1976.
- [Suk07] Suphat Sukamolson. Fundamentals of quantitative research. *Language Institute Chulalongkorn University*, 1(3):1–20, 2007.

-
- [SVS05] Suprateek Sarker, Joseph S. Valacich, and Saonee Sarker. Technology adoption by groups: A valence perspective. *Journal of the Association for Information Systems*, 6(2):37–71, 2005.
- [TF90] Louis G. Tornatzky and Mitchell Fleischer. *The processes of technological innovation*. Lexington, Mass: Lexington Books, 1990.
- [TSK⁺18] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A survey on deep transfer learning. In *Proceedings of the 27th International Conference on Artificial Neural Networks*, volume 3, pages 270–279, Rhodes, Greece, 2018.
- [TSP22] Philip Treleaven, Malgorzata Smietanka, and Hirsh Pithadia. Federated learning: The pioneering distributed machine learning and privacy-preserving data technology. *Computer*, 55(4):20–29, 2022.
- [TWN19] Nattaphol Thanachawengsakul, Panita Wannapiroon, and Prachyanun Nilsook. The knowledge repository management system architecture of digital knowledge engineering using machine learning to promote software engineering competencies. *International Journal of Emerging Technologies in Learning*, 14(12):42–56, 2019.
- [Uni22] European Union. Data act: Commission proposes measures for a fair and innovative data economy, 2022.
- [VBB13] Viswanath Venkatesh, Susan A. Brown, and Hillol Bala. Bridging the qualitative-quantitative divide: Guidelines for conducting mixed methods research in information systems. *Management of Information Systems Quarterly*, 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*, pages 2206–2217, Verona, Italy, 2009.
- [VHC⁺20] Y Venugeetha, BM Harshitha, KP Charitha, K Shwetha, and V Keerthana. Breast cancer prediction and trail using machine learning and image processing. In *Proceedings of the 2nd International Conference on Data Science, Machine Learning and Applications*, pages 957–966, Kolkata, India, 2020. Springer.
- [VPPE⁺14] Willem G. Van Panhuis, Proma Paul, Claudia Emerson, John Grefenstette, Richard Wilder, Abraham J. Herbst, David Heymann, and Donald S. Burke. A systematic review of barriers to data sharing in public health. *BMC Public Health*, 14(1):1–9, 2014.
- [VZB21] Michael Veale and Frederik Zuiderveen Borgesius. Demystifying the draft EU Artificial Intelligence Act — analysing the good, the bad, and the unclear elements of the proposed approach. *Computer Law Review International*, 22(4):97–112, 2021.

- [WH00] Rüdiger Wirth and Jochen Hipp. CRISP-DM: Towards a standard process model for data mining. In *Proceedings of the 4th International Conference on the Practical Applications of Knowledge Discovery and Data Mining*, volume 1, pages 29–39, Manchester, England, 2000.
- [WLD⁺20] Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H. Yang, Farhad Farokhi, Shi Jin, Tony Q. Quek, and Vincent H. Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020.
- [WMOT21] Omar A. Wahab, Azzam Mourad, Hadi Otrok, and Tarik Taleb. Federated machine learning: Survey, multi-level classification, desirable criteria and future directions in communication and networking systems. *IEEE Communications Surveys & Tutorials*, 23(2):1342–1397, 2021.
- [WW02] Jane Webster and Richard T. Watson. Analyzing the past to prepare for the future: Writing a literature review. *Management of Information Systems quarterly*, 26(2):13–23, 2002.
- [WXLM19] Zhiyuan Wan, Xin Xia, David Lo, and Gail C. Murphy. How does machine learning change software development practices? *IEEE Transactions on Software Engineering*, 47(9):1857–1871, 2019.
- [WXLM21] Zhiyuan Wan, Xin Xia, David Lo, and Gail C. Murphy. How does machine learning change software development practices? *IEEE Transactions on Software Engineering*, 47(9):1857–1871, 2021.
- [YJLC22] Joo H. Yoo, Hyejun Jeong, Jaehyeok Lee, and Tai-Myoung Chung. Open problems in medical federated learning. *International Journal of Web Information Systems*, 18(2/3):77–99, 2022.
- [YML⁺22] Bin Yu, Wenjie Mao, Yihan Lv, Chen Zhang, and Yu Xie. A survey on federated learning in data mining. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 12(1):1–20, 2022.
- [ZBO20] Hongyi Zhang, Jan Bosch, and Helena H. Olsson. Federated learning systems: Architecture alternatives. In *Proceedings of the 27th Asia-Pacific Software Engineering Conference*, pages 385–394, Singapore, 2020. IEEE.
- [ZBT11] He Zhang, Muhammad Ali Babar, and Paolo Tell. Identifying relevant studies in software engineering. *Information and Software Technology*, 53(6):625–637, 2011.
- [ZC23] Alex Zarifis and Xusen Cheng. AI is transforming insurance with five emerging business models. In John Wang, editor, *Encyclopedia of Data Science and Machine Learning*, pages 2086–2100. IGI Global, 2023.
- [ZKB⁺21] Nikolas Zolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N. Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz. Advanced technologies adoption and use by us firms: Evidence from the annual business survey. Technical report, National Bureau of Economic Research, 2021.

- [ZLQ⁺20] Yufeng Zhan, Peng Li, Zhihao Qu, Deze Zeng, and Song Guo. A learning-based incentive mechanism for federated learning. *IEEE Internet of Things Journal*, 7(7):6360–6368, 2020.
- [ZLX⁺20] Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, and Yang Liu. {BatchCrypt}: Efficient homomorphic encryption for cross-silo federated learning. In *Proceedings of the 2020 USENIX annual technical conference*, pages 493–506, Virtual, 2020.
- [ZXB⁺21] Chen Zhang, Yu Xie, Hang Bai, Bin Yu, Weihong Li, and Yuan Gao. A survey on federated learning. *Knowledge-Based Systems*, 216:106775, 2021.
- [ZZH⁺21] Yufeng Zhan, Jie Zhang, Zicong Hong, Leijie Wu, Peng Li, and Song Guo. A survey of incentive mechanism design for federated learning. *IEEE Transactions on Emerging Topics in Computing*, 10(2):1035–1044, 2021.

- [ASS⁺23] Philipp Altmann, Leo Sünkel, Jonas Stein, Tobias Müller, Christoph Roch, and Claudia Linnhoff-Popien. SEQUENT: Towards traceable quantum machine learning using sequential quantum enhanced training. In *Proceedings of the 15th International Conference on Agents and Artificial Intelligence*, volume 3, pages 744–751, Lisboa, Portugal, 2023.
- [MGVM22] Tobias Müller, Nadine Gärtner, Nemrude Verzano, and Florian Matthes. Barriers to the practical adoption of federated machine learning in cross-company collaborations. In *Proceedings of the 14th International Conference on Agents and Artificial Intelligence*, volume 3, pages 581–588, Online, 2022.
- [MRSA21] Tobias Müller, Christoph Roch, Kyrill Schmid, and Philipp Altmann. Towards multi-agent reinforcement learning using quantum boltzmann machines. In *Proceedings of the 14th International Conference on Agents and Artificial Intelligence*, volume 1, pages 121–130, Virtual, 2021.
- [MSG⁺23] Tobias Müller, Maximilian Stäbler, Hugo Gascón, Frank Köster, and Florian Matthes. SoK: Assessing the state of applied federated machine learning. *arXiv*, pages 1–9, 2023.
- [MSS⁺21] Tobias Müller, Kyrill Schmid, Daniëlle Schuman, Thomas Gabor, Markus Friedrich, and Marc Geitz. Solving large steiner tree problems in graphs for cost-efficient fiber-to-the-home network expansion. In *Proceedings of the 14th International Conference on Agents and Artificial Intelligence*, volume 3, pages 23–32, Online, 2021.
- [MZM23a] Tobias Müller, Milena Zahn, and Florian Matthes. On the adoption of federated machine learning: Roles, activities and process life cycle. In *Proceedings of the 25th International Conference on Enterprise Information Systems*, volume 1, pages 525–531, Prague, Czech Republic, 2023.
- [MZM23b] Tobias Müller, Milena Zahn, and Florian Matthes. A pathway for the practical adoption of federated machine learning projects. In *Proceedings of the 27th Pacific-Asia Conference on Information Systems*, volume 6, pages 1–16, Nanchang, China, 2023.

- [MZM23c] Tobias Müller, Milena Zahn, and Florian Matthes. A process model for the practical adoption of federated machine learning. In *Proceedings of the 29th Americas Conference on Information Systems*, volume 1, pages 1–10, Panama City, Panama, 2023.
- [MZM23d] Tobias Müller, Milena Zahn, and Florian Matthes. Unlocking the potential of collaborative AI - on the socio-technical challenges of federated machine learning. In *Proceedings of the 31st European Conference on Information Systems*, volume 245, pages 1–14, Kristiansand, Norway, 2023.
- [MZM24] Tobias Müller, Milena Zahn, and Florian Matthes. Revealing the impacting factors for the adoption of federated machine learning in organizations. In *Proceedings of the 57th Hawaii International Conference on System Sciences*, pages 7343–7352, Honolulu, Hawaii, USA, 2024.
- [ZMM24] Milena Zahn, Tobias Müller, and Florian Matthes. Supporting managerial decision-making for federated machine learning: Design of a technology selection tool. In *Proceedings of the 57th Hawaii International Conference on System Sciences*, pages 6738–6747, Honolulu, Hawaii, USA, 2024.

Abbreviations

FedML	Federated Machine Learning
ML	Machine Learning
MLOps	Machine Learning Operations
AI	Artificial Intelligence
SDLC	Software Development Life Cycle
IP	Intellectual Property
CRISP-DM	Cross-Industry Standard Process for Data Mining
PETs	Privay-Enhancing Technologies
TOE	Technology-Organization-Environment
GPT	General Purpose Technology
DOI	Diffusion of Innovation
IEEE	Institute of Electrical and Electronics Engineers
RQ	Research Question
DSR	Design Science Research
IS	Information Systems
SoK	Systematization of Knowledge

APPENDIX A

Embedded Publications in Original Format

Revealing the Impacting Factors for the Adoption of Federated Machine Learning in Organizations

Tobias Müller
Technical University of Munich
and SAP SE
tobias.mueller15@sap.com

Milena Zahn
Technical University of Munich
and SAP SE
milena.zahn@sap.com

Florian Matthes
Technical University of Munich
matthes@tum.de

Abstract

The success of Machine Learning is driven by the ever-increasing wealth of digitized data. Still, a significant amount of the world's data is scattered and locked in data silos, which leaves its full potential and therefore economic value largely untapped. Federated Machine Learning is a novel model-to-data approach that enables the training of Machine Learning models on decentralized, potentially siloed data. Despite its potential, most Federated Machine Learning projects fail to actualize. The current literature lacks an understanding of the crucial factors for the adoption of Federated Machine Learning in organizations. We conducted an interview study with 13 experts from seven organizations to close this research gap. Specifically, we draw on the Technology-Organization-Environment framework and identified a total of 19 influencing factors. Thereby, we intend to facilitate managerial decision-making, aid practitioners in avoiding pitfalls, and thereby ease the successful implementation of Federated Machine Learning projects.

Keywords: Federated Machine Learning, Technology Adoption, TOE Framework, Interview Study

1. Introduction

The ever-increasing wealth of digitized data powers the disruptive potential of Machine Learning (ML) and its immense economic impact. Even though vast amounts of data is freely available, extensive amounts of already generated data is scattered, stored, and locked up in decentralized devices and data silos. Accessing these data silos becomes more difficult with privacy concerns and legal regulations, which leaves the economic potential of the stored data largely untapped.

Federated Machine Learning (FedML) is a novel ML paradigm, with the promise to build prediction models on decentralized data without the need for direct data sharing (McMahan et al., 2016). Through its model-to-data approach, FedML enables the usage of siloed data without disclosing data to third parties. Therefore, FedML has the potential to overcome data silos, enable the usage of currently untapped data and thereby be the catalyst for novel application fields of ML. Despite its advantages, there are only a few production-level applications and most work on FedML comprises prototypes or simulations (Lo et al., 2021). Investigating the challenges, success factors, and influential factors for the adoption of FedML might offer valuable insights into the missing operationalization of FedML. A better understanding of these factors would also aid practitioners to implement FedML projects and thereby support its broader practical adoption.

In contrast to the literature on FedML, research on traditional, centralized Artificial Intelligence (AI) systems already provides relevant insights into the challenges and success factors of AI adoption. For example, research on AI adoption in the financial services industry recognized a lack of AI-related skills, missing top management support, market regulations, and complex implementation as the main challenges (Kruse et al., 2019). Similar studies in the manufacturing and production domain identified leadership support as a crucial success factor (Demlehner and Laumer, 2020). Besides, the complexity of an organization additionally hinders AI adoption in manufacturing firms (Chatterjee et al., 2021). Similar results have also been obtained for AI adoption in public organizations (Neumann et al., 2022).

Organizations that are relatively inexperienced in AI technologies depend on the initiatives of single

employees or are able to successfully implement AI projects with the help of external partners (Bauer et al., 2020). However, top management support is essential to support the allocation of key resources. Once sufficient resources are available to develop AI solutions, the intra-organizational diffusion of AI may increase resistance due to conflicts between different in-house units (Neumann et al., 2022). Further studies confirm that organizational factors such as top management support and thereby organizational readiness are key factors in the adoption of AI in organizations (Alsheibani and Messom, 2019; Dora et al., 2022; Hamm and Klesel, 2021). For small and medium-sized enterprises, the lack of ML know-how poses an additional key challenge (Bauer et al., 2020).

However, FedML introduces another dimension of complexity. Due to its collaborative nature, we argue that FedML projects are additionally subject to collaboration-related challenges. Specifically addressing collaboration challenges in collaborative engineering projects is crucial to projects' efficiency and success (Diirr and Cappelli, 2018; Pauna et al., 2021). Since FedML works at the intersection of AI and collaborative project management, its influential factors for the adoption of FedML in organizations need to be investigated.

The current literature lacks a structured overview of the factors which influence the adoption of collaborative AI paradigms, such as FedML. This work aims towards closing this research gap. Through an expert interview study, we aim to draw on the experiences and expertise of practitioners to investigate the motivations, challenges, and influential factors for the adoption of FedML. Through the structured overview of influential factors, we intend to guide managerial decision-making, help practitioners avoid pitfalls, overcome challenges and overcome risks at an early project stage. We aim to achieve this goal by answering the following research questions (RQs):

RQ1: What are the reasons for the adoption of FedML in organizations and the accompanying main challenges and risks?

RQ2: Which factors influence the practical adoption of FedML in organizations?

2. Theoretical Background

In the following, we will describe the theoretical background of FedML as well as the basis of technology adoption frameworks.

2.1. Federated Machine Learning

FedML is an innovative ML technique that enables the collaborative training of a joint ML model on distributed datasets without the need of sharing data. In traditional ML settings, the data is usually accumulated in a central location, where the ML model is subsequently trained. Hence, data owners need to share their data with a central server and thereby risk losing their Intellectual Property (IP). FedML counteracts this need of sharing datasets through a model-to-data approach.

First introduced by McMahan et al. (2016), FedML can be divided into four distinct steps. These steps are illustrated in Figure 1. The server initially chooses a global model which is suitable for the use case and underlying data structure. In this step, the initial global model can be pre-trained by the server. Secondly, the global model is distributed amongst all participating clients. Thirdly, each client trains the global model on its own local dataset and stores the resulting update gradient. Thereby, each client owns a customized version of the global model based on the clients' individual, local dataset. Lastly, each client sends their stored update gradients back to the server, which are collected and aggregated based on a pre-defined protocol. The aggregate of the individual update gradients is then used to update the global model. These steps can be repeated until a certain accuracy level is reached or until the accuracy converges.

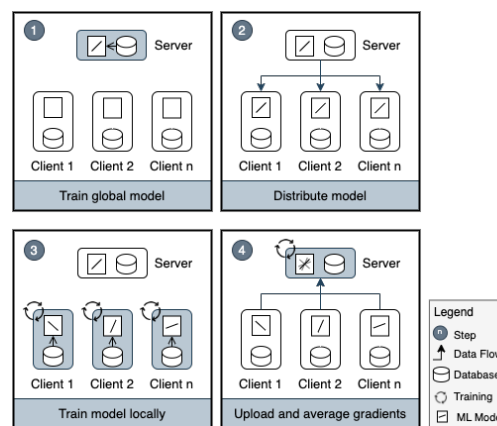


Figure 1. Federated Machine Learning process.

2.2. Technology Adoption Frameworks

The process of adopting innovative technologies in organizations has been a widely studied area within information systems and sparked a multitude of different technology adoption models. These models aim to

identify, predict, and describe the variables that affect adoption behavior in institutions (Dube et al., 2020). Technology adoption frameworks can be classified by their adoption context and categorized into groups of models that aim to study the adoption behavior of groups, individuals, or organizations (Liu et al., 2008).

In this work, we investigate the organizational adoption of FedML and therefore focus on organizational-based technology adoption frameworks. Consequently, models and frameworks which focus on the individual or group level were not considered in our study. On an organizational level, the *Diffusion on Innovations (DOI)* (Rogers, 2003) theory and the *Technology-Organization-Environment (TOE)* (Tornatzky et al., 1990) framework are the two most prominent models to measure the organizational readiness and acceptable use of innovative technologies.

The TOE framework is consistent with the DOI theory. Both emphasize individual, internal, and external characteristics of the organization as influencing factors for the organization's innovativeness. Compared to DOI, TOE additionally considers environmental factors in the technology adoption of organizations. Therefore TOE is considered more complete to explain intra-firm innovation adoption (Oliveira and Martins, 2011). In summary, TOE enables a comprehensive understanding of innovative technology decisions by considering the aspects from a technological, organizational, and environmental perspective.

3. Research Methodology

To explore the influencing factors for the adoption of FedML in organizations, we followed a qualitative research approach by conducting semi-structured interviews and drawing on the experiences of experts. The following sections describe the study design, data collection, and data analysis of our research.

3.1. Study Design

To answer our RQs, we collected data through semi-structured interviews. First, we identified potential interviewees working with FedML either in business or applied research through pre-saved contacts of prior research, referrals, or top search results (e.g., via LinkedIn). We then contacted the identified experts either via email or direct message and scheduled the interview after a positive response. Prior to the scheduled interview, we presented the research purpose, content, and structure to allow for impromptu follow-up questions.

We based our questionnaire on the TOE framework since it provides a solid theoretical and empirically supported structure of the influencing factors for the organizational adoption of innovative technologies, such as FedML. Through leveraging the provided TOE structure, we deem to gather a holistic overview of the influential factors for the practical adoption of FedML in organizations.

In our study, we only included interviewees that had sufficient topic-related knowledge. For the semi-structured interviews, we followed the guidelines proposed by Myers and Newman (2007). Each interview was recorded, transcribed, and coded. The results were iteratively compared with the insights from previous interviews until we reached theoretical saturation, allowing us to close the interview study.

In total, we conducted interviews with 13 experts from seven different organizations. Table 1 provides a codified overview of our resulting sample, including the participants' relevant information, such as their position, organization, and experience in their current position. From hereon, we refer to the experts by their corresponding participant identifier (ID). The interviewee demographic indicates a large variety of "voices" (Myers and Newman, 2007) and comprises a broad spectrum of backgrounds, job roles, and experience, thereby covering multiple viewpoints.

3.2. Data Collection

The interviews were conducted via videotelephony in February 2023, each with two participating researchers to ensure observer triangulation (Runeson and Höst, 2009). We presented a set of pre-defined questions to each interview partner. Each interview had the same outline, and the questions remained unchanged, however, due to the nature and flexibility of semi-structured interviews, slight variations regarding the order of questions or wording occurred. At the start of each interview, the research goal was recalled, and the interview structure was presented to alleviate misunderstandings before we proceeded to the interview questions. The interview guideline was developed based on the RQs and consisted of four different sections. The first section aimed to gather information about the participants' professional backgrounds and experiences with FedML projects. This was followed by general, open questions about their reason for adopting FedML as well as the encountered challenges and risks. Hence, the second section intends to address RQ1. The third section focused on the TOE factors as experienced by the interviewees and aimed towards RQ2. The final section investigated future directions and discussed

Table 1. Overview of expert interviews.

ID	Position	Organization	Experience
I1	Product Manager	Large software enterprise 1	≥10 years
I2	Architect	Large software enterprise 1	≥5 years
I3	Applied Researcher	Industrial software enterprise	≥2 years
I4	Development Expert	Large software enterprise 1	≥19 years
I5	CEO and Founder	FedML Startup 1	≥3 years
I6	Applied Researcher	Research center for AI security	≥1 year
I7	Senior Consultant and Project Lead	Large software enterprise 1	≥6 years
I8	Customer Advisor	Large software enterprise 1	≥2 years
I9	CEO and Founder	FedML startup 2	≥5 years
I10	Product Manager	Large software enterprise 1	≥4 years
I11	Researcher	Research center for software systems	≥4 years
I12	Solution Specialist and Product Manager	Large software enterprise 1	≥12 years
I13	Research Manager	Large software enterprise 2	≥4 years

possible tools that might help overcome challenges in an early project stage.

The results communicated in this work represent the findings of the first three sections. We plan to develop the discussed tool and publish the remaining empirical evidence in a separate work.

3.3. Data Analysis

The transcribed and recorded interviews were coded according to the guidelines of the *Reflexive Thematic Analysis* process (Braun et al., 2018). Consequently, we reviewed the conducted interviews and familiarized ourselves with the content of the collected data. We made notes on the initial insights of each interview and put the insights into the context of the overall data. Additionally, we assigned a unique ID to each expert and dismissed potentially sensitive information to ensure anonymity.

The transcripts were coded and analyzed with the help of MAXQDA2022¹. Each interview was coded according to important features relevant to the RQs. New codes were created whenever new findings could not be assigned to an existing code category, which triggered a re-codification of the previously coded data. Hence, the final coding was created through multiple rounds of coding. The codes were examined and grouped into broader themes. These themes were thereafter named and analyzed in detail to validate if the themes accurately depict the transcript data. Emerging conflicts were discussed by the researchers and resolved by mutual consent. Finally, the interviews with the annotated transcripts and their themes were summarized and contextualized in relation to existing literature.

¹<https://www.maxqda.com>

4. Results

This section presents the summarized results of our interview study. We first address RQ1 by describing the experts' reasons to adopt FedML as well as the main challenges and risks from their experiences. The subsequent sections then answer RQ2 by presenting the identified technological, organizational, and environmental factors.

Reasons of Adoption. We identified three main reasons for adopting FedML, which will be described in more detail in the following. An overview of the aspects and interviewee references can be seen in Table 2.

(1) *Field of Application:* FedML can enable the development of novel use cases and applications that would not have been possible using traditional ML approaches. Furthermore, adopting FedML can improve the performance and capabilities of existing ML products by leveraging additional data from data silos, possibly leading to a competitive advantage.

(2) *Data Privacy:* The potential access to sensitive data and the need for protecting sensitive data also drive FedML adoption. The privacy-enhancing features of FedML help to mitigate privacy concerns by enabling local model training without sharing raw data. This can be particularly important for industries working with sensitive data, such as the healthcare or financial sector. The privacy-enhancing nature of FedML can foster trust between organizations, encouraging collaboration by eliminating the need to share data between organizations and preserving the IP on the data.

(3) *Efficiency:* The improved communication and computation efficiency also motivates FedML adoption. FedML can improve communication

Table 2. Reasons of adoption and main challenges.

Category	Factors	Experts	#Experts (%)
Reasons of Adoption	Field of Application	I4, I5, I9, I10, I11, I13	6 (46.15%)
	Data Privacy	I1, I2, I4, I5, I6, I9, I11	7 (53.84%)
	Efficiency	I3, I5, I9, I10	4 (30.76%)
Challenges and Risks	Uncertainty, Risk Analysis, & Mitigation	I5, I12	2 (15.38%)
	Insufficient Management Support	I5, I7, I12	3 (23.07%)
	Novelty of Technology	I1, I2, I3, I5, I6, I8, I9, I10, I12, I11, I13	11 (84.61%)
	Collaboration	I1, I2, I4, I5, I6, I9, I10	7 (53.84%)
	Complex Implementation	I1, I3, I8, I9, I10, I13	6 (46.15%)
	ML Product	I2, I3, I6, I9, I13	5 (38.46%)

efficiency by minimizing the need for raw data sharing and centralization, resulting in more efficient communication and reduced network overhead. This can be particularly beneficial for organizations with distributed data sources and limited bandwidth. In addition, FedML allows organizations to leverage external expertise and resources without sharing their data, making it an option for outsourcing ML tasks without compromising privacy and security. This also allows companies without sufficient ML in-house expertise to develop such applications.

Challenges and Risks. The interviewees mentioned a total of six main challenges and risks in the adoption of FedML. The list of main challenges and risks including interviewee references can be seen in Table 2.

(1) *Uncertainty, Risk Analysis and Mitigation:* Due to the novelty of FedML, organizations may face uncertainties and challenges related to privacy, security, and compliance. Mitigating these risks requires constant careful analysis and planning. This includes identifying potential risks, assessing their potential impact, and implementing appropriate mitigation measures.

(2) *Insufficient Management Support:* It can be difficult to secure sufficient financial support and investment for FedML initiatives and to gain strategic or tactical buy-in from key decision-makers. Overcoming this challenge may require advocating the value of FedML in terms of its potential impact on business operations, competitive advantage and digitization/data-first strategies.

(3) *Novelty of Technology:* As a relatively new approach, organizations may face challenges as first or early adopters of FedML. These challenges include complex compliance assessments, regulatory and standards uncertainties, and managing the rapid evolution of FedML. Organizations may need to invest in research, collaboration, and proactive monitoring of regulatory and technological developments to

effectively address these challenges.

(4) *Collaboration:* FedML may be applied in collaborative settings with multiple organizations, which may present challenges for managing, coordinating, and achieving critical mass for effective training. To overcome this challenge, organizations may need to establish robust mechanisms for managing responsibilities, suitable communication channels, and ownership frameworks.

(5) *Complex Implementation:* Deploying and managing FedML systems can involve a significant effort to overcome complex technical challenges. Organizations need to carefully plan and execute technical implementation to ensure the effective adoption of FedML, potentially even across company borders.

(6) *ML Product:* The stochastic nature of FedML leads to challenges in managing expectations, evaluating performance, and ensuring reliable results. In addition, organizations may face privacy concerns as FedML involves local model training but does not eliminate privacy concerns.

4.1. Technological Factors

We identified a total of nine technological factors, which can be grouped into four categories. The following presents the categories with the identified factors. Table 3 provides an overview of these factors with references to the interviewees.

Data Considerations. Ensuring high-quality data, sufficient data volume, and data interoperability are critical factors to consider when assessing the feasibility and suitability of adopting FedML in a specific use case.

(1) *Data Quality:* Sufficient data quality is an essential foundation for the implementation of FedML and has a significant impact on the performance and reliability of the resulting models. Organizations

Table 3. Identified technological, organizational and environmental factors.

	Category	Factors	Experts	#Experts (%)
Technology	Data Considerations	Data Quality	I1, I2, I3, I5, I6, I8, I9, I10, I11, I13	10 (76.92%)
		Data Volume and Accessibility	I1, I2, I3, I5, I6	5 (38.46%)
		Data Interoperability	I1, I2, I3, I5, I6, I8, I9, I10, I11, I13	10 (76.92%)
	System Interoperability	FedML System	I3, I5, I13	3 (23.07%)
		Data Integration	I1, I5	2 (15.38%)
	Infrastructure	Compatibility and Accessibility	I1, I5, I9, I13	4 (40.76%)
		Computational Power	I5, I8, I9, I13	4 (30.76%)
	Orchestration	Versioning	I1, I2, I9	3 (23.07%)
Pipelines		I3, I8, I9, I10, I13	5 (38.46%)	
Organization	Organizational Readiness	Management Support	I1, I5, I6, I9, I10, I11, I12, I13	8 (61.53%)
		Knowledge and Expertise	I3, I5, I6, I7, I10, I11, I13	7 (53.84%)
	Federation Considerations	Collaboration Management	I1, I2, I3, I4, I5, I6, I7, I8, I9, I10, I11, I12	13 (100%)
		Co-Creation Management	I1, I9, I13	3 (23.07%)
		IP Management	I1, I7, I10, I11, I13	5 (38.46%)
Environment	Legal Regulations	Cartell Office	I6	1 (7.69%)
		Data Privacy	I3, I9, I11, I12, I13	5 (38.46%)
		Legal Clarity and Unambiguity	I1, I2, I5, I6, I9, I11, I12, I13	8 (61.53%)
	External Pressure	Market Competition	I6, I10	2 (15.38%)
		Regulatory Enforcement	I10	1 (7.69%)

need to consider various aspects of data quality, such as completeness, timeliness, and consistency. Data cleaning may be required to ensure high data quality, which is costly and may outweigh the resulting benefits.

(2) *Data Volume and Accessibility*: An adequate volume of data is essential for the training of accurate and reliable models. Sufficient data availability is a prerequisite for every FedML use case. The availability and accessibility of the data volume, potentially across data silos, is a critical factor for the adoption of FedML. Additional to the training process, organizations also need to ensure that they can also provide an appropriate data sample for the initial feasibility study.

(3) *Data Interoperability*: The data structure and statistical distribution must provide an interoperable basis. For that, organizations need to assess whether the data is homogeneous or can be homogenized. Standardized semantics and industry protocols can help to ensure data interoperability.

System Interoperability. System interoperability is crucial for the seamless training of a joint ML model across multiple clients. It is crucial that the FedML system is implemented appropriately on each side and that the local data storages are integrated and accessible.

(1) *FedML System*: Each participating organization either needs to have the expertise and resources to implement their part of the FedML system, or use an existing FedML platform. Additionally, it needs to be

ensured that the system can be enrolled across all clients.

(2) *Data Integration*: To run the FedML algorithm, data sources need to be integrated. Organizations need to make the data sources accessible so that the FedML system can train on the data sources locally. These data integration tasks must ensure that at each client's side the data from different sources can be combined and used for training in the FedML system.

Infrastructure. The infrastructure for the FedML process is crucial. This includes compatibility with existing infrastructure and ensuring sufficient computing power. Assessing the compatibility of the FedML system with the existing infrastructure on each client's side and evaluating the availability of sufficient computing resources are essential factors to consider in order to ensure a successful implementation of FedML.

(1) *Compatibility and Accessibility*: This relates to the compatibility of the FedML system with the existing IT infrastructure of each client. Organizations must ensure that the FedML system and its components can be implemented within existing infrastructure, including network architecture, hardware, and software. The FedML system may require internet access to enable communication and coordination between the distributed parties. Hence, organizations must ensure that the required connectivity and access are available. Compatibility also includes the integration of FedML with existing IT systems, such as data storage,

processing, and authentication mechanisms.

(2) *Computational Power*: The available computing power required for FedML training needs to be sufficient and depends on several factors, such as the role of the participant (client or aggregator), the specific FedML system, the size and complexity of the ML models, and the amount of data to be trained on. Organizations need to be able to assess whether their existing computational resources are sufficient to support the computational requirements of FedML deployment, or whether additional resources need to be allocated.

Orchestration. This category relates to the deployment of the FedML model including the versioning, training automation, and deployment. Adopting appropriate versioning practices and implementing robust training pipelines with automation can help ensure proper coordination and alignment of ML models between parties in a FedML system.

(1) *Versioning*: Model versioning is critical for managing changes and updates to the ML models used in FedML systems. Organizations need to implement appropriate versioning procedures and mechanisms to ensure accountability and that the FedML models are updated, tracked, and managed efficiently.

(2) *Pipelines*: The Pipelines and automated processes are critical for orchestrating the training process across distributed clients and coordinating gradient exchange, model synchronization, and model serving. Pipeline automation can help streamline the FedML deployment process and reduce manual effort to ensure efficient and scalable training operations.

4.2. Organizational Factors

The organizational factors are divided into two categories with a total of five factors, whereas the federation consideration is only applicable if the project includes a collaboration of different organizations. An overview of these factors with references to the interviewees is provided in Table 3.

Organizational Readiness. The readiness and resources of an organization to adopt emerging technologies and innovative ideas are critical. Organizations need to ensure that there is adequate management support, including awareness, understanding, and willingness to invest in AI projects. Secondly, they need to assess internal knowledge and data science expertise to ensure the successful adoption and implementation in the organization.

(1) *Management Support*: The level of awareness, understanding, and support by management for the

adoption of emerging technologies such as FedML is an important success factor. It includes the management's understanding of the potential of FedML in addressing business challenges, willingness to invest in AI projects, and overall mindset and openness to new technologies. Investments in ML projects are difficult to implement without risk aversion due to the lack of predictability of the outcome. Management support is critical to driving organizational change, providing the necessary resources, and creating a culture that is supportive of innovative ideas. Factors such as company size, risk aversion, and strategic focus on data-driven processes additionally influence organizational readiness.

(2) *Knowledge and Expertise*: The successful implementation of FedML systems also depends on the knowledge and experience of ML (and FedML) of an organization. It includes the organization's existing capabilities and resources for ML projects, digitization maturity, and overall readiness to implement emerging technologies such as FedML. Organizations need to provide internal capabilities and expertise for data science projects, ML model development, data infrastructure, and other necessary resources for their tasks. This factor heavily depends on the role of the participant and the degree of automation within the project. The lack of internal skills could be compensated through the acquisition of external knowledge.

Federation Considerations. If the project is implemented within a setting with different participating organizations, additional challenges arise. These influencing factors relevant to the federated setting include the effective management of collaboration, co-creation, and IP. Organizations need to carefully manage the collaboration between the participants as this is critical to the establishment of the collaboration, the feasibility of the project, and finally the successful implementation of FedML in the use case.

(1) *Collaboration Management*: This factor concerns the management of collaboration between different participants in a FedML environment. Collaboration in FedML can be complex and difficult due to factors such as establishing collaboration, withdrawing from collaboration, and finding and selecting appropriate partners for the use case. The distribution and management of tasks and responsibilities among uneven participants can increase the complexity. Incentive mechanisms may be needed to encourage active participation and collaboration among participants in a FedML environment. Effective collaboration management is critical to the success of a FedML implementation, and organizations need to carefully plan and manage collaboration processes to

ensure smooth and efficient operation.

(2) *Co-Creation Management*: In collaborative settings, the co-creation challenges involve the joint creation and ownership of the FedML model among all collaborators and stakeholders. This may include defining model ownership, data/model contributions, and sharing of results. Co-creation management may also involve defining roles and responsibilities, model usage policies, and governance mechanisms.

(3) *IP Management*: IP management includes issues related to the ownership, use, and protection of the IP. Organizations must carefully define and agree on the ownership and use of IP among participants, which may include legal agreements, contracts, and policies.

4.3. Environmental Factors

We identified two categories with a total of five environmental factors, which will be described in the subsequent paragraphs. All environmental factors and interviewee references can be seen in Table 3.

Legal Regulations. FedML projects in general and especially in collaborative settings need to consider legal regulations to be compliant. These regulations include antitrust compliance, data protection regulations, and ensuring clarity and unambiguity in the legal landscape. Companies need to carefully review and comply with the relevant legal requirements to ensure legally compliant implementation of FedML in their specific use cases and ultimately to be able to use the FedML model. The fast development and emergence of legal regulations which are relevant to AI applications and collaborative projects need to be carefully observed.

(1) *Cartel Office*: In collaborative settings, it needs to be determined whether cooperation with all participants is permissible under local antitrust or competition authority regulations. Depending on the jurisdiction and specific use case, the collaboration between participants may be subject to competition laws and regulations.

(2) *Data Privacy*: Considering legal regulations regarding data privacy is crucial, especially in projects with sensitive data. The type and sensitivity of used data, as well as the jurisdiction in which the FedML system operates, can have significant implications for legal compliance requirements. Organizations must carefully evaluate and understand the privacy implications of using data in a FedML environment, including potential risks associated with data sharing, data use, and privacy. Compliance with data protection regulations, such as the GDPR in the European Union, may be required, and organizations should ensure that

all data processing within the project is compliant with relevant regulations.

(3) *Clarity and Unambiguity*: Legal clarity and unambiguity are complicated topics given the fast pace of FedML's technological advances and the developments of laws. This can lead to uncertainty regarding the legal framework and applicable laws, which is even more complex given the high variability depending on the specific use case. Keeping up to date with the regulations and guidelines is important for regulatory compliance, and currently, there are no certifications or legal frameworks for FedML systems yet.

External Pressure. FedML projects are subject to external factors such as market competition or regulatory enforcement. Organizations must assess and respond to these external pressures by considering the potential benefits, risks, and available resources to make well-informed technology selection decisions about FedML adoption in their specific context.

(1) *Market Competition*: The pressure to collaborate with other organizations influences the development of FedML projects. Organizations may be pressured to engage in collaborations to remain competitive, gain access to sufficient data sources, and generate collective insights. Collaborations may also be needed to meet regulatory requirements, such as transparency along the supply chain in the example of the CO2 footprint.

(2) *Regulatory Enforcement*: The development of FedML systems may be affected by regulatory enforcement or regulatory changes. For example, in certain industries or jurisdictions, privacy-enhancing technologies such as FedML may be required by law to protect sensitive data. As a result, organizations may be forced to allocate additional budgets for implementing privacy-enhancing solutions.

5. Discussion

Through an expert interview study with 13 participants from seven organizations, we investigated the factors influencing the adoption of FedML in organizations. We identified three main reasons for adopting FedML in organizations and a total of six main challenges of practitioners, which are summarized in Table 2. Additionally, we identified a total of 19 factors that impact the adoption of FedML in organizations and summarized our findings in Table 3. We can sum up the results of our RQs as follows:

RQ1: *What are the reasons for the adoption of FedML in organizations and the accompanying main*

challenges and risks?

The main motivation for the usage of FedML revolved around the need to protect sensitive data and thereby its potential to enable the usage of sensitive data. The possibility of using sensitive data enables the training of ML algorithms on currently untapped data. A larger amount of training data yields the possibility to train more sophisticated ML algorithms for complex problem statements, which could not be solved with the current amount of available data. Therefore, the stated motivational driver to use sensitive data and the motivation to tackle novel fields of application are closely interrelated but still differ in the underlying motivation. The results suggest that a better-performing ML model and the increase of training data volume are perceived as more beneficial than improving the communicational and computational efficiency of FedML. However, some experts (I4, I5, I9, I10, I11) were driven by a combination of motivational factors.

As for the challenges and risks, most experts experienced challenges due to the novelty of the technology due to complex compliance assessments, as well as regulatory and standards uncertainties. Additionally, the first and early adopters need to come up with novel business cases, which impedes the allocation of budget and management support for FedML projects. Besides, the experts encountered challenges regarding the collaboration and the complex technical implementation of FedML.

RQ2: *Which factors influence the practical adoption of FedML in organizations?*

We structured the identified factors according to the TOE framework. The most relevant technological factors were aspects regarding data quality and data interoperability. On the organizational side, all experts unanimously agreed that collaboration management is a crucial impacting factor for the adoption of FedML in organizations. This is followed by factors regarding organizational readiness, especially management support as well as knowledge and expertise. The most relevant environmental factors comprised aspects around legal regulations, especially around missing legal clarity. The most relevant TOE factors concur with the identified challenges and risks, which further validates the significance of these aspects.

Contribution. The results of this study contribute to research on the adoption of emerging technologies in organizations. We complement current information systems literature by investigating the influencing factors of FedML adoption. Through systemizing

and presenting the influencing factors of its practical adoption, we intend to provide structured insights into the complex processes of implementing FedML projects. We hope that our insights aid management-oriented as well as technology-oriented audiences in the planning and development process. By knowing the crucial factors for the adoption of FedML, we help to avoid pitfalls, overcome challenges and counteract risks at an early stage. Overall, we intend to facilitate the process of adopting FedML in organizations and thereby help unlock novel fields of applications in the ML domain. In addition, our study provides a basis for further research on challenges and success factors for collaborative AI projects.

Limitations. There are multiple limitations to our work. FedML is an emerging technology and the influencing factors might change with a broader adoption of the technology. More factors might arise and some might be alleviated through the emergence of best practices or changed business understanding towards FedML. Due to its novelty, we were only able to interview first and early adopters, mainly consisting of larger enterprises, research institutes, or start-ups. Middle-sized enterprises were sparsely represented in the interviewee demographic and their experiences might have altered the outcome of our study. Moreover, our study is based on the experiences and expertise of 13 interview participants. Even though we reached theoretical saturation which terminated our interview study, more data from a bigger and more diverse set of interviewees with more perspectives might enrich our results. We encourage researchers and practitioners to further validate our findings in practice, and complement our proposed list of influencing factors.

6. Conclusion

In this paper, we presented a systematized set of critical factors that influence the adoption of FedML in organizations. Through an expert interview study with 13 participants from seven organizations, we identified a total of 19 influencing factors. Additionally, we presented the reasons for the adoption of FedML as well as the main challenges and risks, which were encountered by the interviewed experts. The critical factors with the most occurrences comprised aspects regarding collaboration management, data quality, data interoperability, organizational readiness, and the lack of legal clarity. Due to the novelty of FedML, these factors might change. A broader practical adoption will spark best practices and a change in the business understanding of FedML, which can impact

the landscape of influential factors. We encourage researchers to further extend and improve the list of influencing factors by applying it to various application domains or verifying it in case studies. We hope that our study provides a thorough understanding of the critical factors in the adoption of FedML, aids managerial decision-making, and that it can be used as a basis for further understanding of the challenges and success factors of collaborative AI projects.

7. Acknowledgements

The authors would like to thank SAP SE for supporting this work.

References

- Alsheibani, Y., Sulaiman Abdallah, Cheung, & Messom, C. (2019). Factors inhibiting the adoption of artificial intelligence at organizational-level: A preliminary investigation. *AMCIS 2019 Proceedings*, 2.
- Bauer, M., van Dinther, C., & Kiefer, D. (2020). Machine learning in SME: An empirical study on enablers and success factors. *AMCIS 2022 Proceedings*, 3.
- Braun, V., Clarke, V., Hayfield, N., & Terry, G. (2018). Thematic analysis. In P. Liamputtong (Ed.), *Handbook of research methods in health social sciences* (pp. 1–18). Springer Singapore.
- Chatterjee, S., Rana, N., Dwivedi, Y., & Baabdullah, A. (2021). Understanding AI adoption in manufacturing and production firms using an integrated tam-toe model. *Technological Forecasting and Social Change*, 170(5), 34.
- Demlehner, Q., & Laumer, S. (2020). Shall we use it or not? Explaining the adoption of artificial intelligence for car manufacturing purposes. *Proceedings of the 28th European Conference on Information Systems*.
- Diirr, B., & Cappelli, C. (2018). A systematic literature review to understand cross-organizational relationship management and collaboration. *Hawaii International Conference on System Sciences*.
- Dora, M., Kumar, A., Mangla, S. K., Pant, A., & Kamal, M. M. (2022). Critical success factors influencing artificial intelligence adoption in food supply chains. *International Journal of Production Research*, 60(14), 4621–4640.
- Dube, T., van Eck, R., & Zuva, T. (2020). Review of technology adoption models and theories to measure readiness and acceptable use of technology in a business organization. *Journal of Information Technology and Digital World*, 02, 207–212.
- Hamm, P., & Klesel, M. (2021). Success factors for the adoption of artificial intelligence in organizations: A literature review. *AMCIS 2021 Proceedings*, 1.
- Kruse, L., Wunderlich, N., & Beck, R. (2019). Artificial intelligence for the financial services industry: What challenges organizations to succeed. *Proceedings of the 52nd Hawaii International Conference on System Sciences*, 6408–6417.
- Liu, Z., Min, Q., & Ji, S. (2008). A comprehensive review of research in IT adoption. *2008 4th International Conference on Wireless Communications, Networking and Mobile Computing*, 1–5.
- Lo, S. K., Lu, Q., Wang, C., Paik, H.-Y., & Zhu, L. (2021). A systematic literature review on federated machine learning: From a software engineering perspective. *ACM Comput. Surv.*, 54(5).
- McMahan, H. B., Moore, E., Ramage, D., & y Arcas, B. A. (2016). Federated learning of deep networks using model averaging. *ArXiv*.
- Myers, M., & Newman, M. (2007). The qualitative interview in is research: Examining the craft. *Information and Organization*, 17, 2–26.
- Neumann, O., Guirguis, K., & Steiner, R. (2022). Exploring artificial intelligence adoption in public organizations: A comparative case study. *Public Management Review*, 1–28.
- Oliveira, T., & Martins, M. R. (2011). Literature review of information technology adoption models at firm level. *The Electronic Journal Information Systems Evaluation*, 14(1), 110–121.
- Pauna, T., Lampela, H., Aaltonen, K., & Kujala, J. (2021). Challenges for implementing collaborative practices in industrial engineering projects. *Project Leadership and Society*, 2, 100029.
- Rogers, E. (2003). *Diffusion of innovations*, 5th edition. Free Press.
- Runeson, P., & Höst, M. (2009). Guidelines for conducting and reporting case study research in software engineering. *Empirical Software Engineering*, 14, 131–164.
- Tornatzky, L., Fleischer, M., & Chakrabarti, A. (1990). *The processes of technological innovation*. Lexington Books.

UNLOCKING THE POTENTIAL OF COLLABORATIVE AI – ON THE SOCIO-TECHNICAL CHALLENGES OF FEDERATED MACHINE LEARNING

Research Paper

Tobias Müller, Technical University of Munich, School of Computation, Information and Technology, Department of Computer Science, Germany and SAP SE, Germany, tobias1.mueller@tum.de

Milena Zahn, Technical University of Munich, School of Computation, Information and Technology, Department of Computer Science, Germany and SAP SE, Germany, milena.zahn@tum.de

Florian Matthes, Technical University of Munich, School of Computation, Information and Technology, Department of Computer Science, Germany, matthes@tum.de

Abstract

The disruptive potential of AI systems roots in the emergence of big data. Yet, a significant portion is scattered and locked in data silos, leaving its potential untapped. Federated Machine Learning is a novel AI paradigm enabling the creation of AI models from decentralized, potentially siloed data. Hence, Federated Machine Learning could technically open data silos and therefore unlock economic potential. However, this requires collaboration between multiple parties owning data silos. Setting up collaborative business models is complex and often a reason for failure. Current literature lacks guidelines on which aspects must be considered to successfully realize collaborative AI projects. This research investigates the challenges of prevailing collaborative business models and distinct aspects of Federated Machine Learning. Through a systematic literature review, focus group, and expert interviews, we provide a systemized collection of socio-technical challenges and an extended Business Model Canvas for the initial viability assessment of collaborative AI projects.

Keywords: Federated Machine Learning, Collaborative Data Processing, Business Model, Alliances

1 Introduction

Artificial Intelligence (AI) had an immense economic impact in the last couple of years. In 2021 alone, the market of AI-based services including software, hardware and services exceeded 500\$ billion with a five-year compound annual growth rate of 17.5% (Forradellas and Gallastegui, 2021). The potential profitability raise is currently estimated by an average of 38%, which implies an economic impact of \$14 trillion until 2035¹. Unmistakably, the usage of AI enables new, unprecedented business models with a monumental impact on the industry. The main enabler for this disruptive new market is the emergence of big data, which forms the fundamental basis for AI systems. Even though vast amounts of data is freely available, a considerable amount of the world's data is scattered, stored and locked up in decentralized IoT devices and data silos. Naturally, the siloed data is hardly accessible, leaving a large portion of already generated data, and therefore economic potential, largely untapped. The emergence

¹ https://www.accenture.com/fr-fr/_acnmedia/36dc7f76eab444cab6a7f44017cc3997.pdf

of data silos is strengthened by data protection laws and regulations such as the General Data Protection Regulation (GDPR), California Consumer Privacy Act, Cyber Security Law and the General Principles of the Civil Law. These regulations justifiably aim to protect the privacy of individuals and therefore restrict direct data sharing between different parties (Li et al., 2022). This protection of privacy is an important pursuit but leads to more data silos and therefore unused economic potential.

Federated Machine Learning (FedML) introduced by McMahan et al. (2016) is a novel machine learning (ML) technology with the potential of building prediction models of decentralized and therefore siloed datasets. In contrast to traditional, centralized ML, FedML systems initially train a global ML model which is then distributed to all participants. Then, each participant individually trains the model locally on their own dataset. The clients solely return the update gradient resulting from the local training. Through this model-to-data approach, the data never leaves the client's device, but still enables the development of a joint ML model. Thus, FedML enables tapping the potential of big data without privacy leakage.

FedML technically has the potential to leverage siloed data while still preserving the intellectual property (IP) and privacy of each individuals' dataset. Hence, FedML enables the usage of currently untapped data and therefore brings the potential to be the catalyst for novel, disruptive business model innovation and locking unprecedented value from siloed data. However, this requires the collaboration of multiple parties which own these data silos. Hence, a collaborative business model is needed as a framework for how value can be created, and different parties can be incentivized for participating in such a collaborative network. Setting up collaborative business models is complex and a potential reason for failure. The current literature lacks guidelines for decision-makers on which aspects must be considered for the successful realization of collaborative AI projects.

This work aims towards closing this knowledge gap. More specifically, we investigate the challenges of prevailing collaborative business models through a systematic literature review and identify distinct aspects of collaborative FedML projects by conducting a focus group interview and multiple expert interviews. We work towards a systemized collection of socio-technical challenges and an easily consumable business model canvas (BMC) to aid decision-makers in the initial viability assessment of collaborative AI projects. Summarized, we aim to answer the following research questions (RQs):

RQ1: What are the general challenges of collaborative business models?

RQ2: What are the aspects of inter-organizational FedML business models in relation to prevailing collaborative business models?

RQ3: Which aspects and attributes should be considered for inter-organizational FedML projects and how can these be structured into an extended BMC?

To address these research questions, we first describe the theoretical background of our study by introducing Federated Machine Learning and providing background information on collaborative business models (section 2). Following, we elaborate on our tripartite research methodology, which consists of a systematic literature review, in-depth focus group interviews, and semi-structured expert interviews (section 3). Subsequently, we present the results of our research including a systemized overview of challenges for collaborative business models, a structured list of distinct socio-technical aspects for FedML projects and a proposal for a corresponding extended BMC (section 4). Finally, we discuss our work by reflecting the underlying research problem and research gaps. The discussion is followed by a summary of our contributions, answers to the RQs and limitations of our work. Our study concludes with an outline of future research (section 5).

2 Theoretical Background

The following section presents the theoretical background of our study. We first describe the motivation, terminologies, and the basic concept of FedML as originally proposed by McMahan et al. (2016). Subsequently, we provide general background information on business models to establish a common

understanding for this study. Finally, we elaborate on collaborative business models and corresponding extensions of the BMC by Osterwalder and Pigneur (2010).

2.1 Federated Machine Learning

A classic ML approach requires the collaborating participants to assemble their datasets in a central location and train a unique ML model M_{SUM} , exposing the data to each other and the central server. The participants thereby risk losing their data sovereignty and IP, which inhibits companies to collaborate and share data (Schomakers et al., 2020). Introduced by McMahan et al. (2016), FedML counteracts the need of sharing datasets through a model-to-data approach. As illustrated in Figure 1, a global ML model is chosen, which is distributed amongst all clients. The clients train the model locally on their individual dataset. The update gradients are sent back to the server and used to improve the global model. Thereby, FedML enables data owners to train a joint model M_{FED} without the need to disclose their data.

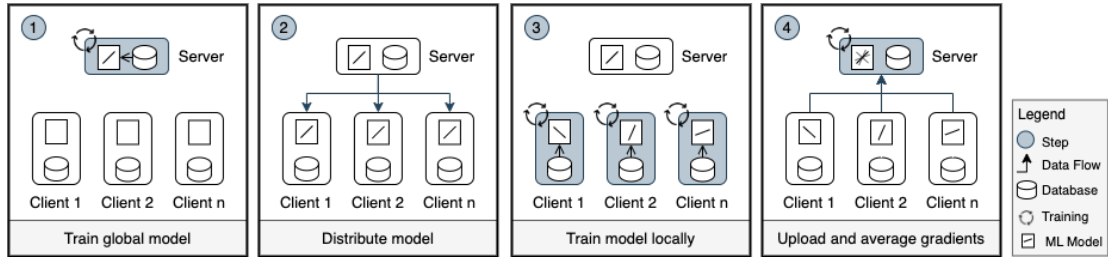


Figure 1. One iteration of the Federated Machine Learning process (source: own work).

In the original *FedAVG* implementation by McMahan et al. (2016) the model is learned through stochastic gradient descent (SGD), where each party k computes the average gradient $g_k = \nabla F_k(w_t)$ on its local data n_k at the current model w_t and iterates multiple times over the update $w_k \leftarrow w_k - \eta g_k$. The party submits the gradients to the central server, which aggregates the updates from all parties as:

$$w_{t+1} \leftarrow w_t - \sum_{k=1}^K \frac{n_k}{n} w_{t+1}^k$$

$$w_{t+1}^k \leftarrow w_t^k - \eta g_k, \forall k$$

While classic FedML operates on a client-server architecture, alternatives that do not rely on a central orchestrating server are also possible. For instance, parties can exchange model updates by establishing a peer-to-peer network, increasing the security of the process at the expense of consuming more bandwidth and resources for encryption (Roy et al., 2019).

Moreover, the distribution of features and samples across datasets may not be homogeneous. *Horizontal Federated Learning (HFL)* refers to the setup in which all datasets $\{D_{i=1}^K\}$ from the K parties contain different samples that share the same feature space. If instead, the same samples are present in all datasets, but feature spaces are disjoint, the setup is known as *Vertical Federated Learning (VFL)*.

Considering the high heterogeneity of data, especially if spread across different organizations, some authors have proposed to overcome the problem of sparse overlapping datasets through *Federated Transfer Learning (FTL)* (Liu et al., 2020). In this scheme, parties may select samples for training that minimizes the distance between their distributions (instance-based FTL) or learn a common feature space collaboratively (feature-based FTL). Alternatively, parties may start by using pre-trained models or by learning models from aligned samples to infer missing features and labels (model-based FTL).

Finally, it is important to note that the performances v_{SUM} and v_{FED} of the respective centralized and federated models, might differ considerably. This performance gap δ is characterized by $v_{SUM} - v_{FED} < \delta$ and will be strongly dependent on the characteristics of the particular application.

Consequently, FedML introduces a potential trade-off between the loss of performance respect to the centralized setup and the privacy guarantees provided by the distributed approach (Yang et al., 2019).

2.2 Collaborative Business Models

A business model describes essential aspects of an organization, explaining how the organization creates, delivers, and captures value (Osterwalder and Pigneur, 2010). In the academic literature, the definition of the term is fragmented, and no consistent boundaries are established. Nevertheless, it can be stated that a business model provides an organizational and strategic design for implementing a business opportunity (George, 2011).

In addition, Osterwalder and Pigneur (2010) argue that a shared understanding of the business model is crucial to its creation and success. Therefore, creating and discussing a business model requires a simple, relevant, and intuitively understandable concept without oversimplifying the complexity of how the organization works. The BMC by Osterwalder and Pigneur (2010) is a tool often used in practice to present a business model structured in nine components.

Business models are not only used for a single company but can also support assessing the feasibility and profitability of collaborations across companies (Kristensen and Ucler, 2016). The trend of an interconnected and dynamic environment encourages organizations to collaborate inter-organizationally and co-create value (Diirr and Cappelli, 2018). In literature, no unified framework exists for collaborations. Still, some approaches utilize Osterwalder and Pigneur (2010) general approach of a business model as a basis and customize it to set a higher focus on specifics (Kristensen and Ucler, 2016). For example, Eppinger and Kamprath (2011) highlight the importance of a partner and customer network in personalized medicine by modifying the canvas components and adding new ones, like intellectual property strategy. The approaches in the literature reach from modifications of business model components (e.g., Eppinger and Kamprath (2011) or Kristensen and Ucler (2016)), to configuration options of the business model (e.g., Curtis (2021) or Man and Luvison (2019)). However, the customizations are mainly application-oriented, tailored to the project to suit the needs and capture unique features influencing the business model and thus decisive for the project's success.

3 Methodology

This research was structured into three distinct parts. After a systematic literature review (SLR) to gain an overview of the challenges of collaborative business models, we organized an in-depth focus group interview to explore the novel field of inter-organizational FedML business models. By this, we aimed to augment the findings from the SLR and identify distinct challenges of business models for collaborative FedML projects. Since focus groups are characterized by their homogeneous group demographic, we pursued more generically applicable results by conducting additional semi-structured expert interviews. The research timeline is displayed in figure 2. The following subsections will go into more detail about the used research methodologies.




Figure 2. Research Timeline

3.1 Systematic Literature Review

To assess and identify the challenges of prevailing collaborative business models, we conducted a systematic literature review. By this, we intend to extract fundamentals and critical attributes of business models for inter-organizational collaborations, which will be collected, structured, and summarized. We followed a search strategy by Zhang et al. (2011) to identify the most relevant literature. Hence, the search is divided into base literature search, main search, and backward search.

The **Base Literature** consists of 13 papers including five publications on business model theory and eight on collaborative projects which were known by the authors prior to the search. Based on the initial literature corpus, we focused on finding keywords related to inter-organizational business models. The resulting search string S is as follows:

String	Query
S1	(collaborate* OR federated OR interorganization* OR inter-organization* OR intercompany OR cross-company OR multi-party OR cross-industr* OR multi-institution* OR shared OR sharing OR alliance OR networked)
S2	("business model*" OR "model canvas" OR "business value model*")
S	S1 AND S2

Table 1. Compiled search string for the database search.

The **Main Search** was conducted from April 2022 to June 2022. The list of searched databases comprises *IEEE Xplore*, *ACM Digital Library*, *ScienceDirect*, *Wiley InterScience* and *SCOPUS*. We only included peer-reviewed English and German publications with full-text access. With the defined search string S , databases, and criteria we collected 262 distinct publications. We only aim to include work in the field of computer science and technology (coarse focus) as well as literature regarding inter-organizational collaborations and business models (narrow focus). Successively, the corpus consisting of 262 distinct publications has been filtered solely by title, abstract and full text regarding the defined coarse and narrow focus. By this, 18 publications remained.

For the **Backward Search**, we scanned the references of the resulting 18 publications from the main search. Again, these referenced publications were filtered by title, abstract and body according to the inclusion and exclusion criteria. After eliminating duplicates, we added one further study resulting in a total of 19 publications.

Finally, relevant information was extracted and synthesized from the final literature corpus. The structured and consolidated output yielded a set of critical attributes and challenges of business models for inter-organizational collaborations in the technological sector. In the following every insight from the SLR is referenced as usual with the corresponding publication.

3.2 In-depth Group Interview

Based on the identified challenges of prevailing collaborative business models from the SLR, we aimed to explore the distinct aspects of collaborative business models for inter-organizational FedML projects. Due to the novelty of the topic and the need for exploration, we organized an in-depth focus group interview (Dilshad and Latif, 2013) to study the business requirements of collaborative ML projects based on the findings from the SLR.

The focus group consisted of five participants and two moderators, where one moderator ensures smooth progress and the other ensures that all topics are covered. All participants worked on a project involving the adoption of FedML in a cross-company use case. The participants were briefed about collaborative business models and were given an overview of the findings from the SLR. Afterwards, the group was

asked about the critical attributes and challenges of business models related to their collaborative FedML project followed by a reflection and lively discussion. Through this, we were able to identify further challenges based on their real-world experiences. The emerging data was coded by two researchers and incorporated into the results of the SLR. In the following, every insight which was gained through the in-depth focus group interview is referenced via the index (*FG*).

3.3 Semi-Structured Expert Interviews

Even though the focus group helped explore the socio-technical challenges of FedML collaborations, the insights might be highly biased due to the homogeneous demographics of the participants. To gain a more generically applicable understanding, we aimed to draw from the experiences of further experts working in the field of applied AI, especially with experience in FedML projects. For these expert interviews, we draw from the Grounded Theory methodology (Hoda et al., 2011). Hence, we confronted the interviewees with a set of pre-defined questions and recorded as well as transcribed the interviews. We successively conducted and compared the results of each interview. After 5 interviews theoretical saturation was reached and consequently, the interview study was closed. The set of interviewees represented a more diverse set of experts from different organizations and domains. Table 1 presents a codified table of our sample. We developed an interview guide based on the research questions and findings from the SLR as well as the focus group interview including open questions about potentially missing attributes, challenges, and further insights. These interviews allowed us to go more in-depth and identify missing aspects and gain more detailed, in-depth individual understanding to develop the guideline questionnaire further.

The interviewees allowed the findings to be published in an anonymized manner but did not agree to disclose the full transcriptions. Therefore, the full transcripts are not included. The findings from the semi-structured interviews are referenced in the following with the participant ID as listed in table 2.

Participant ID	Position	Organization	Duration
E1	AI Business Developer	Large German software enterprise	52
E2	AI Project Lead	Large German software enterprise	44
E3	Principal Data Scientist	Large German software enterprise	45
E4	Applied Researcher	Medium-sized innovation company	35
E5	Scientific Researcher	Research institute for software development	59

Table 2. Interview Study Participants

4 Socio-Technical Challenges of Interorganizational Federated Machine Learning

Applying collaborative models can be challenging in different domains, especially when several companies are involved. When the business is operationalized, complexity increases significantly because the general business model idea needs to balance the interests of all participants (Pauna et al., 2021). Collaborations with multiple participants are complex in nature, and collaboration failure rates are high, leaving much revenue at risk and unrealized value (Man and Luvison, 2019). Moreover, aligning the business model with operational and governance-related aspects is suggested to help position the organization to deliver on its value proposition for a successful implementation of the business model (Curtis, 2021). Hence, early identification of the collaboration challenges is critical for the successful creation of the collaborative business model.

To better understand which specific collaboration challenges should be considered, we first investigate the challenges of prevailing inter-organizational business models and, secondly, which FedML-related

socio-technical aspects are critical for successful implementation and therefore should be considered in a corresponding collaborative business model.

4.1 Challenges of Collaborative Business Models

Joint work of different organizations is complex, and organizations should be prepared to face challenges arising from cooperation. In the following, we give an overview of the systematized results of the SLR on the challenges of inter-organizational business models. We present our key results in a structured manner based on the work of Diirr and Cappelli (2018).

Diirr and Cappelli (2018) divide the challenges of inter-organizational collaborations into three categories: external, internal, and network-related challenges. External and internal challenges are detached from inter-organizational collaboration. External challenges relate to environmental challenges, e.g., natural events, and internal challenges arise from inside the project, for example, infrastructure problems. Network-related challenges focus on the relationships and interactions between organizations and can be further subdivided into management, business process, and collaboration challenges (Diirr and Cappelli, 2018). Based on these categories in conjunction with the findings of the SLR we derived the following network-related challenges as listed in table 3.

Category	Description	Aspects
Management Challenges	Include how organizations create and establish collaboration, compromising the following aspects	Selection of suitable participating actors (Pauna et al., 2021).
		Change management for dynamic collaboration (Caridà et al., 2015; Redlich et al., 2014).
		Cooperation establishment: lack of commitment from participating organizations (Proulx and Gardoni, 2020); building and expanding trust between the parties (Bleja et al., 2020; Diirr and Cappelli, 2018; Redlich et al., 2014).
		Decision-making and coordination slowness within the collaboration (Bleja et al., 2020; Caridà et al., 2015; Diirr and Cappelli, 2018; Redlich et al., 2014).
		Communication with government authorities requires a different approach due to multiple parties' interactions (Pauna et al., 2021).
Business Process Challenges	Addresses the way organization's structure and design partnership operations	Definition of a mutual business goal of the collaboration (Diirr and Cappelli, 2018).
		Co-Creation Management for delivering the value proposition: <ul style="list-style-type: none"> • Distribution of financials, investment (Pauna et al., 2021), costs, and revenues (Bleja et al., 2020; Caridà et al., 2015; Pauna et al., 2021). • Risk allocation (Diirr and Cappelli, 2018). • Ownership structure (Diirr and Cappelli, 2018; Kujala et al., 2020). • Responsibility assignment (Diirr and Cappelli, 2018). • Align on quality of co-creation product (Diirr and Cappelli, 2018). • Intellectual property Management (Eppinger and Kamprath, 2011).

		Slower business strategy and process identification (Berkers et al., 2020; Diirr and Cappelli, 2018).
		Alignment of the structures of heterogeneous organizations with distinct characteristics (Diirr and Cappelli, 2018).
		Infrastructure for managing relationships between multiple collaboration actors (Caridà et al., 2015; Diirr and Cappelli, 2018).
Collaboration Challenges	Describe how organizations jointly work together to achieve the goal of collaboration	Agreement on collaboration and alignment with the organizations' own objectives (Bleja et al., 2020; Costa and Da Cunha, 2015; Diirr and Cappelli, 2018; Man and Luvison, 2019; Pauna et al., 2021).
		Alignment of different organizations: culture and common ethics (Bleja et al., 2020; Diirr and Cappelli, 2018; Kujala et al., 2020).
		Risk of opportunism of participants and consequences of action (Diirr and Cappelli, 2018).

Table 3. Overview of Challenges for Inter-Organizational Collaborations

This overview of challenges is a consolidation of the selected academic sources of the SLR and aims to provide a general understanding of the difficulties of such collaborations. It is important to note that naturally, this list might not be comprehensive and that certain, potentially important, aspects might be missing.

4.2 Aspects of Interorganizational FedML Business Models

To provide initial guidance in the creation of business models for collaborative FedML projects we aim to identify the corresponding critical challenges, which need to be considered at an early stage. These aspects are derived from literature research, in-depth group interviews and expert interviews. The following section presents the aspects in more detail. The basis is formed by the questions catalogue of the BMC by Osterwalder and Pigneur (2010). The extension reflects the specifics of collaboration and technology that can be generalized.

Before setting up a business model for inter-organizational FedML projects, it is necessary to clarify if the underlying problem can be solved by applying FedML. We presuppose a prior feasibility check and task-technology-fit analysis, but still include these two points in our systemized collection of socio-technical challenges.

Osterwalder and Pigneur (2010) use nine components with their associated questions to describe the generic business model clustered in three parts (Create Value, Deliver Value and Capture Value). We augmented these parts with a section that addresses specific aspects of the inter-organizational collaborative environment and FedML. The extensions were obtained from comparable business model extensions retrieved from the literature review, the challenges of section 4.1 and insights from the conducted interviews. Overall, this results in the list as seen in Table 4. The insights from the in-depth focus group interview are referenced by the index *FG* whereas the semi-structured interviews are referenced by the corresponding participant ID (E1 - E5) are described in Table 2.

The complete list with guiding questions including the used literature corpus is provided in a complementary document². The overview of aspects targets important areas of a business model for collaborative AI projects, but due to the nature of the research approach, this list might not be exhaustive.

² Extended Business Model Canvas, Guiding Questions and Literature Corpus: bit.ly/3mXOUQg

Section	Component	Description
Create Value	Value Proposition	Describes the value that can be delivered to a certain customer.
	Customer Relationships	Explains the kinds of relationships an organization makes with particular customer segments.
	Channels	Defines how an organization interacts and reaches its customers to serve the value proposition.
	Customer Segments	Refers to the various groups of people the organization wants to reach and serve.
Deliver Value	Key Partners	Designates the network of suppliers and partners that are essential business model.
	Key Activities	Describes the most important things an organization must do to make its business model work.
	Key Resources	Defines the essential resources required.
Capture Value	Cost Structure	Describes all costs that are decisive for the operation of the business model
	Revenue Streams	Represent the earnings that an organization receives from each customer segment.
Collaboration Management	Collaboration Structure	Describes the negotiation mechanisms for building the network of participants and how decision-making is handled and coordinated within the collaboration. This mainly reflects whether there is a dominant participant in the collaboration (FG, E1, E2, E3, E4).
	Participant Management	Includes the formation regarding the suitability of participants, the change management of the collaboration, and the transparency of the project participants (FG, E1, E4).
	Infrastructure	Includes the management of collaboration- and technology-specific communication channels, as well as the platform to facilitate it (E1, E4).
Co-Creation Management	Distribution of Ownership, Responsibility and Accountability	Encourages distribution mechanisms within the collaboration. The aspects mentioned are crucial for the product created by the ML model (FG, E1, E2, E3, E4, E5).
	Distribution of Revenue and Costs	Describes how participants are rewarded for their participation and how additional effort is compensated (FG, E1, E2, E3, E4).
	Intellectual Property Management	Is crucial for enabling inter-organizational collaboration for joint activities (FG, E1, E2, E3, E4).
Co-Creation Practices	Profit Calculation	Describes the project's estimated cost-effectiveness over the FedML lifecycle with the different participants (E2, E3, E4).
	Risks of Infeasibility	Specify how the feasibility of solving the project with the FedML technology is determined (FG, E1, E4).
	Alignment in Quality	Describes how product quality is defined, ensured, and tracked throughout the FedML lifecycle (E1, E5).
	Implementation of Activities	Explains the extent to which the data-generating parties are involved in operational implementation (FG, E4, E5).
FedML Product	Compliance Data Protections	Include what regulations must be considered to develop a compliant FedML model (FG, E1, E4).
	Versioning	How versioning is handled within the FedML process (E1).
	Retirement	Includes how an ML Model can be recalled and claimed from participating parties and also end customers (E1, E4).

Table 4. Aspects of Inter-Organizational FedML Collaborations

4.3 Extension of the Business Model Canvas

The BMC by Osterwalder and Pigneur (2010) can be used to guide the creation and discussion of a business model. The canvas is a simplification of reality and should also pick up the most critical aspects of the business model already. To provide an easily consumable entry point for the initial viability assessment of collaborative AI projects, we aimed to extend the traditional BMC with the corresponding most critical aspects and challenges. This extended BMC represents another layer of simplification to the provided collection of socio-technical challenges from chapter 4.2. Hence, the presented collection was reduced to the most critical subareas: *Collaboration Management* and *Co-Creation Management*. Both aspects were mentioned as the most critical challenges in the interviews and address the collaborative approach as well as the joint creation of a FedML model. Guiding questions are added to the component titles for ease of use and intuitive comprehension. Figure 3 shows the extension of the canvas with colour-coded support. White tiles represent the business model aspects of a collaboration as a whole and blue tiles the aspects within the collaboration. To spare time and efforts, we suggest that the extended BMC should be used as a first basis to identify potential roadblocks of collaborative FedML projects. Thereupon, decision-makers can use the more detailed and comprehensive collection of socio-technical challenges as a second step of the viability assessment. To develop a concrete collaborative business model further steps (e.g. the identification of value streams between the collaborating parties) are necessary. Our artefacts solely represent a one-stop shop for the early identification of socio-technical aspects, challenges and potential roadblocks in the creation of collaborative AI projects.

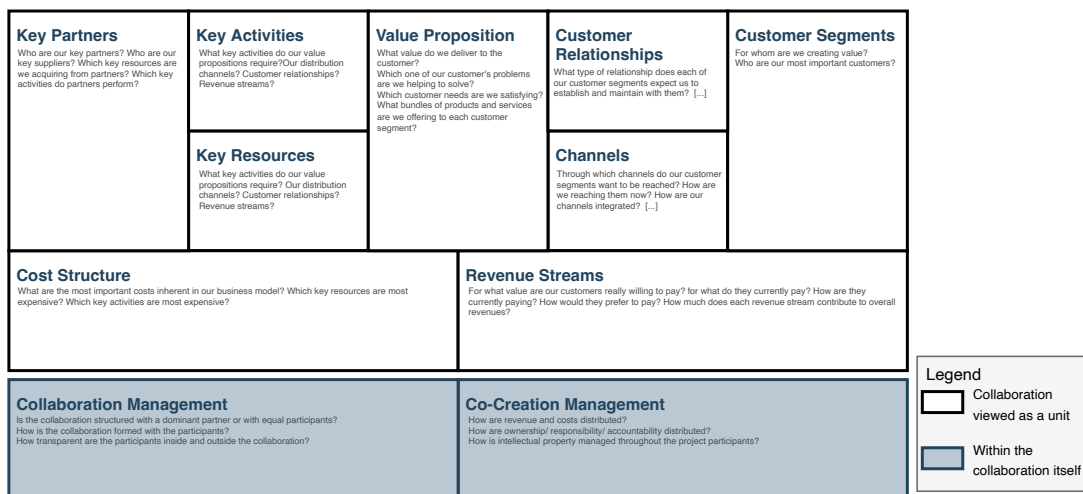


Figure 3. Extended BMC for Inter-Organizational FedML Collaborations

5 Conclusion

A significant amount of the world's data is scattered, stored, and locked up in decentralized data silos. The siloed data is hardly accessible, which leaves a large portion of already generated data and its economic potential largely untapped. The model-to-data approach of FedML technically enables the creation of a joint ML model from decentralized data without the need of data sharing. However, this novel privacy-enhancing ML paradigm requires collaboration of multiple parties which own the data silos. Consequently, a collaborative business model is required to define how value can be created. Setting up these collaborative business models is complex with a high potential of failure.

The information systems research literature offers interesting insights on emerging AI business models and collaborative business models. Current research on the distinct aspects of AI business models resulted in a multitude of relevant contributions. For example, how organizations evaluate AI value creation mechanisms (Alsheiabni et al., 2020) or how the value drivers of AI open new business model

design opportunities (Böttcher et al., 2022). Additionally, studies investigated how AI already transformed the business models of specific domains, such as the insurance sector (Zarifis et al., 2019; Zarifis and Cheng, 2023), Fintech (Zarifis and Cheng, 2022) or education (Renz and Hilbig, 2020; Zarifis and Efthymiou, 2022). Our study complements the current information systems literature by revealing the socio-technical challenges of collaborative AI projects and consequently providing guidance in the creation of a corresponding inter-organizational business model.

5.1 Contributions

Through a systematic literature review, in-depth focus group interview and semi-structured expert interviews, we first investigated the challenges and aspects of business models for inter-organizational FedML projects. These findings were aggregated and structured into a comprehensive set of guiding questions and compressed into an extension of the BMC. The resulted questionnaire represents a set of detailed aspects which need to be considered in the creation of the collaborative business model. Thereby, we aid decision-makers at an early stage of the business model development and prepare them for challenges related to the collaboration. The traditional BMC by Osterwalder and Pigneur (2010) was complemented by two dimensions and should aid decision-makers in the first assessment of value creation, delivery and capturing for inter-organizational FedML projects. We can sum up the results of our research questions as follows:

RQ1: *What are the general challenges of collaborative business models?*

A joint work of multiple organizations is complex, and a multitude of challenges arise from cooperation. Through a systematic literature review on collaborative business models in the technology sector, we aggregated a list of said challenges. Based on the example of Diirr and Cappelli (2018), we structured this list into *Management Challenges*, *Business Process Challenges*, and *Collaboration Challenges*. These categories were filled with challenges from the identified literature corpus. Specifically, the cooperation establishment, slowness within the collaboration, distribution of financials, and agreement on collaboration and alignment with the organizations' own objective seem to be the most critical challenges. However, this list may not be exhaustive but should present the most important challenges.

RQ2: *What are the aspects of inter-organizational FedML business models in relation to prevailing collaborative business models?*

To identify relevant aspects of inter-organizational FedML business models, we organized an in-depth focus group interview to explore this novel topic followed by semi-structured expert interviews to get a more diverse view and to augment the findings from the focus group study. Particularly, the difficulties of allocating rights and responsibilities within the co-creation management appear to be of specific interest in the inter-organizational use of FedML. These results were combined with the identified challenges from RQ1 and structured into groups of aspects for collaborative FedML business models. As a result, we received four aspect clusters: *Collaboration Management*, *Co-Creation Management*, *Co-Creation Practices*, and *FedML Product*. A more detailed specification is listed in the provided collection of socio-technical challenges.

RQ3: *Which aspects and attributes should be considered for inter-organizational FedML projects and how can these be structured into an extended BMC?*

The findings from RQ2 and the resulting questionnaire act as a basis for the extended BMC. Since the canvas should be a simplification of reality but should also consider the most critical aspects of the business model, we selected the most referenced and mentioned subareas. Therefore, *Collaboration Management* and *Co-Creation Management* were selected to expand the original canvas. These subareas seemed to be the most critical challenges of prevailing collaborative business models and were mentioned as the most critical success factors in the focus group interview as well as in the expert interviews. Hence, we argue that these two dimensions are important extensions to the BMC and capture the most critical aspects without losing handiness. This assumption and reasoning needs to be validated in practice but should provide a proper basis.

Overall, the insights of our study advance the understanding of the socio-technical challenges which arise in collaborative AI projects and of the relevant aspects in the development of corresponding inter-organizational business models. The systemized list of socio-technical challenges and resulting extended BMC can be used by decision-makers for the initial viability assessment. The comprehensive set of guiding questions can be used as assistance for the business model development process. By this, we guide decision-makers to make use of FedML and provide support to overcome socio-technical obstacles to the adoption of collaborative AI. Therefore, our study helps to unlock previously inaccessible value from siloed data and contributes to business model innovation.

5.2 Limitations

There are obvious limitations to our work. The participants of the focus group were affiliated with the same company and worked on similar projects within this company. Hence, the findings from this group could be highly biased and might have led to one-sided results. We tried to counteract this by conducting further interviews with experts of different backgrounds and affiliations. However, the theoretical saturation was reached after five interviews which terminated our interview study with a small sample size of five participants. More data and consequently more interesting perspectives from a bigger and even more diverse set of interviewees might enrich our results. Therefore, we encourage researchers and generally interested readers to use our work as a basis to complement, refine and develop our artefacts with their own insights. Moreover, our work poses as a starting point for the development of potential business models of inter-organizational FedML projects. This model only comprises the relevant decisive factors for the success of collaborative FedML projects but does not capture more low-level aspects as the value streams between organizations. This would pose a natural next step in developing a potential business model. We encourage the investigation of a model which captures how actors within the collaborative FedML project might exchange value (e.g., e3-value model). We also assume that a feasibility check was conducted beforehand if FedML is a fitting solution. If FedML can be excluded from the set of reasonable technologies choices for the given problem, there would be no point in going a step further by addressing the multitude of distinct socio-technical challenges of collaborative AI projects and building a concrete business model. Therefore, an a priori task technology analysis would be reasonable.

5.3 Future Research

Generally, the research in the field of FedML is dominated by technical work. Nonetheless, the practical adoption of novel technologies like FedML is dependent on more than the technical dimension. Decision-makers will not consider using FedML if the legal framework is fuzzy, the business model does not provide a proper value proposition or the task and technology contradict. We believe that FedML is a technology with great potential and will open a large variety of possibilities in the era of big data, where huge amounts of data is stored in data silos. To unlock this potential, there needs to be more research on the social and socio-economic challenges of collaborative ML. From legal frameworks and governance concepts to task-technology analyses, the research field is wide open and ready to be explored.

Acknowledgments

The authors would like to thank SAP SE for supporting this work.

References

Alsheibani, S. A., Cheung, Y., Messom, C., and Alhosni, M. (2020). "Winning AI Strategy: Six-Steps to Create Value from Artificial Intelligence," in: Anderson, B. B., Thatcher, J., Meservy, R. D.,

- Chudoba, K., Fadel, K. J., Brown, S. (eds.) *26th Americas Conference on Information Systems, Virtual*.
- Berkers, F., Turetken, O., Ozkan, B., Wilbik, A., Adali, O. E., Gilsing, R., and Grefen, P. (2020). "Deriving Collaborative Business Model Design Requirements from a Digital Platform Business Strategy," *IFIP Advances in Information and Communication Technology* 598 (1), 47–60.
- Bleja, J., Wiewelhoeve, D., Grossmann, U., and Mörz, E. (2020). "Collaborative Business Model Structures for Wireless Ambient Assisted Living Systems," *2020 IEEE 5th International Symposium on Smart and Wireless Systems within the Conferences on Intelligent Data Acquisition and Advanced Computing Systems (IDAACS-SWS)*, Virtual.
- Böttcher, T., Weber, M., Weking, J., Hein, A., and Krcmar, H. (2022). "Value Drivers of Artificial Intelligence". *28th Americas Conference on Information Systems*, Minneapolis, USA.
- Caridà, A., Colurcio, M., and Melia, M. (2015). "Designing a collaborative business model for SMEs," *Sinergie Italian Journal of Management* 33 (1), 233–253.
- Costa, C. C., and Da Cunha, P. R. (2015). "More than a gut feeling: Ensuring your inter-organizational business model works," *28th Bled EConference: #eWellbeing - Proceedings*, 86–99.
- Curtis, S. K. (2021). "Business model patterns in the sharing economy," *Sustainable Production and Consumption* 27, 1650–1671.
- Diirr, B., and Cappelli, C. (2018). "A systematic literature review to understand cross-organizational relationship management and collaboration." *51st Hawaii International Conference on System Sciences*, Hawaii, USA.
- Dilshad, R. M., and Latif, M. I. (2013). "Focus Group Interview as a Tool for Qualitative Research: An Analysis," *Pakistan Journal of Social Sciences* 33, 191-198.
- Eppinger, E., and Kamprath, M. (2011). "Sustainable Business Model Innovation in Personalized Medicine," *R&D Management Conference*, Linköping, Sweden.
- Forradellas, R. F. R., and Gallastegui, L. M. G. (2021). "Digital Transformation and Artificial Intelligence Applied to Business: Legal Regulations, Economic Impact and Perspective," *Laws*, 10 (3), 70.
- George, G. (2011). "The business model in practice and its implications for entrepreneurship research. Entrepreneurship Theory and Practice," *Entrepreneurship Theory and Practice* 35 (1), 83-111.
- Hoda, R., Noble, J., and Marshall, S. (2011). "Grounded theory for geeks," *Proceedings of the 18th Conference on Pattern Languages of Programs - PLoP '11*, Irsee, Germany.
- Kristensen, K., and Ucler, C. (2016). "Collaboration Model Canvas: Using the Business Model Canvas to Model Productive Collaborative Behavior," *2016 International Conference on Engineering, Technology and Innovation/IEEE International Technology Management Conference (ICE/ITMC)*, Wuhan, China.
- Kujala, J., Aaltonen, K., Gotcheva, N., and Lahdenperä, P. (2020). "Dimensions of governance in interorganizational project networks," *International Journal of Managing Projects in Business* 14 (3), 625–651.
- Li, J., Zhang, C., Zhao, Y., Qiu, W., Chen, Q., and Zhang, X. (2022). "Federated learning-based short-term building energy consumption prediction method for solving the data silos problem," *Building Simulation* 15 (6), 1145–1159.
- Liu, Y., Kang, Y., Xing, C., Chen, T., and Yang, Q. (2020). "A Secure Federated Transfer Learning Framework," *IEEE Intelligent Systems* 35 (4), 70–82.
- Man, A.-P. de, and Luvison, D. (2019). "Collaborative business models: Aligning and operationalizing alliances," *Business Horizons* 62 (4), 473–482.
- McMahan, H. B., Moore, E., Ramage, D., and Arcas, B. A. y. (2016). "Federated Learning of Deep Networks using Model Averaging," *CoRR*.
- Osterwalder, A., Pigneur, Y., and Clark, T. (2010). *Business model generation: A handbook for visionaries, game changers, and challengers*, Wiley.
- Pauna, T., Lampela, H., Aaltonen, K., and Kujala, J. (2021). "Challenges for implementing collaborative practices in industrial engineering projects," *Project Leadership and Society* 2, 100029.

- Proulx, M., and Gardoni, M. (2020). "Methodology for Designing a Collaborative Business Model – Case Study Aerospace Cluster," *IFIP Advances in Information and Communication Technology* 594, 387–401.
- Redlich, T., Basmer, S.-V., Buxbaum-Conradi, S., Krenz, P., Wulfsberg, J., and Bruhns, F.-L. (2014). "Openness and trust in value co-creation: Inter-organizational knowledge transfer and new business models," In P. G. Kocaoglu D.F. Anderson T. R., Daim T. U. , Kozanoglu D. C. , Niwa K. (Ed.), *PICMET 2014—Portland International Center for Management of Engineering and Technology, Proceedings: Infrastructure and Service Integration*, Portland, USA.
- Renz, A., and Hilbig, R. (2020). "Prerequisites for artificial intelligence in further education: identification of drivers, barriers, and business models of educational technology companies," *Int J Educ Technol High Educ*, 17(14).
- Roy, A. G., Siddiqui, S., Pölsterl, S., Navab, N., and Wachinger, C. (2019). "BrainTorrent: A Peer-to-Peer Environment for Decentralized Federated Learning," *arXiv*.
- Schomakers, E.-M., Lidynia, C., and Ziefle, M. (2020). "All of me? Users' preferences for privacy-preserving data markets and the importance of anonymity," *Electronic Markets* 30 (3), 649-665.
- Yang, Q., Liu, Y., Cheng, Y., Kang, Y., Chen, T., and Yu, H. (2019). "Federated learning," *Synthesis Lectures on Artificial Intelligence and Machine Learning* 13 (3), 1–207.
- Zarifis, A., and Cheng, X. (2022). "A Model of Trust in Fintech and Trust in Insurtech: How Artificial Intelligence and the Context Influence It," *Journal of Behavioral and Experimental Finance* 36 (1), 1-20.
- Zarifis, A., and Cheng, X. (2023). "AI is Transforming Insurance With Five Emerging Business Models," *Encyclopedia of Data Science and Machine Learning*. IGI Global, 2086-2100.
- Zarifis, A., and Efthymiou, L. (2022). "The four business models for AI adoption in education: Giving leaders a destination for the digital transformation journey," *2022 IEEE Global Engineering Education Conference (EDUCON)*, Tunis, Tunisia.
- Zarifis, A., Holland, C. P., and Milne, A. (2019). "Evaluating the Impact of AI on Insurance: The Four Emerging AI and Data Driven Business Models," *Emerald Open Research*, 1-17.
- Zhang, H., Babar, M. A., and Tell, P. (2011). "Identifying relevant studies in software engineering," *Information and Software Technology* 53 (6), 625–637.

A Process Model for the Practical Adoption of Federated Machine Learning

Completed Research Full Paper

Tobias Müller

Technical University of Munich
and SAP SE
tobias1.mueller@tum.de

Milena Zahn

Technical University of Munich
and SAP SE
milena.zahn@tum.de

Florian Matthes

Technical University of Munich
matthes@tum.de

Abstract

The wealth of digitized data forms the fundamental basis for the disruptive impact of Machine Learning. Yet a significant amount of data is scattered and locked in data silos, leaving its full potential untouched. Federated Machine Learning is a novel Machine Learning paradigm with the ability to overcome data silos by enabling the training of Machine Learning models on decentralized, potentially siloed data. Despite its advantages, most Federated Machine Learning projects fail in the project initiation phase due to their decentralized structure and incomprehensive interrelations. The current literature lacks a comprehensible overview of the complex project structure. Through a Design Science Research approach, we provide a process model of a Federated Machine Learning life cycle including required activities, roles, resources, artifacts, and interrelations. Thereby, we aim to aid practitioners in the project initiation phase by providing transparency and facilitating comprehensibility over the entire project life cycle.

Keywords

Federated Machine Learning, Process Model, Software Engineering, Applied AI, Design Science Research.

Introduction

The disruptive potential of Machine Learning (ML) roots in the emergence of big data and the ever-increasing wealth of digitized data. Yet, the lack of sufficient training data is still a persevering bottleneck in creating sophisticated, data-demanding ML systems. Even though vast amounts of data is freely available, a considerable amount of the world's data is still scattered, stored and locked up in data silos, therefore hardly accessible. This lack of available, suitable training data creates a competitive disadvantage especially for small and medium-sized enterprises (SMEs) (Bauer et al. 2020) leaving their full economic potential unreached. SMEs could overcome data scarcity by breaking up data silos, sharing data and collaborating. However, the companies' willingness to share data is low due to privacy concerns and a potential loss of intellectual property (IP) (Schomakers et al. 2020).

Federated Machine Learning (FedML) is a novel ML technique which allows the creation of a joint ML model on decentralized and therefore siloed datasets (McMahan et al. 2016). Through this model-to-data approach, companies could collaboratively train an ML model without the need for direct data sharing. Technically, FedML has the potential to enable SMEs overcome data silos, use currently untapped data, and thereby leverage the full potential of ML without privacy leakage. Despite its advantages, there are currently only a few production-level applications and most work on FedML comprises prototypes or simulations (Lo, Lu, Wang, et al. 2022). The missing operationalization of FedML may be attributable to

multiple aspects. For example, the persisting discrepancy between engineering traditional, deterministic software and engineering non-deterministic ML systems makes the integration of ML cumbersome (Giray 2021). Additionally, FedML requires the coordination of multiple parties across different life cycle stages due to its decentralized nature. We recognized through focus group discussions and expert interviews, that FedML is currently missing clarity over its complex and multi-faceted process flow. This missing clarity poses a key challenge for practitioners in the project initiation and communication with potential participants. Current literature usually takes a technical perspective when designing FedML systems and lacks a comprehensible overview of the complex project structure. Against this backdrop, we aim to provide transparency by developing a comprehensible process model of an end-to-end FedML project life cycle. The process model intends to break down, structure and illustrate the different project stages including required tasks, resources, roles, artifacts, and their interrelations. In summary, we aim to answer the following research questions (RQs):

RQ 1: What are the most relevant components and aspects needed for a comprehensible overview of a FedML process flow?

RQ 2: How can a generic, structured process model of an end-to-end Federated Machine Learning project life cycle be designed?

Theoretical Background and Related Work

FedML is a novel, disruptive ML paradigm that enables the training of a joint ML model on distributed datasets without the need of sharing data. In traditional ML settings, data is usually accumulated in a central location, where the ML model is subsequently trained. Hence, data owners need to share their data with a central server and risk losing their IP. Introduced by McMahan et al. (2016), FedML counteracts the need of sharing datasets through a model-to-data approach. As visualized in Figure 1, the FedML process can be divided into four steps. First, the server chooses an initial global model which is suitable for the use case and underlying data structure. The global model can be initially trained by the server. Secondly, the server distributes the global model amongst all clients. Thirdly, each client trains the global model on its own local dataset and stores the update gradients. Consequently, each client owns its individually trained ML model based on its local dataset. Finally, the clients send the individually computed update gradients back to the server, which are aggregated based on a pre-defined protocol and used to update the global model. Steps 2-4 can be repeated until a certain accuracy level is reached or until the accuracy converges. Even though FedML usually uses a client-server architecture consisting of a central orchestration server and multiple clients, other architectures have been proposed as well. We refer to client-server architectures with a central orchestrating server since it is the most widely used architectural pattern (Lo, Lu, Zhu, et al 2022).




Figure 1. One Iteration of the Federated Machine Learning Training Process (own work)

IEEE published a reference architecture with generalized information about the structure of FedML model training processes and involved roles (IEEE 2021). The standard can be used as a basis for implementation, but it does not offer information about the entire life cycle. However, considering all life cycle stages is crucial and an existing open problem for an industrial-level implementation of FedML (Zhang et al. 2020). In contrast to the current literature on FedML, such life cycle models can be found in numerous studies on traditional, centralized ML systems. Studer et al. (2021), for example, introduced CRISP-ML(Q), an standard ML process model overarching the entire life cycle from business understanding to monitoring and maintenance with a focus on quality assurance. Kreuzberger et al. (2022) developed an end-to-end workflow including functional components and roles tailored to ML operations (MLOps) principles. Similarly, Kumara et al. (2022) proposed a reference architecture for

MLOps to aid in streamlining the life cycle of ML models in production. Besides, work on software engineering for ML as well as studies on AI governance yielded insights into ML life cycles. Amershi et al. (2019) investigated how software teams at Microsoft developed AI applications and presented a life cycle model including a set of best practices and challenges based on their observations. Laato et al. (2022) explored the incorporation of AI governance into system development life cycle models by conducting expert interviews. These interviews resulted in a set of governance concepts and a visualization of how AI governance stages tie into existing software development life cycles. The studies on ML life cycles yielded valuable insights into existing practices and the different stages of ML project. Based on that, Ritz et al. (2022) presented a process model to additionally illustrate the dependencies and interactions within the different stages and activities. Their model describes the activities and resulting artifacts including the interdependencies throughout the ML software development cycle. The existing life cycle and process models on centralized ML systems can be used as a basis for FedML projects but need to be revisited due to the decentralized nature of FedML. The decentralization introduces additional layers of complexity and coordination, which need to be addressed for a successful industrial-level implementation. In contrast to studies on ML systems, the current literature on FedML lacks such a comprehensible overview of the complex project structure. We aim to close this research gap and enhance the current literature corpus on FedML through a process model which breaks down, structures, and illustrates the different project stages including required tasks, resources, roles, artifacts, and their interrelations.

Research Approach

To investigate why FedML projects fail to actualize, we initially conducted a focus group discussion. We invited three project teams that attempted to realize FedML projects prior and discussed the encountered challenges during the projects. The main challenges have already arisen in the project initiation phase and was two-fold. The first challenge is the structuring of the process flow including the activities, resources, data, and information exchange. Secondly, communicating the division of tasks and coherent interrelations with the collaboration partners posed another main challenge. The focus group participants agreed that a structured end-to-end process model could alleviate both problems. The process model would facilitate the planning phase and communication with external participants by providing an abstract overview of the process flow including its required resources, roles, activities, and their interrelations throughout the project. We recognized that current research on FedML lacked such a holistic process model. To develop a process model for operationalizing FedML projects we used the design science research (DSR) methodology as proposed by (Peffer et al. 2007). We chose the DSR approach since it provides a methodical, rigorous approach for producing and evaluating innovative, purposeful artifacts for a specified problem domain (Hevner et al. 2004). Our research activities according to the DSR steps by Peffer et al. (2007) can be summarized as follows:

(1) Problem Identification and Motivation: Through a focus group, we recognized that structuring the FedML process is complex. Understanding and communicating the division of tasks as well as the coherent interactions with potential partners is challenging. Current literature does not offer a structured, comprehensible overview of the FedML project flow and lacks guidance in the project initiation phase.

(2) Objectives of a Solution: Our objective was the development of a generically applicable end-to-end process model for FedML projects to facilitate the planning phase and communication with external participants. The model should comprise an entire life cycle and provide a clear understanding of interacting entities, their corresponding activities, interactions, and dependencies along with the required resources and resulting artifacts. The content and structure should be closely aligned and consistent with best practices offered by literature on ML life cycles, and software development life cycles. Furthermore, the model should be useful and easily understandable by non-technical stakeholders and practitioners.

(3) Design and Development: We reviewed current literature on ML life cycles, software development life cycles and assessed which practices, procedures and information are applicable for our process model. We gained our knowledge base on the structure of FedML processes from current literature. We conducted an interview study to collect information on the FedML project structure, separation of roles, activities, and interactions between the entities throughout the project life cycle. We describe the interview study in another publication (Müller et al. 2023). According to the described objectives, we designed and structured the elements of our process model. We incorporated the relevant findings from the literature review and interview study. During development, we conducted regular mini focus group discussions with

varying participants to assess the model regularly and to iteratively implement feedback. Additionally, we presented the initial resulting process model to a diverse group of 6 experts with technical backgrounds to gather feedback and change requests on the completeness, comprehensibility, and level of detail.

(4) *Demonstration*. We demonstrated the process model during an expert discussion and within a large industrial lighthouse project as part of a project ideation process to adopt FedML for varying use cases.

(5) *Evaluation*: We conducted two survey-based iterations to evaluate the artifact on the objectives.

(6) *Communication*: Communication is being done through this paper.

Results

A process model should describe an abstract representation of reality and reduce complexity by eliminating details which do not influence relevant behavior (Curtis et al. 1992). We adapted the seven process modeling guidelines (7PMG) by Mendling et al. (2010) wherever possible to design such a process model, which is comprehensible and comprises solely the important details of the underlying process. Following the 7PMG, we use as few elements as possible by clustering activities where possible and only incorporate the most essential elements for the FedML project life cycle. This guideline concurs with our objective to produce an easily understandable process model. To follow the 7PMGs, we also aimed to minimize the routing paths per element by only including the most relevant inputs, outputs, and feedback loops. The full process model is designed so that it can be easily decomposed according to the life cycle stages. Finally, we used verb-object activity labels. Figure 2 displays the resulting process model.

Components of the Process Model

In this section, we answer RQ1 by presenting the components of the process model. These components comprise life cycle *stages*, *activities*, *roles*, *resources*, as well as *artifacts* and are described as follows:

Stages of the life cycle represent interrelated tasks and activities and are structured according to software development life cycle stages in combination with MLOps and FedML-specific stages. Our model comprises five stages: *Project Initiation*, *Project Validation*, *Project Setup*, *System Design and Development*, and *Deployment and Maintenance*. The *System Design and Development* stage contains the FedML training process and can be divided into three sub-stages: *Global Model Design*, *Local Model Training* and *Global Model Aggregation*. The stages are described in the section on the process flow.

Roles describe the behavior and responsibility of an individual or a group. One individual may take multiple roles, but each role describes a distinct set of behaviors and responsibilities attributed to a role. The following gives a short description of the nine different roles and their intended area of responsibility.

(1) *Business Stakeholder* has the business need or problem to be solved and represents the initiator. The Business Stakeholder defines the business requirements, establishes a project strategy and agrees on the project definition with the success criteria (Kreuzberger et al. 2022; Too and Weaver 2014; Zwikael and Meredith 2018). In the process model, the business stakeholder is solely involved in strategic activities.

(2) *Project Manager* is responsible and held accountable by the business stakeholder for the success of the project. The project manager oversees and leads the team to achieve the project's objectives and is responsible for planning and managing the project efficiently (Too and Weaver 2014; Zwikael and Meredith 2018). Therefore, the project manager is involved in strategic activities.

(3) *Subject Matter Expert* (also called domain expert) is an expert in a specific domain and deeply understands the business problem. The Subject Matter Expert aids the Business Stakeholder in the definition of goals during the project initiation stage (Amershi et al. 2019; Studer et al. 2021).

(4) *Solution Architect* analyzes the functional, technical and business requirements and defines the overall solution architecture and technologies to be used (Kreuzberger et al. 2022). In the process model, the Solution Architect is mainly involved in the technical assessment of the ML problem.

(5) *Data Scientist* is an ML expert responsible for analyzing the business problem and building a suitable ML model which solves the business problem. This includes conducting the algorithm selection, feature engineering and performing hyperparameter tuning (Kim et al. 2018; Kreuzberger et al. 2022; Kumara et

al. 2022). Hence the Data Scientist is involved in operational activities such as translating the business problem into an ML problem and further ML-related operational activities such as designing the model.

(6) *Data Engineer* is responsible for pulling and engineering raw data such that the curated data is usable and accessible to the Data Scientist. These activities comprise building and managing data pipelines together with performing data cleaning and feature engineering (Kreuzberger et al. 2022; Kumara et al. 2022). The Data Engineer performs operational activities like preparing data. The Data Engineer is the only role which performs local activities.

(7) *Software Engineer* applies mature software engineering techniques to ensure that the ML system is built in a robust manner. The Software Engineer is responsible for packaging, testing as well as assuring the quality and robustness of the ML model along with the infrastructure (Kreuzberger et al. 2022; Serban et al. 2020). Hence, the Software Engineer is involved in software-related operational activities.

(8) *Development-Operations (DevOps) Engineer* aims to automate and combine the processes of model development and model operations to deploy and serve the ML model. This includes incident management, monitoring, support and delivery of models in production (Kreuzberger et al. 2022; Kumara et al. 2022; Laato et al. 2022). Hence, the DevOps Engineer performs operational activities in the deployment and maintenance stage.

(9) *Legal Representative* checks compliance with legal frameworks (e.g., GDPR or HIPAA) and clarifies the conformity of the potentially established collaboration.

Activities summarize units of work performed by roles. An activity has a clear purpose and usually results in the creation or update of an artifact (Anwar 2014). Specific roles are assigned to each activity. We chose the granularity such that the goal of the activity is comprehensible and such that the flow of activities provides a good understanding of the underlying process. Activities are described with verb-object labels. They either use information or data flows as input or can be triggered through prior activities. Some activities are optional depending on the use case-specific set-up and are indicated through dashed lines. The different activities are arranged horizontally by their chronological order and vertically by the type of activity. Activity types are indicated by swimlanes and are categorized into *strategic activities*, *operational activities*, and *local activities*. Strategic activities comprise management tasks such as planning, analysis, strategy formulation and organization. Operational and local activities cluster technical tasks which are concerned with the development and serving of the software product. We introduced an additional swimlane for local activities specific to FedML. These activities are executed solely in the environment of the local data contributors and accentuate which activities are performed by the decentralized entities. The separation into operational activities and local activities demonstrates the division, interrelations and dependencies between the central entity and local data contributors.

Artifacts are tangible by-products which are produced, modified, or used by a process. Artifacts are used as input to perform an activity or are the result of activities (Anwar 2014). We differentiate between *document artifacts* and *code artifacts*. Documents comprise specifications, descriptions, diagrams, reports, and other documented outputs. Code artifacts include output resulting from implementations such as ML models, infrastructures, or logs. Optional activities are indicated through a dashed outline.

Resources are hardware or software components which are required to fulfil the activities but do not result from the activities of the process. We differ between global and local resources. Global resources can be used by every participant, whereas local resources are used exclusively by the corresponding local data contributor. Resources comprise for example data storage, log storage, model registry and metric storage. We include resources in our process model to fulfill the objective of providing transparency and showcasing the global and local interactions and dependencies within the FedML project life cycle.

Process Flow of a Federated Machine Learning Life Cycle

In this section, we answer RQ 2 by providing a walkthrough of the process model as displayed in Figure 2. We based the stages of our process life cycle on the different phases of software development life cycle models (Akinsola et al. 2020; Alsaqqa et al. 2020; Apoorva and Deepty 2013; Gungur et al. 2020) in combination with ML life cycle models (Amershi et al. 2019; Kreuzberger et al. 2022; Kumara et al. 2022; Laato et al. 2022; Ritz et al. 2022). Finally, we adjusted the stages to the specifics of FedML processes and architectures (Bharti et al. 2022; Bonawitz et al. 2019; IEEE 2021; Lo et al. 2021; Lo, Lu, Zhu, et al. 2022;

Zhang et al. 2020). Subsequently, we will describe the process model according to the life cycle stages. The listed activities represent a high-level abstraction of the needed tasks. The level of abstraction should give non-technical stakeholders and other practitioners a good understanding of the needed activities but should not be too detailed to keep the process model comprehensible and manageable.

Project Initiation stage is the first step in starting a project. The business stakeholder and subject matter expert establish the goal and scope of the project. They define how business value will be delivered and which use cases will be tackled. Hence, a business model, business requirements and use case diagrams will emerge. The goals and scope are communicated with the data scientist, which will translate the business problem into an ML problem. The data scientist produces a technical description of the problem, data descriptions, and a deployment strategy according to the business requirements. Any challenges will be communicated with the business stakeholder to refine the goals and scope if necessary.

Project Validation stage ascertains if the project is feasible and can meet the expected requirements. The ML problem is the input for the validation task which needs to be performed in the strategic and operational dimension. The business stakeholder, project manager, and legal representative conduct the business case analysis resulting in a business case report. The data scientist and solution architect carry out technical assessments. The technical assessment yields data understanding and the infrastructure specifications describing the needed infrastructure for data processing, model training and orchestration.

Project Setup stage lays the basis for the implementation. The strategic activities are concerned with the collaboration creation if the project is implemented in a collaborative setting. Hence, this strategic activity the project setup stage is optional. The business stakeholder and project manager aim to establish and plan the collaboration, whereas the legal representative checks the compliance and legal regulations of the collaboration. The collaboration creation activities result in a collaboration agreement, which also specifies the used communications channels of the collaboration. We provide a comprehensive description of the socio-technical aspects which need to be addressed in the collaboration agreement in a separate publication (Müller et al. 2023). The software engineers set up the infrastructure as specified in the infrastructure specifications and technical description. Therefore, the infrastructure for data processing, model training and orchestration emerges from this stage. Optionally, communication channels are established if the project is implemented within a collaboration.

System Design and Development stage aims to create the software product as well as the ML model and is divided into three substages:

(1) *Global Model Design* substage intends to define the source code for the initial global model which is subsequently trained during the following substages. This requires access to existing data storage and log storage for traceability. The data engineer prepares the data as needed by the specified ML problem. The data scientist builds the corresponding ML model, which is prepared by the software engineer for distribution to the local datasets. Additionally, the data engineer defines data specifications rules, which can be followed by the local data engineers to achieve data homogeneity. The data preparations and features can be iteratively adapted dependent on the ML model performance. Finally, the initial model is stored in a central model registry. Summarized, this stage yields logs to ensure traceability, data preparation specifications to achieve data homogeneity and a global model implementation.

(2) *Local Model Training* substage is conducted by the local participating data contributors. The local data engineers extract their local data from their corresponding data storage and prepare the data according to the defined data preparation specifications. After preparation, they pull the latest model from the model registry and train the local model. This results in a local version of the ML model. The gradients from the training process need to be extracted and sent to the central entity. All training logs need to be stored in a local log storage to ensure traceability. This local data and log storage is not public to guarantee privacy and should only be exchanged in exceptional cases e.g., for accountability matters.

(3) *Global Model Aggregation* substage collects and aggregates the resulting update gradients from the local data contributors. A data scientist builds the global model through the collected update gradients. The model is evaluated and stored in the model registry. Logs are stored in a log storage for traceability. If the model performance is insufficient, another training iteration can be triggered. The evaluation results can also be used to adjust the scope and goals or adapt the data specifications and model design. If the model performs as desired, it is packaged by the software engineer for deployment.

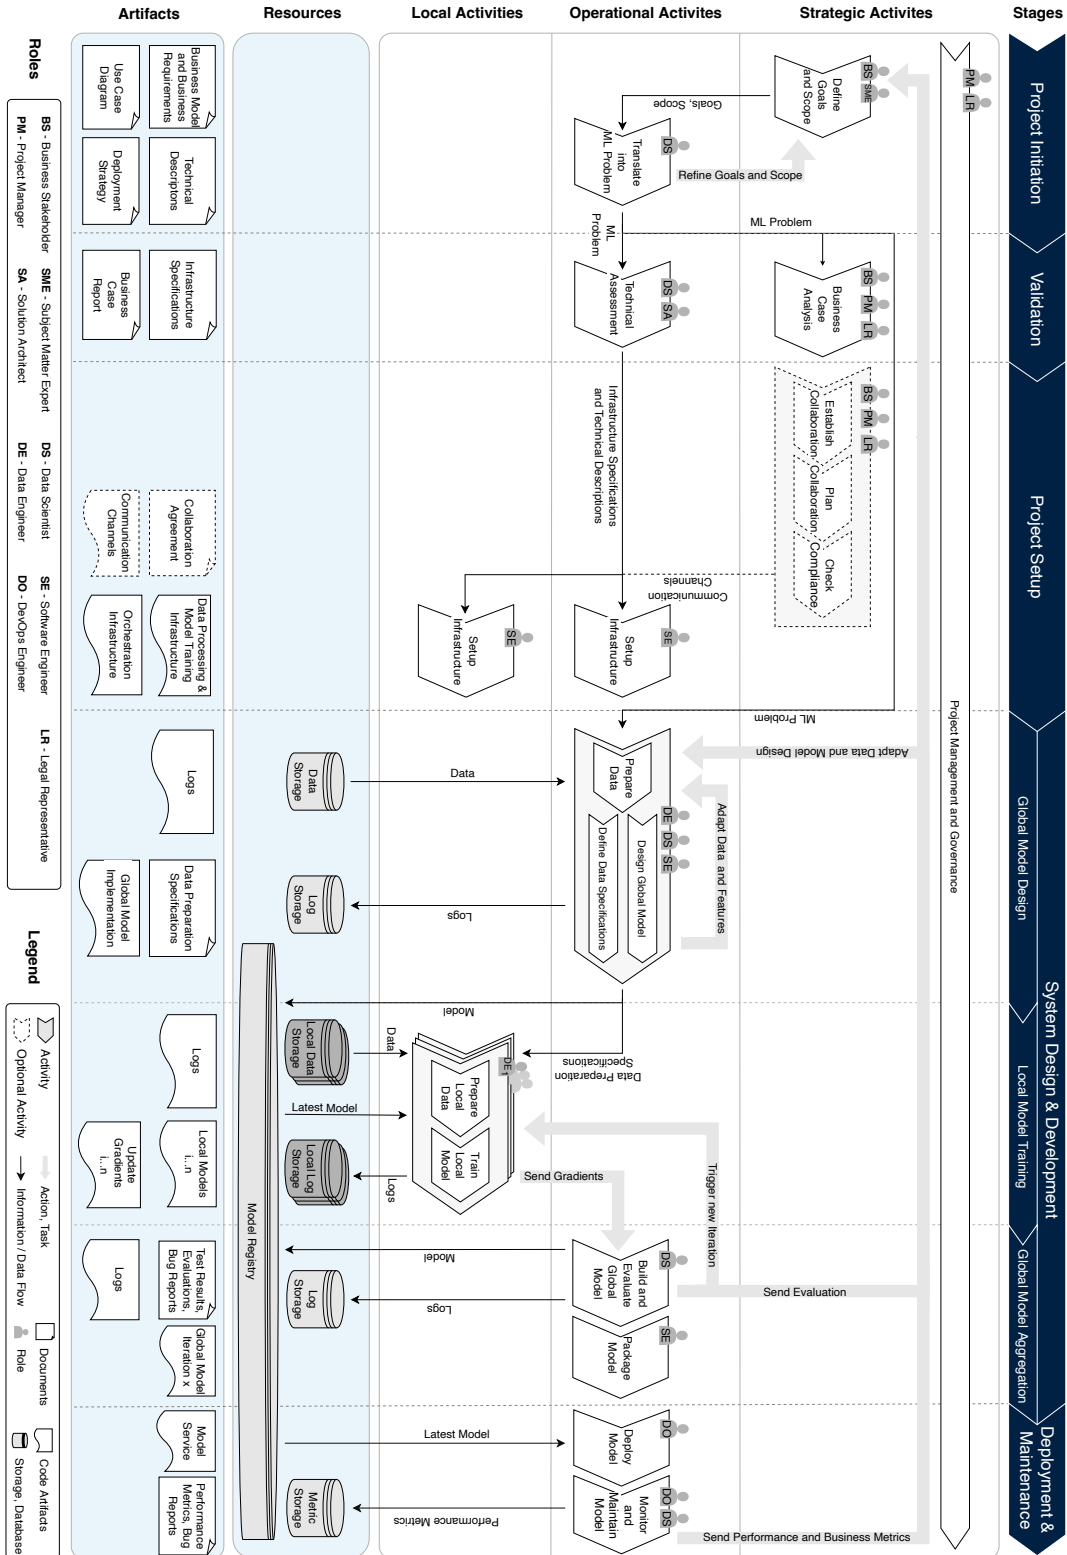


Figure 2. Process Model of an End-to-End FedML Project Life Cycle

Deployment and Maintenance stage intends to serve and maintain the model service. The DevOps engineer deploys the final packaged model. The DevOps engineer and data scientist maintain and monitor the model while storing the recorded metrics in a metric storage. This includes the monitoring report of business metrics. Hence, this stage produces a model service, metrics, and bug reports. If ML service deviates from the desired behavior, the performance metrics can be used to trigger changes in the data preparation scheme and model design as well as trigger an adjustment of the goals and scope.

Demonstration and Evaluation

We conducted two iterations of demonstrations with subsequent survey-based evaluations. The first iteration was part of the *Design and Development* stage as described in the section on the *Research Approach*. We demonstrated the process model and discussed the model with the experts. We received feedback through the discussion and asked all participants to fill out a survey. The survey evaluates the relevance and usefulness of each element with its overall comprehensibility, level of detail, completeness, usefulness, and value. The demonstration was part of a project ideation process to adopt FedML within a large industrial lighthouse project. In total, 14 experts participated in the assessments. Overall, the group of experts comprised a variety of affiliations from startups, big tech companies and research institutions with experience ranging from one to 12 years in their respective field. For each demonstration, all experts confirmed the identified problem and validated the idea, structure, components, and importance of the process model. We incorporated the feedback after each round of discussions.

To evaluate the models' usefulness, we finally presented the refined model to 8 potential users and key stakeholders such as project managers, product specialists, solution architects, and solution advisors. Each participant was involved in prior FedML use case discussions. We aimed to assess if the model can be used in practice, or if further modification is required. One expert stated that the model seems complex at a first glance due to its size. Since the other experts liked that the model displays an entire project life cycle, that the level of detail is fitting, and all components should be included. Based on these statements, we have not made any changes. The involved experts aim to use the model in future discussions.

The overall evaluation results in Figure 3 show that the model was well-received by all participants. A large majority of the participants described the model as detailed enough, complete, comprehensive, and valuable. The model is unanimously considered useful, which implies that the set objective of a useful and understandable model is met. Regarding the components, the stages seem to be relevant for each expert, whereas very few experts deem the artifacts and roles somewhat irrelevant. Three experts stated that the last stage of a project life cycle is missing *Model Maintenance* activities and therefore strongly disagreed that the model covers all important aspects. We agreed and updated our model accordingly. Therefore, we consider the model as presented in Figure 2 as complete. All experts confirmed the structure, comprehensiveness, and level of detail. Thereby, we assume that the objective of a clearly understandable and generically applicable model is met as well. All experts agreed that the process model includes all relevant details and stated that they will use the model in future discussions on FedML with their stakeholders and potential collaborators, which confirms the usefulness for the target user again.

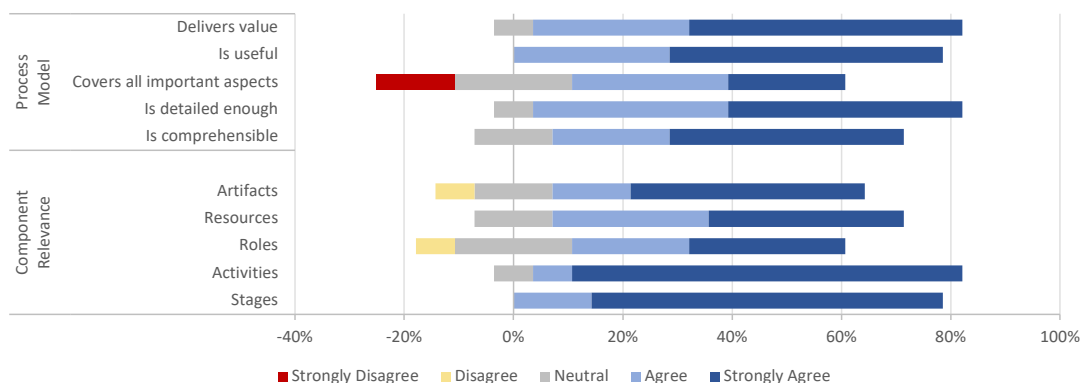


Figure 3. Results of the Evaluation

Conclusion

In this study, we introduced a process model of an end-to-end FedML project to provide transparency and facilitate the comprehensibility of the complex project life cycle. The process model extends current literature on traditional, centralized ML life cycles and complements existing reference architectures through a holistic overview of the FedML project life cycle. By investigating related work, conducting focus group discussions and an interview study, we identified *stages, activities, roles, resources, and artifacts* as the most relevant components for a comprehensible overview of the FedML process (RQ1). Finally, based on our research, we designed and constructed the process model with the corresponding interactions and dependencies of the different components (RQ2). The evaluation results show that potential users deem the model to be useful, deliver value and comprise the most relevant aspects with an appropriate level of detail. The process model can help practitioners to gain an understanding of the structure and interrelations within a FedML project. As stated in the expert evaluations, the model can be used for simplified communication with potential collaboration partners and decision-makers. Additionally, the model can ease the project initiation phase and successful implementation by providing guidance and an overview of the whole project life cycle.

The process model is currently limited to server-client architectures with a central orchestrating server and multiple local data contributors since it is the most widely used architectural pattern. However, more architecture designs are possible and cannot be abstracted by our process model. In the future, adapted process models for novel or different architectures could be established. Also, the model has only been evaluated by 14 experts and within one project ideation phase. We recommend testing and evaluating the model in more scenarios to gather more feedback for improvement. Practical validation might also lead to a more thorough alignment with current MLOps rather than DevOps paradigms to further streamline the model to current best practices. Overall, the process model aids the successful implementation of FedML projects. By that, our model contributes towards more privacy-enhancing ML projects and might help improve the overall performance of ML models through leveraging decentralized and currently untapped dataset. Furthermore, we aim to publish an additional activity model with more detailed activities for each stage of the life cycle to provide further guidance on more in-depth technical aspects of FedML projects.

Acknowledgements

The authors would like to thank SAP SE for supporting this work.

REFERENCES

- Akinsola, J. E. T., Ogunbanwo, A. S., Okesola, O. J., Odun-Ayo, I. J., Ayegbusi, F. D., and Adebisi, A. A. 2020. "Comparative Analysis of Software Development Life Cycle Models (SDLC)," *Intelligent Algorithms in Software Engineering*, In R. Silhavy (Ed.), Springer International Publishing, pp. 310–322.
- Alsaqqa, S., Sawalha, S., and Abdel-Nabi, H. 2020. "Agile Software Development: Methodologies and Trends," *International Journal of Interactive Mobile Technologies (IJIM)* (14:11), pp. 246-270.
- Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B., and Zimmermann, T. 2019. "Software Engineering for Machine Learning: A Case Study," in *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice*, pp. 291–300.
- Anwar, A. 2014. "A Review of RUP (Rational Unified Process)," *International Journal of Software Engineering* (5:2), pp. 8–24.
- Apoorva, M., and Deepty, D. 2013. "A Comparative Study of Different Software Development Life Cycle Models in Different Scenarios," *International Journal of Advance Research in Computer Science and Management Studies* (1:5), pp. 64–69.
- Bauer, M., van Dinther, C., and Kiefer, D. 2020. "Machine Learning in SME: An Empirical Study on Enablers and Success Factors," in *AMCIS 2020 Proceedings*, 3.
- Bharti, S., McGibney, A., and O'gorman, T. 2022. "Design Considerations and Guidelines for Implementing Federated Learning in Smart Manufacturing Applications," *IIC Journal of Innovation* (19:1), pp. 17–35.

- Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konecný, J., Mazzocchi, S., McMahan, H. B., Overveldt, T. V., Petrou, D., Ramage, D., and Roselander, J. 2019. "Towards Federated Learning at Scale: System Design," in *Proceedings of Machine Learning and Systems 1*, Standford, California, pp. 374–388.
- Curtis, B., Kellner, M. I., and Over, J. 1992. "Process modeling," *Communications of the ACM (35:1)*, pp. 75–90.
- Giray, G. (2021). "A Software Engineering Perspective on Engineering Machine Learning Systems: State of the Art and Challenges," *Journal of Systems and Software (180:1)*, pp. 35-97.
- Gurung, G., Shah, R., and Jaiswal, D. 2020. "Software Development Life Cycle Models-A Comparative Study," *International Journal of Scientific Research in Computer Science, Engineering and Information Technology (10:4)*, pp. 30–37.
- Hevner, A. R., March, S. T., Park, J., and Ram, S. 2004. "Design Science in Information Systems Research," *MIS Quarterly (28:1)*, pp. 75–105.
- IEEE. 2021. "IEEE Guide for Architectural Framework and Application of Federated Machine Learning," *IEEE Std 3652.1-2020*, pp. 1–69.
- Kim, M., Zimmermann, T., DeLine, R., and Begel, A. 2018. "Data Scientists in Software Teams: State of the Art and Challenges," *IEEE Transactions on Software Engineering (44:11)*, pp. 1024–1038.
- Kreuzberger, D., Kühn, N., and Hirschl, S. 2022. "Machine Learning Operations (MLOps): Overview, Definition, and Architecture," *IEEE Access (11:1)*, pp. 31866–31879.
- Kumara, I., Arts, R., Di Nucci, D., Heuvel, W. J. V. D., and Tamburri, D. A. 2022. *Requirements and Reference Architecture for MLOps: Insights from Industry*, TechRxiv.
- Laato, S., Birkstedt, T., Mäntymäki, M., Minkinen, M., and Mikkonen, T. 2022. "AI governance in the system development life cycle: Insights on responsible machine learning engineering" in *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*, pp. 113–123.
- Lo, S. K., Lu, Q., Paik, H.-Y., and Zhu, L. 2021. "FLRA: A Reference Architecture for Federated Learning Systems," in *Proceedings of the 15th European Conference on Architecture*, pp. 83–98.
- Lo, S. K., Lu, Q., Wang, C., Paik, H.-Y., and Zhu, L. 2022. "A Systematic Literature Review on Federated Machine Learning: From a Software Engineering Perspective," *ACM Computing Surveys (54:5)*, pp. 1–39.
- Lo, S. K., Lu, Q., Zhu, L., Paik, H., Xu, X., and Wang, C. 2022. "Architectural Patterns for the Design of Federated Learning Systems," *Journal of Systems and Software*, (191:3).
- McMahan, H. B., Moore, E., Ramage, D., and Arcas, B. A. y. 2016. *Federated Learning of Deep Networks using Model Averaging*, CoRR, abs/1602.05629.
- Mendling, J., Reijers, H. A., and van der Aalst, W. M. P. 2010. "Seven process modeling guidelines (7PMG)," *Information and Software Technology (52:2)*, pp. 127–136.
- Müller, T., Zahn, M., and Matthes, F. 2023. "Unlocking the Potential of Collaborative AI - On the Socio-Technical Challenges of Federated Machine Learning," in *ECIS 2023 Proceedings*, Kristiansand.
- Peppers, K., Tuunanen, T., Rothenberger, M., and Chatterjee, S. 2007. "A design science research methodology for information systems research," *Journal of Management Information Systems (24:1)*, pp. 45–77.
- Ritz, F., Phan, T., Sedlmeier, A., Altmann, P., Wiegardt, J., Schmid, R., Sauer, H., Klein, C., Linnhoff-Popien, C., and Gabor, T. 2022. "Capturing Dependencies within Machine Learning via a Formal Process Model," in *Proceedings of the 11th ISO/IEC JTC1/SC32 International Symposium on Software Engineering and Measurement*, Rhodes, pp. 249–265.
- Schomakers, E.-M., Lidynia, C., and Ziefle, M. 2020. "All of me? Users' preferences for privacy-preserving data markets and the importance of anonymity," *Electronic Markets (30:3)*, pp. 649–665.
- Serban, A., van der Blom, K., Hoos, H., and Visser, J. 2020. "Adoption and Effects of Software Engineering Best Practices in Machine Learning" in *Proceedings of the 14th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement*, Torino, pp. 1–12.
- Studer, S., Bui, T. B., Drescher, C., Hanuschkin, A., Winkler, L., Peters, S., and Müller, K.-R. 2021. "Towards CRISP-ML(Q): A Machine Learning Process Model with Quality Assurance Methodology," *Machine Learning and Knowledge Extraction (3:2)*, pp. 392–413.
- Too, E. G., and Weaver, P. 2014. "The management of project management: A conceptual framework for project governance," *International Journal of Project Management (32:8)*, pp. 1382–1394.
- Zhang, H., Bosch, J., and Olsson, H. H. 2020. "Federated Learning Systems: Architecture Alternatives," in *27th Asia-Pacific Software Engineering Conference (APSEC)*, Singapore, pp. 385–394.
- Zwikael, O., and Meredith, J. R. 2018. "Who's who in the project zoo? The ten core project roles," *International Journal of Operations & Production Management (38:2)*, pp. 474–492.

A Pathway for the Practical Adoption of Federated Machine Learning Projects

Completed Research Paper

Tobias Müller

Technical University of Munich
and SAP SE
Munich, Germany
tobias1.mueller@tum.de

Milena Zahn

Technical University of Munich
and SAP SE
Munich, Germany
milena.zahn@tum.de

Florian Matthes

Technical University of Munich
Munich, Germany
matthes@tum.de

Abstract

Big data forms the fundamental basis for the success of Machine Learning. Yet, a large amount of the world's digitized data is locked up in data silos, leaving its potential untapped. Federated Machine Learning is a novel Machine Learning paradigm with the potential to overcome data silos by enabling the decentralized training of Machine Learning models through a model-to-data approach. Despite its potential advantages, most Federated Machine Learning projects fail to actualize due to their decentralized structure and incomprehensive interrelations. Current literature lacks clear guidelines on which steps need to be performed to successfully implement Federated Machine Learning projects. This study aims to close this research gap. Through a design science research approach, we provide three distinct activity models which outline required tasks in the development of Federated Machine Learning systems. Thereby, we aim to reduce complexity and ease the implementation process by guiding practitioners through the project life cycle.

Keywords: Federated Machine Learning, Activity Model, Software Engineering, AI.

Introduction

The success of Machine Learning (ML) systems roots in the emergence of big data and the ever-increasing availability and wealth of digitized information. Even though data forms the fundamental basis for powering ML systems, it also poses ML's major bottleneck. Problem domains become increasingly complex, which results in more sophisticated and therefore data-demanding ML systems. The development of such advanced, intelligent systems is often restricted through a lack of sufficient training data. Especially small and medium-sized businesses (SMEs) experience this problem and suffer from a deficiency of training data (Bauer et al., 2020). Although a considerable amount of data is freely available, vast amounts of the world's data is scattered in decentralized IoT devices and data silos. The siloed data is usually hardly accessible to external prospective parties, which leaves a significant portion of generated data largely untapped. By breaking up these data silos through collaboration and data sharing, SMEs could overcome the persisting problem of data scarcity. Thereby enabling the development of complex ML systems which were not within the realms of possibility before. However, companies are reluctant to share data due to privacy concerns and a potential loss of intellectual property (IP) (Schomakers et al., 2020). Moreover, the constrained usability of these data silos is additionally strengthened by data protection laws and regulations. Important legal regulations such as the General Data Protection Regulation (GDPR), California Consumer Privacy Act,

Cyber Security Law and the General Principles of the Civil Law justifiably aim to protect the privacy of individuals, but also lead to more data silos.

Motivated by this problem McMahan et al. (2016) introduced Federated Machine Learning (FedML), a novel ML approach which enables the training of a joint ML model on distributed datasets without the need for direct sharing training data. Through a model-to-data approach, the ML model is brought to the decentralized data, therefore data never needs to leave the individual's device which enhances privacy by design. By that, companies would be able to collaboratively train a joint ML model without the risk of losing their individual IP and the need of disclosing any data. Hence, FedML technically yields the potential of alleviating the lack of sufficient training data by enabling ML training across company borders and decentralized data silos.

However, despite its advantages, most FedML projects fail to actualize and never leave the prototype or simulation stage (Lo, Lu, Wang, et al., 2022). The reason for the lack of production ready FedML systems may be attributable to multiple aspects. Even integrating centralized ML algorithms into traditional software systems is cumbersome due to the non-deterministic behavior of ML systems, which collides with deterministic software engineering practices (Giray, 2021). Additionally, the decentralized nature of FedML introduces a further dimension of complexity. Among other things, a collaboration may need to be established and multiple parties need to be coordinated throughout the whole project life cycle (Wouters et al., 2017). Through focus group discussions and expert interviews, we recognized that the realization of FedML projects is currently hindered by missing clarity over the multi-faceted process flow. The currently incomprehensible project structure especially impedes product owners and project managers in the project planning and communication with participants. This challenge could be alleviated through a comprehensive step-by-step guide that clearly outlines the needed tasks of implementing an entire FedML project. Current academic and non-academic literature does not address this issue. We aim to close this research gap by providing activity models which describe the sequence of activities needed to successfully implement FedML applications. By this, we aim to provide clarity and guidance for practitioners throughout the project life cycle and thereby aid to facilitate the development of FedML projects. Summarized, we aim to achieve this goal by answering the following research questions (RQs):

RQ 1: What are the required activities for the implementation of Federated Machine Learning projects?

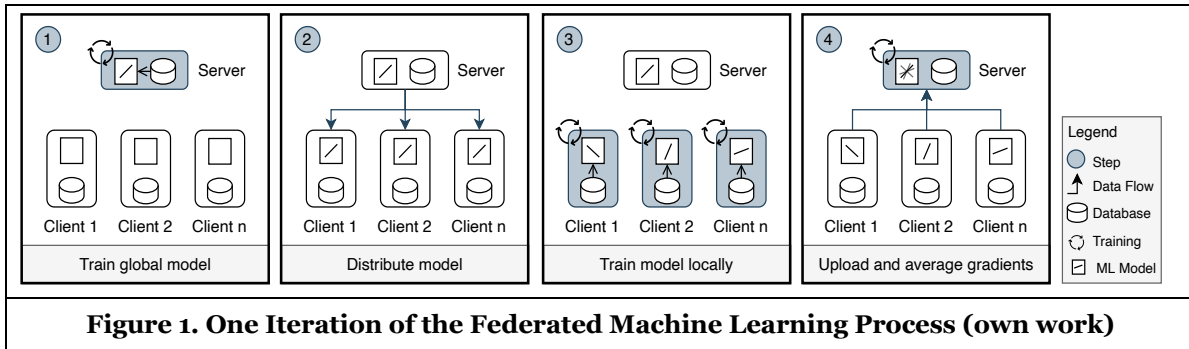
RQ 2: How can a structured model for guiding practitioners through the implementation of Federated Machine Learning projects be designed?

Federated Machine Learning

FedML is a novel, disruptive ML technique that enables the training of a joint ML model on distributed datasets without the need of sharing data. In traditional ML settings, data is usually accumulated in a central location, where the ML model is subsequently trained. Hence, data owners need to share their data with a central server, therefore potentially risk losing their data sovereignty and IP. Introduced by McMahan et al. (2016), FedML leverages a model-to-data approach and thereby alleviated the need of sharing datasets. As visualized in Figure 1, the FedML process can be divided into four steps:

1. The server chooses an initial global model, which is suitable for the use case and underlying data structure. This global model can be initially trained on an initial dataset.
2. The server distributes the global model amongst all clients.
3. Each client trains the global model on its own local dataset and stores the update gradients. After the training process, each client has its individually trained ML model based on its local dataset.
4. The clients send the individually computed update gradients back to the server. These gradients are aggregated based on a pre-defined protocol and used to update the global model.

Steps 2-4 can be repeated over several iterations until a certain accuracy level is reached or until the accuracy converges. A multitude of different FedML architecture have been proposed, however, in this work we solely refer to server-client architectures consisting of a central orchestration server and multiple clients, since this architectural pattern is the most widely used (Lo, Lu, Zhu, et al., 2022).



Related Work

IEEE published a reference architecture with generalized information about the structure of a FedML model training process and high-level descriptions of activities related to defined roles (IEEE, 2021). The reference can be used as a basis for implementation, but it neither provides information about the life cycle nor detailed insights into the activity descriptions and activity sequences. Considering all life cycle stages is crucial and an existing open problem for the implementation of industrial-level FedML systems (Zhang et al., 2020). Also, it is needed to illustrate the chronological sequence as well as the interrelations of the required activities in order to successfully guide practitioners in the development. Opposed to the existing literature on FedML, detailed activity descriptions were proposed in numerous studies on life cycle models of traditional, centralized ML. For example, Studer et al. (2021) introduced CRISP-ML(Q), an ML process model which covers the full life cycle from business understanding to monitoring and maintenance. The authors followed the principles of the Cross-Industry Standard Process Model for Data Mining (CRISP-DM) (Wirth & Hipp, 2000) and constructed their ML process model with a focus on quality assurance. The development activities were organized in six phases and they described the needed activities briefly. Kreuzberger et al. (2022) developed a workflow for ML operations (MLOps) frameworks including functional components and roles. They included activity descriptions and activity sequence as part of their proposed model. Similarly, Kumara et al. (2022) proposed a reference architecture for MLOps to aid in streamlining the life cycle of ML models in production.

Work on software engineering for ML yields important insights into ML life cycles with their corresponding sequence of activities. Amershi et al. (2019) investigated how software teams at Microsoft developed AI applications. Based on their observations, they analyzed and grouped activities which are performed by the teams to construct a life cycle model with a set of best practices. The study describes the activities and illustrates a workflow of the nine stages. Besides, research on AI governance looks into the activities throughout the ML life cycle. Laato et al. (2022) conducted expert interviews to explore the incorporation of AI governance into system development life cycle models. Their study resulted in a set of governance concepts and descriptions of how AI governance stages tie into existing software development life cycles. Even though, the authors did not describe any needed activities per se, the study provides valuable information for the construction of the activity diagrams through their insights on existing best practices during the different stages of an ML project. Additionally, based on that, Ritz et al. (2022) present a process model to illustrate the dependencies and interactions within the different life cycle stages and activities. They described the activities throughout the whole life cycle including the interdependencies throughout the software development life cycle. These existing life cycle models for centralized ML systems can be used as a basis for FedML projects but need to be revisited due to the decentralized nature of FedML.

Some well-defined activity models for traditional, centralized ML models have been proposed for a multitude of application domains. Koshtura et al. (2020) proposed an UML diagram, which illustrates the set of required activities for the implementation of ML-based demand analysis systems to forecast the bicycle use in smart cities. Thanachawengsakul et al. (2019) constructed an activity diagram, which describes implementation of a knowledge repository management system architecture which uses ML, whereas Venugeetha et al. (2022) illustrated an activity diagram for ML-based breast cancer prediction. Again, these models are heavily tailored to their specific use cases and are intended for centralized ML

systems. Since the decentralized FedML process introduces additional layers of complexity, it is needed to revisit activity models and address FedML specifics for a successful industrial-level implementation.

Research Approach

We observed that many FedML projects do not evolve past prototypes. To understand why FedML projects fail to actualize, we conducted a focus group discussion with three project teams that attempted realizing FedML projects in the past. Through this focus group discussion, we aimed to understand the encountered challenges during the project. We noticed that the main barrier lies in the project initiation and planning phase. The project teams reported that it is challenging to structure the complex implementation process due to its decentralized and collaborative nature. The project flow is complicated to understand, and the division of tasks is not clear. Therefore, communicating the tasks throughout the team and potential collaborators is arduous. This barrier of intricated process structure and difficult communication seem to pose the main challenge in building FedML products. The focus group agreed that a process model, which gives a holistic overview of the process and a detailed activity model with a comprehensive sequence of tasks would alleviate the barrier. Both artifacts would help provide transparency and guidance throughout the development of FedML projects.

At this point, is important to emphasize the differences between the activity model and the process model in terms of their target group and intended purpose. The process model aims to provide a highly abstracted overview of the holistic project structure including required resources, role distributions, and resulting artifacts. This high-level overview intends to aid business stakeholder and solution architects in the project initiation phase outline the general project and facilitate communication with potential participants. The activity model, on the other hand, provides a detailed activity sequence through the needed steps of implementing a FedML project. The comprehensive activity descriptions aim to provide guidance for the implementation and therefore intends to support product owner and project manager in the planning phase of the development process. Therefore, while the process model aims to facilitate the project initiation phase for business stakeholder and solution architects, the activity model intends to ease the planning phase for product owner and project manager. In this paper, we will focus on the comprehensive activity model. We plan to communicate the process model¹ in a separate publication since both artifacts are standalone and need to be considered independently for their respective purpose.

Current literature does not offer a structure representation of a FedML project flow and does not provide guidance on the needed sequence of activities required to realize a FedML project in practice. Therefore, we aim to close this research gap by providing a detailed activity model which illustrates the tasks and interrelations of each life cycle stage in a FedML project. To develop such an activity model, we leveraged the design science research (DSR) methodology as proposed by Peffers et al. (2007) since it provides a methodical, rigorous approach for producing and evaluating innovative, purposeful artifacts for a specific problem domain (Hevner et al., 2004). Table 1 provides an overview of our research approach and short description of the conducted activities during the DSR cycle. The remainder of the paper is structured according to the steps of the DSR approach.

Step	Short Description of Activities
(1) Problem identification and motivation	Identified the problem and motivation through focus group discussions. See description above.
(2) Objectives of a solution	Conducted focus group discussions to derive requirements and determine relevant design principles. See chapter on Objectives of a Solution
(3) Design and development	Designed and developed the artifact to provide transparency and guide practitioners through a successful FedML project implementation. See chapter on Design and Development

¹ Process Model: <https://bit.ly/3IHlRAn>

(4) Demonstration	Demonstrated the artifacts in group discussion and during project ideation phase of an industrial lighthouse project. See chapter on Demonstration and Evaluation
(5) Evaluation	Evaluated the comprehensibility, completeness, usability, and value of the artifacts. See chapter on Demonstration and Evaluation.
(6) Communication	Communication is being done through this paper.
Table 1. Design Science Research Steps According to Peffers et al. (2007)	

Results

Objectives of a Solution

The objectives and requirements of the solution were identified through the initial focus group discussion as described in the section on the research approach. The three project teams agreed that the activity model should aid specifically in the planning phase and provide a clear structure over the sequence of activities which are needed to successfully implement FedML projects. Therefore, the activity models should cover the entire end-to-end life cycle of a FedML project. Additionally, the model should help in the communication with non-technical stakeholders and practitioners. Hence, the model should be easily consumable and provide transparency as well as guidance throughout the whole life cycle. The contents should be closely aligned and consistent with best practices of ML life cycles, software development life cycles, and state of the art FedML practices. Lastly, the activity model should be usable for independent of the specific FedML strategy and applicable for all FedML processes with a centralized orchestrator and multiple, distributed clients. The activity model should be applicable independent of the underlying use case, application domain, data structure, and business requirements. The requirements to the artifacts are summarized and indexed in Table 2.

ID	Short Description of the Requirement
R1	Aligned with best practices from ML and software development
R2	Comprise an end-to-end project life cycle
R3	Generically applicable and independent of the use case
R4	Understandable and usable by non-technical stakeholders and practitioners
R5	Provide transparency and structure of the entire project
R6	Provide guidance for the implementation process
Table 2. Identified Requirements on the Activity Models	

Design and Development

To build our knowledge base to meet requirements R1 and R2, we reviewed current literature on ML life cycles (Amershi et al. 2019; Kreuzberger et al. 2022; Kumara et al. 2022; Laato et al. 2022; Ritz et al. 2022), software development life cycles (Akinsola et al., 2020; Alsaqqa et al., 2020; Apoorva & Deepty, 2013; Gurung et al., 2020) and assessed which practices, procedures, and information are applicable for the activity model and its structure. Additionally, we reviewed the state-of-the-art FedML processes through current literature on FedML architectures and algorithms (Bharti et al., 2022; Bonawitz et al., 2019; IEEE, 2021; Lo et al., 2021; Lo, Lu, Zhu, et al., 2022; Zhang et al., 2020). The activity models are partially based on an expert interview study on the socio-technical challenges of FedML projects, which we describe in a separate publication (Mueller et al., 2023). The study especially provided information for strategic activities in a collaborative setting. The relevant findings from the interview study, focus group discussions and

literature review were incrementally combined. During development, we conducted regular mini focus group discussions with varying participants to iteratively assess the activity models and implement feedback.

Per definition, an *activity* represents units of work that are performed by roles. Each activity has a clear purpose and usually results in the creation or update of an artifact (Anwar, 2014). Therefore, the activity models should provide guidance on which units of work need to be followed such that the goals (or artifacts) of the stages can be successfully achieved. However, some activities and their corresponding method of execution may be dependent on technicalities, use case, or business requirements. To ensure requirement R3, we want to leave the choice of the fitting method to the practitioner. Hence, we chose the granularity, such that the goal and required activities is comprehensible, but does not preset specific context-dependent methods. We only included the happy path without exit points and did not incorporate the dependencies and interactions between the stages to enhance comprehensibility. The activity models show one iteration of each stage. It is important to highlight that in ML projects, the stages are usually performed iteratively to refine the process.

The design of the activity models is based on the UML 2.0 notation² since it is a commonly known and normed unified modelling language in the field of software engineering. By leveraging widely known notations, we aim to meet requirement R4. As specified in UML 2.0, every diagram has at least one initial node, which is indicated by a filled circle, and one end point, which is represented by an encircled filled circle. Activity states are illustrated through ellipses, described through verb-object activity labels, and connected with arrows to represent in which order the activities happen. Decision points are modelled through diamonds, and bars represent the start (split) or end (join) of concurrent activities. If multiple actors interact within one activity model, horizontal swimlanes group the activities performed by the same actor. Objects only show critical inputs/outputs and are modelled through rectangles. If similar objects appear multiple times, the life cycle stage is described as state name in rectangular brackets. Shared objects are placed on swimlane separators.

To meet requirement R2 and R5, we structure the activity models according to the life cycle stages of a FedML project, whereas each stage represents a set of interrelated tasks and activities which serve a clear purpose. Figure 2 provides an overview of the different life cycle stages. The first three stages comprise strategic activities, starting with the *project initiation* (1), followed by the *project validation* (2) and *project setup* (3) phase. We clustered these three stages into a single activity diagram since these phases are interconnected and constitute the preparation for the implementation. The resulting activity diagram is illustrated in Figure 3. Thereafter, all development activities are grouped in the *System Design and Development* stage, which can be divided into the *Global Model Design* (4), *Local Model Design* (5), and *Global Model Aggregation* (6) stages. Stage 4 builds the technical basis for the FedML training process, and its corresponding diagram is depicted in Figure 4. The actual FedML training process is conducted in Stage 5 and 6. We clustered both stages into a single activity diagram since these two phases collectively constitute the FedML training process. The activity diagram is illustrated in Figure 5. Finally, the *Deployment and Maintenance* (7) stage serves and maintains the ML product and is shown in Figure 6. The figures display a single iteration of each stage. In practice the stages 4 to 7 are iteratively and continuously performed due to the iterative ML process. The following activity descriptions are structured according to the life cycle stages and describe the accompanying activities in relation to the activity diagrams. The descriptions should help to meet requirement R6 and simultaneously answer RQ1 and RQ2.




Figure 2. Stages of a Federated Machine Learning Project

²<http://www.omg.org/spec/UML/2.4.1/Infrastructure/PDF/>

Stage 1: Project Initiation

The first stage represents the initial step in starting a project and lays the foundation of the project by broadly defining the goal of the project and its corresponding high-level requirements. As seen in Figure 3, this stage contains one main activity (*Initiate Project*) with two sub-activities. The first sub-activity defines the project goals and scope. Depending on the company-specific best practices, the goal definition usually requires the specification of a business model, the corresponding business requirements, and the illustration of use case diagrams to demonstrate the business case. The second sub-activity translates the resulting business problem into an ML problem to define the technical descriptions and deployment strategy.

Stage 2: Project Validation

The second stage aims to investigate the feasibility of the project and includes one main activity (*Check Project Feasibility*) with two concurrent sub-activities (see Figure 3). More specifically, the result of the previous stage is taken as input for a simultaneous business case analysis and technical assessment of the project's feasibility. These assessments determine whether the project's goals and scope are technically executable, deliver value as intended, and which steps should be taken to meet them including the overall FedML strategy. Additionally, it is determined whether collaboration partners are needed to successfully implement the project.

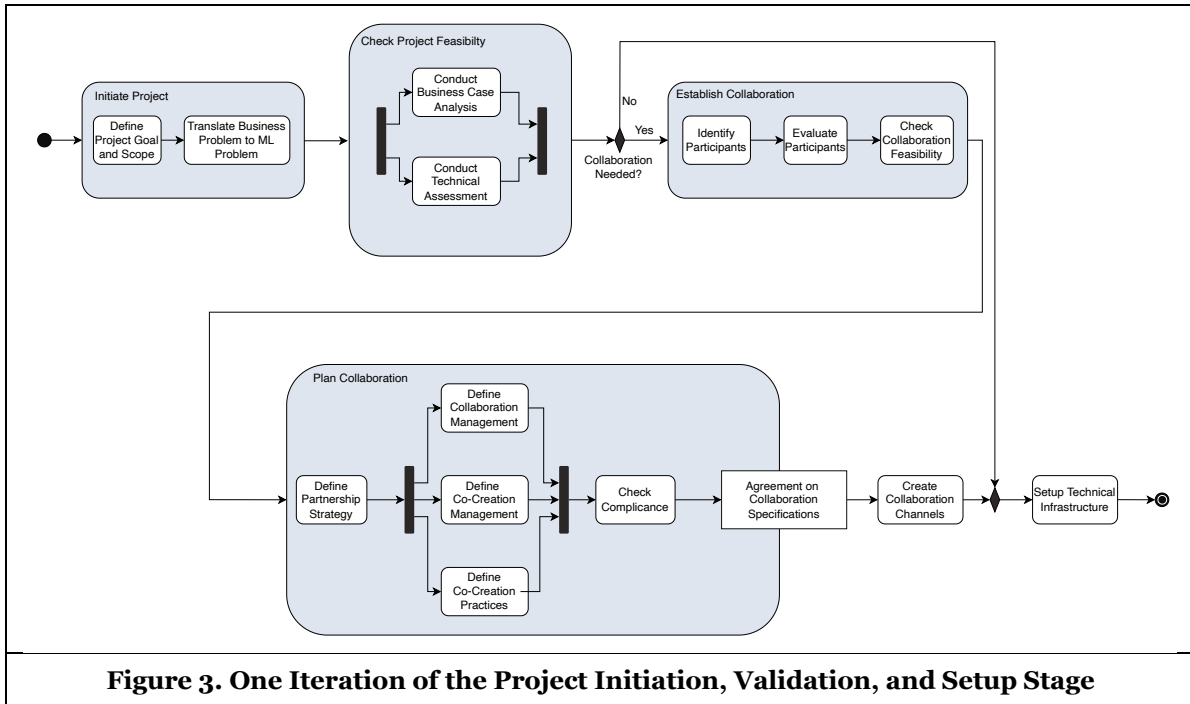
Stage 3: Project Setup

The third stage intends to set up the project by arranging a potential collaboration and setting up the technical foundation for the implementation of the project. Hence, the activities of this stage depend on the decision if collaboration partners are wanted for the project. In a collaborative setup, two additional main activities (*Establish Collaboration* and *Plan Collaboration*) need to be performed, before the technical foundation can be constructed. This stage is depicted in Figure 3.

(3.1) *Establish Collaboration* will create the collaboration and consists of three sub-activities. First, potential collaboration partners need to be identified. Then, it needs to be evaluated if the potential participants would be interested in joining the project. Lastly, the collaboration feasibility needs to be checked, to assess whether the project can be implemented jointly. This step needs to determine if all sides agree on the same objectives and assess on a high level if all participants could deliver the expected value, resources, or expertise. If the collaboration structure cannot perform as intended, all three steps need to be repeated until the collaboration can be established or until a termination decision is made.

(3.2) *Plan Collaboration* will plan the joint project and comprises a total of five sub-activities. To begin the collaboration planning, the overall partnership strategy needs to be described to specify the common goal and underlying partnership structure. Subsequently, the collaboration management, co-creation management, and co-creation practices are defined. The aspects which need to be covered in this step are multi-faceted and comprehensive. We describe the specifics which need to be covered in the collaboration management, co-creation management, and co-creation practices in a separate publication (Mueller et al., 2023). Thereafter, legal compliance needs to be checked. If the compliance check fails, a refinement of the collaboration planning activities is necessary, or termination decision could be made. If all activities were successfully executed, an agreement on the collaboration specifications emerges.

If a collaboration is needed and comes to fruition, all collaboration channels which are needed for the orchestration, communication, and collaboration management are created. Lastly, the required technical infrastructure needs to be set up in all cases. This activity lays the technical basis for the implementation by creating the data processing, model training, and orchestration infrastructure.



Stage 4: Global Model Design

The fourth stage yields the data specifications and source code of the initial ML model which will be the technical foundation of the FedML training process. The activities comprise an ML experimentation task as well as data preparation specifications and can be divided into three distinct main activities. As illustrated in Figure 4, this stage starts with the *Prepare Data* activity and is followed concurrently by the *Design Global Model* and *Define Data Specifications* activities.

(4.1) *Prepare Data* represents the data preparation step of traditional ML pipelines and consists of four sub-activities. It requires access to a data storage which contains a representative dataset for the targeted ML problem. The data is fetched from the data storage and subsequently analyzed on the underlying data distribution, data quality, and overall suitability for the business requirements. Thereafter, the data is preprocessed such that high data quality is ensured and the ML model can interpret the data's features. This task involves data cleaning, data transformation, data reduction, and feature engineering. The final data validation activity ensures the correctness, usefulness, and sufficient quality of the data.

(4.2) *Design Global Model* crafts the initial source code for the initial model which will be used in the FedML training process and consists of four sub-activities. An initial proposal of an ML architecture is trained on the data from step 4.1. The resulting ML model is evaluated on its performance based on the evaluation data and validated on the validation data. Based on the evaluation and validation results, it is determined whether the accuracy suffices for an initial distribution to the local clients for training. Since the model is further trained in the FedML process, this step can be seen as a further feasibility check and should validate the suitability of the ML model for the use case. If the criteria are not met, then the model is analyzed, and another design iteration is triggered. If the criteria are met, the initial ML model is stored in a model registry which is accessible to all participants.

(4.3) *Define Data Specifications* yields the data preparation specifications which are used by every local data contributor (or client) as a data preparation guideline to ensure homogeneous, suitable training data for the FedML process. Therefore, the data cleaning, data transformation, and data reduction rules must be defined and documented. Additionally, the feature engineering rules are defined. The resulting specification document is then communicated with each client.

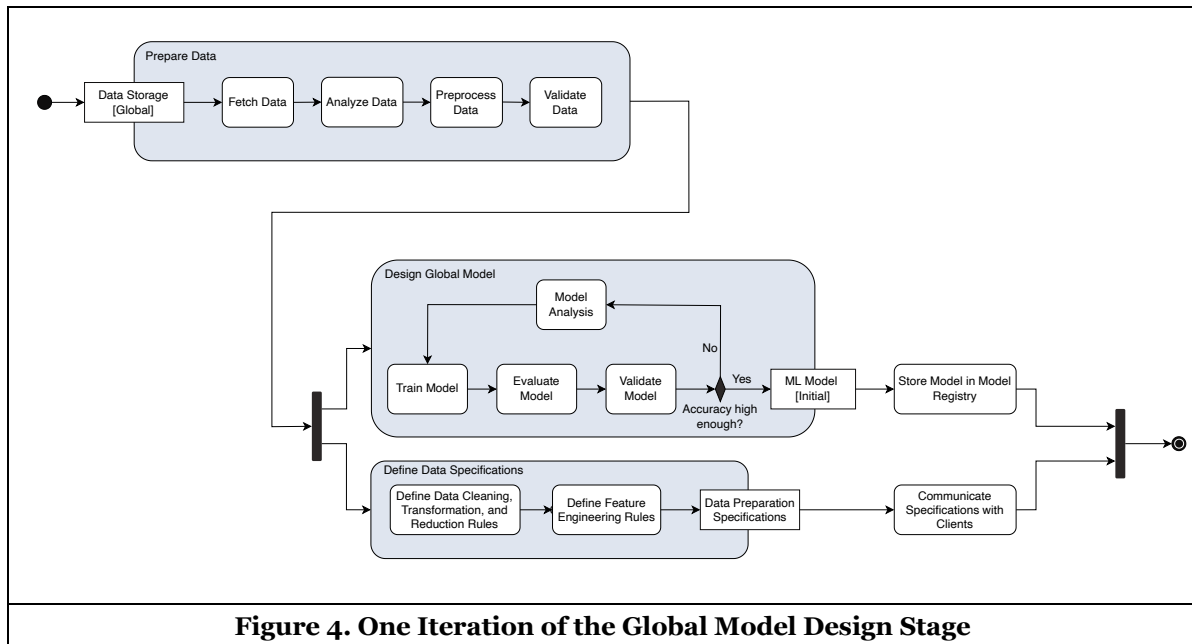


Figure 4. One Iteration of the Global Model Design Stage

Stage 5: Local Model Training

The fifth stage delineates the local training on the data contributor’s (clients) side. As indicated by the swimlanes (*client 1, ..., client n*) in Figure 5, every data contributor performs the same activities simultaneously, which consists of one main activity and a total of five sub-activities. Each data contributor requires access to the model registry and their individual data storage. The different clients fetch the latest model from the model registry, and concurrently prepare the local data. To prepare the local data, it is needed to first fetch the data from their data storage. Thereafter, the data is preprocessed as outlined in the data preparation specifications (see Stage 4). The data preparation specifications ensure homogeneous datasets across all clients. Naturally, in settings that allow heterogeneous data, this step can be neglected. Once preprocessed, the data is then validated on the correctness, usefulness, and sufficient quality of the data. Finally, the latest ML model is trained on the prepared data, resulting in update gradients, which are shared with the aggregator.

Stage 6: Global Model Aggregation

The sixth stage details the aggregation process of the FedML training and is structured into one main activity with six sub-activities (see Figure 5). Once enough update gradients from the local data contributors are collected, the global model can be built and evaluated. The required number of update gradients depends heavily on the use case, business requirements, or scale of the project. The aggregator selects a set of the received upgrade gradients, which are used to train the model. The gradient selection process can either follow a pre-defined sub-sampling scheme or be neglected and simply use every received gradient. This is dependent on the pre-defined FedML strategy. The selected gradients are subsequently aggregated and fused with the latest global ML model. The updated global ML model is then stored in the model registry for traceability and to make the model accessible for the involved clients. Then, the model is validated. Based on the validation results, it is determined whether the model needs to be further trained by triggering a new training iteration or if it can be packaged for deployment.

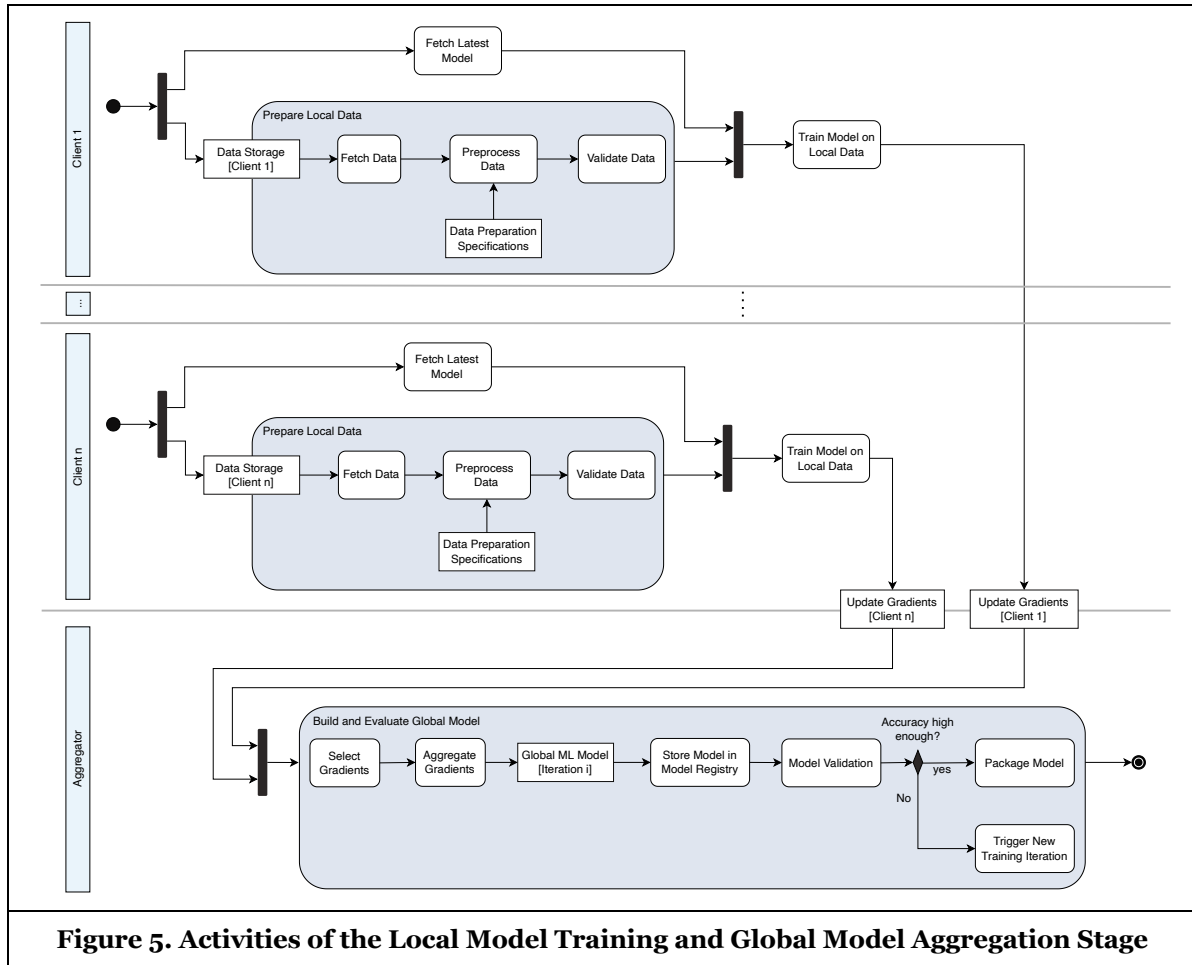


Figure 5. Activities of the Local Model Training and Global Model Aggregation Stage

Stage 7: Deployment and Maintenance

The seventh stage deploys the packaged ML model into production and yields the model service. As illustrated in Figure 6, this stage consists of seven distinct activities. First, the latest ML model is fetched from the model registry and build for deployment. Then, integration testing is performed to ensure that the built code will work as intended. If the tests are successful, the built model will be deployed and served. Again, the specific methods for building and integration testing are dependent on the use case, business requirements, or simply on the developers' preferences. Once the model is served, it is continuously maintained and monitored on its performance, potential deviations from the desired behavior, and pre-defined business key performance indicators. If the model does not perform as desired for example due to concept drift, covariate shift, or simply if an updated ML model version is available, the currently deployed model needs to be sunsetted.




Figure 6. One Iteration of the Deployment and Maintenance Stage

Demonstration and Evaluation

To demonstrate and evaluate our activity model, we performed two iterations of demonstrations with survey-based evaluations. After each demonstration, we received feedback through an expert discussion and surveys, which were additionally filled out by each expert. The survey evaluates if the model is deemed valuable, detailed enough and additionally gathers feedback on the overall usefulness, comprehensibility, and completeness of the activity model. The survey results are shown in Figure 7. We incorporated the feedback after each round of demonstration and evaluation. In total, 14 experts from eight organizations participated in the assessments and provided valuable recommendations. An anonymized list of participants including their current position, their organization, and years of experience in their current position is shown in Table 4.

ID	Position	Organization	Experience
E1	Applied Researcher	Emerging Tech Start-up	> 2 years
E2	Applied Researcher	Industrial Software Enterprise	> 2 years
E3	Applied Researcher	Research Center for AI Security	> 1 year
E4	Research Engineer	Large Engineering Company	> 5 years
E5	ML Engineer and Senior Data Scientist	Large Software Enterprise	> 7 years
E6	Senior Researcher	Large Software Enterprise	> 7 years
E7	Senior Product Manager	Large Software Enterprise	> 6 years
E8	Product Manager	Large Software Enterprise	> 10 years
E9	Senior Consultant and Project Lead	Large Software Enterprise	> 6 years
E10	Solution Advisor	Large Software Enterprise	> 2 years
E11	Development Expert	Large Software Enterprise	> 19 years
E12	Consultant for Emerging Tech	Medium-sized Consultant Company	> 2 years
E13	Solution Architect	Large Software Enterprise	> 5 years
E14	Project Manager	Research Center for AI Security	> 1 year

Table 4. Overview of Evaluation Participants.

In the **first demonstration**, we introduced the activity model to a total of 6 experts (E1-E6). The main focus of this iteration was to evaluate the technical completeness, comprehensibility, and level of detail. Against this backdrop, we invited a diverse group of experts with technical backgrounds. The group comprised three experts on FedML and three applied researchers with expertise in related, emerging technologies but without extensive knowledge on FedML. Therefore, the iteration can be considered as an expert evaluation (Peffer et al., 2012). All experts validated the research problem and confirmed the completeness, comprehensibility, and level of detail. Additionally, the expert round experienced the activity model as a useful and valuable resource for their activities and communications with stakeholders. We received two recommendations to improve the activity model. One expert (E6) mentioned that large-scale FedML projects include an additional gradient selection step, which sub-samples the received gradients. We incorporated the feedback by adding a *Select Gradients* activity as the first task of the *Global Model Aggregation* stage. Another expert (E2) proposed that an explanation of the different steps in the *Plan Collaboration* and *Establish Collaboration* phase would be helpful. We did not adapt the activity model since these activities are dependent on the use case and business requirements.


Overall, the activity model has been received very positively. The participants stated that they will use the activity model for their future activities and discussions. More specifically, the experts remarked that the activity model will facilitate their communication with stakeholders since *“this is very much what decision-*

makers would require from practitioners in order to understand what needs to be done [...]” (E2). Additionally, expert E4 stated that the model “[...] helps [him] to structure a project for privacy-preserving FedML because [he] would know where to apply privacy-enhancing technologies” (E4). Expert E6 validated the usefulness and wished for a similar activity model for other emerging technologies. Summarized, the models aid practitioners in communication with their stakeholders, provide a structure for their projects and thereby facilitate the development of their project. The experts expressed that the models are a “very good resource for FedML projects” (E3).

In the **second demonstration**, we presented the refined activity model to a group of 8 potential users and key stakeholders (E7-E14). The demonstration was part of a project ideation process to adopt FedML within a large industrial lighthouse project. The main focus of this iteration round was to evaluate the models’ usefulness, comprehensibility, and value to potential users. The group of participants included a variety of potential users such as project managers, product owners, solution architects, and solution advisors. Each participant was involved in prior FedML use case discussions and therefore well-suited for this evaluation round. The group of participants reported that the activity model is intuitive to follow, helps to structure the overall process flow, and provides guidance for project implementation. Based on the feedback, we added *Maintain Deployed Model* and *Sunset Deployed Model* tasks in the *Deployment and Maintenance* stage. Besides, the activity model was considered a comprehensible, complete, and comprehensive artifact. The group further agreed that the model is useful and provides value in future activities and discussions.

Overall, the target users from the evaluation group validated the relevancy and effectiveness of the models. Expert E12 specifically emphasized that “[...] the models are intuitive to follow” (E12) and aid in “[...] understanding the FedML process” (E12). The potential users described the models as “a very useful visual tool” (E11) that helps in planning FedML projects due to its “[...] detailed sequence of activities and areas where you can keep focus on” (E11). This makes the complex FedML process for “[...] non-technically proficient stakeholders easier to comprehend” (E14). One target user stated that the models “[...] help make better-informed decisions as a newbie to FedML” (E9). According to the evaluation group, the models provide transparency, and aid non-technical stakeholders to structure and gain an understanding of FedML projects.

As illustrated in Figure 7, the evaluation results show that the model was overall well-received by all 14 experts. The model is considered useful, valuable, and detailed enough. The large majority of the participants experienced the model as comprehensible and that it covers all important aspects.



Discussion

In this study, we designed three distinct activity models which represent the entire life cycle of a FedML project. The problem was identified and motivated through an initial focus group discussion with potential users from three project teams. The relevancy was additionally confirmed in the demonstrations and evaluations. This is also reflected in the survey results since each participant deems the activity models to be useful and valuable (see Figure 7).

Solution Requirements

Through the initial focus group discussion, we were able to formulate a total of six requirements to the activity models (see Table 2). Aiming to fulfill requirement R1, we reviewed and incorporated current literature on ML life cycles, software development life cycles, and state of the art FedML methods. Since technical experts and potential users validated the completeness and usability, we consider that the activity model is aligned with current best practices. Therefore, we consider requirement R1 to be met. The activity models represent the entire life cycle of a FedML project, which was validated on completeness by experts. Hence, requirement R2 should also be fulfilled. Through the two rounds of demonstrations and evaluations, we were able to gather feedback from a large variety of different profiles and backgrounds. The participants ranged from FedML experts and applied researchers to product specialists, projects managers, solution architects, and solution advisors. As reflected in the survey results (see Figure 7), every expert from this diverse set of evaluation participants considers the models valuable, useful, and with a fitting level of detail. Therefore, we consider the activity models as generically applicable and requirement R3 as met. Through the second round of demonstrations, we aimed to evaluate the comprehensibility and usability of potential users and therefore non-technical stakeholders. As reported, the feedback was consistently positive, which leads us to believe, that requirement R4 is fulfilled as well. Combined with the thoroughly positive feedback from the technical experts in the first demonstration, we can assume that the activity models provide transparency and structure to the entire project life cycle (requirement R5). After incorporating the missing aspects suggested by the experts, we consider that the activity models cover the most important aspects. Consequently, the models can be used as guidance for the implementation process and ease the planning process for product owner and project manager, which finally achieves requirement R6. Summarized, we can assume that all objectives have been met to a sufficient extent and that the activity models are useful as well as deliver value to potential users.

Generalizability

The models are currently tailored to server-client architectures with a central orchestrating server and multiple local clients since this is the most widely used architectural pattern (Lo, Lu, Wang et al., 2022). However, more architectural patterns such as completely decentralized processes were suggested (Lo, Lu, Zhu, et al., 2022), which may not be covered by the proposed activity models of stages 5 and 6. Since the architecture choice only influences activities regarding the *local model training* (stage 5) and *global model aggregation* (stage 6), the models of stages 1-4 and stage 7 remain the same and can be used in each FedML project. Additionally, the activities of *local model training* (see Figure 5) are imminent and only the outgoing communication of stage 5 might change depending on the architectural pattern. For partially connected or fully decentralized architectures, clients might communicate amongst themselves, and the aggregation scheme might differ (Lo, Lu, Zhu, et al., 2022). Therefore, the choice of a different architectural pattern might require revisiting the outgoing communication of stage 5 and the activities of stage 6, whereas the remaining activity models can be reused as proposed in this study.

Throughout this study, we focused on designing generically usable activity models that are not specific to use cases or domains. Hence, we only included activities that are required in every FedML project. Moreover, we designed the models such that the need for optional activities is queried (see Figure 3). Thereby, we aimed for generalizable models. To evaluate the generalizability of the models to multiple areas, domains, and target users, we presented the models to a diverse group of potential users. As illustrated in Table 4, the evaluation group consisted of 14 participants, from eight organizations of different sizes and eight different job profiles. Since the large majority of the participants confirmed the comprehensibility and usefulness (see Figure 7), we consider the model to be generically applicable to a large variety of areas. Even though the models were constructed to be generically applicable, some domain-dependent and process-specific activities might be missing. Therefore, the direct application of the models in varying application domains would further validate the generalizability of the models. We encourage practitioners to test our models in their use cases to further develop the artifacts.

Limitations

It should be noted that the activity models were designed to aid in the communication with non-technical stakeholders, provide transparency, and guidance for the implementation process. Consequently, the

models are not intended to be followed down to the smallest detail since the domain, business requirements, or technology specifics might alter certain activities. In this context, we also only included the happy path without exit points throughout the life cycle. The models should be seen as a reference and basis for implementation. Lastly, experiences gathered from practical applications might help to further enhance the models. Even though we gathered insights and feedback for the artifacts from experts with a large variety of different roles and experiences, we conceivably only depict a subset of the potential user demographic. Therefore, we recommend testing the activity models in more diverse settings to receive broader feedback to further improve and develop the models.

One of the main motivational drivers to use FedML is the provided degree of privacy through its model-to-data approach. However, the send gradients could be reverse engineered which might reveal sensitive information (Jere et al., 2020). Further privacy-enhancing technologies could be implemented on top of FedML to mitigate privacy leakages. Our models do not include activities regarding the implementation of privacy-enhancing techniques since these are case-specific and reduce the generalizability of the proposed models. However, as expert E4 stated, the proposed models still aid in the implementation of further privacy-enhancing technologies. Through the detailed step-by-step description of the activity models, expert E4 recognized entry points where privacy-enhancing technologies need to be implemented.

Conclusion

In this study, we introduced three distinct activity models which together depict the sequence of activities that are needed to implement a FedML project. The activity models extend current literature on life cycle models for centralized ML projects and complement FedML reference architectures through activity sequences throughout the whole FedML project life cycle. Thereby, the models provide transparency and structure the entire life cycle of a FedML project to facilitate the comprehensibility of its complex process. The evaluation results show that potential users deem the models to be useful, deliver value and comprise the most relevant aspects with an appropriate level of detail. We showed that the activity models can help practitioners and non-technical stakeholders to gain an understanding of the structure and especially ease the project planning phase by providing guidance over the project life cycle. However, the models are currently limited to server-client architectures with a central orchestrating server and multiple local clients. Other architectural designs may not be covered by our models. Overall, the activity models aid the successful implementation of FedML projects by easing the planning phase for product owners and project managers. Additional to the sequence of activities, current literature on FedML seems to be lacking clarity over the interrelations between the stages, role distribution, and their dependencies throughout the project. We plan to discuss these aspects through a process model in a separate publication.

Acknowledgements

The authors would like to thank SAP SE for supporting this work.

References

- Akinsola, J. E. T., Ogunbanwo, A. S., Okesola, O. J., Odun-Ayo, I. J., Ayegbusi, F. D., & Adebisi, A. A. (2020). Comparative Analysis of Software Development Life Cycle Models (SDLC). In R. Silhavy (Ed.), *Intelligent Algorithms in Software Engineering* (pp. 310–322). Springer International Publishing.
- Alsaqqa, S., Sawalha, S., & Abdel-Nabi, H. (2020). Agile Software Development: Methodologies and Trends. *International Journal of Interactive Mobile Technologies (IJIM)*, 14(11), 246.
- Amershi, S., Begel, A., Bird, C., DeLine, R., Gall, H., Kamar, E., Nagappan, N., Nushi, B., & Zimmermann, T. (2019). Software Engineering for Machine Learning: A Case Study. *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, 291–300.
- Anwar, A. (2014). A Review of RUP (Rational Unified Process). *International Journal of Software Engineering*, 5(2), 8–24.
- Apoorva, M., & Deepty, D. (2013). A Comparative Study of Different Software Development Life Cycle Models in Different Scenarios. *International Journal of Advance Research in Computer Science and Management Studies*, 1(5), 64–69.

- Bauer, M., van Dinther, C., & Kiefer, D. (2020). Machine Learning in SME: An Empirical Study on Enablers and Success Factors. *AMCIS 2020 Proceedings*, 3.
- Bharti, S., McGibney, A., & O'gorman, T. (2022). Design Considerations and Guidelines for Implementing Federated Learning in Smart Manufacturing Applications. *IIC Journal of Innovation*, 19, 17–35.
- Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, H. B., Overveldt, T. V., Petrou, D., Ramage, D., & Roselander, J. (2019). Towards Federated Learning at Scale: System Design. *Proceedings of Machine Learning and Systems*, 1, 374–388.
- Chandrasekaran, V., Jia, H., Thudi, A., Travers, A., Yaghini, M., & Papernot, N. (2021). SoK: Machine Learning Governance. *ArXiv:2109.10870 [Cs]*.
- Giray, G. (2021). A Software Engineering Perspective on Engineering Machine Learning Systems: State of the Art and Challenges. *Journal of Systems and Software*, 180, 35.
- Gurung, G., Shah, R., & Jaiswal, D. (2020). Software Development Life Cycle Models-A Comparative Study. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 30–37.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. *MIS Quarterly*, 28(1), 75–105.
- IEEE. (2021). IEEE Guide for Architectural Framework and Application of Federated Machine Learning. *IEEE Std 3652.1-2020*, 1–69.
- Jere, M. S., Farnan, T., & Koushanfar, F. (2021). A Taxonomy of Attacks on Federated Learning. *IEEE Security & Privacy*, 19(2), 20–28.
- Koshtura, D., Bublyk, M., Matseliukh, Y., Dosyn, D., Chyrun, L., & Lozyn, O. (2020). *Analysis of the Demand for Bicycle Use in a Smart City Based on Machine Learning*.
- Kreuzberger, D., Kühn, N., & Hirschl, S. (2022). *Machine Learning Operations (MLOps): Overview, Definition, and Architecture* (arXiv:2205.02302). arXiv.
- Kumara, I., Arts, R., Di Nucci, D., Heuvel, W. J. V. D., & Tamburri, D. A. (2022). *Requirements and Reference Architecture for MLOps: Insights from Industry*. TechRxiv.
- Laato, S., Birkstedt, T., Määntymäki, M., Minkkinen, M., & Mikkonen, T. (2022). AI governance in the system development life cycle: Insights on responsible machine learning engineering. *Proceedings of the 1st International Conference on AI Engineering: Software Engineering for AI*, 113–123.
- Lo, S. K., Lu, Q., Paik, H.-Y., & Zhu, L. (2021). FLRA: A Reference Architecture for Federated Learning Systems. *Software Architecture*, 12857, 83–98.
- Lo, S. K., Lu, Q., Wang, C., Paik, H.-Y., & Zhu, L. (2022). A Systematic Literature Review on Federated Machine Learning: From a Software Engineering Perspective. *ACM Computing Surveys*, 54(5), 1–39.
- Lo, S. K., Lu, Q., Zhu, L., Paik, H., Xu, X., & Wang, C. (2022). Architectural Patterns for the Design of Federated Learning Systems. *Journal of Systems and Software*, 191(111357).
- McMahan, H. B., Moore, E., Ramage, D., & Arcas, B. A. y. (2016). Federated Learning of Deep Networks using Model Averaging. *CoRR*, abs/1602.05629.
- Mueller, T., Zahn, M., & Matthes, F. (2023). Unlocking the Potential of Collaborative AI - On the Socio-Technical Challenges of Federated Machine Learning. *Proceedings of the 31st European Conference on Information Systems*.
- Peffer, K., Rothenberger, M., Tuunanen, T., & Vaezi, R. (2012). Design Science Research Evaluation. In K. Peffer, M. Rothenberger, & B. Kuechler (Eds.), *Design Science Research in Information Systems. Advances in Theory and Practice* (Vol. 7286, pp. 398–410). Springer Berlin Heidelberg.
- Peffer, K., Tuunanen, T., Rothenberger, M., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77.
- Ritz, F., Phan, T., Sedlmeier, A., Altmann, P., Wieghardt, J., Schmid, R., Sauer, H., Klein, C., Linnhoff-Popien, C., & Gabor, T. (2022). Capturing Dependencies within Machine Learning via a Formal Process Model. *Proceedings of the 11th International Symposium On Leveraging Applications of Formal Methods, Verification and Validation*, 3, 249–265.
- Schomakers, E.-M., Lidynia, C., & Ziefle, M. (2020). All of me? Users' preferences for privacy-preserving data markets and the importance of anonymity. *Electronic Markets*, 30(3), 649–665.
- Studer, S., Bui, T. B., Drescher, C., Hanuschkin, A., Winkler, L., Peters, S., & Müller, K.-R. (2021). Towards CRISP-ML(Q): A Machine Learning Process Model with Quality Assurance Methodology. *Machine Learning and Knowledge Extraction*, 3(2), 392–413.

- Thanachawengsakul, N., Wannapiroon, P., & Nilsook, P. (2019). The Knowledge Repository Management System Architecture of Digital Knowledge Engineering using Machine Learning to Promote Software Engineering Competencies. *International Journal of Emerging Technologies in Learning (IJET)*, 14(12), 42.
- Venugeetha, Y., Harshitha, B. M., Charitha, K. P., Shwetha, K., & Keerthana, V. (2022). Breast Cancer Prediction and Trail Using Machine Learning and Image Processing. In A. Kumar, S. Senatore, & V. K. Gunjan (Eds.), *ICDSMLA 2020* (pp. 957–966). Springer. https://doi.org/10.1007/978-981-16-3690-5_89
- Wirth, R., & Hipp, J. (2000). CRISP-DM: Towards a Standard Process Model for Data Mining. *Proceedings of the 4th International Conference on the Practical Applications of Knowledge Discovery and Data Mining*, 29–40.
- Wouters, L., Creff, S., Bella, E. E., & Koudri, A. (2017). Collaborative systems engineering: Issues & challenges. *2017 IEEE 21st International Conference on Computer Supported Cooperative Work in Design (CSCWD)*, 486–491.
- Zhang, H., Bosch, J., & Olsson, H. H. (2020). Federated Learning Systems: Architecture Alternatives. *2020 27th Asia-Pacific Software Engineering Conference (APSEC)*, 385–394.

APPENDIX B

Guiding Questions for Federated Machine Learning Collaboration Agreements

Disclaimer: This systemized set of guiding questions only includes aspects that should be specifically considered within Federated Machine Learning collaborations.

Collaboration Management

Collaboration Structure

- Does the collaboration consist of one dominant partner or multiple equal partners?
- How is the collaboration established with the participants?
- What are the negotiation mechanisms for building the network of participants?
- How is decision-making handled and coordinated within the collaboration?
- Is the collaboration compliant with antitrust regulations?

Participant Management

- How are project participants e.g., data owners identified, evaluated, and selected?
- What capabilities are required from the different participants?
- How are participants involved and retained in the project?
- Are participants obligated to continue participating in throughout the entire project?
- How transparent are the participants inside and outside the project?

Infrastructure

- When, with whom, and in what context is communication taking place within the project?
- Which communication channel is used for the project?

Co-Creation Management

Distribution of Ownership, Responsibility, and Accountability

- How is the ownership, responsibilities and accountability distributed?
- How is the ownership of the FedML model handled?
- What rights and responsibilities do the project participants, e.g., data owners, have?
- Who takes accountability for the behavior of the FedML model?

Distribution of Revenue and Costs

- What are the different value streams?
- How are revenue and costs distributed?
- Which incentive is created for the project participants?
- Who takes care of the variable and fixed costs?
- How are collaboration costs, e.g., for IT infrastructure on the data owner's side, handled?

Intellectual Property Management

- How is the intellectual property of each project participant's input managed?
- How is intellectual property of the data handled?
- How is the intellectual property of the FedML model architecture and design handled?
- How is contributed intellectual property managed in case a participant leaves the collaboration?

Co-Creation Practices

Profit Calculation

- How is the profit calculation for the project performed?
- How does a potentially dynamic participant network affect the profit calculation?

Risks of Infeasibility

- How is the feasibility of solving the project with FedML determined?

- How is an initial viability assessment of the collaborative FedML project conducted?
- How is the technical assessment conducted?
- How is the business case analysis conducted?

Alignment in Quality

- How is the product quality defined?
- How can the product quality be ensured?
- How can the quality of each contribution be tracked within the collaboration?

Implementation of the FedML Activities

- How are the implementation tasks distributed and managed?
- Who sets up the technical process at each local contributor?
- Is the FedML pipeline automated or manually executed by each local contributor?

FedML Product

Compliance Data Protections

- Which regulations need to be considered to develop a compliant FedML model (e.g., GDPR, HIPAA)?
- How is data privacy for sensitive data assured throughout the FedML process and after aggregating the local contributions?

Versioning

- How are different model versions and libraries managed throughout the entire life cycle?
- How can model versioning be controlled, executed, and ensured for each local contributor?

Recall and Retirement

- How is a model recall enforced for local contributors and end customers?
- How can a model or resulting service be retired?
- How can the collaboration be dissolved and how is the generated output affected?
- How can contribution of a participant be retracted in case the participant decides to leave the collaboration?