

Predictive Compliance Monitoring in Process-Aware Information Systems: State of the Art, Functionalities, Research Directions

Stefanie Rinderle-Ma, Karolin Winter, Janik-Vasily Benzin

Technical University of Munich, Germany
TUM School of Computation, Information and Technology
Department of Informatics
{stefanie.rinderle-ma, karolin.winter, janik.benzin}@tum.de

Abstract

Business process compliance is a key area of business process management and aims at ensuring that processes obey to compliance constraints such as regulatory constraints or business rules imposed on them. Process compliance can be checked during process design time based on verification of process models and at runtime based on monitoring the compliance states of running process instances. For existing compliance monitoring approaches it remains unclear whether and how compliance violations can be predicted, although predictions are crucial in order to prepare and take countermeasures in time. This work, hence, analyzes existing literature from compliance monitoring as well as predictive process monitoring and provides an updated framework of compliance monitoring functionalities. For each compliance monitoring functionality we elicit prediction requirements and analyze their coverage by existing approaches. Based on this analysis, we delimit predictive compliance monitoring as new research area. Afterwards, open challenges and research directions for predictive compliance and process monitoring are elaborated.

Keywords— Predictive Compliance Monitoring; Predictive Process Monitoring; Systematic Literature Review; Research Directions

1 Introduction

The need for predictive and online data analysis is crucial given the highly volatile economic environment in which processes have to constantly adapt to new circumstances, e.g., to COVID-19 circumstances or the Ukraine war and, thus, historical data may be outdated and proactive process management gains importance [113, 140, 13, 119]. In the area of business process management, Predictive Process Monitoring (PPM) [61, 154, 103] has attained tremendous interest recently and several approaches for predicting, for example, the remaining time of cases, the next activity, or the outcome of a process have been

presented. Doing so, PPM can be a valuable means for estimating company-relevant key performance indicators, for example, customer retention. Hence, PPM can enable the proactive management of business processes and operational risks [29].

In addition to PPM, Compliance Monitoring (CM) [97, 144] is an integral part for monitoring and managing business processes in changing, complex regulatory environments such as the financial domain. By combining the respective capabilities of PPM and CM, research can offer companies a means to proactively assess and manage their business processes with respect to future outcomes, compliance status, and risks. Yet, reactive management through auditing is still most prominent in compliance management of companies, explaining why the use of predictive data analysis is an important factor for companies to improve and underlines the need in the current environment to bring both lines of research, i.e., PPM and CM, together [2].

So far, the combination of PPM and CM has not been put to the test, i.e., it has not been systematically analyzed what capabilities of CM are already covered by PPM and what may be missing to fully support *Predictive Compliance Monitoring* (PCM), i.e. predictive compliance management and monitoring in an online setting [97] where compliance violations of process instances are predicted during runtime. For CM, a framework of functionalities called Compliance Monitoring Functionalities (CMF) [97] has been established that serves as means to systematically compare and analyze existing CM approaches. Consequently, the CMFs are well-suited as a starting point for testing the combination of PPM and CM and, therefore, their relationship to PCM. Though some existing PPM approaches claim PCM as prediction goal, e.g., [99], due to the complexity of real-world business processes and their regulatory environments, it remains unclear whether or not PCM is and can be fully supported by PPM.

Example 1 illustrates the complexity of PCM, resulting from a multitude of compliance constraints stemming from different regulatory documents and referring to multiple process perspectives, a process event log/stream that is emitted from multiple, heterogeneous sources/systems, regular and irregular changes of the process, and a tight schedule of up to eleven hours to complete a cycle of the process.

Example 1 Transaction Reporting of Financial Institutions in the EU. *Beginning in 2014, authorities such as the European Commission, the European Parliament, the European Central Bank, the European Banking Authority and the European Securities and Markets Authority imposed regulations on financial institutions conducting money and capital market business in the EU to report money and capital market transactions on a daily basis to the respective authorities as a prerequisite to continue conducting that type of business. The required daily reports are based on the European Markets Infrastructure Regulation (EMIR) [45], Money Markets Statistical Reporting (MMSR) [46], Securities Financing Transactions Regulation (SFTR) [48] and Markets in Financial Instruments (MiFIR/D) [47]. These regulations are typically complemented by multiple addenda and technical specification documents that further clarify the*

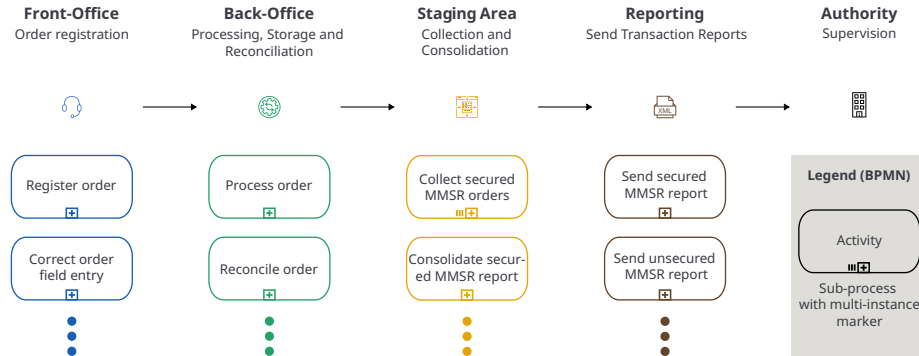


Figure 1: Transaction Reporting Processes of Financial Institutions in the EU [149]

exact requirements the institution has to fulfill, e.g., for MMSR the reporting instructions [44], questions and answers [43], IT appendix [41], data quality checks [42] and further technical specifications documents for the web service and XML schema available for download as a ZIP¹. All in all, the MMSR regulation stipulates the reporting of four different MMSR reports, among others the secured and unsecured market segment reports, that are used to determine the euro short-term rate ($\text{\textsterling}STR$), an important interest rate.

The regulatory documents specify the outcome of the process and necessitate the occurrence of certain activities, but do not specify the actual process model in full detail. Hence, the specification gives financial institutions similar to specifications in the healthcare domain [76, 75] significant flexibility for implementing and executing the process with individual subprocesses for activities [149]. Nevertheless, by abstracting the individual subprocesses of institutions, similar activities for the transaction reporting process can be deduced as depicted in Fig. 1. In the following, the coarse-grained activities that are similar to all financial institutions are described.

Traders or Sales personnel in the front-office agree on a transaction with an external trading party through electronic trading software, e.g., Bloomberg, or over the phone. For each transaction, an order is registered or has to be registered by the trader (“Register order”). If at some point a field of the order has been filled incorrectly, the trader has to correct the entry (“Correct order field entry”). As an agreed transaction needs to be cleared, i.e., payments and the actual transfer of owner/holding rights with respect to the depository need preparation and execution, it is processed in the middle-office (“Process order”). In order to prevent fraud and guarantee clearing according to previously agreed upon terms by both parties of a transaction, orders are reconciled (“Reconcile order”). In order to transform the large number of orders from various electronic order information systems to the specified report structure and format, all secured market segment

¹https://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/money_market/html/index.en.html

orders are collected for the secured MMSR report (“Collect secured MMSR orders”) and consolidated (“Consolidate secured MMSR orders”). Afterwards, the final report can be sent to the respective authority (“Send secured MMSR report”).

According to Annex IV 2. (i-ii) and 3. (iii) the reports have to be accurate, the reporting process must be monitored, and any deviation needs to be explained by the institution within 45 min to 7h depending on the time of occurrence [46, 43]. According to Article 4 1. (a), all MMSR reports have to be sent between 6pm of the same trade day that the transaction was agreed upon and 7am of the following trade day. In [41], the “Send [type] MMSR report” activity is specified. In Fig. 2, an extract from [44] shows a constraint that specifies for fixed-term evergreens as defined by the International Capital Markets Association (ICMA) with collateral having a valid International Security Identification Number (ISIN), i.e. money market loan contracts for which the borrower pledges another, registered capital market product to secure the repayment of the loan, the reporting in the secured market segment and attribute values “T” for both “Trade date” and “Settlement date”. Further exemplary constraints according to [46, 149, 44, 41, 40, 16] are:

- C1 Activity “Send secured MMSR report” needs to be executed once per day.
- C2 Event “Sending MMSR report failed” by the authority must not occur.
- C3 Activity “Collect secured MMSR orders” must be directly followed by “Consolidate secured MMSR orders”.
- C4 Activity “Sending MMSR report succeeded” needs to occur between 6pm of the same day and before 7am of the next trading day.
- C5 Activity “Correct order field entry” with attributes “reason=erroneous entry” is to be sent as an amendment revision in the following report, i.e. the next “Send secured MMSR report” needs to include an XML amendment message for the respective order.
- C6 Activities “Send secured MMSR report” can only be done by a role from a separate department than that of the role that initially registered the order.
- C7 A secured MMSR report can only be sent, if its processing is completed.
- C8 A MMSR report processing starts with checking the MMSR sending channel, i.e. it puts the “Send secured MMSR report” activity into a “ready” state (activation). A successfully sent report event completes the corresponding “Send secured MMSR report” activity.
- C9 The sum of successfully sent reports and failed report sent events needs to equal the number of send report activations.
- C10 For each registered order, the order needs to be processed to be included in the final MMSR report.

C11 All registered and processed orders are simultaneously sent to the authority.

C12 All “Process order” activities must be completed by 7am of the following trade day.

C13 Each report needs to be sent without another send operation in concurrence.

C14 Only four MMSR reports can be sent to the ECB per day.

C15 Unplanned deviations from normal operations (disruptions) and their root causes have to be registered, evaluated and prioritized with respect to their resulting risks and escalated according to predetermined criteria.

All in all, the regulatory documents for MMSR span hundreds of pages containing a multitude of constraints with varying detail. The financial institution is responsible to transform and map all of these requirements to its internal reporting process that adheres on a high-level to the previously described reporting process. Furthermore, all of the documents can change and have changed since the regulation became in force. For example, the regulation itself was amended four times and the reporting instructions at least six times¹. These changes lead to concept drifts in the reporting process, as historically possible activities may or cannot occur anymore and at the same time new activities might have to be executed that have not been observed before; these activities account for unseen behavior [116, 126, 102]. Additionally, interface changes of information systems for electronic trading that are used internally or by external trading parties (e.g. exchanges), as well as other interrelated IT system changes can lead to concept drift in the reporting process.

To sum up, the financial institution has to implement a complex transaction reporting process with a large number of constraints that runs every trade day, has to monitor its execution and explain deviations.

As described in Example 1, the financial institution has to monitor the compliance of the reporting process to the external regulatory constraints and further internal constraints necessary to comply to the governance guideline by the European Banking Authority [40]. To comply to the punctual sending of an accurate report and to proactively manage possible deviations due to IT incidents, wrong data attributes or external factors such as days with exceptionally high order volume, it also needs to ultimately monitor the remaining time prediction to complete the sending of the report and the data attributes predictions and life cycle of that activity with respect to the relevant constraints. In addition, because of concept drifts due to changes in various regulations, various information systems internally and externally, and IT incident and behavior patterns, a PCM system is required that regularly updates both the constraints set and the prediction and is also able to predict unseen behavior, for example data attributes in the report for which the historical data has never recorded a corresponding order. As the financial institution is required to explain deviations, it needs a root-cause analysis functionality and, because of the short period of

Fixed-term evergreens¹⁴ must be reported on an ongoing daily basis as new transactions (NEWT) – both when the transaction is initially conducted and when a rollover occurs – and in principle with the following date structure:

- Trade date: T
- Settlement date: T
- Maturity date: The first date on which the termination of the fixed-term evergreen can occur

Both Trade Date and Settlement Date need to be reported with “T”. However, in case the trade is negotiated with a Settlement Date different from T, for the initial reporting on execution the actual Settlement Date needs to be reported. The Maturity Date of the initial transaction and the subsequent rollovers needs to reflect the first date on which the termination of the fixed-term evergreen can occur and **not** the fixed final termination date of the agreement. Following the Settlement Date of the transaction fixed-term evergreens must be reported each day until they are redeemed (terminated/closed/called).

¹⁴ As defined by ICMA, a *fixed-term evergreen* is a transaction which has a fixed final repurchase date and where both parties have an option to terminate the transaction subject to a notice period. See [A Guide to best practice in the European Repo Market](#), ICMA, March 2021.

Figure 2: Constraint for Case Data of an Order [44]

time for explaining the deviations, ideally some functionality to explain what the results of the system mean.

To further exemplify the transaction reporting process with respect to PPM and CM, in the following, we present two scenarios for Example 1 that represent two days in which violations of the constraint might occur:

Scenario 1 Correct orders. *Consider trade day 29.09.2022, on which for the transaction reporting process of the German BANK the activities for processing, storage and archiving for the MMSR reports (Middle-Office and Back-Office) that usually occur from 8pm to 2am are running one hour longer than they used to run and the data collection and consolidation activities run 4 hours on average such that a violation of the MMSR regulation timeliness constraint is likely. Assume that the longer running time for processing, storage and archiving is due to an exceptionally high number of ”Correct order field entry” activities.*

Scenario 2 Conflicting rules. *Consider trade day 30.09.2022 6pm, at which time the reporting department employee on-call duty has already worked 10 hours. Nevertheless, the report processing breaks at 7pm of the same day such that the employee on-call duty has to intervene. Assume the 10 hour day of the employee was due to the assessment of protection requirements according to [16] and the breaking of the processing due to a hardware anomaly on one of the processing servers.*

In the first scenario, following the aforementioned requirements for monitoring the reporting process, the BANK wants to know and understand in particular how likely it is that the MMSR report is sent in time, whether this likelihood changes over night, what the root cause for the longer running Middle-

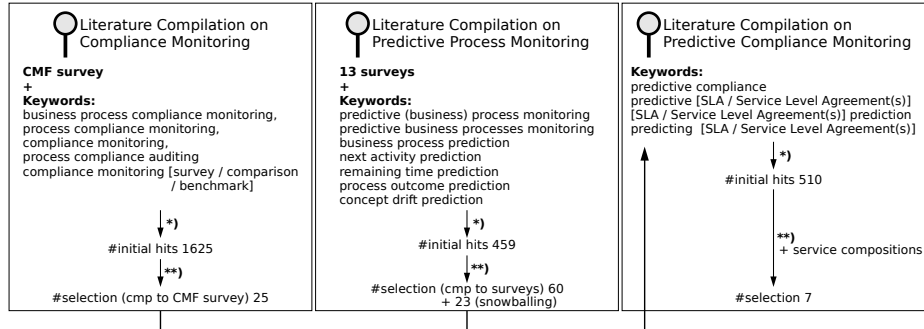
and Back-Office activities is and how to mitigate the potential violation of the timeliness constraint in the short- and long-term. In the second scenario, the BANK has to additionally understand, that it either can meet the requirement of the MMSR reporting process by fixing the broken processing or the requirement for German employees not to work more than 10 hours per day (§3 of the German “Arbeitszeitgesetz”).

This work addresses the question to which extent the PCM system required by the financial institution for monitoring and managing the transaction reporting process in its entirety is addressed and solved by existing PPM and CM approaches (**RQ1**), how the existing PPM and CM approaches are comparable in terms of CMF functionalities, in particular, prediction requirements necessary for the respective functionality (**RQ2**), whether PCM is actually a new research area (**RQ3**) and which challenges and research directions remain still open (**RQ4**).

We start with a compilation of PPM, CM, and PCM literature (\mapsto RQ1). Afterwards, the compiled literature is analyzed along the existing framework on CMFs [97] and for possible extensions to this initial framework (\mapsto RQ2). Analogously to [97] literature is complemented by case studies as source for the CMF extension. Based on the analysis of PPM, CM and PCM literature, we investigate the relation of PPM and PCM, i.e., whether i) PPM can fully encompass PCM by defining compliance violations as prediction goal or ii) PPM is utilized for PCM in predicting, e.g., next activities or remaining time, and interpreting the prediction results, even in combination, over the set of compliance constraints. The analysis of the literature compilation combined with findings from case studies results in an extended CMF framework. The extended CMF framework is then analyzed for the predictive requirements arising for each of the CMFs, illustrated by means of Example 1 (\mapsto RQ2). The CMFs, together with their predictive requirements, are then analyzed for support by existing (mostly PPM) approaches, categorized along their prediction goals, e.g., next activity or outcome (\mapsto RQ3). Then, we discuss the implications of both the extended CMF framework, its prediction requirements and the assessment of existing approaches for the relationship of PPM, CM and PCM as a combination of the former two (\mapsto RQ3). Furthermore, we provide suggestions how each CM functionality can be addressed to elucidate the relationship of the different research areas. Lastly, we provide a set of open challenges and research directions (\mapsto RQ4) that emerge from the assessment and discussion.

The contributions to tackle RQ1 – RQ4 comprise literature reviews for PCM, CM, and PPM, resulting in an extended CMF framework (cf. Sect. 2). The extended CMF framework is described and illustrated in Sect. 3 and serves as basis for deriving requirements for developing a comprehensive PCM solution. Existing PPM approaches are put to the test along these requirements in Sect. 4. The analysis steps finally culminate in open challenges and research directions for PCM (cf. Sect. 5). Limitations of this survey are discussed in Sect. 6.1 and results concluded in Sect. 6.2.

The extended CMF framework and the research directions provide several open research topics from a data, algorithmic, and application perspective for



*) Google Scholar "allintitle" search, January/February 2022
 **) selection criteria: (business) process focus, English language, identifiable/known publication outlet

Figure 3: Search Methodologies for the Literature Compilations on CM, PPM and PCM

PPM and PCM. Overall, this work aims at bridging the gap between online and predictive process analysis techniques and real-world compliance management.

2 Literature Review

Figure 3 depicts the literature search and selection methodology applied in this paper. We start with literature on compliance monitoring (CM), followed by literature on predictive process monitoring (PPM), complemented by a search for predictive compliance monitoring (PCM) approaches. The literature lists are available via <https://www.cs.cit.tum.de/bpm/data/>.

The overall goal of the literature review is to assess whether and to which extent PCM is addressed by existing approaches and to set out a research agenda for PCM. This necessitates building a basis for the assessment, i.e., a set of PCM requirements based on which existing approaches can be evaluated and potential research gaps can be identified. We use the well-established Compliance Monitoring Functionality framework presented in [97] as basis and update and extend it with respect to predictive requirements. The original framework [97] defines the following Compliance Monitoring Functionalities (CMFs):

- *Modeling requirements:* CMF1 (time), CMF2 (data), CMF3 (resources)
 CMF1–3 refer to the modeling capabilities of the compliance constraints. The underlying assumption is that all compliance constraints refer to the control flow of a process, e.g., by referring to the existence of an activity plus a maximal duration of this activity (\rightarrow CMF1).
- *Execution requirements:* CMF4 (non-atomic activities), CMF5 (life cycles), CMF6 (multiple instances constraints)
 CMF4–6 refer to instantiation and execution of the process instances, more precisely the event and life cycle information that is stored in the process

event streams during runtime, and the instantiation of the compliance constraints.

- *User requirements:* CMF7 (reactive management), CMF8 (proactive management), CMF9 (explain root cause of violation), CMF10 (quantify compliance degree)

CMF7–10 refer to support that approaches offer for users to understand and handle compliance violations. CMF8 refers to PCM as proactive management of compliance violations that requires the prediction of such violations.

In the following, we analyze the papers from each literature compilation regarding two aspects, and aggregate the results at the end of this section.

- i) Which existing CMFs from [97] are mentioned/addressed?
- ii) Which possible CMF extensions are mentioned/addressed?

2.1 CM Literature Compilation and Findings

We take the CM functionality framework and systematic literature survey from 2015 [97] as a yardstick, i.e., we assume that CM literature up to 2015 has been mostly covered by [97]. For keywords and selection criteria see Fig. 3. The search resulted in 1625 initial hits. From these 1625 papers, 25 papers were selected as in scope of PCM and not yet analyzed by the CM survey in [97]. The analysis of these 25 papers results in 17 papers that mention or address CMFs in the following way:

- i) *Which CMFs from [97] are mentioned/addressed by CM approaches?*

Most of the existing CM approaches [3, 5, 4, 129, 64, 68, 86, 88, 96, 100, 106, 153, 165, 36] address modeling requirements CMF1–CMF3 fully or partly through their support for the process perspectives time, data, and resources and depending on the employed constraint modeling formalism such as linear temporal logic (LTL), event-based compliance language (ECL), compliance rule graphs (CRG), and variants of event condition action rules (ECA rules), i.e., timed ECA rules or match condition action rules. Fewer approaches mention or address execution requirements CMF4–CMF6, mostly those approaches that support some kind of activity life cycle [5, 4, 129, 86, 88, 106, 165]. Several approaches address user requirements CMF8–CMF10 [3, 5, 4, 129, 86, 86, 100, 106, 36] based on providing reactive and partly proactive management of compliance violations as well as visualization approaches for compliance states, e.g., satisfied or violated. Overall, it seems that the majority of approaches are basic in the sense that they focus on providing support for compliance monitoring in terms of modeling constraints and checking them over event streams, hence addressing modeling requirements CMF1–CMF3. There are also approaches that aim at providing a comprehensive CM solution/framework by addressing modeling and user requirements. For the execution requirements, the “coverage” depends on how approaches are able to deal with different activity life cycles and constraints

that span across multiple instances or processes. The latter (CMF6) has not been covered by existing CM approaches yet.

ii) *Which possible CMF extensions are mentioned by CM approaches?* [3, 96] emphasize the efficiency/performance of the CM approach in order to deal with a large volume of events as well as the aspect of data quality. [25, 86, 96, 106] raise the requirement to support CM in process collaborations, for example, compliance in connection with the dynamic replacement of partners leaving a process collaboration or new partners joining. The support of CM in distributed settings is also emphasized by [88] in supporting the (semantic) aggregation of events from heterogeneous sources. [64] advocate the aggregation of values of multiple events and event correlation for addressing multiple data sources. Moreover, the efficiency of the approaches is put into the spotlight for dealing with a large volume of events. [100] mentions the “*early detection of conflicting constraints*” where only one of the constraints can be fulfilled at a time. [3, 5, 4] emphasize the consistency of the compliance constraint base.

From the analyzed approaches, [89, 141] as well as the survey in [144] cannot be fully assessed due to lack of technical detail.

Conclusion. CMF1 - CMF10 as proposed in [97] are still valid and approaches since its publication in 2015 address several of the outlined CMFs. After 2015, new directions/requirements include:

- Efficiency/performance of compliance monitoring
- Compliance monitoring in distributed processes
- Integration of event streams from multiple data sources
- Consistency of the constraint base

Efficiency and performance of compliance monitoring is motivated in existing literature by the volume of the event data, also in combination with the applied (ML) technique. The assessment of an approach regarding its efficiency and performance is depending on the application and users, e.g., an CM approach taking 5 hours can still be efficient if users expect the results within 12 hours. Hence, the requirement of efficiency/performance will be considered a user requirement. The other requirements refer to *data*. Hence, CMF1–CMF10 will be extended with one user requirement and new *data requirements* accordingly. We will describe and illustrate these extensions in Sect. 2.5.

2.2 PPM Literature Compilation and Findings

For keywords and selection criteria used in the literature search for PPM see Fig. 3. The search results comprise 459 initial hits among which 126 are selected based on the outlined criteria, including 12 survey papers. The publication dates of the surveys include 2 surveys in 2018 [61, 103], 3 surveys in 2019 [154, 146, 142], 3 surveys in 2020 [138, 66, 112], and 4 surveys in 2021 [163, 78, 140, 121]. We add the survey in [110] due to snowballing.

The 13 surveys provide classifications for PPM approaches based on prediction goals [61, 103, 163], techniques [66, 109, 121], goals and technique [112], and use cases [138] as well as benchmarks regarding specific techniques [143], a specific prediction goal [147, 154], or specific data requirements [138]. [140] provides a classification based on the explainability of PPM approaches. The classifications provided by existing surveys are partly utilized for the analysis conducted in Sect. 4 and comprise prediction goals relevant in the context of PCM, e.g., LTL rules, and for CMF extensions, e.g., inter-case metrics. Also small data sets and explainability point to possibly relevant CMF extensions.

At first, the analysis of the 13 surveys results in 10 papers (partly contained in the 126 papers in the literature selection) that claim to provide a solution to PCM, i.e., compliance monitoring through PPM. Out of these 10 approaches, 4 do not address any constraint definition, 3 address constraint definition through SLA [20, 92, 28], and 3 address constraint definition in the form of predicates [99, 60, 130], e.g., based on LTL constraints. The latter three approaches, in particular, define predicates as prediction goals, i.e., it becomes directly possible to predict violations. [108] outlines three basic PPM approaches based on machine learning, constraint satisfaction, and QoS aggregation which are all relevant for PCM. We consider SLA to be in scope for the subsequent analysis on PPM and the search on PCM, as the majority of relevant PPM papers targets SLAs. In this context, we also consider *prescriptive process monitoring* approaches, which are concerned with finding the optimal execution time of interventions to optimize relevant KPIs or meet SLAs of the process such as the cycle time of process instances through predictions at runtime [11, 107]. These approaches can be also interesting for possible actions after violations have been predicted.

Similar to the CM literature review, we use the 13 surveys as yardstick to distinguish “non-survey” papers into papers that have already been analyzed and contribute to the conclusions of one or several of the 13 surveys, and into papers that have not been analyzed by a survey yet. The latter class of papers comprises 60 papers. Snowballing yields 23 PPM papers to be potentially relevant for PCM. From these 60+23 papers, 14+8 mention/address “compliance” or “SLA” and go into the following discussion of questions i) and ii) set out in the previous section for CM.

i) *Which CMFs from [97] are mentioned/addressed by PPM approaches?* Existing PPM approaches [92, 32, 70, 152, 117, 157] address modeling requirements CMF1–CMF3 fully or partly as PPM perspectives time, data, and resource. Control flow, by contrast to [97], is explicitly mentioned as PPM perspective, especially in the context of next activity/event prediction approaches such as [30, 49, 115] as well as pattern [7] and predicate prediction approaches [60]. The support of execution requirements CMF4–CMF5 is missing as existing approaches have not supported any kind of (activity) life cycle yet. CMF6 is addressed by [132] through predictions considering intra-case and inter-case features. User requirement CMF8 is addressed [126, 92] by, e.g., providing BPI cockpits [20] or suggesting mitigation actions [35]. Several PPM approaches focus on the explainability of the prediction results [8, 104, 155]. They can be mapped onto CMF9 on explaining root causes for compliance violations as

proposed in [97], but CMF9 can be refined into more precise CMFs, i.e., 1) root cause analysis and 2) effective communication of root cause as in [97], as well as additionally in 3) explaining and visualizing prediction results, 4) explaining and visualizing the set of future violations, and 5) explaining and visualizing the effects of mitigation actions on predicted/future violations. Moreover, we advocate to rename CMF9 into *CMF9': Explainability*.

ii) *Which possible CMF extensions are mentioned by PPM approaches?* [28, 14, 10, 57] mention the requirement to consider external (process) context data. This can be underpinned by other recent approaches such as [139, 38] showing that context can provide useful information for root cause analysis and explainability in PPM. Prediction in distributed processes is addressed in [10]. Finally, PPM approaches address the properties and quality of the data, i.e., the event streams, including the size of the input data [79], and the output e.g., the reliability of the predictions [26]. This will be reflected in an additional CMF on data properties and quality.

Conclusion. An additional requirement reflecting the control flow perspective of compliance constraints will be added to the CMF framework. Moreover, the CMF framework will be extended by a CMF on the ability to exploit external (process) context data and data properties and quality. These additional CMFs can be added to the new group *Data requirements*. Further on, a refinement of CMF8 and CMF9 will reflect the work on explaining and visualizing results of prediction.

Note that the PPM approaches from the literature review will be analyzed with respect to how they meet PCM requirements in Section 4.

2.3 PCM Literature Compilation and Findings

For keywords and selection criteria used in the literature search for PCM see Fig. 3. The number of initial hits accounts to 510 among which we selected papers based upon the outlined criteria additionally extended to include papers on SLA predictions in the context of service compositions. Papers from the medical domain predicting whether a medical treatment would result in the desired effects or whether patients are likely to follow the medical advice are considered out of scope due to a missing connection to business processes. Moreover, we excluded theses and finally came up with 7 papers [92, 80, 24, 127, 27, 73, 94]. Out of these, [92] was also found in the compilation of PPM literature and has therefore already been assessed.

PCM literature yields the following insights with respect to *i) existing CMFs from [97]* and *ii) CMF extensions*. [80] can be classified as remaining time PPM approach and hence addresses CMF1. [27] mentions prediction across multiple process cases, i.e., instance spanning predictions, but no concrete solutions are provided. The approach is directed towards explainability by providing a measure for reliability of predictions, but just for individual cases. Those aspects are covered by CMF6 and CMF9. [73] uses an abstract notation for service orchestrations, i.e., “*compositions with a centralized control flow*” and “*predict*

possible situations of SLA conformance and violation, and to obtain information on the internal parameters of the orchestration (branch conditions, loop iterations) that may occur in these situation". These aspects are covered by CMF8 and CMF9 and touch distributed processes which was already identified as an additional CMF in the previous sections. [94] predicts SLAs and adapt service compositions in order to avoid a violation of SLAs. Mitigation actions and adaptations are part of CMF8. [24] targets the problem of state space explosion which addresses the newly added requirement on efficiency of compliance monitoring. An analysis of non-compliance to prevent compliance violations in the future with only limited prediction capabilities is presented in [127] and addresses CMF8.

Conclusion. PCM approaches confirm the findings of the CM and PPM literature reviews regarding CMF framework extensions.

2.4 Findings Based on Case Studies

Analogous to [97], we analyze case studies and real-world compliance constraint collections to identify CMFs with respect to *i) existing CMFs [97]* and *ii) possible CMF extensions*. Case studies can be found in various domains including data protection [160], finance [156], and manufacturing [63, 162]. As discussed in [156], real-world compliance constraints refer to the modeling requirements CMF1-3 plus control flow patterns such as existence, absence, and ordering. A collection of real-world constraints that span multiple process instances and processes can be found in [123]. Constraints spanning multiple instances are referred to by CMF6 in the original CMF framework [97]. When looking into literature and the real-world constraints, CMF6 can be refined into constraints that reflect *i) the simultaneous execution of events, ii) constrained execution, iii) order of events, iv) non-concurrent execution of events, and v) constrained start of following instances [162, 161]*.

Conclusion. Case studies and collections of real-world compliance constraints confirm modeling requirements CMF1-3. Additional requirements can be specifically identified in real-world constraints that span multiple process instances or processes and will be included as refinement of CMF6.

2.5 Extended Compliance Monitoring Functionality Framework

This section summarizes the extended CMF framework based on the findings from literature reviews and case studies in Sect. 2.1 – 2.4). We deliberately propose an extended CMF framework rather than, for example, a predictive CMF framework as the extensions refer to prediction and monitoring. Following [97], each of the CMFs is partitioned into sub CMFs reflecting specific requirements on the expressiveness of the CMF. Each of the CMFs is illustrated by an example in Sect. 3.

The first extension refers to the modeling requirements by explication of CMF0 on control flow. Following control flow patterns for compliance constraints [97], we opt for the basic building blocks **existence (CMF0.1)**, **absence (CMF0.2)**, and **ordering (CMF0.3)**. Note that CMF 1.1 time qualitative becomes obsolete due to adding CMF0.3 ordering.

For the execution requirements, the extensions comprise the refinement of CMF6 on multiple instance constraints, following the categorization for instance-spanning constraints proposed in [162], i.e., constraints on **simultaneous (CMF6.2)**, **constrained (CMF6.3)**, **order (CMF6.4)**, and **non-concurrent (CMF6.5)** execution of tasks across process instances/processes as well as **constrained start of following instances (CMF6.6)**.

The user requirements are extended by refinement of CMF8 and CMF9. For CMF8, the **update of the set of possible and future violations (CMF8.3)** of compliance is added. CMF9 is renamed to **CMF9': Explainability** and refined by **explain and visualize the prediction results (CMF9.3)**, **explain and visualize the set of possible and future violations (CMF9.4)**, and **explain and visualize the effects of mitigation actions (CMF9.5)**. Moreover, **CMF11 on efficiency/performance of CM** is added as user requirement.

Finally, the group of **Data requirements** on process event data such as logs and streams as input for solving the PCM problem is added. The consistency of the constraint base as also mentioned in the literature is considered beyond the scope of this work. In detail, the extensions comprise **integration of data from multiple sources (CMF12)**, **distributed processes (CMF13)**, **context data (CMF14)**, and **data properties and quality (CMF15)**.

3 Extended CMF Framework and Prediction Requirements

The extended and updated CMFs when compared to the original framework presented in [97] are described by means of the template also provided in [97] in order to equip them with a more precise meaning. The template contains for each CMF, its *name*, a brief *overview*, a *description*, *evaluation criteria* with particular focus on how the CMF could be verified through PPM in terms of prediction requirements, *examples*, and clues on the *implementation* [97]. In case of refined CMFs, the templates are centering around the refinements. Moreover, the original CMFs are also illustrated based on Ex. 1 and evaluation criteria are listed with the focus on prediction requirements. This illustration is understood as an addition, i.e., the *overview*, *description* etc. from still hold [97] for the original CMFs. The list of evaluation criteria results in a list of requirements for a holistic view that serves as an input for the assessment of existing PPM and CM approaches in covering PCM, i.e. both PPM and CM, (cf. Sect. 4) and lays the foundation for the subsequent discussion on the relationship between PPM, CM and PCM (cf. Sect. 4.5).

For the illustration and description of the extended CMF Framework and elicitation of prediction requirements it is important to distinguish two prediction goals to implement the prediction ability of PCM:

- i) formulate the prediction goal as the truth value of the compliance constraint (*predicate prediction*)
- ii) formulate the prediction goal as next (set of) events/activities potentially in combination with time or data related prediction goals and subsequently check compliance constraints (*next activity prediction “plus”*).

Other common prediction goals such as remaining time by themselves are not expressive enough to capture all relevant compliance constraints (e.g. *C1* of Ex. 1). For predicate prediction, the expressiveness of the chosen logic for compliance constraints determines the assessment of an approach for modeling and execution requirement CMFs. For next activity prediction, the assessment of modeling and execution requirement CMFs is also dependent on the prediction method. Therefore, the evaluation criteria for modeling and execution requirement CMFs in Sect. 3.1 and 3.2 focus on prediction requirements for next activity prediction. With the exception of CMF9.3 and CMF9.4 that coincide for predicate prediction, the user and data requirement CMFs are both similarly relevant to predicate and next activity prediction.

Furthermore, it is important to note that the prediction goal of an approach is independent of the ability to *match* reality observed through recorded/received events/activities with the respective events/activities stated in the compliance constraints. In predicate prediction, the approach must match before training the prediction method in order to compute the target for prediction, whereas in next activity prediction the approach must match after the training, even after the prediction. Consequently, the next activity prediction separates the task of predicting from the task of matching and checking compliance. Also note that matching requires an equivalence notion that formalizes the case when we consider two events/activities as equal. Since the often assumed label equivalence requires the usage of a controlled vocabulary throughout the design, execution and change of monitored processes and respective information systems, this equivalence notion is a non-valid simplification [124]. Hence, further equivalence notions such as attribute equivalence are necessary for proper matching.

3.1 Modeling requirements

CMF0 to CMF3 pertain to the ability of CM approaches to deal with control-flow, time, data and resource constraints.

CMF0.x: existence, absence, ordering

Overview: As basic building blocks of control-flow patterns, compliance rules need the ability to express that an activity must occur or exists in a process instance, that it is absent from a process instance and in what order two or more activities occur in the process instance.

Description: The existence or occurrence of an activity in a process instance can either be a primitive condition that a certain activity exists in the instance or be more advanced in the sense that the condition additionally carries information on how often the activity occurs in the instance. The former is stated either explicitly as a must occur or implicitly as part of another condition, for example, an activity data condition that implies the activity to occur. The latter typically adds some pattern for quantities such as "at least twice". To check conditions on absence of an activity in a process instance in particular requires the ability to decide whether a process instance is completed. The ordering conditions surpass the original time qualitative conditions in [97] in expressiveness. Whereas time qualitative conditions can only express eventually follows, ordering conditions also capture conditions on the directly follows of events/activities or statements such as "follows after three intermediate activities".

Evaluation criteria: To fully support CMF0.x, a PCM approach must be able to predict the set of next activities/events ranked by probability of occurrence and with a distinction of directly/eventually follows or the complete order of activities/events per current process instance, the approach must be able to search the complete process instance for activities/events and match recorded activities/events with the activity/event specified in the condition.

Examples (cf. Ex. 1): C1 (Existence), C2 (Absence), C3 (Ordering)

Implementation: For training the prediction method for predicate prediction (i) or after prediction of the next activities/events for next activity prediction (ii) the compliance monitoring approach must be able to check the truth of the existence and order conditions, i.e., it requires some equivalence notion for matching activities and events such as label equivalence or attribute equivalence [88] (e.g., in case events of conditions have to be matched with events contained in console logs that do not come with easily separable labels), a threshold for the probability of events/activities to decide whether they will actually occur, a mechanism to correlate events to process instances and a means to infer the order of the predicted events per process instance. For next activity prediction, the implementation needs to decide how to infer the absence of activities: Not predicting the occurrence of a certain activity is interpreted as being absent (only valid for approaches able to predict unseen behavior) or predicting it with a below threshold probability can be interpreted as being absent (valid interpretation for approaches able to predict unseen behavior and for approaches without that ability).

CMF1: time quantitative

For *Overview* and *Description*, and *Implementation* see [97].

Evaluation criteria: To fully support the quantitative time functionality, the approach must be able to additionally predict the remaining time to either complete the process, to complete the activity, the remaining time until the next event happens per process instance or timestamps for all predicted events/activities.

Examples (cf. Ex. 1): C4 (quantitative)

CMF2.x: activity data, case data

For *Overview* and *Description*, and *Implementation* see [97].

Evaluation criteria: To fully support these data-related functionalities, the approach must be able to additionally predict relevant event attributes potentially depending on the prediction of case data over time.

Examples (cf. Ex. 1): C5 (activity data), constraint extracted from Fig. 2 (case data)

CMF3.x: unary resource condition, extended resource condition

For *Overview* and *Description*, and *Implementation* see [97].

Evaluation criteria: To fully support resource-related functionalities, the approach must be able to additionally predict the associated resource of activities/events based on resource predictions and potentially further event attributes necessary to monitor and check the conditions.

Examples (cf. Ex. 1): Orders without collateral need to be registered by a senior trader (unary resource condition); C6 (extended resource condition)

Summary of PCM system (prediction) requirement list based on modeling functionalities: Compliance prediction with respect to modeling CMFs requires next activity / event prediction, including fine-granular probabilities. The prediction of activity absence is especially interesting, as it requires to infer it from the predicted activities. Moreover, temporal and resource prediction as well as the prediction of data values is required.

3.2 Execution requirements

CMF4 to CMF6 pertain to the ability of CM approaches to deal with execution-based constraints. Furthermore, these CMFs enable the assessment of approaches to deal with domain-dependent information that is only available during execution of process instances.

CMF4: non-atomic activities

For *Overview* and *Description*, and *Implementation* see [97].

Evaluation criteria: To fully support non-atomic activities, the approach must be able to additionally predict and distinguish different event types and their relation to the respective activity.

Examples (cf. Ex. 1): C7 (explicit)

CMF5: life cycle

For *Overview* and *Description*, and *Implementation* see [97].

Evaluation criteria: To fully support life cycles, the approach must be able to additionally predict and distinguish life cycle states/transitions of next events/activities.

Examples (cf. Ex. 1): C8 (activation+completion); C9 (balance start/complete events)

CMF6.x: Multiple instantiation of compliance constraints, simultaneous execution, constrained execution, order, non-concurrent execution, constrained start

Overview: Multiple instances of compliance constraints may not only be necessary for a single process instance due to multiple occurrences of related activities, but in the case of instance-spanning constraints also for multiple instances of one or multiple process types.

Description: As constraints may impose requirements over multiple process instances or even processes and the constraint’s instantiation trigger may occur multiple times during execution, each instantiation needs to be monitored simultaneously. Consider, for example, the instance-spanning constraint ”The centrifugation may only be started when at least five samples have arrived” [159]. If there exist five centrifuges in the lab, the instance-spanning constraint may be instantiated five times in parallel to monitor the arriving of samples at each centrifuge individually before the centrifugation may be started. If for four of the five centrifuges more than five samples have arrived and for the fifth centrifuge only three, but all centrifugations have started, the monitoring framework should only identify the respective constraint for the fifth centrifuge as violated.

Evaluation criteria: To fully support multiple instances functionalities, the approach must be able to predict various combinations of previous prediction goals such as events with attributes, life cycles and/or timestamps in conjunction with a notion of process instances and processes to predict events across multiple instances/processes. In addition, the approach must be able to be flexible in handling various data granularities in constraints with respect to event attributes in received events, as the aggregation of data values for specific events may be necessary.

Examples (cf. Ex. 1): *C10* (multiple instantiation); *C11* (simultaneous execution); *C12* (constrained execution); *C13* (non-concurrent activities); *C14* (constrained start)

Implementation: Encoding inter-case features [133] can be a solution to predict, for example, the utilization of shared resources. Compliance prediction of constraints that span across multiple instances or processes pose requirements on the ability of a PCM solution to a) model *context* in the constraints, e.g., the set of instances the constraint refers to [52]; b) differentiate process instances/processes; c) track the contexts of rule activations and their compliance status separately; d) transfer the separate handling of rule activations and their contexts to instance-spanning constraints.

PCM system (prediction) requirement list based on execution functionalities: Compliance prediction with respect to execution CMFs requires approaches to distinguish the semantics of different event types and life cycle states/transitions and to predict different event types and life cycle states/transitions. Moreover, predictions of next activity/event plus prediction of time, data, and resources should be possible across multiple process instances and processes.

3.3 User requirements

CMF7 to CMF11 pertain to the ability of CM approaches to support the user in understanding and managing compliance violations in time.

CMF7: reactive management

For *Overview* and *Description*, and *Implementation* see [97].

Evaluation criteria: To fully support the reactive management of compliance violations, the approach must be able to continue predicting events for

process instances that have already violated constraints and to continue the monitoring of further constraints for these instances. On the one hand, continuous monitoring of violations gives a more complete and fine-grained view on the compliance status of a process/system. On the other hand, it supports and enriches further functionalities of reactive management such as recovery, compensation mechanisms and reporting/documentation features for auditing.

Examples (cf. Ex. 1): C15 (continuous monitoring)

CMF8.x: early detection of conflicting rules, possible/future violation, update set of possible and future violations, recommendations for users to avoid violations

Overview: The ability for pro-active detection and management of compliance violations does not only include the detection of conflicting rules that lead to implicit violations, the detection of possible and inevitable future violations and subsequent recommendation of mitigation actions, but also includes the capacity of an approach to react to evolving or high velocity processes in real-world environments by updating the set of possible and future violations.

Description: The typical training of a prediction model on ex-post process instances may be too inflexible to capture the dynamics and variability of evolving processes in the real-world [126]. Hence, CM approaches should be able to flexibly deal with processes exhibiting concept drift or a multitude of variants over time through mechanisms to update the prediction model and/or the set of possible and future violations as the process evolves.

Evaluation criteria: To fully support the pro-active detection and management of compliance violations, the approach must be able to detect conflicting rules as soon as possible with precise probability/likelihood, to continuously update the set of conflicting rules as the event stream evolves, to detect and predict the set of compliance violations as soon as possible and as complete as possible with precise probability/likelihood and to continuously update compliance prediction for all compliance constraints and events as the event stream evolves. Furthermore, the approach has to determine and provide mitigation actions based on compliance predictions as soon as possible, with precise assessment of risk and impact of the mitigation actions. Additionally, it has to continuously update the mitigation actions based on updates of compliance predictions.

Examples (cf. Ex. 1):

- Consider *Scenario 2*. The accuracy constraint of the MMRS regulation is in conflict with a constraint set by the German “Arbeitszeitgesetz” (early detection of conflicting rules);
- Consider *Scenario 1*. Missing the deadline for sending the report is a possible violation to be predicted (possible/future violation);
- Consider *Scenario 1*. Each time an activity of the data collection and consolidation in the staging area is completed, the likelihood of possibly violating the MMSR timeliness constraint may change (update set of possible and future violations);
- Consider *Scenario 1*. Recommendations to avoid the possible violation may be to provision additional server resources and register them with the execution engine in the staging area; or to avoid all future possible violations may be to

reduce the number of correcting order field entries by putting attention to the topic in the front-office or by aligning their incentives with a correcting order field activity metric (recommendations for users);

Implementation: For updating the set of possible and future violations as the event stream evolves, the CM system can either take a brute-force approach by always updating the predictions as new events arrive, which may be computationally intensive or even infeasible for high velocity streams, or optimize the points in time to update the predictions as well as the prediction model according to new events. Consequently, the approach needs both a strategy for updating the prediction and for updating the prediction model in light of concept drift. These strategies can either be optimized for the particularities of the overall PCM system or generic and configurable for various settings.

CMF9.x: root cause analysis, effective communication of root cause, explain and visualize prediction results, explain and visualize set of possible and future violations, explain and visualize effects of mitigation actions

Overview: Next to root cause analysis and the effective communication of identified root causes, the prediction results, set of possible and future violations and the mitigation effects need to be communicated to the end user in a meaningful way by appropriate visualizations and explanatory presentations. These additional abilities increase the usability of the CM approach and facilitate the efficient use of the same in practice.

Description: Depending on the prediction goal(s) of the CM approach, the prediction results and the set of possible and future violations may coincide (in case of predicting truth values of constraints) or be separate from another (in case of predicting the sequence of events before evaluating the constraints on the sequence of events). Nevertheless, for both and the effects of mitigation actions, the CM approach should provide the user with explanations and visualizations in support of a quick assimilation of the key factors determining the results, what they mean and - for mitigation actions - what their execution will lead to.

Evaluation criteria: To fully support root cause analysis, explainability and visualization functionalities, the approach must be able to precisely determine root causes for predicted compliance violations as soon as possible, provide root cause analysis and continuously visualize root causes for predicted compliance violations to users. Furthermore, the approach has to continuously provide explanations for compliance predictions at algorithmic level (i.e., which input leads to which output) and continuously visualize the prediction results in their context, possibly together with providing post hoc explanations (together with CMF9.1). The predicted compliance violations have to be continuously visualized together with their root causes and effects (cf. CMF9.1/CMF9.3) and with their mitigation actions and the effects of applying the mitigation actions. Each of these functionalities requires the ability to handle single and multiple instances (the latter also in an aggregated manner) and multiple process perspectives and views.

Examples (cf. Ex. 1):

- Consider *Scenario 1*. The root cause analysis for the longer Middle- and

Back-office activities should reveal the exceptionally high number of correct order activities as the root cause (root cause analysis);

- Consider *Scenario 1*. Effective communication of the exceptionally high number of correct order activities both entails a visualization for each MMSR reporting process showing the relation of the correct order activities to the respective, longer running activities in the Middle- and Back-Office with the respective correlations and statistical tests (effective communication);
- Consider *Scenario 1*. The prediction of the later completion times for the Staging area should explain that this is due to the longer running Middle- and Back-Office activities and to what extent this is the case. The quantifying metrics for the extent are visualized in a graph showing all the features and to what extent they determine the current prediction and their likelihood (prediction results);
- Consider *Scenario 1*. The content of CMF9.2 in conjunction with a visualization of the predicted overtime by which the staging area will violate the timeliness requirement assembles all required information for the MMSR reporter to understand the possible violation and decide on how to proceed (set of possible/future violations);
- Consider *Scenario 1*. The type of provisioned servers and their effect is visualized to show to what extent they affect the performance of the staging area. The properties of the execution engine that determine the effect and its limits is explained on top of the visualization (mitigation actions);

Implementation: To streamline the result presentation of the root cause analysis, prediction (and compliance checking) and mitigation actions with their respective effect to the end user, the implementation can choose to support tools for *Business Intelligence* [65] or implement a standalone graphical user interface. The former option transfers the functionalities with respect to communicating, explaining and visualizing (CMF9.2-CMF9.5) to an external tool and the end user. Since the end user has to know the technical details of acquiring relevant information from the CM approach, this option may entail a considerable barrier for using the system. The latter option does not come with these limitations. Here, the CM approach can leverage *Visual Analytics* [148] to enable the interaction of the end user with the results. By summarizing the key results and their implications in a dashboard (as proposed, for example, in [136]), the CM approach can support the user in understanding the various outputs of the system, as presenting information through dashboards has proven beneficial in practice [37, 81, 91].

CMF10.x: compliance degree of single traces, compliance degree of an entire process/system

For *Overview* and *Description*, and *Implementation* see [97].

Evaluation criteria: To fully support compliance degree functionalities, the approach must be able to continuously exploit the predicted probabilities/likelihoods of compliance violations for continuously determining and updating the compliance degree of single process instances and across all process instances and processes. Depending on the complexity of the compliance constraint base,

the prediction goal in the case of predicate prediction can be tweaked to predict compliance degree classes instead of truth values for compliance constraints.

Examples (cf. Ex. 1):

- The respective regulations are different for the four reporting processes and the institution is required to monitor and explain deviations in the reporting processes leading to the need for individual compliance degrees of each reporting process instance (single traces);
- To assess the overall risk that is imposed through the transaction-based reporting system, the bank must measure and monitor the compliance degree of the entire process/system [16] (entire process/system);

CMF11: efficiency/performance of CM

Overview: To realize pro-active detection and management of compliance violations, an appropriate performance of the CM approach is a prerequisite.

Description: As the efficiency/performance of the CM approach determines the point in time the end user can react to compliance violations by executing proposed compensation/recovery mechanisms, reporting the violation and documenting it; or proactively act to mitigate the possible violations by executing proposed mitigation actions, it is crucial that this point in time adheres to the performance requirements set for the monitored system and it comes before the possible violations occurs. Depending on how challenging the performance requirements are, a particular CM approach may be unsuitable.

Evaluation criteria: To fully support the performance functionality, the approach must be able to provide performance optimization strategies for compliance prediction and its continuous update based on, e.g., delta approaches. Furthermore, it needs to provide benchmarks with respect to compliance prediction performance in offline and online settings.

Examples (cf. Ex. 1): If the CM approach takes more than 11 hours to update its predictions, it can never capture the dynamics inherent to the reporting process at runtime, rendering its capabilities as only reactive or for auditing purposes;

Implementation: Due to the trade-off between performance and accuracy [135], CM approaches with the highest accuracy values in offline settings may render themselves unsuitable in online settings. As performance requirements on the CM approach may greatly vary, the implementation can either specialize on a certain performance threshold in online settings and optimize its results accordingly or introduce configuration parameters/options that determine the effective performance/efficiency of the approach in practice.

PCM system (prediction) requirement list based on user functionalities: Compliance prediction with respect to user CMFs requires approaches to continuously predict and update conflicting compliance constraints and compliance violations as soon as possible with precise probability/likelihood and to exploit this information for precise user feedback. This user feedback comprises explanations at the algorithmic level as well as the visualization of compliance violation predictions and their probabilities. Moreover, the root cause for predicted compliance violations has to be investigated and presented to users through visualization. Furthermore, mitigation actions based on the compli-

ance violation prediction analysis are to be determined, continuously updated, and their effects and updates are to be visualized for users, as well. Finally, performance optimization strategies for prediction and continuous update and benchmarks with respect to prediction performance are to be supported.

3.4 Data requirements

CMF12 to CMF15 pertain to the ability of CM approaches to deal with properties of data that are determined by the means of recording and storing the data and the properties of data sources.

CMF12: integration of data from multiple sources

Overview: Real-world processes can be supported by multiple information systems at once thereby requiring the CM approach to integrate the data from each of the information systems to construct the complete processes. If context data is considered, it needs to also be integrated with the event data.

Description: When business processes are supported by multiple, distributed information systems, observing events by extracting them from the various information systems as a starting point for CM approaches can become challenging, since there exists a high degree of heterogeneity among the systems [111] leading to a need for subsequent integration. Reasons for companies to support their processes with multiple information systems can be specialized software or external systems that support the process, e.g. after outsourcing [77]. If these information systems do not follow the same data governance, integration becomes inevitable [21].

Evaluation criteria: To fully support the data integration functionality, the approach must be able to support various event extraction techniques and be able to transition from basing predictions on label equivalence to equivalence notions based on activity semantics, e.g., attribute equivalence [124] and integration of other process perspectives and case ids.

Examples (cf. Ex. 1):

- The activities of the transaction reporting process can either be explicitly recorded through the underlying events in information systems such as Murex² or Bloomberg³ or be implicitly recorded as database transactions in ERP systems such as SAP⁴. The various data sources together with additional data sources for context data have to be integrated.
- “If the loan request is greater or equal to one million, the solvency level of the customer needs to be at least A, a manager needs to process the request, and the solvency information must not be older than two days. [...] the information necessary to check this rule is distributed across multiple systems [88].

Implementation: The CM approach can address the challenge of multiple, potentially heterogeneous sources by supporting various event extraction techniques, equivalence notions and event abstraction mechanisms.

²<https://www.murex.com/>

³<https://www.bloomberg.com/>

⁴<https://www.sap.com/>

CMF13: distributed processes

Overview: Compliance in process choreographies refers to the fulfillment or violation of constraints at different levels, i.e., the global choreography level, the local private process level, a mixture of both, and assertions [53], i.e., constraints might span across multiple partners of the choreography. In addition, constraints might refer to one or several choreography instances (and span across multiple partners at the same time).

Description: Constraints in distributed settings span across several partners in a process choreography and additionally might span one or several choreography instances. Due to the distribution of the partners, the events of process choreographies are typically recorded in multiple information systems. Hence, CMF13 is related to CMF12 and predicting compliance in distributed settings is a task that typically cannot be performed at one partner’s side, but across several partners in the choreography. Additionally, compliance prediction has to deal with confidentiality issues if private processes of partners are affected. Overall, dealing with process choreographies poses new challenges for the CM approach, e.g. the prediction of and checking of constraints under privacy and confidentiality issues [53] and the realization of the PCM system with respect to how the prediction is actually performed (at one partner’s side, in a distributed way?).

Evaluation criteria: To fully support the distributed processes functionality, the approach must be able to transition from basing predictions on label equivalence to equivalence notions based on activity semantics, e.g., attribute equivalence [124] and integration of other process perspectives, case ids, and message ids. Additionally, the approach must be able to provide compliance predictions on the event streams/compliance constraints with confidentiality requirements constituted by, e.g., hidden private process information.

Examples (cf. Ex. 1): All eight reporting processes run concurrently. To assess and predict the overall report data collection and consolidation performance and completeness for internal reporting purposes, equivalence notions based on activity semantics have to be employed.

Implementation: So far only few approaches have addressed the challenges that come with compliance in process choreographies, e.g., [53], and approaches for predicting compliance in distributed processes are missing. Future research efforts should particularly target settings with private information, as these settings can be common in competitive markets.

CMF14: context data

Overview: Context data, i.e., data that is external to the process data, has proven useful in enabling or enriching CMFs such as root cause analysis and explainability (cf. Sect. 2), so CM approaches should come with the ability to deal with such data.

Description: Taking context data for pro-active detection and management of compliance violations into account includes the ability to identify relevant context data, relating the context data to the events/activities that are actually affected or in that context represented by the data, leveraging context data for improving predictions (e.g. potentially enabling predictions for processes whose

events are private) and considering them as further evidence for root causes and for explaining what the compliance violation means. By definition, it is likely that context data is not recorded in the same information system as the process, which relates this functionality with CMF12.

Evaluation criteria: To fully support the data integration functionality, the approach must be able to continuously exploit context data for compliance predictions, particularly for predictions at the presence of unseen process and data behavior and for predicting unseen context data behavior. Furthermore, the approach must be able to continuously exploit context information for explaining prediction results.

Examples (cf. Ex. 1): In the case of major external events such as the beginning of the Ukraine war in 2022, trading activity in financial markets can be higher than usual leading to an increase number of orders that need to be reported, putting a significantly higher burden on the reporting process (cf. Ex. 1). In this case, the prediction of timestamps and durations is affected and therefore, the possible violation of the timeliness constraints.

Implementation: For each ability of dealing with context data, there exist proposed approaches, e.g. context ontologies for processes can guide the identification of relevant context [151, 33, 137, 90, 114]. Hence, the CM approach may either connect existing approaches in a meaningful way or opt to develop new methods for each required step.

CMF15: data properties and quality

Overview: Properties of the process and context data may require data transformations or suitable relation mechanisms (cf. CMF 12) as part of data preprocessing [51], whereas the quality of data can lead to the exclusion of data attributes during preprocessing. In the context of CM, excluding data for quality reasons may limit the monitoring to events of process instances that are more likely to be compliant.

Description: For CM, a crucial data property is transparency, i.e., to understand which properties of data lead to which effects during execution of the CM approach to facilitate the understanding of the results (cf. CMF9.x). As the goal of CM is to reactively and pro-actively detect and manage compliance violations, events with data quality issues can signify compliance violation. Consequently, CM approaches should avoid preprocessing to improve data quality, in particular during runtime, but rather have the ability to flexibly deal with the recorded data quality.

Evaluation criteria: To fully support the data properties/quality functionality, the approach must be able to consider and exploit properties and quality of the input event streams, interpret data (quality) properties with respect to prediction results and elaborate strategies for dealing with data quality properties and problems under prediction result quality guarantees.

Examples:

- Since data quality issues such as erroneous values or wrong formats either lead to more "Correct order field entry" activities or to amendment/correction messages to be included in the next report [44], they affect both what activities in the reporting process occur and their performance.

- “For example, although the average number of activities in the EnvLog dataset is only 44 (compared to 20 for BPI’12), the dataset only provides 787 instances for 331 possible activities resulting in a high sparsity of 0.42, whereas BPI’12 has a comparably low sparsity of 0.0028.” [69].

Implementation: So far research has yet to address how to exploit data properties and quality for CM.

PCM system (prediction) requirement list based on data functionalities: Compliance prediction with respect to data CMFs requires approaches for addressing and optimizing the performance of compliance predictions, especially for online predictions and at the presence of a multitude of compliance constraints and event data (\mapsto volume and velocity of the input data). Moreover, the heterogeneity of the input data, i.e., compliance constraints and event data from multiple sources, has to be addressed by novel data integration methods (\mapsto variety). This challenge is aggravated for distributed processes as the input data might also be subject to confidentiality requirements, resulting in partly hidden, invisible data. Finally, compliance prediction results are to be examined with respect to the properties and quality of the input data (\mapsto veracity).

4 Assessment of Predictive Compliance and Process Monitoring

In this section, we analyze the 122 papers from the PPM literature compilation plus the 23 papers added by snowballing for their coverage of the extended CMF framework presented in Sect. 2.5. We structure the section along the extended CMFs and discuss the relation between supporting CMF groups and their prediction requirements (cf. Sect. 3) and the prediction goal of the PPM approaches. The goal is to understand the relation between CM, PPM, and PCM (cf. Sect. 4.5) as well as to identify open research challenges (cf. Sect. 5).

The assessment of the extended CMFs and their prediction requirements is done as follows: + indicates that the prediction requirements for the CMF are fully met by existing approaches, \sim means the prediction requirements are partly met and $-$ that the prediction requirements are not met. We use c to express that the requirements can be met by combining different existing PPM approaches, e.g., predicting resources is connected with predicting next activities. An assessment summary is provided at the end in Table 1.

Note that we present an assessment that considers seen behavior, i.e., we assume that we have observed all behavior already in a historic log/stream. For the case of unseen behavior the assessment would for almost all CMFs evaluate to $-$ as only few approaches such as [116] can deal with unseen next activity prediction in the form of updating the prediction model.

4.1 Assessment of modeling requirements CMF0.x–CMF3.x

As discussed in Sect. 3, compliance violations can be predicted basically in the following two ways:

- i) formulate the prediction goal as the truth value of the compliance constraint (*predicate prediction*)
- ii) formulate the prediction goal as next (set of) events/activities potentially in combination with time or data related prediction goals and subsequently check compliance constraints (*next activity prediction “plus”*).

The assessment of *i) predicate prediction* approaches with respect to modeling requirements in general depends on the expressiveness of the predicate language, i.e., which process perspectives (control flow, time, data, and resources) can be expressed as predicates. In the literature compilation, we identified 3 distinct approaches that deal with predicting possible violations of predicates that range from simple SLAs [93] to LTL based formulae [99], and first-order event expression (FOE) based formulae [130]. The predicates refer to control flow, time and event attributes. Hence, CMF0.1–0.3 are covered (+) if the activities have been already observed. [99] is also able to deal with quantitative time prediction (CMF1). [130] explicitly deals with event attributes, i.e., activity data (CMF2.1, +), and resources (CMF3.x, +)

ii) next activity prediction “plus”: Next activity / event prediction means to make statements about upcoming activities / events that are referred to by one or several compliance constraints. Take, for example, constraint “*If Correct order field entry activity is executed for a an order of the previous or even older day, then it must never be reported as a new order, but as an amendment.*” (cf. Ex. 1), which refers to activities *Correct order field entry* and *Send secured MMSR report*. Existing approaches predict next activities if the activities in the compliance constraint have already been observed so far. Absence of an activity can then also be implicitly predicted, based on probabilities. Consider for example a compliance constraint stating “*b must not directly occur after a*”. In this case, we would expect that for a trace in which we have observed *a* the probability of *b* as next activity should be 0. If it is not, we could end up with a compliance violation. Hence, next activity prediction approaches support CMF0.1, CMF0.2, and CMF0.3 (+).

Next activity / event prediction could also serve as “anchor” for predicting the modeling requirements CMF1–3.x referring to time, data, and resources by combining next event/activity prediction with remaining time and resource prediction. Except for [122] that predicts the future path of a trace and then the delay to the next event starting from the current event, combined approaches are missing. Instead time, data, and resources are used by existing approaches as features to improve next activity prediction.

Finally, there are dedicated approaches for predicting remaining time, numeric indicators, and resources. Existing approaches for remaining time/delay in combination with next activity / event prediction cover CMF1 on quantitative time (*c/+*). In [28], activity data (CMF2.1) is used for prediction, i.e., cost for executing tasks is seen as a risk parameter. Resource prediction (CMF3.1 and CMF3.2) in connection with temporal prediction is mostly seen from a scheduling perspective, i.e., how to determine and avoid potential temporal problems such as bottlenecks by assigning resources [128, 134]. Other approaches utilize

resources as features for temporal predictions [58, 83]. [19] predicts the resource / resource pool an upcoming event will be assigned to. Hence, CMF3.1 and CMF3.2 can be assessed with $c/+$.

4.2 Assessment of execution requirements CMF4–CMF6.x

If the input event stream contains different event types (CMF4, CMF5) such as **start**, **complete**, or **running** (cf. life cycle model for XES [1]), the corresponding event labels are encoded as features, but a distinction of any kind of semantics of the event types is missing (–).

Existing PPM approaches to address CMF6.1 on the multiple instantiation of compliance constraints are missing.

For the simultaneous execution of process instances (CMF6.2), inter-case features for batching (i.e., executing process instances in one batch) are used in order to improve remaining time predictions [133, 84, 120] (c/\sim). Aggregated PPIs that might include data constraints can be predicted based on [31] as well as aggregated risks over multiple instances by [28] (CMF6.3, c/\sim). [56] propose a probabilistic approach for remaining time prediction taking into account hidden dependencies between process instances (CMF6.4 \sim). CMF6.3 and CMF6.5 are not covered by temporal prediction approaches (–). CMF6.6 is only touched upon, i.e., [58]s predict how many instances will start in a particular time window (c/\sim).

4.3 Assessment of user requirements CMF7–CMF11

In general, it is often difficult or impossible to identify the output of existing prediction approaches. However, the prediction requirements state to at least predict the set of next activities / events, ranked by probability of occurrence and a distinction between immediately/eventually occurs. If we look at more complex compliance constraints referring to several activities and their occurrence/absence and order, possibly in combination with time, data, and resources, a fine-granular prediction feedback with probabilities would be desired, which is basically possible, but not explicitly provided by any of the approaches.

Regarding CMF7, descriptive monitoring approaches [11, 50] constitute a means to propose recovery or compensation in the form of interventions to the user and provide feedback (+).

Regarding CMF8.1 and CMF8.2, predicate prediction approaches [99, 93, 130] can basically provide early detection of (future) compliance violations (+). Updates of prediction results, especially compliance violations (CMF8.3), is addressed in a preliminary way by incremental learning approaches that focus on updating the prediction model when drifts occur, e.g., [116, 101] as well as based on continuous task monitoring through sliding windows as proposed in, e.g., [31] (c/\sim). For CMF8.4, approaches mention that predictions can provide recommendations for users [99]. Alerts [18] and dashboards [118, 55, 136] provide information on predictions and compliance to users and can help to

avoid violations. Risk predictions [28] are provided to users as recommendations (CMF8.4) which can also partly serve as mitigation actions for lowering risk for specific risk types. However, all of these recommendations do not target compliance violations (–). Prescriptive monitoring approaches can foster the early detection of compliance violations (CMF8.1) and the preparation of mitigation actions (CMF8.4). [11, 50], for example, enable the generation of alarms that trigger interventions to prevent an undesired outcome or mitigate its effect. Similarly, [158] “*supports the proactive handling of deviations, i.e. inserted and missing events in process instances, to reduce their potential harm*”. However, both approaches do not provide any recommendations or mitigation action, in particular, with respect to a set of constraints and not KPIs or SLAs.

Root cause analysis (CMF9.1) can be implicitly based on probabilities and feature vectors, i.e., by answering the question whether certain data elements influence the prediction of the next activity/event [74, 62, 67] (c/\sim). However, the effective communication of root causes to users (CMF9.2) is missing (–), although explainability (CMF9.3), e.g., based on features, is targeted by several approaches recently [140, 62, 31, 105, 67] (c/\sim). CMF9.3 on explaining and visualizing results is implicitly supported via helping to choose parameters by [6] and by [84] in the context of inter-case features for batching. Quality metrics for the prediction results are provided, including stability [82] and reliability [26]. However, existing approaches lack the ability to explain the prediction results in natural language and by means of visualizations that both are comprehensible by domain experts without knowledge of the prediction method or without a mathematical background (c/\sim). In particular, visualization approaches for explaining prediction results, predicted violations (CMF9.4), and the effects of mitigation actions (CMF9.5) are missing (–).

For the assessment of compliance degrees, there is no approach for single instances (CMF10.1, –). By providing temporal predictions for inter-case features, [133] can contribute partly to CMF10.2, i.e., the prediction of the compliance degree (‘healthiness’) of an entire process (c/\sim).

Regarding efficiency and performance (CMF11), first approaches contribute by applying, for example, scalable online learning algorithms, [122], hyperparameter optimization [59], and temporal predictions “*in a parallel and distributed manner, on top of a cloud-based service-oriented infrastructure*” [22], yet fail to develop a case study that elicits user requirements on the performance/efficiency of the system and comprehensively benchmarks existing approaches based on the case study; only for outcome-oriented approaches a benchmark exists for execution times [146] (\sim).

4.4 Assessment of data requirements CMF12–CMF15

For CMF12 on the integration of data from multiple sources, [117, 145] deal with structured and unstructured data, i.e., textual data, as input for PPM (\sim). There are no approaches for predictions in distributed processes (CMF13, –).

The potential of contextual data (CMF14) is mentioned by several approaches. Internal contextual data is exploited by existing approaches such as [15, 74] by encoding them as features for next activity/event prediction. Internal context data is also utilized by [57] for temporal predictions ($c/+$), but none of the existing approaches exploits external context data ($-$).

The influence of data quality and properties (CMF15) is considered by first approaches that consider data properties [69] and deal with small data sets [78] (\sim) as well as by [85] with respect to reliability of the predictions (\sim).

Class	CMF	Coverage
Modeling req.	CMF0.1 existence	+
	CMF0.2 absence	+
	CMF0.3 ordering	+
	CMF1 time quantitative	$c/+$
	CMF2.1 activity data	+
	CMF2.2 case data	-
	CMF3.1 unary resource condition	$c/+$
	CMF3.2 extended resource condition	$c/+$
Execution req.	CMF4 non-atomic activities	-
	CMF5 life cycle	-
	CMF6.1 multiple instantiation of compliance constraints	-
	CMF6.2 simultaneous execution	c/\sim
	CMF6.3 constrained execution	c/\sim
	CMF6.4 order	\sim
	CMF6.5 non-concurrent execution	-
	CMF6.6 constrained start	c/\sim
User req.	CMF7 reactive management	+
	CMF8.1 early detection of conflicting rules	+
	CMF8.2 possible/future violation	+
	CMF8.3 update set of possible and future violations	c/\sim
	CMF8.4 recommendations for users to avoid violations	-
	CMF9.1 root cause analysis	c/\sim
	CMF9.2 effective communication of root cause	-
	CMF9.3 explain and visualize prediction results	c/\sim
	CMF9.4 explain and visualize set of possible and future violations	-
	CMF9.5 explain and visualize effects of mitigation actions	-
	CMF10.1 compliance degree of single traces	-
CMF10.2 compliance degree of an entire process/system	c/\sim	
CMF11 efficiency/performance of CM	\sim	
Data req.	CMF12 integration of data from multiple sources	\sim
	CMF13 distributed processes	-
	CMF14 context data (internal external)	$c/+ -$
	CMF15 data properties and quality	\sim

Table 1: Coverage assessment of extended CMFs by existing literature; CMF extensions in bold; +:covered, \sim :partly covered, -:not covered, c :combination necessary

4.5 Setting the PCM problem space

Table 1 summarizes the CMF assessment by existing PPM approaches. Modeling requirements CMF0.x–CMF3.x are supported by existing PPM approaches except for CMF2.2 regarding case data, through predicate prediction or approaches for next activity prediction (CMF0.x), temporal prediction (CMF1), indicator prediction (CMF2.1), and resource prediction (CMF3.x). While predicate prediction can target several modeling requirements at the same time depending on the expressiveness of the predicate language, a combination of

next activity prediction with, for example, temporal prediction becomes necessary and has been only addressed partly. Execution requirements CMF4 and CMF5 have not been addressed yet due to the missing exploitation of activity and instance life cycle information for prediction. CMF6.1 is also missing as the constraints and their instantiation has not been of interest for PPM approaches yet. CMF6.2, CMF6.3, CMF6.4, and CMF6.6 can be partly supported, e.g., by inter-case data, and through combination of the approaches. However, PPM approaches have not dealt with CMF6.5 so far. User requirements CMF7 and CMF8.1–CMF8.3 are supported by existing approaches by providing reactive measures in case of compliance violations and being able to predict compliance violations. There is at least partial support for root cause analysis (CMF9.1) and the explanation of prediction results, mainly based on visualizations (CMF9.3). CMF8.4, CMF9.2, CMF9.4, CMF9.5 are not supported, i.e., no approaches for determining recommendations and mitigation actions as well as the explanation of future violations exist. Regarding the compliance degree, approaches partly support statements for the entire system (CMF10.2), but not for single traces (CMF10.1). CMF11 is partly supported regarding the management of data volume. Regarding data requirements, existing approaches deal with internal context data (CMF14.1) and consider first aspects of input streams from multiple data sources and the data quality (CMF12 and CMF15). External context data (CMF14.2) as well as compliance prediction in the context of distributed processes (CMF13) has not been considered so far.

We can conclude from the analysis that several CMFs are not supported by existing approaches yet or require at least a combination of existing PPM approaches. Supporting only one or a small subset of the CMFs is not sufficient for a comprehensive PCM solution/system. At first, real-world PCM problems are *complex* as demonstrated by the example in Sect. 1 and hence require the support of a multitude of CMFs in their combination. Moreover, PCM is a *continuous problem*. As a consequence, a comprehensive PCM system has to support CMFs in a continuous way by continuous compliance prediction and re-evaluation of previous results.

Another conclusion is that there are basically two approaches to tackle PCM through PPM, i.e., *i) predication prediction* and *ii) next activity prediction “plus”*. So far, PPM research has mainly considered i) for compliance prediction. However, i) bears several disadvantages when put to test in complex and continuous real-world PCM problems. In a real-world PCM problem, multiple constraints and multiple instances of these constraints are monitored simultaneously. Predicting compliance violations then requires multiple models, one per constraint/predicate. Moreover, as PCM is a continuous problem, whenever a compliance constraint changes we need to immediately update the corresponding model. If a new constraint comes into effect we need to formalize the corresponding predicate and train a new model for that specific constraint. Moreover, root cause analysis is limited for predicate prediction because knowledge on related events and data cannot be created and exploited.

Compared to i), option ii), i.e., using PPM approaches as enabler for PCM by combining them with actual monitoring capabilities, is preferable in many

ways. Though model updates cannot be avoided in this case as well, we can still reduce the number of models that need to be trained. In particular, consider we have multiple control flow related constraints that we need to check. In this case we could just train one next activity/event prediction model and apply different monitors for each of the constraints. This is way more feasible than having to manage different predicate prediction models at once. Moreover, the results of the predictions can be used for other goals that a company needs to achieve, i.e., we can reuse prediction models for other purposes.

Overall, we understand PCM as a system that helps a company to understand its future compliance status which encompasses and supersedes all existing approaches. Understanding PCM as a system results in approaches claiming to develop or implement PCM that these approaches must cover all the extended CMF functionalities and, ideally, also come with a real-world case study that is not simulated using event logs demonstrating the system in practice. Existing event logs such as logs from the Business Process Intelligence Challenge recently organized by the IEEE Task Force on Process Mining⁵ do not challenge the PCM system in many of the required functionalities, e.g. BPIC logs are already extracted from the information systems. To implement a PCM system, no or almost no abstractions or simplifying assumptions should be made about reality, as these assumptions may decrease the applicability and value in practice. Hence, existing PPM and CM approaches are methods for PCM and only cover fractions of the required functionality (as depicted in Tab. 1). It remains yet to be shown, how the existing approaches can be combined and applied in a case study (cf. Sect. 5). If a method is developed that falls into PPM or CM and can be used for PCM, we advocate to phrase it as such that the method helps to realize a PCM system, but in itself is neither PCM nor a PCM system.

5 Open Challenges and Research Directions

Section 4 shows the potential of existing PPM approaches for a comprehensive PCM support. This section summarizes the open PCM challenges and set out research directions for PCM and PPM along the four classes of the extended CMF framework.

Modeling requirements

1. Holistic modeling and prediction

Challenge: The modeling and prediction of time, resources and data is tied in with next activity prediction “plus” or is implicitly predicted through the predicate in predicate prediction and is mostly covered by existing PPM approaches. What is still missing is the modeling and prediction of case data and a holistic approach that encompasses all modeling functionalities in a semi-automatic way, thereby taking the burden of stating constraints in some logic from the user.

⁵<https://www.tf-pm.org/competitions-awards/bpi-challenge>

Research direction: For compliance constraints, the simultaneous prediction of events plus time, resources and data is crucial. Assume, for example, a compliance constraint from the logistics domain on a transportation process that states if the transport takes up to 5 hours, the destination of the transport is 'London', if the transport takes longer than 5 hours, the destination is 'Berlin'. Predictions regarding this constraint necessitate the prediction of remaining time plus location. Additionally, data value prediction might not only refer to process data, but also to contextual data such as time series, for example, if decision rules are based on time series data such as temperature [131]. Here, the combination of PCM with time series prediction approaches constitutes a promising research direction [72]. Altogether, future work should focus on complex examples from reality (cf. Ex. 1) and try to avoid introducing simplifying assumptions or abstractions as much as possible, as these might limit the research to only a subset of modeling requirement CMF combinations. In particular, research should avoid assuming the constraints to be readily available and stated in some logic.

Execution requirements

2. Life cycle handling

Challenge: None of the existing PPM approaches exploits the life cycle of activities, i.e., exploits the semantics of distinct life cycle states/transitions of activities in the event stream. However, these life cycle states/transitions might contribute to predict, for example, the activity duration or might indicate exceptional behavior, resulting in unseen behavior or drift, and subsequently necessitating adequate mitigation actions.

Research direction: Incorporating and exploiting life cycle states and transitions into PPM and PCM might tremendously increase prediction quality and applicability in real-world settings. Consider, for example, the transportation of goods to different locations. By distinguishing the **start** and **complete** events of activities, activity duration can be considered in the prediction. By exploiting more life cycle states such as **suspend** and **abort**, upcoming exceptions might be predicted and exception handling actions be defined and taken. Assume that, for example, for activity 'transport' a **suspend** event occurs. This might result in delay which should be incorporated in a temporal prediction, e.g., of the remaining time of the affected instance. If an **abort** event occurs for 'transport', we can conclude that the transport will not be completed (i.e., event **complete** will not occur for 'transport') and this might result in a compliance violation.

3. Instance and process spanning constraints

Challenge: Predicting the compliance of constraints that span multiple process instances and/or multiple processes has not been explicitly addressed by existing approaches, except for using inter-case features for predicting next activities or predicting performance indicators across multiple instances. However, many application domains crave for compliance

support in instance and process spanning settings, e.g., logistics, medicine, and manufacturing.

Research direction: State-of-the-art PPM approaches mostly focus on predicting next activity/event within the context of single instances. However, when considering compliance constraints that span across multiple processes and process instances, it becomes necessary to predict interactions between the affected process instances and their behavior, as well. In particular, such compliance constraints refer to data and/or resources shared by processes/instances. Take as an example the compliance constraint: “Each clerk is allowed to issue approve loan as long as a threshold (around \$1M) is not reached. Otherwise he has to delay this event to the following day” [159]. First of all, we can see that the constraint imposes a condition across several process instances that refers to data element ‘threshold’ and implicitly to time (‘within one day’). Hence, compliance predictions across multiple processes and process instances have to consider a combination of the control flow, data, time, and resource perspective. Moreover, constraints spanning multiple instances and processes typically state and trigger actions, e.g., delaying instance execution until the following day. Incorporating the effects of this behavior imposed by the constraints into the prediction is of utmost importance.

User requirements

4. Provision of mitigation actions

Challenge: Though recommendations are used to support users in taking counteractions regarding delays or other risks, none of the approaches suggests mitigation actions to overcome compliance violations. In particular, approaches are missing that provide mitigation actions at different granularity levels, analyze and visualize the effects of applying mitigation actions, and provide users with estimations on their significance.

Research direction: Based on compliance violation predictions combined with root cause analysis, the effects of mitigation actions can be assessed; either by simulating what will happen if a user applies a specific countermeasure or by determining and suggesting mitigation actions for avoiding the compliance violation. Consider again the transportation example provided in research direction *life cycle handling*. Assume that based on data gathered before and during transportation, the transportation is predicted to be aborted for a certain process instance. Based on the prediction, we can immediately start to define countermeasures for avoiding the compliance violation of not arriving at the destination, together with predicting their effects. One possible mitigation action in this case is to start another transportation process arriving on time. The prediction can then estimate whether or not the application of this countermeasure compensates the failure. By taking the results of the estimation and an ontology or a knowledge graph either on the domain, i.e. in this case logistics, or on the company executing the process (cf. research direction on explainable, supervised machine learning [17]), future work may come up with novel

ways to explain the mitigation actions and their effects in an understandable natural language or interpret them as explainable events (cf. event management in [23]).

5. Visualization and explanation of predictions and violations

Challenge: Explainability of prediction results has gained attention. However, visualization approaches for prediction results, especially future compliance violations are mostly missing. Moreover, root cause analysis has to be extended in order to deal with predicting violations of real-world compliance constraints. The challenge becomes even more difficult, when the approach does not stop at some metric quantifying the individual impact of features on the result and visualizing results, but start to tell a story beginning at the root cause and ending at the compliance violation, as coherent, logical storytelling is the means of explaining for humans [17]. *Research direction:* PCM requires an aggregated view on several perspectives, including the compliance constraints and the process respectively process instance perspective, i.e., a view on the current event stream combined with continuous updates. Though some approaches already provide visualizations for simple SLA violations including color coding, e.g., red means the SLA is violated, green the SLA is not violated, there is still room for improvement and extensions when considering complex compliance constraints. Therefore, visualization approaches are required to depict possible complex information at once, i.e., all compliance states for all processes and process instances captured by the event stream, as well as the definition and visualization of single views, e.g., visualization of compliance predictions for one constraint, one particular instance, or one perspective such as time. Moreover, information on root causes for compliance violations, the current prediction model in use, and mitigation actions together with their effects should be conveyed to users based on visualization approaches. Through visualization, the PCM system can start to connect trained classifiers with relevant constraints that facilitates fine-grained root cause analysis. Consider PPM approaches that train classifiers based on available data and use them to, e.g., predict the next event/activity. Consider, e.g., a compliance constraint c stating that “for premium customers any loan request below € 100 will in any case be granted without performing an additional check”. If a decision tree algorithm is trained to predict the next activity based on event attributes *customer type* and *requested loan amount*, we would be able to determine the underlying decision rule, i.e., $(customer\ type = premium \ \& \ requested\ loan\ amount < 100) \rightarrow grant\ loan\ immediately$, otherwise *perform additional check*. In order to enable a fine-grained root cause analysis, we need to link this decision rule to compliance constraint c . At last, the PCM system helps the user in understanding the visualized results by supporting natural language explanations that connect the results in a logical way.

6. **Update:** To implement a PCM system, various functionalities for updating parts of the system have to be developed.

(a) **Treatment of unseen behavior and update**

Challenge: Activity occurrence, absence, and ordering of activities in compliance constraints is covered by existing PPM approaches for those activities that have been already observed. Unseen behavior remains largely uncovered. Unseen behavior can occur in event streams if the underlying process model is not or only partly known or due to concept drift. In combination with compliance constraints, the requirement to predict unseen process behavior becomes even more likely as the compliance constraints do not have to be part of either an underlying process model or the observed behavior in the event stream. Similar observations also hold for the treatment of unseen data and unseen data values (internal and external data).

Research direction: Recently, strategies on how to update the prediction model in case of unseen process behavior in the context of next activity prediction have been proposed, including “do nothing”, “retrain without hyperparameter optimization”, “full retrain”, and “incremental update” [116, 126, 102, 101]. In addition, we need to consider not just updating the model whenever unseen behavior has occurred, but also how to predict the unseen behavior as such, e.g., by considering available context data.

(b) **Online PCM, online (re-)training of prediction models**

Challenge: One situation is to train the prediction model based on available historical information and then constantly update it whenever new information in terms of events or constraints occurs. The more difficult situation is to start from scratch and having to learn and train without any previous knowledge. Both cases of learning/updating the prediction model are challenges for PCM. When considering context data, the decision on when it is necessary to update may be different.

Research direction: As both cases have been addressed by online process mining approaches with respect to new events, research on PCM has to investigate the challenges based on existing work in online process mining. However, approaches to consider new constraints or constraint evolution are missing. Moreover, strategies for updating prediction models at the presence of continuous context data have to be elaborated, e.g., constant updates versus updates if significant changes in the context data occurs.

(c) **Continuous update of prediction results and compliance violations**

Challenge: The prediction results and compliance violations have to be updated continuously as new arriving events reveal information about the actual progress of the monitored process(es). Here, the challenge is to investigate whether the performance of the PCM system allows to update all results for each new incoming event or whether batching of incoming events is necessary. These considera-

tions have to take the (end) user requirement on the performance of the approach into account. The challenge becomes more difficult, if the evolution of the constraints base through, for example, changing regulatory documents (cf. Ex. 1) is also taken into account.

Research direction: This challenge requires thorough case studies that result in a comprehensive benchmarking environment for PCM systems. The case studies should come with time series data on requirements (simulating the evolution of requirements during development and operations of a PCM system), regulatory documents, relevant and irrelevant context and the raw events from various information systems, possibly also from various processes even spanning multiple partners. Only then can proposed PCM systems be properly compared.

(d) **Compliance degree and update**

Challenge: First approaches for predicting the compliance degree across multiple process instances have been presented. Yet, especially in combination with updating compliance violations, an open challenge remains how to define and update the compliance degree while new events arrive throughout the event stream and to predict compliance states of single instances. Furthermore, it remains unclear how these instance-level compliance states and degrees can be lifted to the process or choreography level. As prediction results include probabilities, e.g., how likely is the occurrence of a certain activity, these probabilities should be considered for compliance states and degrees.

Research direction: Following [100], currently we distinguish compliance states *possibly violated/satisfied* and *violated/satisfied* for single instances. *Possibly violated* means that the violation can be still healed, i.e., by the occurrence of an activity that is mandatory according to the compliance constraint (cf. [98]). A (final) violation, by contrast, states that the constraint cannot be healed anymore, e.g., if the constraint is *possibly violated* and then the end event of a process instance or all end events occur. There is also a distinction between *fully* and *partly violated/satisfied* where full violation of a compliance constraint means that this constraint is violated for all process instances, and a partly violation means that the constraint is violated for at least one constraint [150]. It is unclear whether the proposed compliance states are actually sufficient in light of the prediction and continuous update performed by PCM systems. In particular, the presence of probabilities on prediction and, thus, compliance state results expressing the reliability and certainty of the PCM system on the results may necessitate further compliance states that communicate relevant aspects of the results to the end user in a compact and understandable way. For more complex compliance constraints referring to activities, data, time, and resources, the probabilities

of satisfying/violating these constraints have to be calculated in an adequate manner, e.g., how likely is the occurrence of a certain activity producing a certain data value in a given time span? If these probabilities can be determined, in turn, the risk of violations can be assessed and the compliance state properly identified. Compliance states may also need to express boundary considerations (based on the information we have, what has happened in the best and worst case at the partner?) when considering private processes in process choreographies. In connection with the challenges on data quality, a compliance state may need to also reflect the uncertainty on the process or context data, as data sources such as sensors can carry information on their accuracy. For uncertain process data, PCM system research should investigate based on recent work on data quality in the data preparation phase [164] or in conformance checking [54]. By defining a suitable aggregation mechanism, work on compliance states and degrees can be lifted to the process or choreography level. Somewhat orthogonal to the aforementioned lines of research is the question on how to update the compliance degrees and states as prior parts of the PCM systems (see above) are updated. This remains unsolved so far.

Data requirements

6. Heterogeneous data from distributed processes and (contextual) data sources

Challenge: First approaches start addressing data requirements in PPM and PCM. However, approaches for event and constraint data from distributed and heterogeneous sources and processes are missing. Moreover, the ongoing exploitation of contextual data is promising, but under-researched yet.

Research direction: Future research on heterogeneous data for PCM may have to borrow techniques from data integration research that deal with the problems of, for example, entity recognition, data fusion or schema alignment [34] before the line of research that tackles the problems of event extraction, event abstraction and handling event data that does not come with a single case identifier can be followed. These problems are likely to occur, if relevant data on processes is distributed across multiple information systems either within a company or among partners. These problems are aggravated when compliance constraints are not stated in a machine-readable format carrying semantics of a logic such as LTL. Typically, constraints are stated in natural language and scattered across multiple regulatory documents (cf. Ex. 1). Then, the regulatory documents are further data sources that are challenging to include, as constraints have to be extracted in a way that we can match events/activities with them. Including context data potentially coming from other sources than the process data into PCM can significantly increase prediction capabilities and quality. In manufacturing processes, for example, several sensor data streams are

measured continuously that report the environment state/context of the process, e.g., the temperature of a room/machine or the fluid level in the machine. Detecting deviations in the context data can increase prediction effectiveness, e.g., concept drifts might be predicted early [139]. This can be furthered by augmenting PCM by predictions of the context data.

7. Prediction and compliance in process choreographies

Challenge: As there can exist constraints for the whole process choreography spanning multiple partners, predicting compliance in this distributed setting is a task that typically cannot be performed at one partner's side. Hence, process choreographies pose new challenges for PCM such as how the prediction and checking of constraints is actually performed (in a distributed way?). If private processes of partners come into play, the prediction of and checking of constraints has to adhere to privacy and confidentiality requirements.

Research direction: Although first approaches enable to check compliance in process choreographies (e.g., [53]), the problem of predicting compliance in process choreographies has not been addressed so far. One challenge arises from the confidentiality requirements of the private partner processes: if they are affected by compliance predictions, it might become necessary to distribute the prediction among the partners. This endeavour may be realized by means of *secure multi-party computation* [95].

8. Data properties and quality

Challenge: The exploitation of data properties and quality is promising. One drawback of existing PPM approaches with respect to the input data is the assumption of label equivalence, i.e., the prediction are based on labels of events. Label equivalence is not sufficient, particularly when merging event streams from heterogeneous input sources (variety). Another challenge is the size of the input data which can be too small or too big (volume). For distributed processes, event streams might contain information on message exchanges between partners –how can they be exploited for prediction or being predicted themselves?– and might also contain hidden/invisible parts due to confidentiality requirements of the partners. Here, an initial hurdle is the lack of data sets. When considering data quality, the data values of low quality may signify a compliance violation or be used as a feature for prediction. Determining whether the former is the case and the latter is beneficial, is yet another challenge.

Research direction: To tackle variety challenges, PPM and PCM can benefit from equivalence notions that aim at the functionality of activities, e.g., attribute equivalence [88]. First approaches for dealing with volume challenges boost small data sets [79, 71]. Other approaches aiming at efficiency and performance of PPM and PCM with respect to both, volume and high velocity event streams are missing (note that for PCM also a large set of compliance constraints might exist). To deal with PCM

in distributed settings, the collection, provision, and preparation of (real-world) data for different process scenarios containing multiple perspectives remains an ongoing direction of research. With respect to data quality and its exploitation in a PCM system, approaches are missing. Here, PCM can investigate based on work considering data quality as uncertain data values [164, 54].

Overall

9. Systematic assessment of data mining/machine learning techniques

Challenge: In the light of a multitude of challenges and research directions and the lack of a single prediction technique implementing or supporting all extended CMFs, it remains a challenge to apply an existing or develop a new prediction technique that addresses all of the challenges and implements or supports all extended CMFs. Due to the many facets of the problem and functionalities required, it is likely that not each problem can be tackled by existing data mining/machine learning techniques.

Research direction: First, a systematic assessment of (existing) prediction techniques to address these challenges is required. As it is likely, that not all of the challenges for PCM can be solved by existing prediction techniques, it might be necessary to develop new techniques. Another possible direction could be to decompose the problem and reformulate it either as a *multi-view learning* problem or a *multi-task learning* problem [166].

6 Discussion and Conclusion

6.1 Discussion

The survey aims at shedding light on the current position of existing CM and PPM approaches with respect to the extended CMF framework and whether their combination PCM is actually a new research area to be followed to address the challenges derived from the assessment of existing CM and PPM approaches. We followed selected principles of conducting a systematic literature review [12] and adapted them in terms of incorporating existing surveys as basis whenever possible. In this spirit, we took the established CMF framework [97] and extended it based on more recent findings. Despite this careful method design, the following limitations can be identified.

- New approaches on PPM are published constantly. Therefore, one limitation of this work is that new approaches since the literature compilation in February might have been published which are consequently not covered within this paper yet. A search on Google Scholar with `allintitle:predictive process monitoring` and selecting papers after 2022 results in 31 hits⁶. From these 31 hits, this survey covers [125, 74,

⁶accessed 2022-10-20

83]. 28 papers are not covered, out of which 2 are out of scope following our search methodology in Fig. 3, 15 papers have been published as technical reports and 1 as PhD thesis. Looking at this most recent work, the majority of the approaches is concerned with explainability, some combined with data issues such as [39] (cf. CMF9, CMF14, and CMF15) and with updating the prediction model (cf. CMF8).

- The main focus of this paper is on prediction tasks and compliance monitoring. Hence, further related areas such as online process mining, concept drift detection, and anomaly detection approaches have only been considered if papers from these areas were detected during the systematic literature review. Online process mining is often geared towards concept drift detection. This work covers several concept drift detection approaches [101, 139], also in connection with updating prediction models at the presence of concept drift [126, 116]. Anomaly detection can provide insights to PCM. However, most anomaly detection approaches work offline, some can be applied on event streams, e.g., [87], but process anomaly predictions beyond the approaches studied in this work such as [9] are missing.
- If surveys exist for the investigated research areas, i.e., for CM and PPM, we used these surveys as a basis for our further literature analysis. Doing so might result in missing papers that have been published prior to the existing surveys and not having been treated by them. However, we conducted a full search without restricting the publication dates first and then compared the identified set of papers to existing surveys. Doing so, we limit the risk of missing out relevant prior work.
- The aim of this paper is not to assess PPM approaches in terms of machine learning or data mining techniques in detail, i.e., the goal is not to identify the PPM approach currently performing best. Instead, the goal is to provide a comprehensive outline and analysis of the PCM problem and how it is addressed by the current literature. Hence, at this point, we do not investigate or propose particular techniques from a technical point of view. This endeavour is left as a challenge for future work (cf. Challenge 9).
- Though there are case studies for compliance monitoring available (\mapsto Sect. 2.4), we still need to have a detailed look and investigate whether these are suitable for evaluating approaches tackling the mentioned research directions, i.e. they meet the requirements set in Challenge 6c for a benchmark case study.

6.2 Conclusion

This work provides a comprehensive analysis of existing CM and PPM research in light of a real-world example and in terms of the extended CMF framework and their combination PCM that entails a new research area. It has tackled

research questions RQ1 – RQ4 (cf. Sect. 1) as summarized in the following. In addition to findings on the relationship between CM, PPM and PCM, the study particularly provides findings on PPM and its capabilities.

RQ1: To what extent is the PCM system required by the financial institution in Ex. 1 addressed and solved by existing PPM and CM approaches? RQ1 is addressed by an extensive compilation of literature on PPM, CM and PCM. Based on analyzing the literature, we conclude that the PCM system in general has not been developed by now, i.e., no specific PCM approaches exist and no existing PPM or CM approach constitutes a PCM system.

RQ2: How are the existing PPM and CM approaches comparable in terms of CMF functionalities, in particular, in terms of prediction requirements necessary for the respective functionality? The selected literature from PPM, CM and PCM research emphasizes that the compliance monitoring functionalities as originally proposed by [97] in 2015 are still valid and can serve as requirements for the PCM system. The CMF framework is extended based on analyzing the literature compilation regarding PPM and CM directions after 2015, most prominently, towards the explainability of the prediction results and input data requirements. To further the predictive capabilities necessary for PCM, the extended CMF framework is analyzed for predictive requirements arising for each of the CMFs, illustrated by means of Ex. 1, and is subsequently used to assess existing mostly PPM approaches.

RQ3: Is PCM a new research area? Papers explicitly addressing PCM are missing. In general, PPM holds the capabilities to tackle PCM challenges and can be seen as an enabler for PCM system development. These capabilities are derived for each CMF functionality of the extended CMF framework and used to assess existing approaches. The assessment finds some capabilities to be (partly) supported, but there is no comprehensive solution for all PCM challenges, i.e., the assessment eventually results in a list of open PCM challenges. Hence, we conclude that PCM is a new research area that, on the one hand, puts more focus on predictive capabilities necessary within CM research and, on the other hand, brings necessary CM functionalities such as integrating heterogeneous data sources, handling distributed processes and more complex constraints and their monitoring to the forefront of PPM. In both existing lines of research, the system development aspect of PCM is also missing to a large extent, underlining the need to understand PCM as a new research area.

RQ4: Which open challenges and research directions remain for full PCM support? Based on the identified open PCM challenges, together with the PCM system (prediction) requirement list, research directions for PCM are elaborated. These research directions comprise the holistic prediction of activities with data, time and resources, an appropriate life cycle handling, a support for instance- and process-spanning constraints, the provision of mitigation actions, the explainability of compliance violation predictions, multiple functionalities

to update various parts of the PCM system containing a call for a thorough case study necessary to benchmark approaches, the capability to deal with process choreographies and the explicit treatment of data quality and properties, also for heterogeneous processes and data sources. All of these open challenges constitute key success factors for predictive compliance monitoring.

The research directions point to several future research opportunities. Working on the research directions will necessitate a comprehensive assessment of existing machine learning and data mining techniques and might result in the development of extended or even new techniques. Moreover, this work assumes that compliance constraints are formalized using some notion. In future work, we will incorporate the sources, e.g., regulatory documents, into predictive compliance monitoring. In order to evaluate and compare new techniques and approaches, appropriate data sets are crucial, i.e., event streams, contextual data, data from different sources and processes, and unseen data.

Acknowledgments

This work has been supported by Deutsche Forschungsgemeinschaft (DFG), GRK 2201 and by the Austrian Research Promotion Agency (FFG) via the Austrian Competence Center for Digital Production (CDP) under the contract number 881843.

References

- [1] IEEE standard for extensible event stream (xes) for achieving interoperability in event logs and event streams. *IEEE Std 1849-2016*, pages 1–50, 2016. doi:10.1109/IEEESTD.2016.7740858.
- [2] Amy Matsuo, Dan Click, Jeff Garfield, Juan Gonzalez, Mike Lamberth, Brent McDaniel, Anthony Monaco, Jaime Pego, Todd Semanco, Jennifer Shimek, Elizabeth Hammer, Karen Staines, and Joe Slaninka. KPMG 2021 CCO Survey. Technical report, KPMG LLP, 2021. URL: <https://assets.kpmg/content/dam/kpmg/cn/pdf/en/2021/10/kpmg-2021-cco-survey.pdf>.
- [3] Ahmed Awad, Sherif Sakr, and Amal Elgammal. Compliance Monitoring as a Service: Requirements, Architecture and Implementation. In *Cloud Computing*, pages 1–7, 2015. doi:10.1109/CLOUDCOMP.2015.7149636.
- [4] Ahmed Barnawi, Ahmed Awad, Amal Elgammal, Radwa Elshawi, Abdullah Almalaise, and Sherif Sakr. An Anti-Pattern-based Runtime Business Process Compliance Monitoring Framework. *Advanced Computer Science and Applications*, 7, 2016.
- [5] Ahmed Barnawi, Ahmed Awad, Amal Elgammal, Radwa El Shawi, Abdullah Almalaise, and Sherif Sakr. Runtime self-monitoring approach of

- business process compliance in cloud environments. *Cluster Computing*, 18(4):1503–1526, December 2015. doi:10.1007/s10586-015-0494-0.
- [6] Nico Bartmann, Stefan Hill, Carl Corea, Christoph Drodts, and Patrick Delfmann. Applied predictive process monitoring and hyper parameter optimization in camunda. In *CAiSE Forum 2021*, pages 129–136, 2021. doi:10.1007/978-3-030-79108-7_15.
- [7] Antonio Bevacqua, Marco Carnuccio, Francesco Folino, Massimo Guarascio, and Luigi Pontieri. A data-adaptive trace abstraction approach to the prediction of business process performances. In *Enterprise Information Systems*, pages 56–65, 2013. doi:10.5220/0004448700560065.
- [8] Kristof Böhmer and Stefanie Rinderle-Ma. Probability based heuristic for predictive business process monitoring. In *On the Move to Meaningful Internet Systems.*, pages 78–96, 2018. doi:10.1007/978-3-030-02610-3_5.
- [9] Kristof Böhmer and Stefanie Rinderle-Ma. Mining association rules for anomaly detection in dynamic process runtime behavior and explaining the root cause to users. *Inf. Syst.*, 90:101438, 2020. doi:10.1016/j.is.2019.101438.
- [10] Michael Borkowski, Walid Fdhila, Matteo Nardelli, Stefanie Rinderle-Ma, and Stefan Schulte. Event-based failure prediction in distributed business processes. *Inf. Syst.*, 81:220–235, 2019. doi:10.1016/j.is.2017.12.005.
- [11] Zahra Dasht Bozorgi, Irene Teinemaa, Marlon Dumas, Marcello La Rosa, and Artem Polyvyanyy. Prescriptive process monitoring for cost-aware cycle time reduction. In *Process Mining*, pages 96–103, 2021. doi:10.1109/ICPM53251.2021.9576853.
- [12] Pearl Brereton, Barbara A. Kitchenham, David Budgen, Mark Turner, and Mohamed Khalil. Lessons from applying the systematic literature review process within the software engineering domain. *J. Syst. Softw.*, 80(4):571–583, 2007. doi:10.1016/j.jss.2006.07.009.
- [13] Dominic Breuker, Martin Matzner, Patrick Delfmann, and Jörg Becker. Comprehensible Predictive Models for Business Processes. *MIS Quarterly*, 40(4):1009–1034, 2016. URL: <https://www.jstor.org/stable/26629686>.
- [14] Jens Brunk. Structuring business process context information for process monitoring and prediction. In *Business Informatics*, pages 39–48, 2020. doi:10.1109/CBI49978.2020.00012.
- [15] Jens Brunk, Matthias Stierle, Leon Papke, Kate Revoredo, Martin Matzner, and Jörg Becker. Cause vs. effect in context-sensitive prediction of business process instances. *Inf. Syst.*, 95:101635, 2021. doi:10.1016/j.is.2020.101635.

- [16] Bundesanstalt für Finanzdienstleistungsaufsicht. Bankaufsichtliche Anforderungen an die IT (BAIT), August 2021. URL: <https://www.bundesbank.de/resource/blob/863264/6e6c87e41efd6d70b48b587ec899b89b/mL/2017-10-bait-2021-data.pdf>.
- [17] Nadia Burkart and Marco F. Huber. A Survey on the Explainability of Supervised Machine Learning. *Journal of Artificial Intelligence Research*, 70:245–317, January 2021. doi:10.1613/jair.1.12228.
- [18] Cristina Cabanillas, Claudio Di Ciccio, Jan Mendling, and Anne Baumgrass. Predictive task monitoring for business processes. In *Business Process Management*, pages 424–432, 2014. doi:10.1007/978-3-319-10172-9\31.
- [19] Manuel Camargo, Marlon Dumas, and Oscar González Rojas. Learning accurate LSTM models of business processes. In *Business Process Management*, pages 286–302, 2019. doi:10.1007/978-3-030-26619-6\19.
- [20] Malu Castellanos, Fabio Casati, Umeshwar Dayal, and Ming-Chien Shan. A Comprehensive and Automated Approach to Intelligent Business Processes Execution Analysis. *Distributed and Parallel Databases*, 16(3):239–273, November 2004. doi:10.1023/B:DAPD.0000031635.88567.65.
- [21] Alfonso Castro, Víctor A. Villagrà, Paula García, Diego Rivera, and David Toledo. An Ontological-Based Model to Data Governance for Big Data. *IEEE Access*, 9:109943–109959, 2021. doi:10.1109/ACCESS.2021.3101938.
- [22] Eugenio Cesario, Francesco Folino, Massimo Guarascio, and Luigi Pontieri. A cloud-based prediction framework for analyzing business process performances. In *Workshop on Privacy Aware Machine Learning for Health Data Science*, pages 63–80, 2016. doi:10.1007/978-3-319-45507-5\5.
- [23] Xinhong Chen and Qing Li. Event modeling and mining: a long journey toward explainable events. *The VLDB Journal*, 29(1):459–482, January 2020. doi:10.1007/s00778-019-00545-0.
- [24] Giuseppe Cicotti, Luigi Coppolino, Salvatore D’Antonio, and Luigi Romano. Runtime model checking for SLA compliance monitoring and qos prediction. *J. Wirel. Mob. Networks Ubiquitous Comput. Dependable Appl.*, 6(2):4–20, 2015. doi:10.22667/JOWUA.2015.06.31.004.
- [25] Marco Comuzzi. Alignment of process compliance and monitoring requirements in dynamic business collaborations. *Enterprise Information Systems*, 11(6):884–908, 2017. doi:10.1080/17517575.2015.1135482.

- [26] Marco Comuzzi, Alfonso E. Márquez Chamorro, and Manuel Resinas. Does your accurate process predictive monitoring model give reliable predictions? In *Service-Oriented Computing*, pages 367–373, 2018. doi:10.1007/978-3-030-17642-6_30.
- [27] Marco Comuzzi, Alfonso E. Márquez Chamorro, and Manuel Resinas. A hybrid reliability metric for SLA predictive monitoring. In *Symposium on Applied Computing*, pages 32–39, 2019. doi:10.1145/3297280.3297285.
- [28] Raffaele Conforti, Massimiliano de Leoni, Marcello La Rosa, Wil M. P. van der Aalst, and Arthur H. M. ter Hofstede. A recommendation system for predicting risks across multiple business process instances. *Decis. Support Syst.*, 69:1–19, 2015. doi:10.1016/j.dss.2014.10.006.
- [29] Raffaele Conforti, Sven Fink, Jonas Manderscheid, and Maximilian Röglinger. PRISM – A Predictive Risk Monitoring Approach for Business Processes. In *Business Process Management*, pages 383–400, 2016. doi:10.1007/978-3-319-45348-4_22.
- [30] Alfredo Cuzzocrea, Francesco Folino, Massimo Guarascio, and Luigi Pontieri. A multi-view multi-dimensional ensemble learning approach to mining business process deviances. In *Neural Networks*, pages 3809–3816, 2016. doi:10.1109/IJCNN.2016.7727691.
- [31] Alfredo Cuzzocrea, Francesco Folino, Massimo Guarascio, and Luigi Pontieri. A predictive learning framework for monitoring aggregated performance indicators over business process events. In *International Database Engineering & Applications Symposium*, pages 165–174, 2018. doi:10.1145/3216122.3216143.
- [32] Massimiliano de Leoni, Wil M. P. van der Aalst, and Marcus Dees. A general process mining framework for correlating, predicting and clustering dynamic behavior based on event logs. *Inf. Syst.*, 56:235–257, 2016. doi:10.1016/j.is.2015.07.003.
- [33] Patricia Dockhorn Costa, João Paulo A. Almeida, Luís Ferreira Pires, and Marten van Sinderen. Situation specification and realization in rule-based context-aware applications. In *Distributed Applications and Interoperable Systems*, pages 32–47, 2007. doi:10.1007/978-3-540-72883-2_3.
- [34] Xin Luna Dong and Theodoros Rekatsinas. Data Integration and Machine Learning: A Natural Synergy. In *Management of Data*, pages 1645–1650, May 2018. doi:10.1145/3183713.3197387.
- [35] Marlon Dumas and Fabrizio Maria Maggi. Enabling Process Innovation via Deviance Mining and Predictive Monitoring. In *BPM - Driving Innovation in a Digital World*, pages 145–154. 2015. doi:10.1007/978-3-319-14430-6_10.

- [36] Gregorio Díaz and Luis Llana. Contract Compliance Monitoring of Web Services. In *Service-Oriented and Cloud Computing*, 2013. doi:10.1007/978-3-642-40651-5_10.
- [37] Wayne W Eckerson. *Performance dashboards: measuring, monitoring, and managing your business*. John Wiley & Sons, 2010.
- [38] Matthias Ehrendorfer, Juergen Mangler, and Stefanie Rinderle-Ma. Assessing the impact of context data on process outcomes during runtime. In *Service-Oriented Computing*, pages 3–18, 2021. doi:10.1007/978-3-030-91431-8_1.
- [39] Ghada El-Khawaga, Mervat Abu-Elkheir, and Manfred Reichert. Explainability of predictive process monitoring results: Can you see my data issues? *CoRR*, abs/2202.08041, 2022. URL: <https://arxiv.org/abs/2202.08041>.
- [40] European Banking Authority. Guidelines on Internal Governance, March 2018. URL: https://www.eba.europa.eu/sites/default/documents/files/documents/10180/2164689/531e7d72-d8ff-4a24-a69a-c7884fa3e476/Guidelines%20on%20Internal%20Governance%20%28EBA-GL-2017-11%29_EN.pdf?retry=1.
- [41] European Central Bank. MMSR - IT Appendix for reporting agents - version 3.0, December 2017.
- [42] European Central Bank. Money Market Statistical Reporting - Data Quality Checks, July 2021. URL: https://www.ecb.europa.eu/stats/money/mmss/shared/files/MMSR-Data_Quality_Checks.pdf.
- [43] European Central Bank. Money Market Statistical Reporting (MMSR) - Questions and Answers - version 3.5, December 2021.
- [44] European Central Bank. Reporting Instructions for the Electronic Transmission of Money Market Statistical Reporting (MMSR) version 3.5, December 2021. URL: https://www.ecb.europa.eu/stats/money/mmss/shared/files/MMSR-Reporting_instructions.pdf.
- [45] European Parliament and European Council. Regulation (EU) No 648/2012 of the European Parliament and of the Council of 4 July 2012 on OTC derivatives, central counterparties and trade repositories, August 2014.
- [46] European Parliament and European Council. Regulation (EU) No 1333/2014 of the European Central Bank of 26 November 2014 concerning statistics on the money markets (ECB/2014/48), January 2015.
- [47] European Parliament and European Council. Regulation (EU) No 600/2014 of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Regulation (EU) No 648/2012Text with EEA relevance, January 2018.

- [48] European Parliament and European Council. Regulation (EU) No 2015/2365 of the European Parliament and of the Council of 25 November 2015 on transparency of securities financing transactions and of reuse and amending Regulation (EU) No 648/2012, April 2019.
- [49] Joerg Evermann, Jana-Rebecca Rehse, and Peter Fettke. Predicting process behaviour using deep learning. *Decis. Support Syst.*, 100:129–140, 2017. doi:10.1016/j.dss.2017.04.003.
- [50] Stephan A. Fahrenkrog-Petersen, Niek Tax, Irene Teinemaa, Marlon Dumas, Massimiliano de Leoni, Fabrizio Maria Maggi, and Matthias Weidlich. Fire now, fire later: alarm-based systems for prescriptive process monitoring. *Knowl. Inf. Syst.*, 64(2):559–587, 2022. doi:10.1007/s10115-021-01633-w.
- [51] Fazel Famili, Wei-Min Shen, Richard Weber, and Evangelos Simoudis. Data preprocessing and intelligent data analysis. *Intell. Data Anal.*, 1(1-4):3–23, 1997. doi:10.1016/S1088-467X(98)00007-9.
- [52] Walid Fdhila, Manuel Gall, Stefanie Rinderle-Ma, Juergen Mangler, and Conrad Indiono. Classification and formalization of instance-spanning constraints in process-driven applications. In *Business Process Management*, pages 348–364, 2016. doi:10.1007/978-3-319-45348-4_20.
- [53] Walid Fdhila, David Knuplesch, Stefanie Rinderle-Ma, and Manfred Reichert. Verifying compliance in process choreographies: Foundations, algorithms, and implementation. *Information Systems*, page 101983, 2022. doi:https://doi.org/10.1016/j.is.2022.101983.
- [54] Paolo Felli, Alessandro Gianola, Marco Montali, Andrey Rivkin, and Sarah Winkler. Conformance Checking with Uncertainty via SMT. In *Business Process Management*, pages 199–216, 2022. doi:10.1007/978-3-031-16103-2_15.
- [55] Filipe Ferreira, Ahm Shamsuzzoha, Américo Lopes Azevedo, and Petri Helo. Virtual enterprise process monitoring: An approach towards predictive industrial maintenance. In *Progress in Systems Engineering*, pages 285–291, 2014. doi:10.1007/978-3-319-08422-0_43.
- [56] Iman Firouzian, Morteza Zahedi, and Hamid Hassanpour. Real-time Prediction and Synchronization of Business Process Instances using Data and Control Perspective. *International Journal of Nonlinear Analysis and Applications*, 10(1), 2019. doi:10.22075/ijnaa.2019.4065.
- [57] Francesco Folino, Massimo Guarascio, and Luigi Pontieri. Discovering context-aware models for predicting business process performances. In *On the Move to Meaningful Internet Systems*, pages 287–304, 2012. doi:10.1007/978-3-642-33606-5_18.

- [58] Francesco Folino, Massimo Guarascio, and Luigi Pontieri. A prediction framework for proactively monitoring aggregate process-performance indicators. In *Enterprise Distributed Object Computing*, pages 128–133, 2015. doi:10.1109/EDOC.2015.27.
- [59] Chiara Di Francescomarino, Marlon Dumas, Marco Federici, Chiara Ghidini, Fabrizio Maria Maggi, Williams Rizzi, and Luca Simonetto. Genetic algorithms for hyperparameter optimization in predictive business process monitoring. *Inf. Syst.*, 74:67–83, 2018. doi:10.1016/j.is.2018.01.003.
- [60] Chiara Di Francescomarino, Marlon Dumas, Fabrizio Maria Maggi, and Irene Teinemaa. Clustering-Based Predictive Process Monitoring. *IEEE Transactions on Services Computing*, 12(6):896–909, November 2019. Conference Name: IEEE Transactions on Services Computing. doi:10.1109/TSC.2016.2645153.
- [61] Chiara Di Francescomarino, Chiara Ghidini, Fabrizio Maria Maggi, and Fredrik Milani. Predictive process monitoring methods: Which one suits me best? In *Business Process Management*, pages 462–479, 2018. doi:10.1007/978-3-319-98648-7_27.
- [62] Riccardo Galanti, Bernat Coma-Puig, Massimiliano de Leoni, Josep Carmona, and Nicolò Navarin. Explainable predictive process monitoring. In *Conference on Process Mining*, pages 1–8, 2020. doi:10.1109/ICPM49681.2020.00012.
- [63] Manuel Gall and Stefanie Rinderle-Ma. Evaluating compliance state visualizations for multiple process models and instances. In *Business Process Management Forum*, volume 427, pages 126–142, 2021. doi:10.1007/978-3-030-85440-9_8.
- [64] Ping Gong, David Knuplesch, Zaiwen Feng, and Jianmin Jiang. bpC-Mon: A Rule-Based Monitoring Framework for Business Processes Compliance. *Web Services Research*, 14(2):81–103, 2017. doi:10.4018/IJWSR.2017040105.
- [65] Wilfried Grossmann and Stefanie Rinderle-Ma. *Fundamentals of Business Intelligence*. Data-Centric Systems and Applications. Springer, 2015. doi:10.1007/978-3-662-46531-8.
- [66] Nitin Harane and Sheetal Rathi. Comprehensive Survey on Deep Learning Approaches in Predictive Business Process Monitoring. In *Modern Approaches in Machine Learning and Cognitive Science: A Walkthrough: Latest Trends in AI*. 2020. doi:10.1007/978-3-030-38445-6_9.
- [67] Maximilian Harl, Sven Weinzierl, Mathias Stierle, and Martin Matzner. Explainable predictive business process monitoring using gated graph neural networks. *Journal of Decision Systems*, 29:312–327, 2020. doi:10.1080/12460125.2020.1780780.

- [68] Mouna Hedjeres, Abdellah Boukerram, Samir Sebahi, and Mohand-Said Hacid. Temporal Event based Compliance Monitoring. In *Advances in Computing and Communication Engineering*, pages 17–22, 2018. doi: 10.1109/ICACCE.2018.8441700.
- [69] Kai Heinrich, Patrick Zschech, Christian Janiesch, and Markus Bonin. Process data properties matter: Introducing gated convolutional neural networks (GCNN) and key-value-predict attention networks (KVP) for next event prediction with deep learning. *Decis. Support Syst.*, 143:113494, 2021. doi:10.1016/j.dss.2021.113494.
- [70] Robert Heinrich, Philipp Merkle, Jörg Henss, and Barbara Paech. Integrating business process simulation and information system simulation for performance prediction. *Softw. Syst. Model.*, 16(1):257–277, 2017. doi:10.1007/s10270-015-0457-1.
- [71] Tobias Herbert, Juergen Mangler, and Stefanie Rinderle-Ma. Generating reliable process event streams and time series data based on neural networks. In *Enterprise, Business-Process and Information Systems Modeling*, pages 81–95, 2021. doi:10.1007/978-3-030-79186-5_6.
- [72] Yuxiu Hua, Zhifeng Zhao, Rongpeng Li, Xianfu Chen, Zhiming Liu, and Honggang Zhang. Deep learning with long short-term memory for time series prediction. *IEEE Commun. Mag.*, 57(6):114–119, 2019. doi:10.1109/MCOM.2019.1800155.
- [73] Dragan Ivanovic, Manuel Carro, and Manuel V. Hermenegildo. Constraint-based runtime prediction of SLA violations in service orchestrations. In *Service-Oriented Computing*, pages 62–76, 2011. doi: 10.1007/978-3-642-25535-9_5.
- [74] Abdulrahman Jalayer, Mohsen Kahani, Asef Pourmasoumi, and Amin Beheshti. Ham-net: Predictive business process monitoring with a hierarchical attention mechanism. *Knowl. Based Syst.*, 236:107722, 2022. doi:10.1016/j.knosys.2021.107722.
- [75] Georg Kaes, Jürgen Mangler, Florian Stertz, Ralph Vigne, and Stefanie Rinderle-Ma. ACaPlan - Adaptive Care Planning. In *BPM Demos*, volume 1418, pages 11–15. CEUR-WS.org, 2015. URL: <http://ceur-ws.org/Vol-1418/paper3.pdf>.
- [76] Georg Kaes, Stefanie Rinderle-Ma, Ralph Vigne, and Juergen Mangler. Flexibility Requirements in Real-World Process Scenarios and Prototypical Realization in the Care Domain. In *OTM Workshops*, pages 55–64, 2014. doi:10.1007/978-3-662-45550-0_8.
- [77] Andrew Kakabadse and Nada Kakabadse. Outsourcing: current and future trends. *Thunderbird international business review*, 47(2):183–204, 2005.

- [78] Martin Käppel, Stefan Jablonski, and Stefan Schönig. Evaluating predictive business process monitoring approaches on small event logs. In *Quality of Information and Communications Technology*, pages 167–182, 2021. doi:10.1007/978-3-030-85347-1_13.
- [79] Martin Käppel, Stefan Schönig, and Stefan Jablonski. Leveraging small sample learning for business process management. *Inf. Softw. Technol.*, 132:106472, 2021. doi:10.1016/j.infsof.2020.106472.
- [80] Naveed Khan, Zulfiqar Ali, Aftab Ali, Sally I. McClean, Darryl Charles, Paul N. Taylor, and Detlef D. Nauck. A generic model for end state prediction of business processes towards target compliance. In *Artificial Intelligence*, volume 11927, pages 325–335, 2019. doi:10.1007/978-3-030-34885-4_25.
- [81] Vishnu Singh Khatuwal and Digvijay Puri. Business Intelligence Tools for Dashboard Development. In *Intelligent Engineering and Management*, pages 128–131, April 2022. doi:10.1109/ICIEM54221.2022.9853086.
- [82] Jongchan Kim and Marco Comuzzi. Stability metrics for enhancing the evaluation of outcome-based business process predictive monitoring. *IEEE Access*, 9:133461–133471, 2021. doi:10.1109/ACCESS.2021.3115759.
- [83] Jongchan Kim, Marco Comuzzi, Marlon Dumas, Fabrizio Maria Maggi, and Irene Teinemaa. Encoding resource experience for predictive process monitoring. *Decis. Support Syst.*, 153:113669, 2022. doi:10.1016/j.dss.2021.113669.
- [84] Eva L. Klijn and Dirk Fahland. Identifying and reducing errors in remaining time prediction due to inter-case dynamics. In *Process Mining*, pages 25–32, 2020. doi:10.1109/ICPM49681.2020.00015.
- [85] Christopher Klinkmüller, N. R. T. P. van Beest, and Ingo Weber. Towards reliable predictive process monitoring. In *CAiSE Forum*, pages 163–181, 2018. doi:10.1007/978-3-319-92901-9_15.
- [86] David Knuplesch, Manfred Reichert, and Akhil Kumar. A framework for visually monitoring business process compliance. *Inf. Systems*, 64:381–409, 2017. doi:10.1016/j.is.2016.10.006.
- [87] Jonghyeon Ko and Marco Comuzzi. Keeping our rivers clean: Information-theoretic online anomaly detection for streaming business process events. *Inf. Syst.*, 104:101894, 2022. doi:10.1016/j.is.2021.101894.
- [88] Patrik Koenig, Juergen Mangler, and Stefanie Rinderle-Ma. Compliance Monitoring on Process Event Streams from Multiple Sources. In *Process Mining*, 2019. doi:10.1109/ICPM.2019.00026.

- [89] Kala Kotamarthi, Xianzhi Wang, Georg Grossmann, Quan Z. Sheng, and Sarath Indrakanti. A Framework Towards Model Driven Business Process Compliance and Monitoring. In *Enterprise Distributed Object Computing Workshop*, pages 24–32, 2015. doi:10.1109/EDOCW.2015.27.
- [90] D. Kronsbein, D. Meiser, and M. Leyer. Conceptualisation of contextual factors for business process performance. *Lecture Notes in Engineering and Computer Science*, 2210, 2014.
- [91] S M Kumar and Meena Belwal. Performance dashboard: Cutting-edge business intelligence and data visualization. In *Smart Technologies For Smart Nation*, pages 1201–1207, August 2017. doi:10.1109/SmartTechCon.2017.8358558.
- [92] Philipp Leitner, Johannes Ferner, Waldemar Hummer, and Schahram Dustdar. Data-driven and automated prediction of service level agreement violations in service compositions. *Distributed Parallel Databases*, 31(3):447–470, 2013. doi:10.1007/s10619-013-7125-7.
- [93] Philipp Leitner, Johannes Ferner, Waldemar Hummer, and Schahram Dustdar. Data-driven and automated prediction of service level agreement violations in service compositions. *Distributed Parallel Databases*, 31(3):447–470, 2013. doi:10.1007/s10619-013-7125-7.
- [94] Philipp Leitner, Anton Michlmayr, Florian Rosenberg, and Schahram Dustdar. Monitoring, prediction and prevention of SLA violations in composite services. In *Web Services*, pages 369–376, 2010. doi:10.1109/ICWS.2010.21.
- [95] Yehuda Lindell. Secure multiparty computation. *Communications of the ACM*, 64(1):86–96, January 2021. doi:10.1145/3387108.
- [96] Daniela Loreti, Federico Chesani, Anna Ciampolini, and Paola Mello. A distributed approach to compliance monitoring of business process event streams. *Future Generation Computer Systems*, 82:104–118, 2018. doi:10.1016/j.future.2017.12.043.
- [97] Linh Thao Ly, Fabrizio Maria Maggi, Marco Montali, Stefanie Rinderle-Ma, and Wil M. P. van der Aalst. Compliance monitoring in business processes: Functionalities, application, and tool-support. *Inf. Syst.*, 54:209–234, 2015. doi:10.1016/j.is.2015.02.007.
- [98] Linh Thao Ly, Stefanie Rinderle-Ma, Kevin Göser, and Peter Dadam. On enabling integrated process compliance with semantic constraints in process management systems - requirements, challenges, solutions. *Inf. Syst. Frontiers*, pages 195–219, 2012. doi:10.1007/s10796-009-9185-9.
- [99] Fabrizio Maria Maggi, Chiara Di Francescomarino, Marlon Dumas, and Chiara Ghidini. Predictive monitoring of business processes. In *Advanced*

- Information Systems Engineering*, pages 457–472, 2014. doi:10.1007/978-3-319-07881-6_31.
- [100] Fabrizio Maria Maggi, Marco Montali, and Ubaier Bhat. Compliance Monitoring of Multi-Perspective Declarative Process Models. In *Enterprise Distributed Object Computing Conference*, pages 151–160, 2019. doi:10.1109/EDOC.2019.00027.
- [101] Marco Maisenbacher and Matthias Weidlich. Handling concept drift in predictive process monitoring. In *Services Computing*, pages 1–8, 2017. doi:10.1109/SCC.2017.10.
- [102] Amolkirat Singh Mangat and Stefanie Rinderle-Ma. Next-activity prediction for non-stationary processes with unseen data variability. In *Enterprise Design, Operations, and Computing*, pages 145–161, 2022. doi:10.1007/978-3-031-17604-3_9.
- [103] Alfonso Eduardo Márquez-Chamorro, Manuel Resinas, and Antonio Ruiz-Cortés. Predictive monitoring of business processes: A survey. *IEEE Trans. Serv. Comput.*, 11(6):962–977, 2018. doi:10.1109/TSC.2017.2772256.
- [104] Nijat Mehdiyev, Joerg Evermann, and Peter Fettke. A Novel Business Process Prediction Model Using a Deep Learning Method. *Business & Information Systems Engineering*, 62(2):143–157, April 2020. doi:10.1007/s12599-018-0551-3.
- [105] Nijat Mehdiyev and Peter Fettke. Explainable artificial intelligence for process mining: A general overview and application of a novel local explanation approach for predictive process monitoring. In *Interpretable Artificial Intelligence: A Perspective of Granular Computing*, pages 1–28, 2021. doi:10.1007/978-3-030-64949-4_1.
- [106] Giovanni Meroni, Luciano Baresi, Marco Montali, and Pierluigi Plebani. Multi-party business process compliance monitoring through IoT-enabled artifacts. *Information Systems*, 73:61–78, 2018. doi:10.1016/j.is.2017.12.009.
- [107] Andreas Metzger, Tristan Kley, and Alexander Palm. Triggering Proactive Business Process Adaptations via Online Reinforcement Learning. In *Business Process Management*, pages 273–290, 2020. doi:10.1007/978-3-030-58666-9_16.
- [108] Andreas Metzger, Philipp Leitner, Dragan Ivanovic, Eric Schmieders, Rod Franklin, Manuel Carro, Schahram Dustdar, and Klaus Pohl. Comparing and combining predictive business process monitoring techniques. *IEEE Trans. Syst. Man Cybern. Syst.*, 45(2):276–290, 2015. doi:10.1109/TSMC.2014.2347265.

- [109] Dominic A. Neu, Johannes Lahann, and Peter Fettke. A systematic literature review on state-of-the-art deep learning methods for process prediction. *Artificial Intelligence Review*, 2021. doi:10.1007/s10462-021-09960-8.
- [110] Dominic A. Neu, Johannes Lahann, and Peter Fettke. A systematic literature review on state-of-the-art deep learning methods for process prediction. *Artif. Intell. Rev.*, 55(2):801–827, 2022. doi:10.1007/s10462-021-09960-8.
- [111] Sina Niedermaier, Falko Koetter, Andreas Freymann, and Stefan Wagner. On Observability and Monitoring of Distributed Systems – An Industry Interview Study. In *Service-Oriented Computing*, pages 36–52, 2019. doi:10.1007/978-3-030-33702-5_3.
- [112] O. Ogunbiyi, A. Basukoski, and T. J. Chausalet. Comparative analysis of clustering-based remaining-time predictive process monitoring approaches. *International Journal of Business Process Integration and Management*, 2020.
- [113] Kasey Panetta. Gartner top 10 data and analytics trends for 2021, 2021. (accessed 2022-01-13). URL: <https://www.gartner.com/smarterwithgartner/gartner-top-10-data-and-analytics-trends-for-2021>.
- [114] Gyunam Park, Janik-Vasily Benzin, and Wil M. P. van der Aalst. Detecting context-aware deviations in process executions. In *Business Process Management Forum*, pages 190–206, 2022. doi:10.1007/978-3-031-16171-1_12.
- [115] Vincenzo Pasquadibisceglie, Annalisa Appice, Giovanna Castellano, and Donato Malerba. A multi-view deep learning approach for predictive business process monitoring. *IEEE Trans. Serv. Comput.*, 15(4):2382–2395, 2022. doi:10.1109/TSC.2021.3051771.
- [116] Stephen Pauwels and Toon Calders. Incremental predictive process monitoring: The next activity case. In *Business Process Management*, volume 12875, pages 123–140, 2021. doi:10.1007/978-3-030-85469-0_10.
- [117] Marco Pegoraro, Merih Seran Uysal, David Benedikt Georgi, and Wil M. P. van der Aalst. Text-aware predictive monitoring of business processes. In *Business Information Systems*, pages 221–232, 2021. doi:10.52825/bis.v1i.62.
- [118] Mirko Polato, Alessandro Sperduti, Andrea Burattin, and Massimiliano de Leoni. Time and activity sequence prediction of business process instances. *Computing*, 100(9):1005–1031, 2018. doi:10.1007/s00607-018-0593-x.

- [119] Rouven Poll, Artem Polyvyanyy, Michael Rosemann, Maximilian Röglinger, and Lea Rupprecht. Process Forecasting: Towards Proactive Business Process Management. In *Business Process Management*, pages 496–512. doi:10.1007/978-3-319-98648-7_29.
- [120] Mahsa Pourbafrani, Shreya Kar, Sebastian Kaiser, and Wil M. P. van der Aalst. Remaining time prediction for processes with inter-case dynamics. In *Process Mining Workshops*, pages 140–153, 2021. doi:10.1007/978-3-030-98581-3_11.
- [121] Efrén Rama-Maneiro, Juan C. Vidal, and Manuel Lama. Deep Learning for Predictive Business Process Monitoring: Review and Benchmark. *arXiv:2009.13251 [cs]*, 2021.
- [122] Pedro Rico, Félix Cuadrado, Juan C. Dueñas, Javier Andión, and Hugo A. Parada G. Business process event prediction through scalable online learning. *IEEE Access*, 9:136313–136333, 2021. doi:10.1109/ACCESS.2021.3117147.
- [123] Stefanie Rinderle-Ma, Manuel Gall, Walid Fdhila, Jürgen Mangler, and Conrad Indiono. Collecting examples for instance-spanning constraints. *CoRR*, abs/1603.01523, 2016. URL: <http://arxiv.org/abs/1603.01523>.
- [124] Stefanie Rinderle-Ma, Manfred Reichert, and Martin Jurisch. On utilizing web service equivalence for supporting the composition life cycle. *Int. J. Web Serv. Res.*, 8(1):41–67, 2011. doi:10.4018/jwsr.2011010103.
- [125] Williams Rizzi, Chiara Di Francescomarino, and Fabrizio Maria Maggi. Explainability in Predictive Process Monitoring: When Understanding Helps Improving. In *Business Process Management Forum*, pages 141–158, 2020. doi:10.1007/978-3-030-58638-6_9.
- [126] Williams Rizzi, Chiara Di Francescomarino, Chiara Ghidini, and Fabrizio Maria Maggi. How do I update my model? on the resilience of predictive process monitoring models to change. *Knowl Inf Syst*, 2022. doi:<https://doi.org/10.1007/s10115-022-01666-9>.
- [127] Carlos Rodríguez, Patrícia Silveira, Florian Daniel, and Fabio Casati. Analyzing compliance of service-based business processes for root-cause analysis and prediction. In *Current Trends in Web Engineering*, volume 6385, pages 277–288, 2010. doi:10.1007/978-3-642-16985-4_25.
- [128] Andreas Rogge-Solti and Mathias Weske. Prediction of business process durations using non-markovian stochastic petri nets. *Inf. Syst.*, 54:1–14, 2015. doi:10.1016/j.is.2015.04.004.
- [129] Sherif Sakr, Ahmed Awad, and Amal Elgammal. Compliance Monitoring as a Service: Requirements, Architecture and Implementation. In *Cloud Computing*, 2015. doi:10.1109/CLOUDCOMP.2015.7149636.

- [130] Ario Santoso and Michael Felderer. Specification-driven predictive business process monitoring. *Software and Systems Modeling*, 19(6):1307–1343, November 2020. doi:10.1007/s10270-019-00761-w.
- [131] Beate Scheibel and Stefanie Rinderle-Ma. Decision mining with time series data based on automatic feature generation. In *Advanced Information Systems Engineering*, pages 3–18, 2022. doi:10.1007/978-3-031-07472-1_1.
- [132] Arik Senderovich, Chiara Di Francescomarino, Chiara Ghidini, Kerwin Jorbina, and Fabrizio Maria Maggi. Intra and inter-case features in predictive process monitoring: A tale of two dimensions. In *Business Process Management*, pages 306–323, 2017. doi:10.1007/978-3-319-65000-5_18.
- [133] Arik Senderovich, Chiara Di Francescomarino, and Fabrizio Maria Maggi. From knowledge-driven to data-driven inter-case feature encoding in predictive process monitoring. *Inf. Syst.*, 84:255–264, 2019. doi:10.1016/j.is.2019.01.007.
- [134] Arik Senderovich, Matthias Weidlich, Avigdor Gal, and Avishai Mandelbaum. Queue mining for delay prediction in multi-class service processes. *Inf. Syst.*, 53:278–295, 2015. doi:10.1016/j.is.2015.03.010.
- [135] Stelios Sidiroglou-Douskos, Sasa Misailovic, Henry Hoffmann, and Martin Rinard. Managing performance vs. accuracy trade-offs with loop perforation. In *SIGSOFT symposium and Foundations of software engineering*, pages 124–134, September 2011. doi:10.1145/2025113.2025133.
- [136] Patrícia Silveira, Carlos Rodríguez, Fabio Casati, Florian Daniel, Vincenzo D’Andrea, Claire Worledge, and Zouhair Taheri. On the Design of Compliance Governance Dashboards for Effective Compliance and Audit Management. In *Service-Oriented Computing*, pages 208–217. 2010. doi:10.1007/978-3-642-16132-2_20.
- [137] Rongjia Song, Jan Vanthienen, Weiping Cui, Ying Wang, and Lei Huang. Towards a comprehensive understanding of the context concepts in context-aware business processes. In *Subject-Oriented Business Process Management*, pages 5:1–5:10, 2019. doi:10.1145/3329007.3329020.
- [138] Florian Spree. Predictive Process Monitoring-A Use-Case-Driven Literature Review. In *EMISA Forum: Vol. 40, No. 1*. De Gruyter, 2020.
- [139] Florian Stertz, Stefanie Rinderle-Ma, and Juergen Mangler. Analyzing process concept drifts based on sensor event streams during runtime. In *Business Process Management*, pages 202–219, 2020. doi:10.1007/978-3-030-58666-9_12.

- [140] Matthias Stierle, Jens Brunk, Sven Weinzierl, Sandra Zilker, Martin Matzner, and Jörg Becker. Bringing light into the darkness - A systematic literature review on explainable predictive business process monitoring techniques. In *European Conference on Information Systems*, 2021. URL: https://aisel.aisnet.org/ecis2021_rip/8.
- [141] Oleg Svatoš. Requirements for Business Process Legal Compliance Monitoring. *Journal of Systems Integration*, 8(2), 2017. doi:10.20470/jsi.v8i2.301.
- [142] Bayu Adhi Tama and Marco Comuzzi. An empirical comparison of classification techniques for next event prediction using business process event logs. *Expert Systems with Applications*, 129:233–245, 2019. doi:10.1016/j.eswa.2019.04.016.
- [143] Bayu Adhi Tama and Marco Comuzzi. An empirical comparison of classification techniques for next event prediction using business process event logs. *Expert Syst. Appl.*, 129:233–245, 2019. doi:10.1016/j.eswa.2019.04.016.
- [144] Johan J C Tambotuh, Harjanto Prabowo, Sani M Isa, and Wahyu Pudjianto. Process mining in governance, risk management, compliance (GRC), and auditing: A systematic literature review. *JATIT*, (18), 2021.
- [145] Irene Teinemaa, Marlon Dumas, Fabrizio Maria Maggi, and Chiara Di Francescomarino. Predictive business process monitoring with structured and unstructured data. In *Business Process Management*, pages 401–417, 2016. doi:10.1007/978-3-319-45348-4_23.
- [146] Irene Teinemaa, Marlon Dumas, Marcello La Rosa, and Fabrizio Maria Maggi. Outcome-Oriented Predictive Process Monitoring: Review and Benchmark. *ACM Trans. Knowl. Discov. Data*, 13(2):17:1–17:57, 2019. doi:10.1145/3301300.
- [147] Irene Teinemaa, Marlon Dumas, Marcello La Rosa, and Fabrizio Maria Maggi. Outcome-oriented predictive process monitoring: Review and benchmark. *ACM Trans. Knowl. Discov. Data*, 13(2):17:1–17:57, 2019. doi:10.1145/3301300.
- [148] J.J. Thomas and K.A. Cook. A visual analytics agenda. *IEEE Computer Graphics and Applications*, 26(1):10–13, January 2006. Conference Name: IEEE Computer Graphics and Applications. doi:10.1109/MCG.2006.5.
- [149] Thomas Wenzel and Christian Boeth. Transaction Reporting. Technical report, Deloitte Touche Tohmatsu Limited, December 2020. URL: <https://www2.deloitte.com/content/dam/Deloitte/de/Documents/risk/Deloitte-Transaction-Reporting.pdf>.

- [150] Silvano Colombo Tosatto, Guido Governatori, and Nick van Beest. Checking regulatory compliance: Will we live to see it? In *Business Process Management*, pages 119–138, 2019. doi:10.1007/978-3-030-26619-6\10.
- [151] Wil M. P. van der Aalst and Schahram Dustdar. Process mining put into context. *IEEE Internet Computing*, 16(1):82–86, 2012.
- [152] Wil M. P. van der Aalst, M. H. Schonenberg, and Minseok Song. Time prediction based on process mining. *Inf. Syst.*, 36(2):450–475, 2011. doi:10.1016/j.is.2010.09.001.
- [153] J. M. E. M. van der Werf and H. M. W. Verbeek. Online Compliance Monitoring of Service Landscapes. In *Business Process Management Workshops*, pages 89–95, 2015. doi:10.1007/978-3-319-15895-2\8.
- [154] Ilya Verenich, Marlon Dumas, Marcello La Rosa, Fabrizio Maria Maggi, and Irene Teinemaa. Survey and cross-benchmark comparison of remaining time prediction methods in business process monitoring. *ACM Trans. Intell. Syst. Technol.*, 10(4):34:1–34:34, 2019. doi:10.1145/3331449.
- [155] Ilya Verenich, Hoang Nguyen, Marcello La Rosa, and Marlon Dumas. White-box prediction of process performance indicators via flow analysis. In *Software and System Process*, pages 85–94, 2017. doi:10.1145/3084100.3084110.
- [156] Thomas Voglhofer and Stefanie Rinderle-Ma. Collection and elicitation of business process compliance patterns with focus on data aspects. *Bus. Inf. Syst. Eng.*, 62(4):361–377, 2020. doi:10.1007/s12599-019-00594-3.
- [157] Chi Wang and Jian Cao. Interval-based remaining time prediction for business processes. In *Service-Oriented Computing*, pages 34–48, 2021. doi:10.1007/978-3-030-91431-8\3.
- [158] Sven Weinzierl, Sebastian Dunzer, Johannes Christian Tenschert, Sandra Zilker, and Martin Matzner. Predictive business process deviation monitoring. In *European Conference on Information Systems*, 2021.
- [159] Karolin Winter and Stefanie Rinderle-Ma. Discovering instance-spanning constraints from process execution logs based on classification techniques. In *Enterprise Distributed Object Computing Conference*, pages 79–88, 2017. doi:10.1109/EDOC.2017.20.
- [160] Karolin Winter and Stefanie Rinderle-Ma. Untangling the GDPR using conrelminer. *CoRR*, abs/1811.03399, 2018.
- [161] Karolin Winter and Stefanie Rinderle-Ma. Defining instance spanning constraint patterns for business processes based on procllets. In *Conceptual Modeling*, pages 149–163, 2020. doi:10.1007/978-3-030-62522-1\11.

- [162] Karolin Winter, Florian Stertz, and Stefanie Rinderle-Ma. Discovering instance and process spanning constraints from process execution logs. *Inf. Syst.*, 89:101484, 2020. doi:10.1016/j.is.2019.101484.
- [163] Frederik Wolf, Jens Brunk, and Jörg Becker. A Framework of Business Process Monitoring and Prediction Techniques. In *Innovation Through Information Systems*, 2021. doi:10.1007/978-3-030-86797-3_47.
- [164] Moe Thandar Wynn and Shazia Sadiq. Responsible Process Mining - A Data Quality Perspective. In *Business Process Management*, pages 10–15, 2019. doi:10.1007/978-3-030-26619-6_2.
- [165] Marwa Hussein Zaki, Ahmed Awad, and Osman Hegazy. Enabling Compliance Monitoring for Process Execution Engines. *BPMDS 2017 RADAR*, page 8, 2017.
- [166] Yu Zhang and Qiang Yang. A Survey on Multi-Task Learning. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–1, 2021. doi:10.1109/TKDE.2021.3070203.