Computers in Industry 146 (2023) 103837

Contents lists available at ScienceDirect

Computers in Industry

journal homepage: www.sciencedirect.com/journal/computers-in-industry

The biggest business process management problems to solve before we die

Iris Beerepoot [a,*], Claudio Di Ciccio [b,*], Hajo A. Reijers [a,*], Stefanie Rinderle-Ma [c,*], Wasana Bandara [d], Andrea Burattin [e], Diego Calvanese [f], Tianwa Chen [g], Izack Cohen [h], Benoît Depaire [v], Gemma Di Federico [e], Marlon Dumas [i], Christopher van Dun [j], Tobias Fehrer [k], Dominik A. Fischer [j], Avigdor Gal [l], Marta Indulska [g], Vatche Isahagian [m], Christopher Klinkmüller [n], Wolfgang Kratsch [o], Henrik Leopold [p], Amy Van Looy [q], Hugo Lopez [e], Sanja Lukumbuzya [r], Jan Mendling [s], Lara Meyers [z], Linda Moder [j], Marco Montali [f], Vinod Muthusamy [m], Manfred Reichert [w], Yara Rizk [m], Michael Rosemann [d], Maximilian Röglinger [j], Shazia Sadiq [g], Ronny Seiger [t], Tijs Slaats [x], Mantas Simkus [r], Ida Asadi Someh [g], Barbara Weber [t], Ingo Weber [c,u], Mathias Weske [y], Francesca Zerbato [t]

[a] *Utrecht University, Heidelberglaan 8, Utrecht, 3584 CS, Netherlands*
[b] *Sapienza University of Rome, Viale Regina Elena 295, Rome, 00161, Italy*
[c] *Technical University of Munich, Arcisstraße 21, Munich, 80333, Germany*
[d] *Queensland University of Technology, 2 George Street, Brisbane, 4000, Queensland, Australia*
[e] *Technical University of Denmark, Anker Engelunds Vej 1, Lyngby, 2800, Denmark*
[f] *Free University of Bozen-Bolzano, Piazza Università, 1, Bozen-Bolzano, 39100, Italy*
[g] *The University of Queensland, St Lucia, Brisbane, 4072, Queensland, Australia*
[h] *Bar-Ilan University, Ramat Gan, 590002, Israel*
[i] *University of Tartu, Ülikooli 18, Tartu, 50090, Estonia*
[j] *University of Bayreuth, Universitätsstraße 30, Bayreuth, 95447, Germany*
[k] *University of Augsburg, Universitätsstraße 2, Augsburg, 86159, Germany*
[l] *Technion – Israel Institute of Technology, Haifa, 3200003, Israel*
[m] *IBM Research, Cambridge, MA, United States*
[n] *Capgemini Invent, Sydney, New South Wales, Australia*
[o] *Augsburg University of Applied Sciences, An d. Hochschule 1, Augsburg, 86161, Germany*
[p] *Kühne Logistics University, Großer Grasbrook 17, Hamburg, 20457, Germany*
[q] *Ghent University, Tweekerkenstraat 2, Ghent, 9000, Belgium*
[r] *Vienna University of Technology, Wiedner Hauptstraße, Vienna, 1040, Austria*
[s] *Humboldt-Universität zu Berlin, Unter den Linden 6, Berlin, 10117, Germany*
[t] *University of St. Gallen, Dufourstrasse 50, St. Gallen, 9000, Switzerland*
[u] *Fraunhofer-Gesellschaft, Hansastraße 27c, Munich, 80686, Germany*
[v] *Hasselt University, Martelarenlaan 42, Hasselt, 3500, Belgium*
[w] *Ulm University, Helmholtzstraße 16, Ulm, 89081, Germany*
[x] *Copenhagen University, Nørregade 10, Copenhagen, 1165, Denmark*
[y] *Hasso Plattner Institute, Prof.-Dr.-Helmert-Straße 2-3, Potsdam, 14482, Germany*
[z] *Powerlink Queensland, Virginia, Queensland, Australia*

ARTICLE INFO

Keywords:
Business process management
Grand challenges
Process mining
Business process redesign
Digital twins
Process modeling

ABSTRACT

It may be tempting for researchers to stick to incremental extensions of their current work to plan future research activities. Yet there is also merit in realizing the grand challenges in one's field. This paper presents an overview of the nine major research problems for the Business Process Management discipline. These challenges have been collected by an open call to the community, discussed and refined in a workshop setting, and described here in detail, including a motivation why these problems are worth investigating. This overview may serve the purpose of inspiring both novice and advanced scholars who are interested in the radical new ideas for the analysis, design, and management of work processes using information technology.

* Corresponding authors.
E-mail addresses: i.m.beerepoot@uu.nl (I. Beerepoot), diciccio@di.uniroma1.it (C. Di Ciccio), h.a.reijers@uu.nl (H.A. Reijers), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma).
https://doi.org/10.1016/j.compind.2022.103837
Received 13 October 2022; Received in revised form 16 December 2022; Accepted 17 December 2022
Available online 5 January 2023
0166-3615/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Science is a complex and dynamic network in which scientists, institutions, concepts, physical entities and forces "knit, weave and knot" together into an overarching scientific fabric (Latour, 1987). A commonly accepted view is that science has evolved in a nonlinear way: big discoveries as well as gradual advances occur in unpredictable patterns. A rather new development, however, is the "incrementalization of science" (Sverdlov, 2018). Due to perverse incentives, including shortages of long-term funding, scientists try to minimize the risk of having their papers rejected. As Bruce Alberts, the former editor-in-chief of Science, once formulated it (Gitschier, 2012): "If a scientist has to quickly produce a lot of publications, he or she will not do anything new". It is, therefore, likely that many contemporary scientists have little interest in identifying the big problems in their fields, let alone devoting their precious time to solving these.

This paper is meant as a countermove to the incrementalization of science, at least as far as it pertains to the discipline of Business Process Management (BPM). BPM is concerned with the analysis, design, and management of work processes within and across organizations, which often involves the use of Information Technology.[1] Clearly, our work does not neutralize the underlying incentives that lead to incremental research. However, by identifying and characterizing the major open problems in the BPM field, we aim to provide a perspective on the more challenging and important topics that do exist. Hopefully, this may inspire researchers to set time aside to pick up on these challenges and contribute to their solutions, potentially in long-term and collaborative community efforts. We also hope to encourage an intensified partnership between academia and industry to collectively tackle the problems.

The problems that are presented in this paper have been collected through a workshop, an established means to identify the grand challenges for a field. For example, the National Science Foundation organizes workshops from time to time for this purpose across a range of fields; also, various academic conferences have started organizing dedicated side events with a similar purpose, such as the European Semantic Web Conference (ESWC), Academy of Management (AoM) and European Group for Organizational Studies (EGOS). The workshop, the first International Workshop on BPM Problems to Solve before We Die (Beerepoot et al., 2021), was organized as an event directly preceding the 2021 edition of the Business Process Management conference, the flagship event of the discipline.[2] We widely distributed a so-called 'Call for Problems', which invited submissions on the following topics:

- The identification of challenges, wicked problems, and brain-breakers;
- Suggestions for new research fields or paradigms;
- Descriptions of fundamental trade-offs and limitations;
- Suggestions for innovative applications;
- Ideas that are fundamentally different from what we have seen before in the field of BPM;
- Suggestions for the syntheses of contributions from BPM and other domains to solve new or existing problems.

The call attracted 25 submissions from research teams all across the world. A board of senior scholars in the field screened the submissions with respect to their novelty, criticality, magnitude, and appeal. Based on this screening, nine papers were selected for the event. Registration for the event was open to anyone to stimulate discussion of the problems as well as their refinement. In total, more than 50 scholars participated in this event, representing a wide range of research backgrounds. Additionally, a number of them have industrial affiliations or advisory experience in industry. The author teams of the nine submissions were invited to present their problems, after which the audience was encouraged to ask critical questions. The audience then ranked the questions using a voting tool. After the workshop day, the author teams were asked to revise their submissions taking into account the three highest-ranked questions about their problem. They were also asked to reflect on their problem in relation to the process life-cycle, the level of automation that they expected a solution to entail, and the application domains that would be impacted by a solution.

As a result of this endeavor, this paper presents the nine problems, edited and enriched as a result of the discussions during the workshop day. In addition, we analyze some of the patterns that lie beneath these problems. In the following sections, we discuss each of the identified problems, after which we reflect on the implications for science and society of getting these problems solved.

2. The problems

The first problem relates to digital innovation, more in particular the **BPM-driven value creation from data**. The data deluge and associated technological proliferations have changed the landscape of how businesses are run considerably. These changes, in turn, necessitate profound changes in how business processes are managed. Yet, as organizations aspire towards embracing data-driven approaches both technically and culturally, the socio-technical barriers for value creation from data are becoming increasingly evident.

Second is the problem of **expansive BPM**. Despite large investments in BPM over decades, organizations are still left with process fragments by seeing 'process-trees' rather than the entire 'BPM-forest'. This was evident during the COVID-19 pandemic with organizations struggling with a plethora of ad-hoc, and often uncoordinated process changes (Van Looy, 2021). BPM approaches that put individual processes at the center of their attention are unlikely to be able to address 'big processes', i.e., processes that stretch far beyond the boundaries of an enterprise, are closely intertwined with other processes and impacted by various management disciplines.

Third is the problem of **automated process (re-)design**. Driven by the recent "hyperautomation" trend (Panetta, 2020) and the widespread adoption of process-aware information systems, organizations increasingly aspire to leverage automation potential in the context of process operations (Beverungen, 2014). Despite all automation efforts, process (re-)design has remained a manual, cognitively demanding task, making it time-consuming, labor-intensive and error-prone. Several characteristics of business processes and the complexity of process (re-)design itself prevent or at least complicate its automation, such as the creativity required for (re-)design, the socio-technical nature of processes as well as the complex process context.

The fourth problem is related to **constructing digital twins**. Business processes are subject to frequent changes due to various factors (Reichert and Weber, 2012). These changes may be *organic* — resulting from actors in a process adjusting their behavior to respond to emergent situations — or *planned* — agreed-upon changes to the norms, guidelines, policies, or IT systems. This particular problem arises in the context of planned changes, which we refer to as *interventions*. Examples of interventions include reordering two or more tasks, adding a task, adding a resource, changing the decision logic of a branching point, or automating a task. An intervention may have a positive or negative impact on one or more performance measures, e.g. cycle time, activity processing time, or resource utilization.

Fifth is the problem described as **the lack of objectivity in process descriptions**. Process models provide a foundation for many business process management activities including process design, documentation, analysis, and automation (Malinova and Mendling, 2018). Available methods, tools, and notations propose elements for model creation and for capturing aspects related to the functional, behavioral, organizational, and informational process perspectives (Jablonski and Bussler, 1996). Yet, they largely fail to provide guidance with respect

[1] For an overview of how the discipline evolved, see Reijers (2021).
[2] See https://bpm-conference.org/.

to how the semantics of a process can be conveyed through a consistent vocabulary and how to obtain a consistent level of granularity. This is not only a problem for manual model creation, but extends to the use of process discovery algorithms (Klinkmüller et al., 2021). Note that we further discuss implications for process mining in the context of the sixth problem. Hence, model creation is more art than science and the resulting freedom can exacerbate the effective utilization of models. Process models are concise, selective and arguably subjective representations, because there is a lack of objectivity regarding terminology, perspectives and granularity.

Granularity is also discussed in relation to another problem, namely that of **fixed granularity levels for process analysis**. Since process mining techniques do not work well with fine- and mixed-grained events (van Zelst et al., 2020) event logs are typically created by abstracting the raw, fine-grained events into coarse-grained ones. However, finding an appropriate granularity level of process activities during pre-processing is challenging as multiple granularity levels can be relevant for the analysis, and the concept of "activity" often emerges as the data is explored (Zerbato et al., 2021). Thus, analysts must iterate between the Event Abstraction, Granularity Selection, and Mining & Analysis phases until they find a suitable granularity level. Still, they can select only one specific level, which limits what they can observe and control.

The role of the process analyst and, in addition, the domain expert, is also evident in the identified problem of **augmenting process mining with common sense and domain knowledge**. Often, event logs are of medium quality: they might miss events and contain wrong or redundant events. Even when they are of good quality, they might lack the explicit representation of facts that are relevant for the analysis. Humans handle such incorrect and incomplete information by applying their common sense and domain knowledge. However, this knowledge is not transferred to process mining algorithms: it either stays in the head of domain experts, or it gets embedded into ad-hoc pieces of code used in data pre-processing.

The need for complementing event data with common sense and domain knowledge may in part be attributed to work not being recorded as it is executed in real life, which is identified within the problem of **worker-centric process management**. BPM tools manage workers' tasks by imposing varying degrees of structure while tracking progress and ensuring traceability (especially in highly regulated industries). As modern software tools continue to infiltrate the workplace, dynamic and ad-hoc work outside BPM tools (or "behind the system's back") has become the norm (van der Aalst et al., 2005). Not only do current BPM tools fail to capture this type of dynamically changing ad-hoc work, but also create an additional overhead on the workers by forcing them to redo their work inside these BPM tools.

Work that is executed outside of BPM systems may be captured from other sources, but this brings its own challenges, as outlined in the problem of **mining processes using stochastic data**. Current times are characterized by increasing amounts of event data that are generated from multiple sources including physical devices and sensors. The quality of such sources may be low and questionable due to many factors, among them the quality of data capturing devices, quality reduction following data processing and the use of probabilistic models. The end result of collecting such data into process logs is described in the literature as uncertain sensor data. Consider, for example, the use of a machine learning algorithm to detect activities in video clips. Such an algorithm typically offers, as a last stage before decision making, a probability distribution over a space of alternative activities. The probabilistic information can be utilized to quantify the uncertainty associated with event data, and propagate it to the log to create a stochastic, rather than deterministic, log. However, performing process mining tasks using the stochastic data remains a challenge.

3. Towards fixing the problems

A major element of big research problems is that there are no solutions available as of yet. For some problems, we are very far from a solution. For others, recent works and developments in the BPM field may provide potential paths toward a solution. Perhaps the solution to some problems does not lie in research at all, but in industrial practice. In the following sections, we reflect on the implications of the problems, as well as on possible starting points for solutions where they are available. In addition, we describe what the resolution of these problems will open up in the future.

3.1. BPM-driven value creation from data

The technical advancements in data science and machine learning, as well as the third wave of AI, have raised expectations of business transformation. However, while Information Systems research provides many insights in the context of value creation from IT assets, organizations still struggle with value creation from data. Those who are able to overcome the barriers, which are often socio-technical in nature, will gain a competitive advantage over those whose data remains an untapped asset. BPM-related research has a role to play in helping organizations in the private as well as the public sector overcome these barriers, as depicted in Fig. 1. In particular, there are opportunities for:

1. Process orientation research to inform organizational structures and processes to better support data-driven work;
2. Process design research to facilitate agile use of data-driven insights;
3. Process mining research to assist with identifying the right balance between automation and human involvement in data-driven processes;
4. Reference model research to assist with capturing and sharing value creation best practices.

Addressing the challenges relating to value creation from data will enable more effective structuring of organizations and teams in a way that transforms the organization into a data-driven entity. It will result in data analytics being pervasively embedded in an enterprise-wide approach, and closer ties between data analytics teams and business users and thus improved data-driven decision making. While these organizational and cultural changes will not address technical challenges, process mining has the potential to alleviate data quality challenges, identify key data assets, and also help organizations measure the value of their data. The use of reference models can be furthermore employed to share best practices to create and capture value from data.

3.2. Expansive BPM

Business processes are critical organizational assets because they influence the performance, rules and even culture of how work is conducted (Dumas et al., 2018). They are complex because they can overlap, parts can be reused, and they can vary (e.g., in region, division, regulatory compliance application, product, and customer profile). However, contemporary BPM practices often create 'silos' due to a focus on single processes, resulting in potential conflicts between individual process performance and overall enterprise performance. Furthermore, BPM efforts are often sponsored by divisions or functional managers. Even for enterprise-wide BPM efforts, KPIs or motivational metrics often remain divisionally focused. Instead, one needs to know the operating model to understand how a business process stretches across the organization and beyond. Improving a process (or a portfolio of processes) is an interdisciplinary challenge; where the typical process-centric view needs to be expanded, to involve strategy, metrics, risks, data, a business architecture, etc. Moreover, various disciplines recognize processes and different cross-disciplinary perspectives provide value to address a business problem or to align to strategy. Hence,

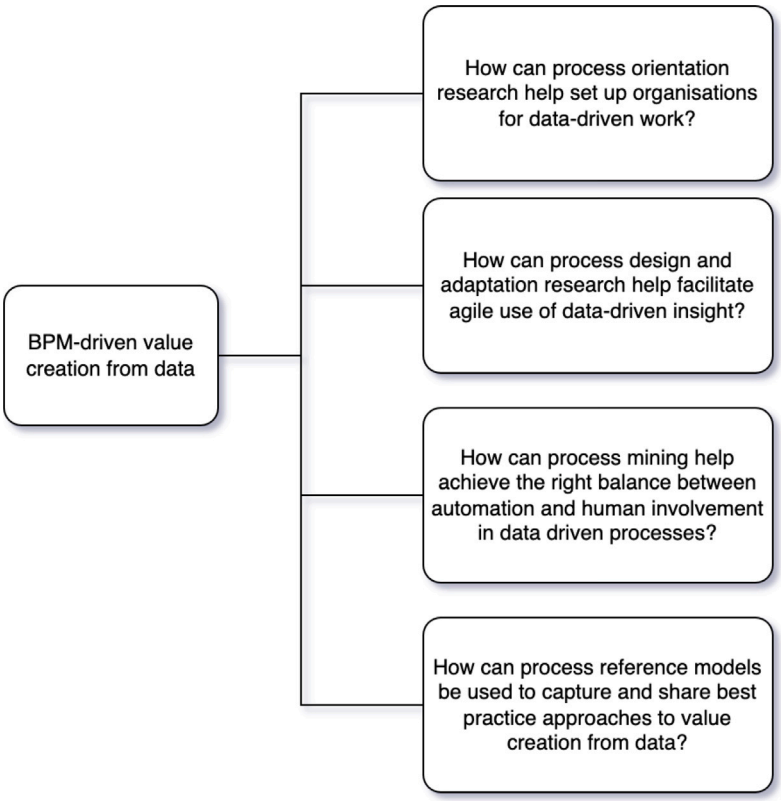

Fig. 1. BPM-driven value creation from data.

interdisciplinarity should be at the center of BPM practice. Current BPM practice is far from this holistic state, with businesses being incapable to manage processes pre and post their corporate sphere of control (Merideth et al., 2020).

BPM will need to expand its unit of analysis and move from the study of single processes to the study of 'big processes'. As Fig. 2 depicts, we envision the use of Process Architectures (PAs) with expanded capabilities as a core mechanism to reach Expansive BPM. Architectural perspectives can provide the modules to show how the building blocks interact and integrate into end-to-end flows, indicating what stakeholders and customers may be involved in, having different outcomes and expectations for the processes. These inter-relationships will differ across diverse scenarios, and needs to be swiftly, and robustly tackled. Strategic and tactical demands may drive the enterprise portfolio of process priorities. PAs host affordances to form the missing holistic view of the enterprise, and the on-going growth of process digitalization will enable the proposed expansions of PA practices (e.g., process data can be more efficiently maintained and acted upon with real-time data flows).

Expansive BPM would extend current BPM practices by the enactment of three critical lenses, by:

1. Expanding the end-to-end process view: capturing the 'true' beginnings and ends of processes;
2. Incorporating a multi-process view: recognizing that multiple, interrelated processes need to be looked at to solve a critical problem or capitalize on an opportunity;
3. Integrating a multi-disciplinary view: multiple disciplines matter to and ongoingrested in business processes. This calls for a coordinated effort across related disciplines, to simultaneously provide input to the re-design and continuous management of business processes.

3.3. Automated redesign

BPM continuously attracts the attention of academia and practice, as it is known to drive organizational performance (Beverungen, 2014; Kerpedzhiev et al., 2021). Especially process (re-)design entails significant business value by introducing adjusted process designs or completely new processes to address existing issues or to seize opportunities. Process (re-)design introduces innovation, reduces costs, and improves quality, productivity, efficiency as well as customer experience (Teinemaa et al., 2019; Kreuzer et al., 2020). Thus, it is considered an essential phase in the BPM lifecycle (Gross et al., 2021). Today, organizations in the private as well as the public sector must overthink their business processes at an increasingly fast pace, consider continuously rising customer needs and expectations, create novel process-based value propositions, and engage in innovation to stay successful (Beverungen, 2014; Grisold et al., 2019; Gross et al., 2021). Technological developments are rapidly gaining momentum, processes are at drift, and ever more players enter the global market, resulting in a strong competitive landscape and in the organizational environment continuously becoming more volatile, uncertain, complex, and ambiguous (VUCA, in the words of Bennett and Lemoine, 2014). Even though this poses pressure on organizations, it also offers a wide range of opportunities. While automation is prevalent in other BPM lifecycle phases (e.g., in process execution — see van der Aalst, 2013a), process (re-)design commonly requires manual activities such as traditional creativity techniques (Gross et al., 2021; Vanwersch et al., 2016), making it time-consuming, expensive and labor-intensive. Automated process (re-)design (Fig. 3) holds high yet unexploited potential for long-term corporate success, since it could accelerate process (re-)design and make it more efficient and effective as well as less dependent on human creativity. Thereby, it would enable organizations to exploit emerging opportunities more rapidly and automatically to secure a competitive advantage.

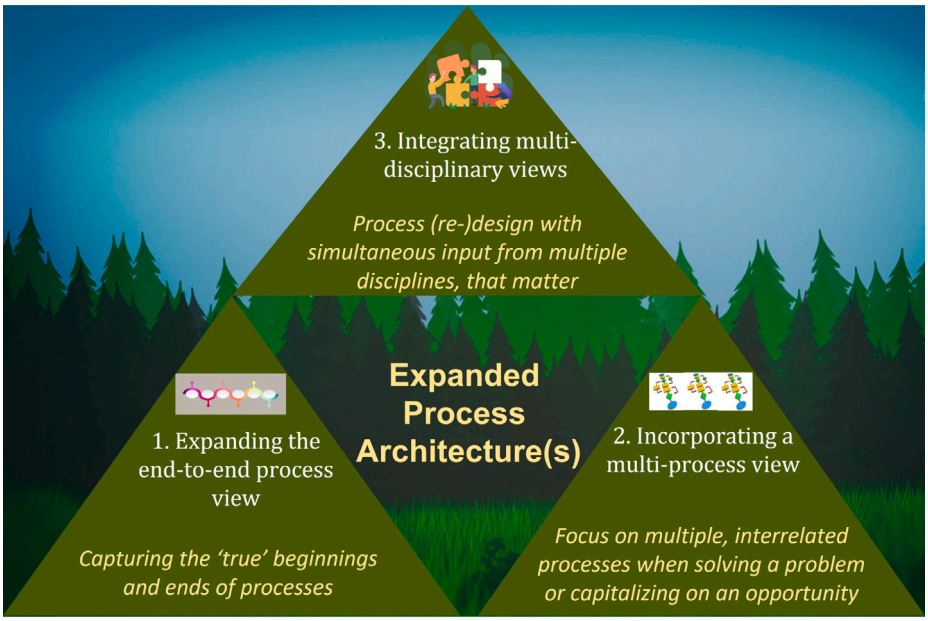

Fig. 2. Expansive BPM: Three lenses for a holistic approach.

The automation of process (re-)design promises a flexible response to changing conditions, competitive advantages, the unburdening of human decision-makers and, not least, more sustainable use of resources. First, the introduction of a new class of automated process improvement systems (APIS) will shift the tasks of traditional BPM more toward the process participants. Today's descriptive process mining methods become accessible by evolving into prescriptive recommendation services, democratizing the technology. Thus, process improvement becomes part of daily business. With increasing degrees of autonomy, knowledge, and intelligence, processes improve without further intervention, in increments, but ultimately also through radical redesigns. These developments will allow the automatic design of entirely new processes and thus have the potential to create entire business models. The question will then be, who will prevail against their market competitors in the future? Whether it is access to the most advanced models and methods, the computing power to simulate all possible scenarios, or simply the ability to successfully manage constant change, with the possible advantages, also other issues will arise: legal and ethical concerns have been much debated in the field of artificial intelligence, such as who is responsible for decisions made by self-learning algorithms and whether the decisions can be trusted. In addition to these external constraints, from a business perspective, it will be necessary for organizations to set strategic guidelines and goals in such a way that processes always behave within these guidelines. Thus, altogether, solving this problem will create new opportunities as well as new challenges.

3.4. Constructing digital twins

Digital Process Twins (DPTs) would allow decision makers to estimate the performance improvements that a given process intervention is likely to bring about (Fig. 4), and hence to build a business case for intervention by comparing its expected benefits against its expected costs or drawbacks. This capability would enable managers to improve the quality of their decisions during the redesign phase of the business process lifecycle (Dumas et al., 2018). Importantly though, the use of a DPT as a decision-making tool for business process redesign requires that it fulfills the following two non-functional requirements:

1. The predictions about the effects of interventions should be accurate. Here, accuracy may be measured in terms of the error between the predicted and the actual performance after intervention.
2. Sometimes, it might not be possible to predict the effect of a given intervention, for example because no similar interventions have been observed in the past, and thus there is no information about how the actors in the process will react. If this is the case, either no prediction should be made or the prediction should be accompanied by a reliability estimate.

As for the second requirement, we remark that the reliability estimate would enable business process managers to put in place appropriate measures to validate the predictions made by the DPT, such as undertaking preliminary checks – e.g. via A/B testing (Satyal et al., 2019).

In addition to its direct application as a decision-making tool for business process redesign, a DPT would enable the development of advanced techniques for automated process optimization. Here, automated process optimization refers to the ability to identify optimal interventions on a process to maximize or minimize a given objective function (defined in terms of one or multiple performance measures) under certain constraints (e.g., resource utilization should remain below 80% — see Si et al., 2018). Fundamentally, automated process optimization algorithms operate by exploring a large number of possible process interventions and evaluating them in order to determine which intervention (or combination of interventions) yields the highest gain. An accurate and reliable DPT would allow such algorithms to evaluate the gain of each candidate intervention in order to navigate through the search space. To be useful in this setting, the DPT needs to be sufficiently efficient (computationally) to be executed thousands or millions of times in a short period of time. This use case brings an additional requirement: the DPT should be computationally efficient, particularly when it comes to generating predictions.

3.5. Lack of objectivity in process descriptions

The current lack of objectivity in process descriptions has implications on business process management in diverse stages of its lifecycle. During the modeling phase, for example, analysts cannot draw on objective guidelines to create unambiguous, holistic, and universally usable models. Instead, they follow their own intuition when deciding which process aspects to include and how to express

	Incremental process improvement	Radical process innovation
Automated process (re-)design	**Automated Improvement** Incrementally improve an existing process automatically	**Automated Innovation** Radically create a new process automatically
Manual process (re-)design	**Manual Improvement** Incrementally improve an existing process manually	**Manual Innovation** Radically create a new process manually

Fig. 3. The problem of automated process (re-)design.

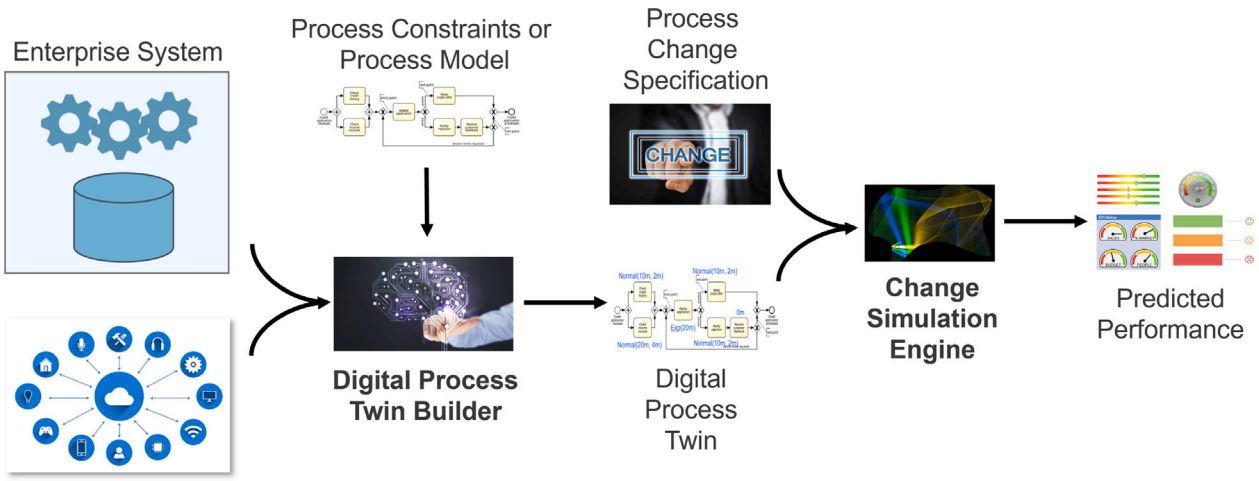

Fig. 4. Constructing digital twins for what-if process analysis.

them. In this regard, they need to make assumptions about opaque future model use cases, typically leading to usage intentions that differ from the actual usage requirements. Ultimately, this negatively affects model comprehension and application, as can be observed in various application scenarios of business process models where the semantics of the model are important (Mendling et al., 2015). For example, process model matching algorithms are designed to automatically identify activities that represent similar functionality. Despite the substantial attention that process model matching received, existing solution approaches have not yet yielded satisfactory and practically usable performance, as prominently demonstrated in the process model matching contests (Cayoglu et al., 2013; Antunes et al., 2015). This is a direct result of the lack of objectivity: matching approaches rely on general-purpose, off-the-shelf knowledge bases and techniques, but need to interpret less objective process models with, e.g., heterogeneous terminology (i.e., differences in labeling styles, domain terminology, etc., as observed by Klinkmüller et al., 2013) and different structures that express similar control flow relationships (Klinkmüller and Weber, 2017). A promising direction to improve approaches is to learn from user feedback (Klinkmüller and Weber, 2021) and let human model users interpret the models. However, several studies have demonstrated that humans often arrive at diverging views due to interpreting model aspects in different ways (see, e.g., Rodríguez et al., 2016; Kuss et al., 2018). Note that similar problems can be observed in other domains. For example, the field of semantic interoperability is concerned with the communication between distributed systems based on a common understanding of the meaning of exchanged services and data (Heiler, 1995). While many ontologies have been developed to establish such a common understanding, they are generally use case and domain-specific, requiring further normalization to bridge the different views expressed by the ontologies (Ganzha et al., 2017). The existence of similar challenges in other domains further demonstrates the severity of the underlying issues.

Increasing our ability to create objective process models will alleviate many problems in business process management, as depicted in Fig. 5. First, the creation of process models, be it for documentation, analysis, implementation or other purposes, will be eased and unified. This will reduce the cognitive load for model creators, as they can focus on understanding the process and do not need to take into account for which purposes a model might be used and how relevant information for these potentially conflicting use cases is best conveyed. Second, tool support for model creation can be extended by simplifying steps in the model creation process, e.g., via templates, and by providing advanced tools for verification and validation that go beyond control-flow semantics. While this will above all affect manual model creation, the field of process mining and in particular the development of process discovery algorithms can benefit from objective model creation guidelines as well, especially when those guidelines can constrain the space of models that can explain observed behavior. Section 3.6 provides more details on the implications for process mining, especially focusing on the lack of fixed levels of granularity. Third, model analysis and interpretation is eased, as models provide comprehensive information and express process aspects clearly and unambiguously. This reduces the human effort required to understand a process model and potentially to collect additional information. Similar to tool support for model creation, it would be easier to develop intelligent approaches that can interpret process models to, e.g., provide improvement recommendations.

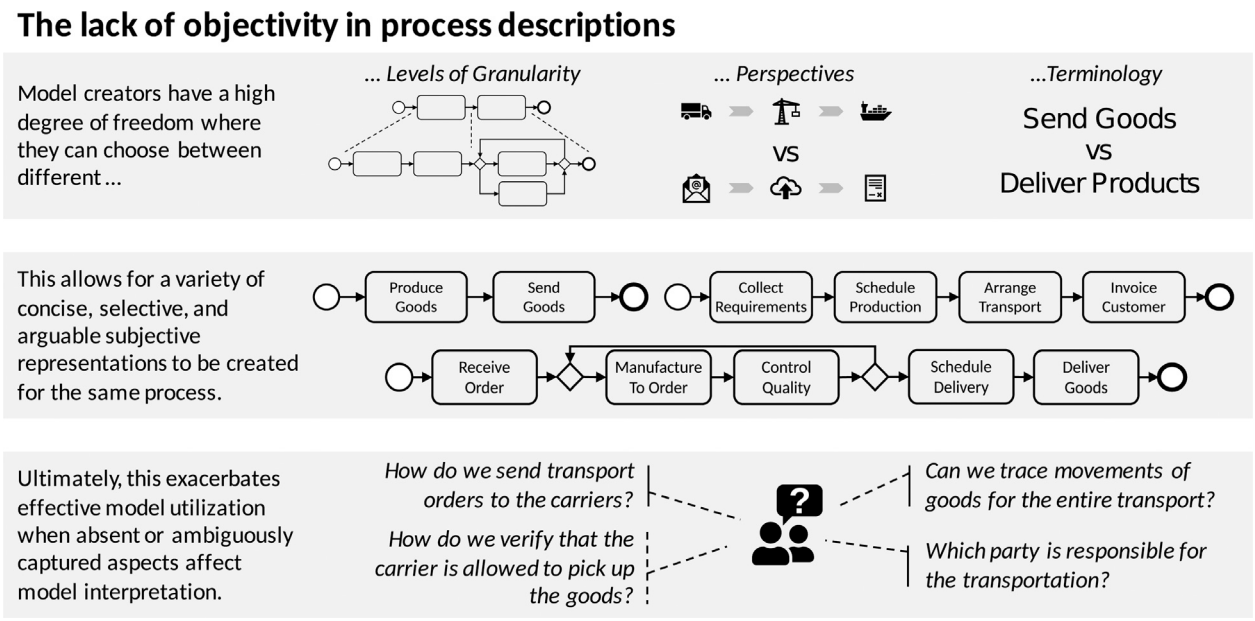

Fig. 5. The lack of objectivity in process descriptions.

3.6. Fixed granularity levels for process analysis

Finding an appropriate granularity level for process activities poses challenges for the pre-processing and analysis of event data in various application areas of process mining. In particular, this is the case for knowledge-intensive scenarios where fine- and mixed-grained event data is derived from executing loosely specified processes (Reichert and Weber, 2012). Indeed, since fine-grained events do not convey business-relevant meaning and are unsuitable for most process mining techniques, they need to be abstracted (cf. the Event Abstraction phase in Fig. 6). This requires knowing the most appropriate granularity level to represent business-relevant process activities to be able to select it before the Mining & Analysis phase.

Granularity selection is challenging for several reasons. First, it is often impossible to know which granularity level will reveal meaningful analysis patterns in advance. Second, it is not trivial to establish what constitutes a "process activity", especially in the context of loosely specified processes. Third, since only one granularity level can be selected at a time, analysts have to iterate among many possible granularity levels without having the possibility of integrating them. Finally, the "fixed view" imposed on the data by the granularity selection limits what process analysts can do in the Mining & Analysis phase. Process analysts cannot:

1. Observe multi-granular patterns within the same analysis;
2. Trace which raw events have been aggregated into process activities since the original events are often lost in the abstraction phase;
3. Adjust and control the granularity level during the analysis — they shall go back to the Granularity Selection or Event Abstraction phases.

Solving the problem of having fixed granularity levels for process analysis would foster the application of process mining techniques to knowledge-intensive and loosely specified processes. We envision a solution that overcomes the above challenges by bridging the pre-processing (specifically, the Event Abstraction) and the Mining & Analysis phases of typical process mining workflows. This can be realized in different ways. A first approach could be to develop an interactive system allowing analysts to reselect the granularity level or recompute the abstraction on the fly. Another method would be to devise human-in-the-loop approaches enabling process analysts to run and control process mining algorithms on multi-granular event logs. Regardless of how the solution could be implemented, enabling users to explore event data at different granularity levels in an integrated manner brings several advantages:

1. It would ease the granularity selection, enabling analysts to explore the data and identify an appropriate granularity level for process activities in domains where they typically emerge during the analysis;
2. It will allow for a relaxation of the notion of process, allowing analysts to look at the process from different perspectives and analyze events and their relationships at different granularity levels and from different dimensions (Esser and Fahland, 2021; Ghahfarokhi et al., 2021);
3. It will give analysts more "control" over the granularity level of the data, allowing them to choose where to zoom in and out to obtain details-on-demand (Shneiderman, 1996) or freeze interesting sub-models (Schuster et al., 2021);
4. It would make the event abstraction more transparent, enabling analysts to trace what abstractions have been applied to which events.

3.7. Augmenting process mining with common sense and domain knowledge

This problem relates to the grand challenge of *garbage-in garbage-out* in process mining and, more in general, in data science. While the community has extensively worked on data quality targeting the pieces of information explicitly contained in an event log, no attention has been devoted to the insights that can be obtained through reasoning on such data. At the same time, the elicitation and usage of common sense and domain knowledge constitute central open problems in (general) artificial intelligence (Davis and Marcus, 2015). Focusing on a concrete setting, such as that of process mining, grounds this problem in a concrete context, paving the way to more accessible results that could in turn provide insights on how to advance with the problem in its full generality. Within the BPM community, this problem is largely unexplored. The literature has so far mainly targeted the problem of integrating structural knowledge with process models (see, e.g., Artale et al., 2019), or in providing richer ontological foundations for process models themselves (see, e.g., Adamo et al., 2021). The huge body of

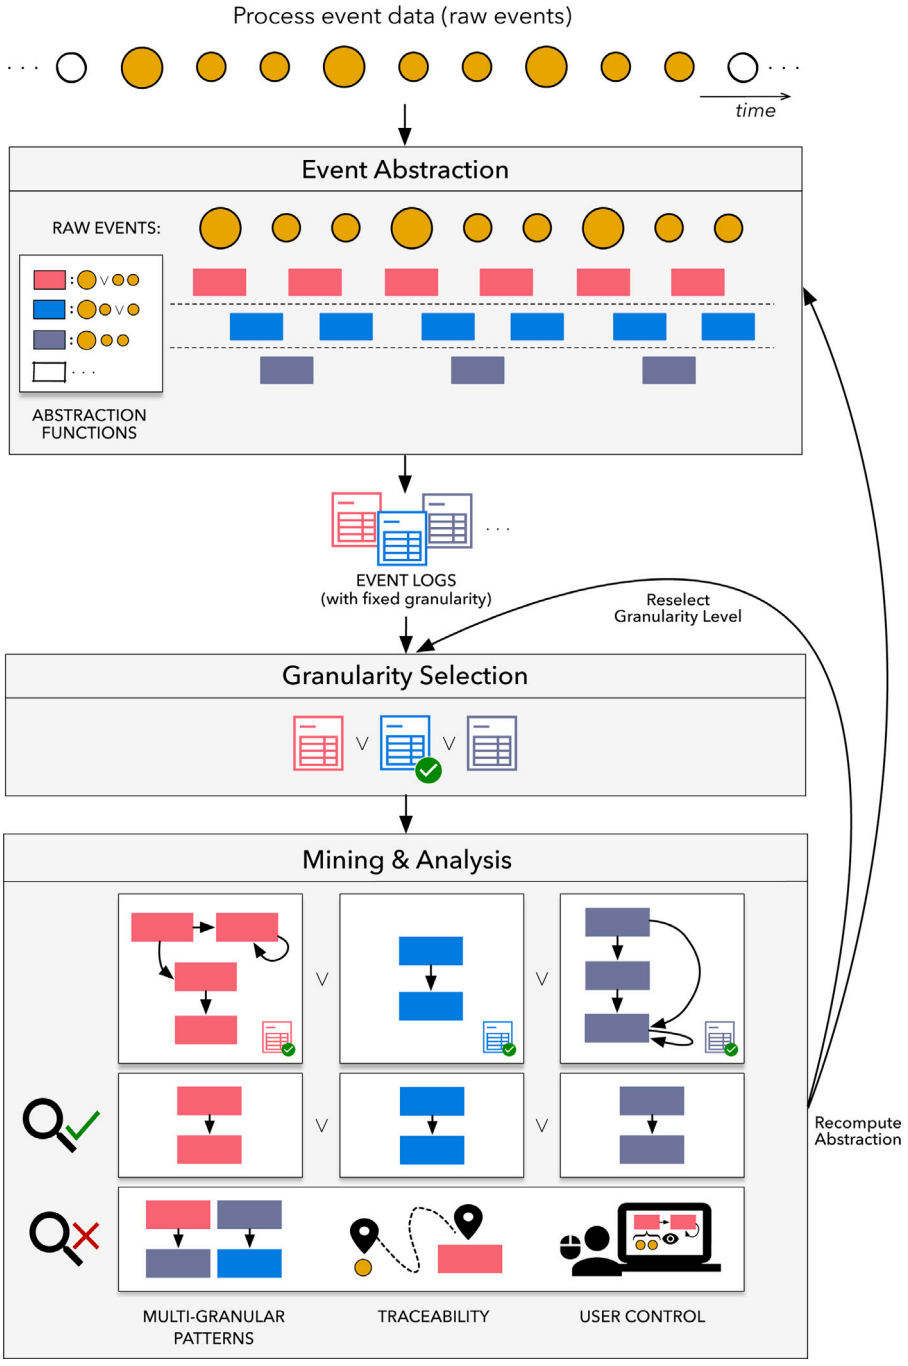

Fig. 6. Fixed granularity levels for process analysis.

work on data quality for event data is complementary and in synergy with the problem.

In the majority of cases, process mining is applied on event logs with missing/wrong/redundant events, or that do not explicitly contain all relevant facts for the analysis. The application of conventional process mining techniques on such data inevitably produces low-quality results. This, in turn, leads to two possible, undesirable scenarios. If the process mining analyst does not realize that the results are of poor quality, such results lead to unjustified corrective actions. If (s)he does, extensive manual effort is required to curate, if possible, the input data, and to restart the entire process mining pipeline from scratch. Differently from this problematic scenario, knowledge-augmented process mining techniques (Fig. 7) would be able to exploit commonsense and domain knowledge to automatically detect and compensate errors in the events, as well as inferring missing events and relevant implicit facts. This, in turn, has the potential to handle data quality and trustworthiness issues directly within the process mining pipeline. In addition, knowledge may be used to infer and return justifications and explanations about the produced results, and to provide estimates about their degree of uncertainty. The process mining analyst could then immediately exploit the so-produced results and, together with domain experts, refine the analysis and take well-justified corrective actions. All in all, augmenting process mining with common sense and domain knowledge is essential to enable a successful process mining lifecycle including a virtuous and effective continuous interaction between algorithms and human stakeholders.

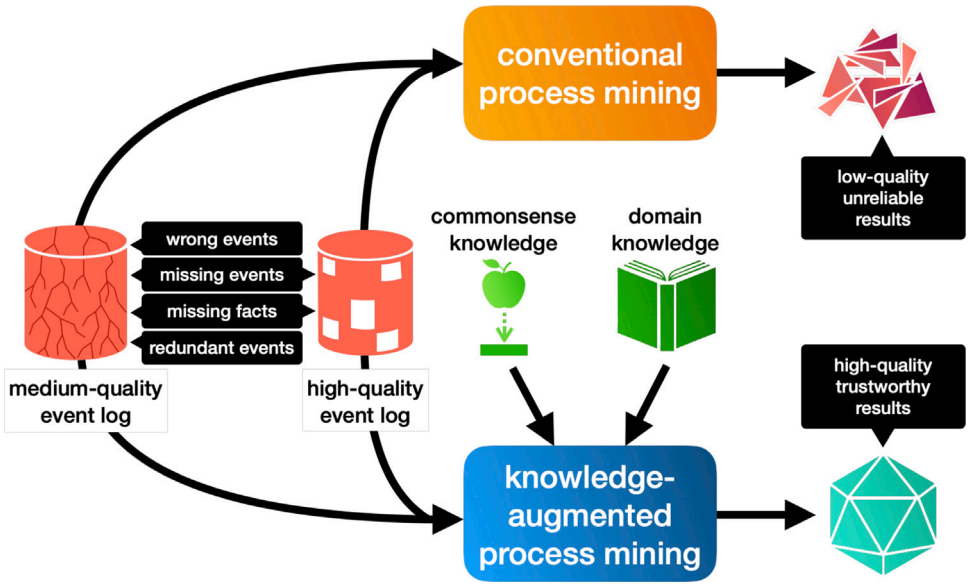

Fig. 7. Augmenting Process Mining with Common Sense and Domain Knowledge.

3.8. Worker-centric process management

Business processes are considered the backbone of enterprises. Through processes, companies track work, enforce policies, and ensure compliance. There was a pre-conceived notion of how work needs to be done, and tools were designed and developed in part to exemplify that notion. However, the recent adoption of digital transformation enabled a new agile (and ad-hoc) way of working that is outside the scope of business process tools. This adoption was accelerated by the recent pandemic and advancements in AI. While the technology provided the flexibility and agility in the way workers completed tasks, tools are still reflecting the older pre-conceived notions. Workers interfacing with these tools are now dealing with overhead of duplicate work, mundane set of tasks, and inefficient processes that are tailored to the tools rather than the workers. This not only affects their productivity but leads to enterprise-wide inefficiencies and increased costs. Recent work focused on enhancing these BPM tools and expanding their scope to accommodate this changing work landscape. While that is encouraged, the main focus of these tools would still be the business process. We envision worker-centric business processes and imagine that recent advancements in AI (including computer vision, natural language processing and machine learning) would facilitate the creation of a new generation of tools that meet workers where they are.

Once more, worker-centric BPM tooling (leveraging the latest AI advancements) is available to engage knowledge workers in the channels and modalities where they are already most comfortable and productive. As a result, workers would only concern themselves with performing their business tasks in the appropriate channels, and the BPM tooling takes care of the record-keeping. Interactions with the BPM tools to examine case history would also happen in the channels where workers operate. For example, a conversational assistant would allow a knowledge worker to query for overdue work items through a "more natural" free-form chat interface (i.e., direct interactions with the BPM tools should be deprecated in favor of more channel-specific ones). This is illustrated in Fig. 8 where the record-keeping step, denoted by dotted arrows, is no longer a manual step in the worker-centric solution. A worker-centric BPM solution brings several benefits. First, knowledge workers no longer need to perform the redundant step of recording their work in the BPM tool after already having performed it in the channel of their choice, such as a voice call. Second, workers that need to interact with the BPM artifacts have more choice in the modality (Rizk et al., 2020). Third, process owners and workers get a more complete and more accurate picture of the end-to-end "as-is" process (van der Aalst, 2013b) since directly instrumenting the channels where work is happening and automating the record-keeping means less work getting done "behind the system's back". This, in turn, provides a more accurate dataset for any process analysis to find opportunities for optimization or automation.

3.9. Mining processes using stochastic data

The increasingly growing amounts of uncertain process data, sourced from sensors or models with stochastic outputs, calls for process mining models over stochastically known event data. Consider the use of a machine learning algorithm to detect activities in video clips, such as video cameras on food preparation processes as in video-monitored restaurant kitchens. The cook prepares drinks and foods according to recipes (process models). Given a known (or discovered) set of models (cookbook recipes or historical supervised datasets), we wish to automatically identify the prepared dish based on video clips. Such identification can serve in confirming that a dish is prepared according to its recipe, informing diners regarding expected dish arrival time, or providing kitchen performance improvement. Other use cases exist in environments such as healthcare, security and manufacturing. The predicted trace, which is the result of data processing and learning techniques, is probabilistic (e.g., a softmax layer of a neural network). In practice, we expect such a problem to include a large number of events, depending on the length of the overall process and the sampling resolution. In the video clip example, the sampling may result in a large number of video frames. Also, whenever sampling is performed at a predetermined frequency, time points should be grouped into higher level activities. Therefore, the challenge lies in the magnitude of possible traces that follows from the uncertain trace representation and the fact that, to date, no conformance technique has been proposed to handle this type of stochastic uncertain traces.

Fig. 9 accommodates the spectrum of conformance checking tasks that may originate from stochastic data starting with a deterministic setting in Cases 1 and 3. In both these cases the process realizations and models are deterministic, where Case 1 is the standard conformance checking and Case 3 uses conformance checking to classify an observation as one of several process models based on the best conformance. Cases 5 and 7 are the counterparts of Cases 1 and 3 only the observation is stochastic. Case 5 may represent a setting in which one wants to check, for example, the conformance of a surgical procedure with its

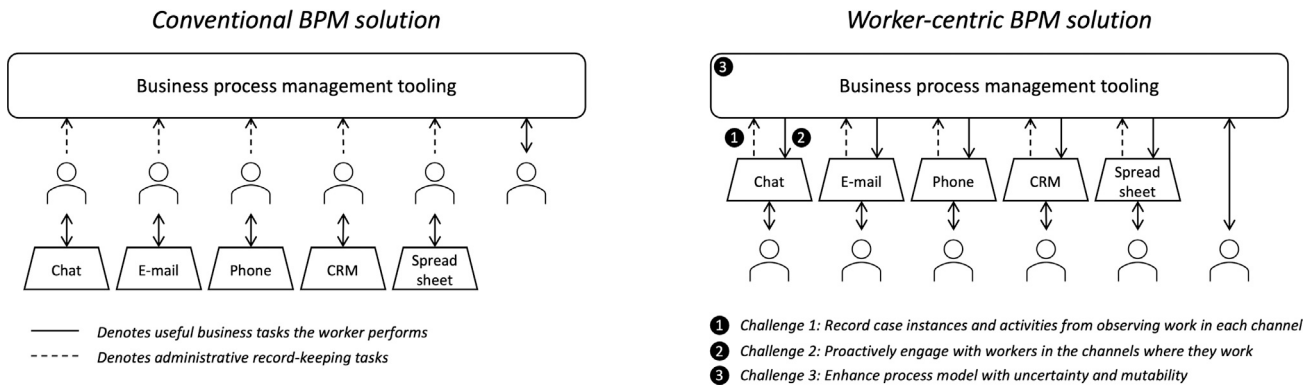

Fig. 8. Worker-centric process management.

model (e.g., for educating surgeons or debriefing purposes). Such a case poses the challenge of developing a conformance technique that explicitly accommodates the probabilistic information. In such a case, an example observation may be modeled using a probability matrix in which the rows correspond to activities, columns to timestamps, and entries represent the probability of an activity to occur in a time point. The matrix can be the outcome of a softmax layer of a neural network. In Case 7, a stochastically observed process needs to be classified into one of the process models using a conformance measure. A representative use-case may include a dataset of food preparation models (e.g., latte, tea, scrambled eggs, and cheese sandwich) and a stochastic log based on a video recorded dish preparation that needs to be automatically classified as one of the models. In such a case, we suggest conformance checking of the observation with respect to each of the models-the best conforming model is selected as the prepared dish.

In Cases 2, 4, 6 and 8, the models are stochastic. Such settings may arise when creating a fully supervised dataset is too costly. A natural way to discover the models is to apply neural network techniques on videos of known dishes, which would result in a stochastic trace for each historical video with a deterministically known label (i.e., the dish name is known). Cases 6 and 8 in which both models and the log are stochastic, are the most challenging since it is hard to distinguish between two types of stochasticity. The first reflects variations across process realizations (e.g., in 60). Resolving the problem would open several options. First, we would be able to perform conformance checking using stochastic data. This, in turn, will enable to analyze processes that were monitored by sensors and to fuse deterministic event data with stochastic data. Consider, for example, checking the conformance of processes that are monitored by cameras. This may facilitate the option of triggering real-time notifications about process deviations in critical settings such as aircraft maintenance processes and complex medical procedures. At times, offline notifications following conformance checking may suffice. In other cases, the problem is resolved by classifying noisy process observations into one of several process classes. In such a case, we would use stochastic alignment cost between a noisy observation and each process class as a loss function. This may facilitate automatic identification of observed processes. Finally, a stochastically known process trace 'hides' a multitude of possible trace realizations from which only one is the actual, observed trace. Solving the problem would allow to denoise the stochastic trace and reveal the actual trace. There are many additional implications, including correcting errors in uncertain event logs.

4. Cross-problem analysis

When comparing the problems, we make several observations. In the following sections, we reflect on the phases of the process life-cycle in which the problems are believed to reside, as well as on the balance between automation and human involvement in solving the problems. Moreover, we reflect on the application domains or scenarios where the solution is expected to have a major impact.

4.1. The problems and the process life-cycle

For each of the problems, we determined their importance in relation to the different phases of the process life-cycle: design, execute, monitor, and optimize. In Fig. 10, we illustrate the results. Every phase is identified as the most important one at least once, and the combinations are highly varied. Design, monitor, and optimize, are almost equally decisive over all problems. Interestingly, the execute phase is the exception, as it is overall regarded as less critical across the set of identified problems.

4.2. The level of automation and general themes

Each problem illustrates in its own way that there is a tension between the role that humans play in work processes and the level of automation that is being applied within such processes. This tension relates to both the presence and the absence of humans. In extension, it is evident that the problems cannot be solved through research alone. Solving some of the problems requires a different way of working in organizations.

We also observe, based on the discussions during the workshop, that the problems are important for different reasons. A major portion of the problems is difficult to solve because of their complexity, i.e., the problem is so complex to deal with in its entirety. Two problems relate to abstraction, or granularity, because the problem manifests itself on many different levels. The last theme that we identified is the environmental or contextual aspect, where the relation with the environment or context is crucial for a particular problem.

In Fig. 11, we illustrate the problems in relation to the extent to which we expect their solution to lean towards automation or human-driven tasks and the themes in which they belong. From the figure, we can observe that the expected level of automation is highly varied, especially for the problems dealing with a high amount of complexity. It seems that we cannot ignore either of automation and human involvement in those situations. For the problems that are mostly about abstraction, granularity, or environmental factors, the focus is mostly on the involvement of humans. Full automation is believed to be unrealistic there.

What is also interesting to note is a possible connection between the level of automation and the phases of the process life-cycle. The problems for which the solutions are believed to be highly automatic, i.e., **mining processes using stochastic data**, **automated process redesign**, and **augmenting process mining with common sense and domain knowledge**, happen to be the ones where the optimize phase scores high, but the execute phase is believed to be the least important.

	main task: conformance		main task: classification	
Model (Dataset) →	Single process		Multiple processes	
↓ Observation (Log)	DK	SK	DK	SK
Deterministically Known (DK)	1	2	3	4
Stochastically Known (SK)	5	6	7	8

Fig. 9. Mining processes using stochastic data. In the picture, eight cases are classified according to the characteristics of the process and observed log.

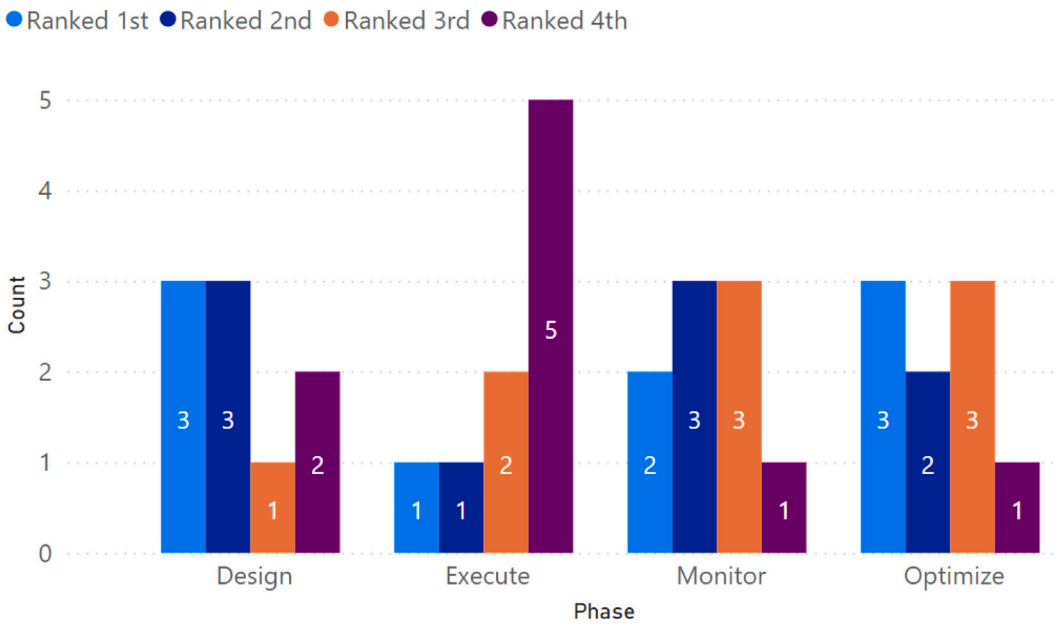

Fig. 10. Ranking the problems' importance in relation to the phases of the process life-cycle.

Fig. 11. Expected level of automation of the solution (100 = fully automated, 0 = fully human-driven).

Fig. 12. The Future Graveyard of BPM Problems.

One possible interpretation for this is that automation can refer to different phases, i.e., automation in the sense of a non-interactive way of executing processes, or automation in design and analysis tasks — which is part of the described problems themselves. The former was apparently not an essential part of the discussed problems.

4.3. Application domains and scenarios

Solutions to the identified problems are expected to predominantly impact knowledge-intensive business processes across industries. Especially in healthcare, many scenarios can be thought of that will be impacted by these solutions. Consider a concrete example around the problem of **value creation from data**. A state-wide public health service is striving to blend internal process-orientation with an external patient-focused approach, and through this strike a balance between efficiencies and quality of patient care. The use of data from internal and external sources can have a tremendous impact on building patient-oriented processes. The healthcare context, however, is complex with traditional function-oriented structures that have siloed data and prevent its consistent and pervasive use across the health system and by its various stakeholders. The management team recognizes that there is an array of processes and patient conditions that could benefit from use of data and algorithms, and is seeking a process oriented approach to implement an organization-wide change to unleash the value of data. Process orientation research that informs organizational structures and processes to better support data-driven work can develop the guidance required by organizations in such a scenario.

Manufacturing is another application domain in which solutions to the mentioned problems are expected to be valuable. Consider as an example the problem of **expansive BPM**. Processes can become more agile and capable of swiftly designing and delivering products based on rapid demands thanks to an overall process view and the simplified detection of change-related implications. In manufacturing, there is a large demand for solutions to data-related problem such as data value creation and providing process data at different granularity levels. Worker assistance is also a highly relevant problem in manufacturing. Moreover, digital twins or digital shadows play a pivotal role for the simulation and the testing of manufacturing processes.

The high complexity of many of the discussed problems may be due to the difference between theory and practice of business processes. Single processes may be modeled accurately, but in practice, they are part of a large network of processes, as outlined in the problem of **expansive BPM**. Because of the influence of outside factors, the execution in practice may be different than predicted. The wide range of factors that need to be taken into account in practice also complicate the achievement of goals such as **automated redesign** and **digital twins**. Techniques such as process mining, and more specifically, conformance checking, provide a means to compare the designed procedures to evidence from practice. However, giving insight into process executions in practice requires high-quality data. Many of the solutions to the identified problems promise to provide better data quality and as a result, a more truthful representation of work that is being executed in real life. Perhaps healthcare is expected to be impacted most as this is known to be a domain where data quality issues are widespread and much work takes place outside information systems. For example, a solution to the problem of **worker-centric process management** is believed to record data on process executions in the very same place as where it is executed. As such, the recording is done more naturally and previously hidden activities are now included. A solution to the problem of **augmenting process mining with common sense and**

domain knowledge will allow event data to be complemented to fill in the blanks and make sense of the data. Similarly, the solutions to the problems of **mining processes using stochastic data** and **fixed granularity levels for process analysis** are aimed at improving the view on processes from the data, to allow process analysts to draw more truthful insights from it.

When the problems are solved, we expect to see no more untapped value from data, no more process silos and no more manual process (re-)designs. In relation to event data and process models, we foresee a scenario without inaccurate what-if models, subjective models, fixed granularity levels, missing knowledge behind event data, separated work and recording, or low quality sensor data. Fig. 12 illustrates the future graveyard of BPM problems that we hope to see buried in the coming decades.

5. Conclusion

In this paper, we distinguish nine important problems that we currently face in BPM research. We discuss what challenges will need to be overcome to solve them and outline potential starting points for developing those solutions. In terms of digital and process innovation, we expect to get rid of untapped value from data, process silos, and manual process re-designs in the coming decades. We expect improvements in terms of more accurate prediction models and turning subjective models into objective ones. As for process mining, we hope to achieve a more faithful representation of reality through tackling fixed granularity, separated work and recording, missing knowledge, and low quality sensor data.

As we hope to have clarified, the future will bring many opportunities, if only we can eliminate the identified problems. Therefore, we end this paper with a call to pick up on these challenges in the interest of further developing the BPM discipline and serving society with research that reaches beyond the next increment.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

Adamo, G., Di Francescomarino, C., Ghidini, C., Maggi, F.M., 2021. Beyond arrows in process models: A user study on activity dependences and their rationales. Inf. Syst. 100, 101762.

Antunes, G., Bakhshandeh, M., Borbinha, J., et al., 2015. The process model matching contest 2015. In: Proc. of EMISA. pp. 127–155.

Artale, A., Kovtunova, A., Montali, M., van der Aalst, W.M.P., 2019. Modeling and reasoning over declarative data-aware processes with object-centric behavioral constraints. In: Proc. of BPM. In: LNCS, vol. 11675, Springer, pp. 139–156.

Beerepoot, I., Di Ciccio, C., Marrella, A., Reijers, H.A., Rinderle-Ma, S., Weber, B. (Eds.), 2021. Proceedings of the International Workshop on BPM Problems to Solve Before We Die (PROBLEMS 2021) co-located with the 19th International Conference on Business Process Management. BPM 2021, Rome, Italy, September 6-10, 2021, In: CEUR Workshop Proceedings, vol. 2938, CEUR-WS.org, URL http://ceur-ws.org/Vol-2938.

Bennett, N., Lemoine, G.J., 2014. What a difference a word makes: Understanding threats to performance in a VUCA world. Bus. Horiz. 57 (3), 311–317. http://dx.doi.org/10.1016/j.bushor.2014.01.001.

Beverungen, D., 2014. Exploring the interplay of the design and emergence of business processes as organizational routines. Bus. Inf. Syst. Eng. 6 (4), 191–202. http://dx.doi.org/10.1007/s12599-014-0335-3.

Cayoglu, U., Dijkman, R., Dumas, M., et al., 2013. The process model matching contest 2013. In: Proc. of BPM Workshops. pp. 442–463.

Davis, E., Marcus, G., 2015. Commonsense reasoning and commonsense knowledge in artificial intelligence. Commun. ACM 58 (9), 92–103.

Dumas, M., La Rosa, M., Mendling, J., Reijers, H.A., 2018. Fundamentals of Business Process Management. Springer, http://dx.doi.org/10.1007/978-3-662-56509-4.

Esser, S., Fahland, D., 2021. Multi-dimensional event data in graph databases. J. Data Semant. 10 (1), 109–141. http://dx.doi.org/10.1007/s13740-021-00122-1.

Ganzha, M., Paprzycki, M., Pawłowski, W., Szmeja, P., Wasielewska, K., 2017. Semantic interoperability in the internet of things: An overview from the INTER-IoT perspective. J. Netw. Comput. Appl. 81, 111–124. http://dx.doi.org/10.1016/j.jnca.2016.08.007, URL https://www.sciencedirect.com/science/article/pii/S1084804516301618.

Ghahfarokhi, A.F., Berti, A., van der Aalst, W.M.P., 2021. Process comparison using object-centric process cubes. arXiv preprint arXiv:2103.07184.

Gitschier, J., 2012. Scientist citizen: An interview with Bruce Alberts. PLoS Genet. 8 (5), e1002743.

Grisold, T., Gross, S., Röglinger, M., Stelzl, K., vom Brocke, J., 2019. Exploring explorative BPM – Setting the ground for future research. In: Proc. of BPM. Springer, pp. 23–31. http://dx.doi.org/10.1007/978-3-030-26619-6_4.

Gross, S., Stelzl, K., Grisold, T., Mendling, J., Röglinger, M., vom Brocke, J., 2021. The business process design space for exploring process redesign alternatives. Bus. Process Manag. J. http://dx.doi.org/10.1108/BPMJ-03-2020-0116.

Heiler, S., 1995. Semantic interoperability. ACM Comput. Surv. 27 (2), 271–273. http://dx.doi.org/10.1145/210376.210392.

Jablonski, S., Bussler, C., 1996. Workflow Management - Modeling Concepts, Architecture and Implementation. International Thomson.

Kerpedzhiev, G.D., König, U.M., Röglinger, M., Rosemann, M., 2021. An exploration into future business process management capabilities in view of digitalization. Bus. Inf. Syst. Eng. 63, 83–96. http://dx.doi.org/10.1007/s12599-020-00637-0.

Klinkmüller, C., Seeliger, A., Müller, R., Pufahl, L., Weber, I., 2021. A method for debugging process discovery pipelines to analyze the consistency of model properties. In: Proc. of BPM.

Klinkmüller, C., Weber, I., 2017. Analyzing control flow information to improve the effectiveness of process model matching techniques. Decis. Support Syst. 100, 6–14.

Klinkmüller, C., Weber, I., 2021. Every apprentice needs a master: Feedback-based effectiveness improvements for process model matching. Inf. Syst. 95, 101612.

Klinkmüller, C., Weber, I., Mendling, J., Leopold, H., Ludwig, A., 2013. Increasing recall of process model matching by improved activity label matching. In: Proc. of BPM. pp. 211–218.

Kreuzer, T., Röglinger, M., Rupprecht, L., 2020. Customer-centric prioritization of process improvement projects. Decis. Support Syst. 133, http://dx.doi.org/10.1016/j.dss.2020.113286.

Kuss, E., Leopold, H., van der Aa, H., Stuckenschmidt, H., Reijers, H.A., 2018. A probabilistic evaluation procedure for process model matching techniques. Data Knowl. Eng. 117, 393–406.

Latour, B., 1987. Science in Action: How to Follow Scientists and Engineers Through Society. Harvard University Press.

Malinova, M., Mendling, J., 2018. Identifying do's and don'ts using the integrated business process management framework. Bus. Process Manag. J..

Mendling, J., Leopold, H., Pittke, F., 2015. 25 Challenges of semantic process modeling. Int. J. Inf. Syst. Softw. Eng. Big Companies (IJISEBC) 1 (1), 78–94.

Merideth, J.C., Bandara, W., O'Neill, D., 2020. Process portfolio management for enhanced digital readiness: Insights from a large Australian bank. In: Proc. of ICIS. Association for Information Systems.

Panetta, K., 2020. Gartner top strategic technology trends for 2021. URL https://www.gartner.com/smarterwithgartner/gartner-top-strategic-technology-trends-for-2021/.

Reichert, M., Weber, B., 2012. Enabling Flexibility in Process-Aware Information Systems: Challenges, Methods, Technologies, Vol. 54. Springer Science & Business Media.

Reijers, H.A., 2021. Business process management: The evolution of a discipline. Comput. Ind. 126, 103404.

Rizk, Y., Isahagian, V., Boag, S., Khazaeni, Y., Unuvar, M., Muthusamy, V., Khalaf, R., 2020. A conversational digital assistant for intelligent process automation. In: Proc. of BPM. pp. 85–100.

Rodríguez, C., Klinkmüller, C., Weber, I., Daniel, F., Casati, F., 2016. Activity matching with human intelligence. In: Proc. of BPM Forum. pp. 124–140.

Satyal, S., Weber, I., Paik, H., Di Ciccio, C., Mendling, J., 2019. Business process improvement with the AB-BPM methodology. Inf. Syst. 84, 283–298. http://dx.doi.org/10.1016/j.is.2018.06.007.

Schuster, D., Zelst, S.J.v., van der Aalst, W.M.P., 2021. Freezing sub-models during incremental process discovery. In: Proc. of ER. Springer, pp. 14–24. http://dx.doi.org/10.1007/978-3-030-89022-3_2.

Shneiderman, B., 1996. The eyes have it: A task by data type taxonomy for information visualizations. In: IEEE Symposium on Visual Languages. IEEE, pp. 336–343. http://dx.doi.org/10.1109/VL.1996.545307.

Si, Y., Chan, V., Dumas, M., Zhang, D., 2018. A Petri nets based generic genetic algorithm framework for resource optimization in business processes. Simul. Model. Pract. Theory 86, 72–101. http://dx.doi.org/10.1016/j.simpat.2018.05.004.

Sverdlov, E., 2018. Incremental science: Papers and grants, yes; Discoveries, no. Mol. Genet., Microbiol. Virol. 33 (4), 207–216.

Teinemaa, I., Dumas, M., La Rosa, M., Maggi, F.M., 2019. Outcome-oriented predictive process monitoring: Review and benchmark. ACM Trans. Knowl. Discov. Data 13 (2), 1–57. http://dx.doi.org/10.1145/3301300.

van der Aalst, W.M.P., 2013a. Business process management: A comprehensive survey. ISRN Softw. Eng. 2013, 1–37. http://dx.doi.org/10.1155/2013/507984.

van der Aalst, W.M.P., 2013b. "Mine your own business": Using process mining to turn big data into real value. In: Proc. of ECIS, (no. 225), URL https://aisel.aisnet.org/ecis2013_cr/225.

van der Aalst, W.M.P., Weske, M., Grünbauer, D., 2005. Case handling: A new paradigm for business process support. Data Knowl. Eng. 53 (2), 129–162.

Van Looy, A., 2021. How the COVID-19 pandemic can stimulate more radical business process improvements: Using the metaphor of a tree. Knowl. Process Manag. 28 (2), 107–116.

Vanwersch, R.J., Vanderfeesten, I., Rietzschel, E., Reijers, H.A., 2016. Improving business processes. In: Proc. of BPM. Springer, pp. 3–18. http://dx.doi.org/10.1007/978-3-319-23063-4_1.

van Zelst, S.J., Mannhardt, F., de Leoni, M., Koschmider, A., 2020. Event abstraction in process mining: Literature review and taxonomy. Granul. Comput. 1–18.

Zerbato, F., Seiger, R., Di Federico, G., Burattin, A., Weber, B., 2021. Granularity in process mining: Can we fix it? pp. 40–44, In Beerepoot et al. (2021), URL http://ceur-ws.org/Vol-2938.