



From process mining to augmented process execution

David Chapela-Campa¹ · Marlon Dumas¹

Received: 20 July 2023 / Revised: 4 October 2023 / Accepted: 6 October 2023 / Published online: 4 November 2023
© The Author(s) 2023

Abstract

Business process management (BPM) is a well-established discipline comprising a set of principles, methods, techniques, and tools to continuously improve the performance of business processes. Traditionally, most BPM decisions and activities are undertaken by business stakeholders based on manual data collection and analysis techniques. This is time-consuming and potentially leads to suboptimal decisions, as only a restricted subset of data and options are considered. Over the past decades, a rich set of data-driven techniques has emerged to support and automate various activities and decisions across the BPM lifecycle, particularly within the process mining field. More recently, the uptake of artificial intelligence (AI) methods for BPM has led to a range of approaches for proactive business process monitoring. Given their common data requirements and overlapping goals, process mining and AI-driven approaches to business process optimization are converging. This convergence is leading to a promising emerging concept, which we call *(AI-)augmented process execution*: a collection of data analytics and artificial intelligence methods for continuous and automated improvement and adaptation of business processes. This article gives an outline of research at the intersection between process mining and AI-driven process optimization, classifies the researched techniques based on their scope and objectives, and positions augmented process execution as an additional layer on top of this stack.

Keywords Business process management · Predictive analytics · Prescriptive analytics · Autonomous systems

1 Introduction

Business process management (BPM) is an established discipline comprising a set of principles, methods, techniques, and tools to continuously improve the performance of the business processes of an organization [15]. The core activities covered by the BPM discipline are often conceptualized as a lifecycle, which starts from the identification of business processes in the organization and proceeds with their discovery, analysis, redesign, implementation, execution, and monitoring. In this lifecycle, the output of each phase feeds into the next one, and the output of the monitoring phase feeds back into the discovery phase [15].

In traditional BPM approaches, most of the activities in the BPM lifecycle are performed manually by managers, analysts, developers, process workers, and other business stakeholders. For example, traditional approaches to discover business processes involve interviews, focus groups, field observation, and other time-consuming data collection methods, followed by manual analysis, synthesis, and validation of the collected data, leading to models describing the business process in question.

In the past years, we have seen the rise of a body of techniques to automate the activities in the BPM lifecycle based on data available in systems of records, such as Enterprise Resource Planning (ERP) or Customer Relationship Management (CRM) systems. These techniques are known under the umbrella term of *process mining* [48].

More recently, the uptake of artificial intelligence (AI) methods for business process management has led to a range of approaches for monitoring business processes proactively, most notably via predictive machine learning models. These techniques usually rely on data similar to (or sometimes identical to) the data used as input by process mining techniques. Unavoidably, given their common data requirements and

Communicated by Ladjel Bellatreche, Christian Attiogbe, and Sadok Ben Yahia.

✉ David Chapela-Campa
david.chapela@ut.ee

Marlon Dumas
marlon.dumas@ut.ee

¹ University of Tartu, Tartu, Estonia

their similar goals, process mining techniques and AI-driven approaches to business process optimization are converging, leading to a promising emerging concept, which we call *(AI-)augmented process Execution* [16]. Broadly speaking, augmented process execution is a collection of data analytics and artificial intelligence methods for continuous and automated improvement and adaptation of business processes.

This paper gives an outline of research at the intersection between process mining and AI-driven process optimization, classifies the researched techniques into layers based on their scope and objectives, and positions augmented process execution as an additional layer on top of this stack.

2 From process mining to augmented process execution

Process mining (PM) emerged in the late 1990s and early 2000s as a collection of techniques to discover and monitor business processes [48]. Earlier techniques for process mining focused on descriptive use-cases, such as analyzing so-called *event logs* extracted from systems of records to discover graphical representations of “as is” processes executed over these systems. More recently, research in the field of process mining has expanded to cover predictive and prescriptive use-cases driven by machine learning techniques, e.g., predicting if a case of a process will end up in a negative outcome, or recommending actions at runtime to optimize one or more performance indicators [21, 45]. Below, we use the term *data-driven BPM* to refer to the entire body of approaches that use data extracted from any system recording the execution of activities, tasks, events, or steps in a process, in order to support one or more phases of the BPM lifecycle.

Figure 1 presents a conceptualization of the ongoing evolution of data-driven BPM approaches in the form of a pyramid of capabilities. In this pyramid, each layer builds on top of the layers below, and serves as the basis for the ones above. Each layer comprises techniques pertaining to two types of use-cases: (i) tactical use-cases, where the goal is to inform managers in their business process change decisions, typically with timeframes of a few weeks to a few months between decision and change implementation; and (ii) operational use-cases, where the goal is to provide information, issue recommendations, or trigger actions in the context of running cases of one or more processes, to improve their performance on a day-to-day basis.

Perhaps not coincidentally, the chronology of research in the field of data-driven BPM, generally reflects the structure depicted in this pyramid. The first layer, *Descriptive Process Analytics*, mostly comprises research done in the early years of process mining. In this layer, the developed techniques focus on describing the current state of the process, e.g., the control-flow structure of the process, or the distribution of




Fig. 1 The augmented BPM pyramid: a classification of data-driven BPM approaches

performance metrics. The second layer, *predictive process analytics*, moved beyond describing the current state to predicting what will happen in the future, e.g., predicting the outcome of a running case, or predicting what will happen if the number of cases created per day increases by 10%. Although predicting the future state of the process may provide useful insights, it only creates value when actions are triggered based on these predictions. This naturally led to the next layer of the pyramid, viz. *Prescriptive Process Optimization*, which focuses on techniques to prescribe actions that, based on the current state and future predictions, might increase (or decrease) the probability of certain undesired events occurring, e.g., a loan offer being rejected by the customer. More recently, we observe that data-driven BPM research is going one step forward, as captured in the top layer of the pyramid, namely *augmented process execution*. This layer goes beyond techniques to inform or recommend decisions to improve one or more business processes. The aim of augmented process execution is to autonomously manage and optimize the process to achieve the desired business outcomes, within constraints and boundaries set by managers.

As we go up in the pyramid levels, the need for human interaction with the system decreases. To make an analogy with the levels of business process execution autonomy presented in [49], predictive process analytics corresponds to level 0, where the system assists the human but does not intervene. Prescriptive process optimization corresponds to levels 1 and 2, where the system actively assists the human in certain actions. Finally, levels 3, 4, and 5 are represented by augmented process execution, where the system drives on its own, and the role of the human ranges from intervening in certain actions, to merely supervising the autonomous system.




Fig. 2 Generic architecture of descriptive process analytics systems

3 Descriptive process analytics

At the lower layer of the pyramid, establishing the ground base for all the others, we find *Descriptive Process Analytics*. The broad purpose of techniques in this layer is to describe the current state of one or more business processes. This includes, for example, describing the most frequent sequences of activities in a process, summarizing the performance of the process via performance indicators, charts, and other visualizations, or detecting, quantifying, and describing deviations between the intended behavior of a process, captured as a process model, and its actual behavior as recorded in an event log. Research on descriptive process analytics reached its pinnacle in the 2000s and early 2010s, and is currently mature and consolidated.

As depicted in Fig. 2, the input of the techniques in this layer is extracted from systems of records, be it via extractions of records of events occurring during a past time period (historical data) or via continuous data extraction (e.g., real-time data streams). The data collected in these event logs is typically pre-processed and stored in an intermediate form, known as an *event log*. An event log is a collection of records, each of which captures the execution of an activity (or a step within an activity) in the context of a business process. Each record in an event log refers to an activity and has at least one timestamp (sometimes multiple). Each record also refers to one instance of a process (a case) and/or to one or more objects manipulated by one or more business processes [38]. A record in an event log may contain other attributes, such as the resource who performed the activity in question, or other attributes describing the case where the event occurred. Using this input, process mining algorithms produce different types of insights about the current state of the process.

Descriptive process analytics techniques developed in the field of process mining generally tackle tactical use-cases. In other words, the insights extracted from event logs by process mining techniques are used by business stakeholders (managers, analysts) to make decisions that, once implemented,

will affect the execution of all or a subset of cases in a process. Typically, the timeframe for the design and implementation of these change decisions ranges from a few weeks to several months. In this layer, we find four main capabilities that address primarily tactical use-cases:

- **Automated process discovery.** The ability to discover process models—i.e., a diagrammatic representation of the process—from data [1, 23, 50, 55] in order to put into evidence the main pathways in the control flow of the process. These diagrammatic representations depict the main structure of the behavior happening in the process, and are useful not only to analyze this behavior, but also to identify unexpected exceptions and highlight potential wastes (e.g., rework, overprocessing).
- **Conformance checking.** The ability to analyze the similarities and discrepancies between the modeled behavior and the pathways executed in the process [2, 33, 52]. Conformance-checking techniques analyze the behavior recorded in the event log against the behavior modeled by the process model, and enable both (i) the identification of deviations in the process, including violations of compliance rules—e.g., purchase orders without invoices—or deviations between the observed execution flows and normative pathways; and (ii) the assessment of the resemblance between the designed and the real process.
- **Performance mining.** The ability to analyze one or more processes in terms of performance measures, such as cycle time, cost per case, defect rate, etc. The insights of these techniques are commonly displayed as enhanced process models or as performance dashboards that help analysts and project managers when making decisions about the process by providing an image of the current state of the project.
- **Variant analysis.** The ability to identify positive and negative deviance in a process by comparing how the process is performed for different subsets of cases, e.g., in different regions. This is achieved, for example, via enhanced process models (capturing commonalities and differences between two or more variants of a process) or via comparative performance dashboards (i.e., dashboards that display the performance of multiple variants) [13, 34, 43].

There are also a number of descriptive process analytics techniques to support operational management use-cases. In particular, runtime performance dashboards, such as those produced by Business Activity Monitoring (BAM) tools [15], allow managers and workers to monitor the performance of ongoing cases of a process, helping them to plan their daily work and/or to detect potential short-term issues in the process, e.g., an increased rate of service-level agreement (SLA) violations.




Fig. 3 Generic architecture of predictive process analytics systems

4 Predictive process analytics

The capabilities offered by the descriptive process analytics layer allow managers to detect and investigate issues or improvement opportunities in a process. However, they do not help managers to preempt issues before they occur, or to estimate the future performance of a process. These latter use-cases are supported by the second layer of the pyramid: the *predictive process analytics* layer. The focus of this layer is on building predictive models capable of estimating the future state of the process, be it at a macro-level or at the level of individual process instances. A wide range of predictive process analytics methods have been proposed over the past decade [12, 20, 30, 42, 51]. Figure 3 depicts a generic architecture of this body of techniques. These techniques use event logs to train predictive models that, using the data of ongoing cases as input, predict future states.

The techniques in this layer can be broadly classified into two categories:

- **What-if digital process twins.** The ability to predict the impact of a (potential) process change. For example, consider an Order-to-Cash process executed on an ERP system. By applying descriptive process mining, we may find that a bottleneck in the packaging step of the process is causing many delays. By using process mining, statistical analysis, and machine learning techniques, we can build high-fidelity business simulation models from event logs [9–11, 29, 57]. These simulation models serve as a replica of a process, and for this reason, they are sometimes referred to as Digital Process Twins (DPT) [14].¹ For example, we can then use a DPT learned from one or more event logs, to simulate what will happen if we add more resources to the packaging step in an order-to-cash

process. The simulation of a DPT leads to an estimate of the impact of this and other possible changes on the percentage of late deliveries.

- **Predictive process monitoring.** The ability to predict future states of a process [28, 39]. Typically, predictive process monitoring is implemented using machine learning [45, 53] or deep learning [39] techniques. These techniques have been applied to use-cases in the field of credit collection [44] or logistics [18], among others. A predictive model is first trained on historical data, and later applied to a stream of events to produce a stream of predictions. Predictive process monitoring techniques can operate at two levels:

- At the *case level* where the purpose is to make predictions about individual cases in a process, e.g., predicting if a running case will end in a positive or negative outcome [20, 27, 45], predicting the remaining time until completion of a case [12, 51, 53], or predicting the next activities in a case and their timing [30, 39, 42].
- At the *process level*, where the purpose is to predict the value(s) of one or more process performance metrics (calculated across all or a subset of cases) [37], e.g., predicting the on-time delivery rate in three days time.

A key difference between these two categories of techniques is that DPTs allow us to predict the future state of the process after a change, e.g., after automating some of the steps of the process, while the predictive process monitoring techniques focus on predicting the future state of one or more cases of the process, assuming that we do not alter the process. Also, the first category of techniques focuses on tactical use-cases, while the second category targets operational use-cases.

An additional challenge that emerges in predictive process monitoring systems, as well as systems from the upper layers of the pyramid, is in the post-deployment phase, specifically, how to maintain the performance of a model after its deployment given that processes are prone to changes? To address this challenge, mechanisms need to be put in place to monitor the performance of the model and to retrain the model when needed. These mechanisms need to strike a tradeoff between the accuracy and reliability of the predictions and the cost of retraining the predictive model. Repeatedly (re-)training a predictive model helps to maintain its performance, but it may be impractical from a cost perspective. Drift detection algorithms [36] can be used to monitor changes in the process that may affect the performance of the predictive model, while incremental learning techniques can be used to efficiently update it [36, 40].

¹ Arguably, a more accurate term would be *Digital Process Shadow* since changes to the simulation model do not affect the real process, which is one of the characteristics usually expected from a full digital twin.




Fig. 4 Generic architecture of prescriptive process optimization systems

5 Prescriptive process optimization

Predicting potential future threats in a process is informative. However, predictions only create value when they are followed by actions. This brings us to *Prescriptive Process Optimization*, the third layer of the Augmented BPM pyramid. Prescriptive process optimization aims to turn predictions into actions, optimally directed and timed, in order to optimize the performance of the process. Unlike the two previous layers, research regarding prescriptive process optimization is recent, and the field is rapidly evolving. For the most, the techniques in this layer have yet to be empirically validated in real-life applications [21].

Figure 4 depicts a generic architecture of prescriptive process optimization systems. Similar to the techniques in the previous layer, a predictive model is first trained with historical data (complete cases) of the process. Then, these predictive models take data from ongoing cases and generate predictions, such as predictions of the remaining time until completion of a case, or predictions of case outcomes. However, instead of directly reporting these predictions to the managers, they are given as input to a recommender system, which produces recommendations to perform certain actions (a.k.a. treatments or interventions), taking into account a given cost model, which captures the cost and the benefits of these actions. These recommendations are then either automatically triggered by the prescriptive monitoring system, or presented to managers, who may decide whether to follow these recommendations or not. In some techniques, the effect of the recommended actions may be fed back into the prescriptive system, which uses this feedback to update the model used to generate recommendations.

Not unlike the previous layers, the capabilities provided by prescriptive process optimization techniques cover tactical as well as operational use-cases. Accordingly, these capabilities can be sub-categorized as follows:

- Automated process optimization.** The ability to recommend changes to a process to strike a tradeoff between competing performance indicators, for example, lowering costs while minimizing defect rate and cycle times [22]. For example, an automated process optimization system may recommend to a process owner to change the allocation rules and work schedules of some resources, to alleviate certain bottlenecks at the beginning of each week. It may also recommend performing additional verification steps for some purchase order types to prevent mishandled orders and associated rework. A recommender system with these characteristics can be achieved, for example, by using search-based optimization algorithms (e.g., genetic algorithms), combined with a DPT. The search-based algorithm generates possible configurations of the process, which are then evaluated using the DPT, to determine which configuration provides the best tradeoff between the relevant process indicators (e.g., cycle time and cost-per-instance) [26].
- Prescriptive process monitoring.** The ability to recommend actions to optimize the performance of a process with respect to one or more performance indicators, in real-time or near-real-time [21]. Existing techniques for prescriptive process monitoring systems are designed to recommend actions to improve the performance at the *case level*. For example, a prescriptive process monitoring technique may detect that a loan application is likely to be rejected, and recommend to the loan officer to send a second alternative offer to the applicant to maximize the probability that the applicant takes a loan [7, 17, 31, 41, 56]. Other prescriptive process monitoring techniques focus on triggering actions to reduce the cycle time of a case [6]. We observe that there are currently no prescriptive process monitoring techniques that operate at a more macro-level, i.e., that recommend actions affecting not one case at a time, but all or a subset of cases in a process. For example, such techniques could detect that

there are currently too few resources assigned to one task, and too many assigned to another one, and recommend redeploying some of those resources to prevent a surge in the number of delayed cases. The challenge here is how to reliably evaluate the short-term impact of such actions, when they affect many cases, within a process or even across multiple processes.

6 Augmented process execution

In prescriptive process optimization, the machine recommends possible actions to human actors. The human actors decide whether to apply these recommendations or ignore them. Although prescriptive systems are helpful for process optimization, they still require human interaction. Conceptually, the human is the one executing the process, and the system is assisting them (i.e., human-driving systems). The last layer of the pyramid, *augmented process execution*, inverts these roles. In this layer, the objective is to develop systems that execute the process while the human assists them, either intervening or supervising. Also, in this layer, we shift from “process optimization” to “process execution.” Indeed, augmented process execution is not only about discovering patterns, or generating process re-design recommendations. It is about driving the execution of cases in the process.

Figure 5 depicts the overall architecture of an augmented process execution system. In these systems, the human defines a frame—i.e., a set of restrictions—inside which the system triggers and allocates actions on its own. The augmented system may be composed of different subsystems, building on techniques from the lower layers of the pyramid, to extract information needed to execute and optimize the process. The managers are part of the process, but only interact with the augmented system when it needs to take actions out of the frame, or when the system is uncertain about what to do next. We can distinguish two types of augmented systems, depending on the approach followed in such situations:

- **Autonomic process execution systems.** The system works on its own inside the frame. When it detects a situation out of this frame, or a situation for which it cannot make a decision due to high uncertainty, it reverts to a human stakeholder, so they act accordingly. For example, a system may determine the verifications that should be done when a purchase order is received, based on historical execution data. However, when the system detects a new type of purchase order that it has never seen before, it escalates to the human operator, who determines which verifications should be done for this new type of order. The system records the decision of the human opera-

tor and applies it when a purchase order of this type is received again.

- **Autonomous process execution systems.** The system has complete control of the process. It can freely operate within the frame, and even make changes to it, provided that these changes will contribute to achieving a set of business goals. In other words, the system pursues a business goal, and uses the frame as a reference (i.e., as a recommendation on how the business goal is usually pursued). In an autonomous process execution system, the human acts as a supervisor, intervening only to prevent undesired consequences. For example, when such a system sees a type of purchase order it has not seen before, it may derive a plan for handling it by extrapolating from similar types of previously seen purchase orders. The system may then flag this case to a human actor. Depending on the control policy, the system may need to wait for input from the human actor (supervised autonomy), or it may be allowed to proceed without such input (full autonomy). Fully autonomous (“zero-touch”) process execution may only be desirable in situations where the implications of the actions the system takes are well-scoped.² Otherwise, the autonomous system needs to be framed with hard constraints (must be fulfilled), while other constraints in the frame may be tagged as being “soft constraints.”

Augmented process execution research is still a nascent concept [16, 32, 47]. However, it is a promising field, and we expect to experience its rise in the following years.


7 Conversational process management

With the coming to maturity of large language models (LLMs), we are witnessing the emergence of conversational systems that assist managers through natural language [4, 5, 54]. Research on this topic is emerging across all layers of the pyramid in Fig. 1. The main purpose of these techniques is to provide access to such systems to non-expert workers [3]. With the use of such techniques, workers no longer need to have expert knowledge about specific notations—e.g., BPMN or Petri nets—or analytics or query languages—e.g., SQL—to interact with the system. The conversational systems perform this translation.

Importantly, these techniques are orthogonal to the pyramid’s layers. Conversational systems (including LLMs) support both tactical and operational BPM use-cases and provide capabilities that are relevant across the four layers of the pyramid:

² Zero-touch paradigms have gained prominence, for example, in the field of network and service management [24].

Fig. 5 Overview of the architecture of augmented process execution systems



- Descriptive process analytics.** A direct application of conversational systems within the descriptive process analytics layer is to translate business questions and hypotheses expressed in natural language by business stakeholders, into executable specifications (e.g., SQL queries, analysis scripts) to generate relevant metrics, charts, and reports [5]. For example, a business stakeholder may inquire if there are certain cases for which the cycle time is much lower than for others. This may result in a series of charts displaying variations in cycle times across countries, product types, and time periods. Another application of conversational systems is to present findings from process mining in a user-friendly way, for example, using a natural language generation (NLG) system to create the description of a process from an event log, or to provide a concise report outlining how the process is currently performing [19, 25, 35, 46].
- Predictive process analytics.** In this layer, conversational systems may act as a translation layer to facilitate the interaction between a business stakeholder and a DPT. For example, the authors in [4] present a tool that translates questions such as “what would be the impact of automating activity X in the process” or “what would be the impact of removing two resources from task X and assigning them to task Y.” The statements given to the chatbot are then translated into simulation scenarios, specified in the language of a business process simulator, to estimate this impact, and the result is used to produce natural language statements. Another application of conversational systems in this layer is to provide explanations for the predictions generated by a machine learning model, e.g., if the model predicts that a case will end up late, the user could ask why, and the conversational system could provide a summary of relevant factors affecting this decision as well as counterfactual statements and examples of similar past cases.
- Prescriptive process optimization.** In this layer, a conversational prescriptive system could translate the recommendations and/or warnings of prescriptive systems to natural language, in order to ease the understanding of the managers. In the case of simple prescriptions, these systems could get closer to the augmented layer, and accept answers in natural language to trigger the recommendations. For example, the system recommends sending an email to inform the client about a certain event, the client answers affirmatively, and the system performs the recommendation on its own. Note that the difference with augmented techniques is that there is no inversion of roles. In this case, the worker is running the process, only transferring certain actions to the system. Another application of conversational systems, specifically LLM-based systems, at the prescriptive layer, is to generate recommendations on how to improve a given business process. For example, given a prompt including a process model discovered from an event log of a purchase-to-pay process (P2P), as well as a summary of relevant statistics, a user could ask questions such as “How could I improve the on-time completion rate of this process, given that my performance target is X%.” The LLM could use its knowledge of P2P process optimization to generate a number of recommendations, such as reallocating resources to reduce certain bottlenecks, or inserting additional controls early on during the process. The resulting recommendations could then be assessed using a DPT, and the conversational system could then return a summary of recommendations validated by the DPT.
- Augmented process execution.** In this layer, conversational systems could complement the system-driven management of the process by interacting or giving feedback to the worker through natural language processing (NLP) [46] and NLG systems. A system with which you can interact and automatically performs the corrections.

For example, the analyst asks for a report of the current state during the morning, or for the predictions for the week. The system outputs the information in natural language, and when the analyst asks to perform a specific optimization or action (either prescribed or original by the analyst), the system automatically performs it. Another application of conversational interfaces in this layer is to enable users to specify the “process frame” in natural language, and to subsequently interact with the process execution system to refine this frame.

8 Conclusion

The last two and a half decades have witnessed the emergence of a rich field of techniques that support and automate the activities comprised under the BPM lifecycle based on data extracted from systems of records. This trend started with the emergence of process mining and business activity monitoring in the early 2000s, which resulted in a body of techniques and tools for discovery, analysis, and monitoring of as-is processes based on event logs or event streams.

More recently, a number of techniques have been proposed to train and use machine learning models to predict future states of a process. In parallel, we have seen the emergence of process mining techniques, sometimes combined with machine learning, to discover DPTs for what-if analysis. These DPTs allow analysts to assess and compare business process redesign options before their implementation.

As a natural step forward, a number of techniques have been proposed for prescriptively recommending process redesigns or recommending actions, at runtime, to enhance the performance of a process with respect to a collection of performance indicators.

We foresee that the next step in this evolution will be the emergence of techniques and tools for autonomic and autonomous execution of business processes, which we place under the umbrella term of augmented process execution. These systems will leverage process mining and AI techniques to achieve continuous and automated improvement and adaptation of business processes. However, transitioning to such systems presents a set of still open challenges. While hyper-automation—i.e., increasing the automation of business processes by introducing artificial intelligence, machine learning, and robotic process automation—holds the potential to revolutionize the way businesses operate and compete in the digital age, it simultaneously introduces numerous complexities. One of the main challenges derived from such a transition, is composability, which concerns how to design these systems in such a way that they can be safely composed, in an environment prone to frequent changes [8].

Orthogonally, we foresee that disruptions in the field of conversational systems, particularly those driven by LLMs,

will lead to richer interactions between managers, analysts, and process workers, on the one hand, and process execution systems, on the other. Several challenges will, however, have to be tackled along the way, most notably, the challenge of ensuring that the process execution system does not perform actions that, while technically falling under the frame specified by the business stakeholders, lead to undesirable business outcomes. The black-box nature and the inherent unpredictability of LLMs will require carefully designed guardrails and new paradigms for managing business process execution risk.

Acknowledgements Research funded by the European Research Council (PIX Project 834141).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Augusto, A., Conforti, R., Dumas, M., et al.: Split miner: automated discovery of accurate and simple business process models from event logs. *Knowl. Inf. Syst.* **59**(2), 251–284 (2019)
2. Augusto, A., Armas-Cervantes, A., Conforti, R., et al.: Measuring fitness and precision of automatically discovered process models: a principled and scalable approach. *IEEE Trans. Knowl. Data Eng.* **34**(4), 1870–1888 (2022)
3. Barbieri, L., Madeira, E.R.M., Stroeh, K., et al.: Towards a natural language conversational interface for process mining. In: *Workshops of the 3rd International Conference on Process Mining (ICPM Workshops 2021)*, LNBIP, vol. 433. Springer, pp. 268–280 (2021)
4. Barón-Espitia, D., Dumas, M., Rojas, O.G.: Coral: Conversational what-if process analysis (extended abstract). In: *Proceedings of the ICPM Doctoral Consortium and Demo Track 2022*, *CEUR Workshop Proceedings*, vol. 3299. CEUR-WS.org, pp. 118–122 (2022)
5. Berti, A., Schuster, D., van der Aalst, W.M.P.: Abstractions, scenarios, and prompt definitions for process mining with llms: a case study. *CoRR*. [arXiv:2307.02194](https://arxiv.org/abs/2307.02194) (2023)
6. Bozorgi, Z.D., Kopolov, A., Dumas, M., et al.: Prolift: automated discovery of causal treatment rules from event logs (extended abstract). In: *Proceedings of the ICPM Doctoral Consortium and Demo Track 2022*, *CEUR Workshop Proceedings*, vol. 3299. CEUR-WS.org, pp. 108–112 (2022)
7. Bozorgi, Z.D., Dumas, M., Rosa, M.L., et al.: Learning when to treat business processes: Prescriptive process monitoring with causal inference and reinforcement learning. In: *Proceedings of the 35th International Conference on Advanced Information Systems Engineering (CAiSE 2023)*, LNCS, vol. 13901. Springer, pp. 364–380 (2023)

8. Bucchiarone, A., Marconi, A., Pistore, M., et al.: Dynamic adaptation of fragment-based and context-aware business processes. In: Goble, C.A., Chen, P.P., Zhang, J. (eds.) *Proceedings of the 19th IEEE International Conference on Web Services (ICWS 2012)*. IEEE Computer Society, pp. 33–41 (2012)
9. Camargo, M., Dumas, M., González, O.: Automated discovery of business process simulation models from event logs. *Decis. Support Syst.* **134**, 113284 (2020)
10. Camargo, M., Dumas, M., González-Rojas, O.: Learning business process simulation models: a hybrid process mining and deep learning approach. *Inf. Syst.* **117**, 102248 (2023)
11. Chapela-Campa, D., Dumas, M.: Modeling extraneous activity delays in business process simulation. In: *Proceedings of the 4th International Conference on Process Mining (ICPM 2022)*. IEEE, pp. 72–79 (2022)
12. van Dongen, B.F., Crooy, R.A., van der Aalst, W.M.P.: Cycle time prediction: When will this case finally be finished? In: *Proceedings of the 2008 OTM Confederated International Conferences “On the Move to Meaningful Internet Systems” (OTM 2008)*, LNCS, vol. 5331. Springer, pp. 319–336 (2008)
13. Dumas, M.: Consolidated management of business process variants. In: *Workshops of the 9th International Conference on Business Process Management (BPM Workshops 2011)*, LNBIP, vol. 100. Springer, p. 1 (2011)
14. Dumas, M.: Constructing digital twins for accurate and reliable what-if business process analysis. In: *International Workshop on BPM Problems to Solve Before We Die (PROBLEMS 2021)*, CEUR Workshop Proceedings, vol. 2938. CEUR-WS.org, pp. 23–27 (2021)
15. Dumas, M., Rosa, M.L., Mendling, J., et al.: *Fundamentals of Business Process Management*, 2nd edn. Springer, Berlin (2018)
16. Dumas, M., Fournier, F., Limonad, L., et al.: Ai-augmented business process management systems: a research manifesto. *ACM Trans Manag Inf Syst* **14**(1), 111–1119 (2023)
17. Fahrenkrog-Petersen, S.A., Tax, N., Teinemaa, I., et al.: Fire now, fire later: alarm-based systems for prescriptive process monitoring. *Knowl. Inf. Syst.* **64**(2), 559–587 (2022)
18. Feldman, Z., Fournier, F., Franklin, R., et al.: Proactive event processing in action: a case study on the proactive management of transport processes (industry article). In: *Proceedings of the 7th ACM International Conference on Distributed Event-Based Systems, (DEBS 2013)*. ACM, pp. 97–106 (2013)
19. Fontenla-Seco, Y., Lama, M., Bugarín, A.: Process-to-text: A framework for the quantitative description of processes in natural language. In: *1st International Workshop on Trustworthy AI - Integrating Learning, Optimization and Reasoning (TAILOR 2020)*, LNCS, vol. 12641. Springer, pp. 212–219 (2020)
20. Francescomarino, C.D., Dumas, M., Maggi, F.M., et al.: Clustering-based predictive process monitoring. *IEEE Trans. Serv. Comput.* **12**(6), 896–909 (2019)
21. Kubrak, K., Milani, F., Nolte, A., et al.: Prescriptive process monitoring: Quo vadis? *Peer J. Comput. Sci.* **8**, e1097 (2022)
22. Lashkevich, K., Ponce, L.M.M., Camargo, M., et al.: Discovery of improvement opportunities in knock-out checks of business processes. In: *Proceedings of the 17th international conference on research challenges in information science (RCIS 2023)*, LNBIP, vol. 476. Springer, pp. 381–397 (2023)
23. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs—a constructive approach. In: *Proceedings of the 34th International Conference on Application and Theory of Petri Nets and Concurrency (PETRI NETS 2013)*, LNCS, vol. 7927. Springer, pp. 311–329 (2013)
24. Liyanage, M., Pham, Q., Dev, K., et al.: A survey on zero touch network and service management (ZSM) for 5g and beyond networks. *J. Netw. Comput. Appl.* **203**, 103362 (2022)
25. López, A., Sánchez-Ferreres, J., Carmona, J., et al.: From process models to chatbots. In: *Proceedings of the 31st International Conference on Advanced Information Systems Engineering (CAiSE 2019)*, LNCS, vol. 11483. Springer, pp. 383–398 (2019)
26. López-Pintado, O., Dumas, M., Yerokhin, M., et al.: Silhouetting the cost-time front: multi-objective resource optimization in business processes. In: *Forum of the 19th International Conference on Business Process Management (BPM Forum 2021)*, LNBIP, vol. 427. Springer, pp. 92–108 (2021)
27. Maggi, F.M., Francescomarino, C.D., Dumas, M., et al.: Predictive monitoring of business processes. In: *Proceedings of the 26th International Conference on Advanced Information Systems Engineering (CAiSE 2014)*, LNCS, vol. 8484. Springer, pp. 457–472 (2014)
28. Márquez-Chamorro, A.E., Resinas, M., Ruiz-Cortés, A.: Predictive monitoring of business processes: a survey. *IEEE Trans. Serv. Comput.* **11**(6), 962–977 (2018)
29. Martin, N., Depaire, B., Caris, A.: The use of process mining in business process simulation model construction—structuring the field. *Bus. Inf. Syst. Eng.* **58**(1), 73–87 (2016)
30. Mehdiyev, N., Evermann, J., Fettke, P.: A novel business process prediction model using a deep learning method. *Bus. Inf. Syst. Eng.* **62**(2), 143–157 (2020)
31. Metzger, A., Kley, T., Palm, A.: Triggering proactive business process adaptations via online reinforcement learning. In: *Proceedings of the 18th International Conference on Business Process Management (BPM)*, LNCS, vol. 12168. Springer, pp. 273–290 (2020)
32. Montali, M.: Constraints for process framing in ai-augmented BPM. In: *Workshops of the 20th International Conference on Business Process Management (BPM Workshops 2022)*, LNBIP, vol. 460. Springer, pp. 5–12 (2022)
33. Muñoz-Gama, J.: Conformance Checking and Diagnosis in Process Mining—Comparing Observed and Modeled Processes, LNBIP, vol. 270. Springer, Berlin (2016)
34. Nguyen, H., Dumas, M., Rosa, M.L., et al.: Multi-perspective comparison of business process variants based on event logs. In: *Proceedings of the 37th International Conference on Conceptual Modeling (ER 2018)*, LNCS, vol. 11157. Springer, pp. 449–459 (2018)
35. Ocaña, M., Chapela-Campa, D., Álvarez, P., et al.: Automatic linguistic reporting of customer activity patterns in open malls. *Multimed. Tools Appl.* **81**(3), 3369–3395 (2022)
36. Pauwels, S., Calders, T.: Incremental predictive process monitoring: The next activity case. In: *Proceedings of the 19th International Conference on Business Process Management (BPM 2021)*, LNCS, vol. 12875. Springer, pp. 123–140 (2021)
37. Poll, R., Polyvyanyy, A., Rosemann, M., et al.: Process forecasting: towards proactive business process management. In: *Proceedings of the 16th International Conference on Business Process Management (BPM 2018)*, LNCS, vol. 11080. Springer, pp. 496–512 (2018)
38. Popova, V., Fahland, D., Dumas, M.: Artifact lifecycle discovery. *Int. J. Cooperative Inf. Syst.* **24**(1), 1550001:1–1550001:44 (2015)
39. Rama-Maneiro, E., Vidal, J.C., Lama, M.: Deep learning for predictive business process monitoring: review and benchmark. *IEEE Trans. Serv. Comput.* **16**(1), 739–756 (2023)
40. Rizzi, W., Francescomarino, C.D., Ghidini, C., et al.: How do I update my model? On the resilience of predictive process monitoring models to change. *Knowl. Inf. Syst.* **64**(5), 1385–1416 (2022)
41. Shoush, M., Dumas, M.: When to intervene? Prescriptive process monitoring under uncertainty and resource constraints. In: *Forum of the 20th International Conference on Business Process Management (BPM Forum 2022)*, LNBIP, vol. 458. Springer, pp. 207–223 (2022)
42. Tax, N., Verenich, I., Rosa, M.L., et al.: Predictive business process monitoring with LSTM neural networks. In: *Proceedings of the*

- 29th International Conference on Advanced Information Systems Engineering (CAiSE 2017), LNCS, vol. 10253. Springer, pp. 477–492 (2017)
43. Taymouri, F., Rosa, M.L., Dumas, M., et al.: Business process variant analysis: survey and classification. *Knowl. Based Syst.* **211**, 106557 (2021)
 44. Teinemaa, I., Dumas, M., Maggi, F.M., et al.: Predictive business process monitoring with structured and unstructured data. In: *Proceedings of the 14th International Conference on Business Process Management (BPM 2016)*, LNCS, vol. 9850. Springer, pp. 401–417 (2016)
 45. Teinemaa, I., Dumas, M., Rosa, M.L., et al.: Outcome-oriented predictive process monitoring: review and benchmark. *ACM Trans. Knowl. Discov. Data* **13**(2), 17:1–17:57 (2019)
 46. van der Aa, H., Carmona, J., Leopold, H., et al.: Challenges and opportunities of applying natural language processing in business process management. In: *Proceedings of the 27th International Conference on Computational Linguistics (COLING 2018)*. Association for Computational Linguistics, pp. 2791–2801 (2018)
 47. van der Aalst, W.M.: Hybrid intelligence: to automate or not to automate, that is the question. *Int. J. Inf. Syst. Proj. Manag.* **9**(2), 5–20 (2021)
 48. van der Aalst, W.M.P.: *Process Mining - Data Science in Action*, 2nd edn. Springer, Berlin (2016)
 49. van der Aalst, W.M.P.: Six levels of autonomous process execution management (APEM). *CoRR*. [arXiv: 2204.11328](https://arxiv.org/abs/2204.11328) (2022)
 50. van der Aalst, W.M.P., Weijters, T., Maruster, L.: Workflow mining: discovering process models from event logs. *IEEE Trans. Knowl. Data Eng.* **16**(9), 1128–1142 (2004)
 51. van der Aalst, W.M.P., Schonenberg, M.H., Song, M.: Time prediction based on process mining. *Inf. Syst.* **36**(2), 450–475 (2011)
 52. van der Aalst, W.M.P., Adriansyah, A., van Dongen, B.F.: Replaying history on process models for conformance checking and performance analysis. *WIREs Data Min. Knowl. Discov.* **2**(2), 182–192 (2012)
 53. Verenich, I., Dumas, M., Rosa, M.L., et al.: 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)
 54. Vidgof, M., Bachhofner, S., Mendling, J.: Large language models for business process management: opportunities and challenges. In: Francescomarino CD, Burattin A, Janiesch C, et al (eds) *Proceedings of the 2023 Business Process Management Forum (BPM Forum 2023)*, LNBIP, vol. 490. Springer, pp. 107–123 (2023)
 55. Weijters, A.J.M.M., Ribeiro, J.T.S.: Flexible heuristics miner (FHM). In: *Proceedings of the 2011 IEEE Symposium on Computational Intelligence and Data Mining (CIDM 2011)*. IEEE, pp. 310–317 (2011)
 56. Weinzierl, S., Dunzer, S., Zilker, S., et al.: Prescriptive business process monitoring for recommending next best actions. In: *Forum of the 18th International Conference on Business Process Management (BPM Forum 2020)*, LNBIP, vol. 392. Springer, pp. 193–209 (2020)
 57. Wynn, M.T., Dumas, M., Fidge, C.J., et al.: Business process simulation for operational decision support. In: *Workshops of the 5th International Conference on Business Process Management (BPM Workshops 2007)*, LNCS, vol. 4928. Springer, pp. 66–77 (2007)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



David Chapela-Campa obtained his Ph.D. in Computer Science from the University of Santiago de Compostela in 2021. Currently, he is a postdoctoral researcher in the Information Systems group at the University of Tartu, Estonia. His research focuses on process mining (PM) and business process management (BPM), specifically on data-driven techniques for business process simulation and business process optimization.



Marlon Dumas is Professor of Information Systems at University of Tartu, Estonia and co-founder of Apomore – a company dedicated to developing and commercializing open-source process mining solutions. His research focuses on data-driven methods for business process management, including process mining and predictive process monitoring.