Knowledge Graph Reasoning for Intelligent and Explainable Business Process Technologies

Leon Bein

Technical University of Munich, Bildungscampus 2, 74076 Heilbronn, Germany

Abstract

Supporting process executions via software systems is a key concern of the Business Process Management field. With increasing digitization and technological advances, the potential scope of this support shifted and extended, from simply orchestrating tasks following pre-defined routines toward providing comprehensive decision support. However, modern intelligent process support approaches suffer from scattered knowledge bases and black-box technologies, leading to intransparent, hard-to-adapt automated decision making. Knowledge graph technologies appear promising to address this gap. These enriched graph representations of the real world allow for integrative encoding of knowledge of different dimensions and performing explainable reasoning thereupon. However, their usage as key driver for business process management systems has not yet been investigated. In this doctoral project, we will investigate the usefulness of knowledge graph technologies for Business Process Management Systems (BPMSs), and design and implement a reference architecture for knowledge-graph-based BPMSs. This research proposal outlines the potentials for using Knowledge Graph technologies for the design of an explainable, deep-knowledge-driven Business Process Management System, provides an overview of existing work, and presents a research outline and the current progress therein.

Keywords

Business Process Management, Knowledge Graphs, Prescriptive Process Monitoring, Ontology-based Business Process Modeling

1. Introduction

A key part of the Business Process Management (BPM) field is the support of business process design and execution using dedicated software systems. These *Business Process Management Systems (BPMSs)* support organizations by providing process modeling and monitoring capabilities, as well as orchestrating their execution by determining, prioritizing, and assigning tasks, managing handovers between them, and detecting and escalating business errors [1, 2]. With an increasing digitization of processes and technological advances, particularly in the umbrella field of Artificial Intelligence (AI), the scope and applicability of these systems have been steadily widening. A shift in focus, from simply orchestrating tasks to providing comprehensive decision support hat been observed [3, 4]. Recent visions of *AI-augmented* business process management systems (ABPMS) promise smarter, more autonomous, and more adaptive orchestration. Among other things, this promises to extend software support to *Knowledge-intensive Processes* (KiP) – processes that are driven by expert decisions based on domain knowledge, exhibiting high variability and need for flexibility [2, 5, 3].

To realize this vision, *Knowledge Graph* (KG) technologies have been identified as a potential key component [5, 6, 7]. Knowledge Graphs can be defined as data graphs that represent entities and their relations in a domain of interest, and are extended by context in the form of ontologies, metadata, and more [8]. Associated with this notion are a field of research and a toolbox of methods and technologies, concerned with the elicitation of domain knowledge, its encoding into graphs, and (software-driven) reasoning thereupon [8].

Using knowledge graph technologies for business processes can be motivated from multiple angles. Our motivation is mainly based, first, on the ability of KGs to encode and consequently integrate

Proceedings of the Best BPM Dissertation Award, Doctoral Consortium, and Demonstrations & Resources Forum co-located with 23rd International Conference on Business Process Management (BPM 2025), Seville, Spain, August 31st to September 5th, 2025.
leon.bein@tum.de (L. Bein)
0000-0001-9064-7905 (L. Bein)

knowledge of different process perspectives, notably including domain-specific concepts that do not fall under any common BPM category. This enables a unified management, retrieval, and processing of the knowledge from these different perspectives. Second, the ability of knowledge graph reasoning methods to jointly consider that knowledge and provide comprehensive explanations [8]. These two capabilities address two key gaps in state-of-the-art process support systems, namely the limiting focus on knowledge about activities and their relations, and the lack of explainability for decisions made by current intelligent (neural-net-based) approaches, resulting from this lack of semantic underpinning and the black-box nature of the employed technologies [9].

To realize these potentials, changes to all phases of the interaction with classical BPMS can be made: For instance, (i) the knowledge stored needs to be extended and shifted from non-uniform separated representations of activities, goal, organization etc. to a unified process knowledge graph, (ii) process modeling needs to be extended to more general process knowledge elicitation, and (iii) strictly following process models or black-box decision making on a shallow knowledge base needs to be shifted to process-knowledge-graph-based explainable reasoning. We refer to such systems as *Knowledge-graph-based BPMS* (KG-BPMS).

While the potentials of KG-BPMS have been hinted at in previous works (s.a.), the concept has not been investigated in depth so far. Investigating the realization of KG-BPMSs can shed light on their true potentials and provide a puzzle piece for advancing the design of ABPMSs. Consequently, this doctoral project aims to investigate the following overarching questions:

RQ1: How can a business process management system be designed that utilizes knowledge graphs as primary process knowledge representation and associated technologies for process design and execution?

RQ2: How do systems following the design perform? What are advantages and challenges of such a design?

Using knowledge graphs as main process knowledge representation and associated technologies as main driver for process design and execution potentially touches upon, and consequently affects, most aspects of a BPMS. These different application potentials synergize, e.g., when using KGs as main knowledge representation, KG-based reasoning can directly be applied without the need to first translate the present knowledge. Thus, they can hardly be investigated in full isolation from each other. Consequently, to adequately assess KG potentials for BPMSs, we plan to investigate a full cross-cut, i.e., from domain knowledge to decision support. To structure our research, we break this cross-cut down into several sub-problems to be tackled one after another, informed by the identified potentials of KG-based BPMS and the common structure of BPMSs: (i) the encoding of process knowledge as (knowledge) graph, (ii) its elicitation and import into said graph, and (iii) its utilization for process support.

In the remainder of this proposal, we first outline the potentials for KG-BPMSs in more detail to identify the subproblems of interest, outline relevant existing works, and finally present the resulting work packages and research plan for this project and discuss challenges for the research design.

2. Potentials for Knowledge Graphs in Business Process Management Systems

Rich Integrated Process Encoding A KG-BPMS can integrate the various dimensions of process knowledge into one uniform data structure, a *Process Knowledge Graph* (PKG). These dimensions include (i) common process elements, such as activities and decision points, and their relations (as usually depicted in, e.g., BPMN or Declare process models), (ii) organizational structure, such as (human) resources, their roles, permissions, and skills, and resource requirements on activities, (iii) goals and constraints, often implicitly captured or expressed declaratively (cf. [3] for KiP modeling requirements). Further, notably, a PKG can also integrate knowledge of and about domain-specific concepts that do not fall under common classifications from the BPM field. For instance, in medical processes, knowledge about diseases, drugs, symptoms, etc., and their relations is relevant for decision-making. In these cases, connecting to and reusing existing domain-specific ontologies is promising (cf. [10, 11] for examples of medical ontologies). This way, subsequent querying on and processing is able to consider a *rich* width of process knowledge, that is encoded into one *integrated* structure.

In order to achieve this rich, integrated process encoding, an adequate format needs to be found and respective ontologies (to be extended) engineered. Existing formalisms for the common process knowledge dimensions can be reused here, e.g., for activity relations and execution history [12, 13], for organizational knowledge [14], and for goals and constraints [15]. Further, ways to elicit the process knowledge need to be investigated. Existing sources of process knowledge, such as human input, event logs, textual process documentations, process models, regulatory documents, etc., as well as existing methods, e.g., from the process mining and knowledge graph construction fields, appear promising to be reused here. Challenges herein lie with transforming the different formats into graph form and ensuring consistency with existing encoded knowledge. We want to highlight here the potential for *Large Language Models* (LLMs) for translating textual process descriptions as well as for conversational knowledge input and conflict resolution.

Explainable Prescriptive Process Support Having a process knowledge graph at the ready allows applying KG-based knowledge retrieval and reasoning techniques on an integrated representation of the different process knowledge dimensions. This yields several promises for the design of BPMSs. First, to address requirements for designing ABPMSs (cf. [5]) and supporting KiPs (cf. [3]). KiP looseness and variability can be realized by encoding relevant entities and constraints in one structure rather than relying mainly on activity relationships, such as in BPMN or even Declare, allowing systems to act more autonomously while considering a well-defined frame. When execution is inherently derived from goals and constraints rather than routines, adaptation to unexpected events and errors becomes easier. Process evolution is easier and more comprehensive, as changes only have to be made to the relevant entities instead of solely to activity relationships.

Second, a KG-BPMS promises more comprehensive decision-making and prescriptive support. This stems from the availability of explainable graph reasoning approaches (as opposed to black-box process prediction models) that can utilize the width of process knowledge for their reasoning and recommendations (as opposed to only providing, e.g., counterfactuals on attributes of an event log).

Finally, a KG-BPMS can support process participants and owners in providing insight about the process(es) and their execution. Applying graph knowledge retrieval approaches, allows for explaining, e.g., where the dependencies of two activities stem from, *based on* the background process knowledge encoded in the PKG.

To achieve explainable prescriptive support first requires the adaptation of existing graph reasoning approaches to inform a novel kind of process engine. Approaches need to be able to interpret the width of process knowledge, while at the same time producing comprehensive explanations.

3. Related Work

Ontology-based Business Process Modeling One of the most important existing streams of BPM for this project is the field of *Ontology-based Business Process Modeling* (OBPM), also called *Semantic Business Process Modeling* (cp., e.g., [16]). OBPM encompasses approaches to extend classical process models with instances of ontologies that represent additional background process knowledge, with the explicit intent to make this knowledge usable for machines [17, 18]. Existing works already allow expressing BPMN and execution traces as graphs [12, 13], enabling, e.g., checking compliance rules on process models *considering background knowledge* [17, 16]. However, these works are limited to reasoning at process design-time, and the major focus lies on integrating with existing process modeling approaches rather than reasoning [18].

Knowledge-driven Process Execution Visions and Frameworks Further single key works stand out to be mentioned because of their relevance for and/or similarity to this work: In their ABPMS manifesto [5], leading figures of the BPM community introduce the concept and the promised capabilities. These include framed autonomous execution and adaptation, continuous self-improvement, and proactive explaining, partially motivating the KG-BPMS potentials presented in this proposal. The paper names knowledge graphs as one technology to use for internal knowledge representation and process framing in ABPMSs, without going into detail on how to realize such usage, thus providing a perfect gap for our work.

Similar to our work, Beheshti et al. [7] propose a transformer-architecture-based approach for using a process knowledge graph to dynamically determine process flow and perform next-activity-prescription. While their work already provides proposals on how to build the graph (by crowdsourcing) and perform reasoning on it (using a transformer architecture), it stays at a proposal level.

In [4], Kir and Erdogan present an intelligent BPM framework based on multi-agent systems, and, notably, ontological modeling. They define different aspects to be modeled in the ontology, suggest existing ontologies to be adopted, and propose to use graph rules for expressing preconditions and effects of actions, i.e., activities, and for expressing process rules. While their work already gives insight into what might be relevant aspects of process knowledge and proposes one method for reasoning on process knowledge graphs (multi-agent systems), it does not contribute to explainable decision support or give guidance on the elicitation of that knowledge.

To sum up, existing works already suggest and motivate using knowledge graphs for BPM, and solutions for single problems have been proposed. However, there is a lack of an implemented, end-to-end, knowledge-graph-based BPMS approach that integrates rich process knowledge and provides explainable execution support.

4. Research Design & Current Progress

The intended main artifact of this research is a design of a knowledge-graph-based BPMS, intended to give guidance on the usage of KGs for business process technologies and for evaluating potentials and challenges. Consequently, as overarching method, we adopt a design-science research (DSR) approach, as proposed, e.g., by Peffers et al. [19] or Hevner et al. [20]. The subchallenges raised in section 2 hereby serve to structure design iterations, each with their own (sub-)objectives and (sub-)evaluation, feeding into the overall objectives and evaluation and contributing to the overall KG-BPMS design. Consequently, we define the following work packages:

WP1: Base Architecture As a base and reference for future iteration, the first step in this project is to develop a base architecture for KG-BPMSs and create a respective base prototype. So far, groundwork has been done by identifying the potentials of the knowledge graph usage for BPMSs and generating an initial overview of the dimensions of process knowledge. A respective paper presenting the general vision has been accepted and presented at the AI4BPM workshop in fall 2024 [6]. A simplified draft for a base architecture is further displayed in Fig. 1, showcasing the envisioned components and interaction across system boundaries. The development and investigation of the base prototype is in progress, with a rule-based approach for next-activity prediction as well as a base process knowledge ontology already in place and a base version of knowledge elicitation yet to be developed.

WP2: Encoding of Process Knowledge An adequate graph format needs to be found and respective ontologies engineered. Following guidelines from ontology-based process modeling [18], the goal of this WP is to formalize an initial set of standard process concepts into a base process knowledge graph design, which organizations can then extend. To this end, we started and will continue to scan and relate the dimensions of process knowledge described in BPM literature. Rather than iterating, we plan the insights from the subsequent work packages to continuously seep into this baseline, extending it as necessary.

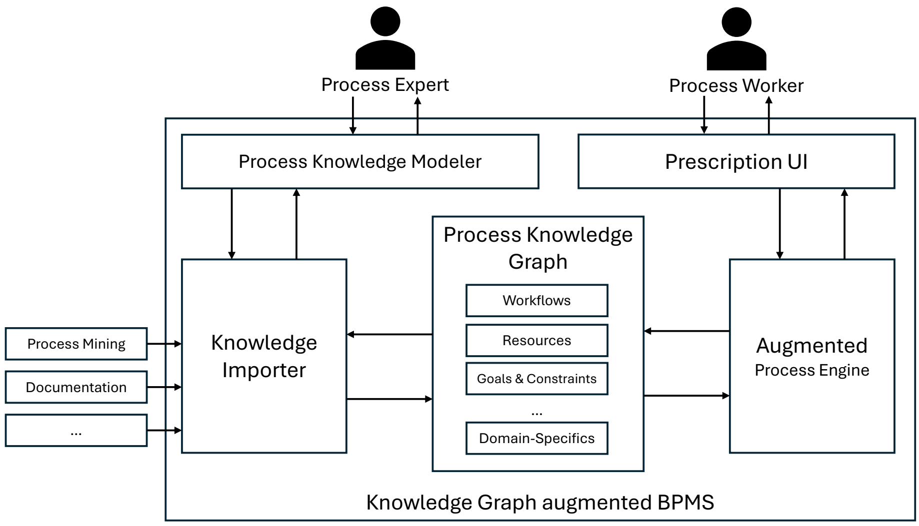

Figure 1: Simplified draft architecture of a KG-BPMS

WP3: KG-based Process Execution Support Goal of this WP is to investigate the usage of knowledge graph reasoning techniques for supporting process execution, especially focusing on the consideration of process knowledge for system decisions and explanations. Specifically, we want to investigate KGR for (i) *next activity prescription* and (ii) *resource allocation*. For this, we plan to transfer and adapt techniques from Knowledge Graph Reasoning and integrate them with existing BPM paradigms. Multiple DSR iterations of this work package can be done to improve performance or investigate different avenues of application. A paper outlining a base approach for using a KG-BPMS for resource allocation has been handed in for review as of writing this proposal.

WP4: Process Knowledge Elicitation Goal of this WP is to devise methods of "extracting" the necessary knowledge to fill the process knowledge graph and allow for execution support. Again, we plan to investigate multiple directions, namely (i) from textual process descriptions, for which we plan to apply and adapt existing knowledge extraction methods, potentially based on LLMs, (ii) from event logs and execution data, for which we plan to utilize existing process mining and graph rules mining methods, and (iii) from process experts directly, for which we envision a conversational user interface. So far, we have investigated existing graph rules mining methods. Initial insights, however, indicated poor applicability, motivating the development of tailored solutions.

WP5: Overall evaluation All of the previous WPs contribute to and inform the design of a final KG-BPMS and a prototypical implementation thereof. In a final step, we plan to evaluate this overall artifact, ideally using a real-world industry case, strengthening overall validity by showing feasibility and utility in organizational settings.

4.1. Research Challenges & Limitations

Due to the cross-cut nature of the planned research, a fine line has to be walked between trying to be too general, which runs the risk of investigating every area too shallowly, and delving too deeply into the single sub-problems, which runs the risk of the scope of the project becoming too large. To address this challenge, we have structured the WPs as described above, focusing on one hand on finishing the cross-cut before iterating on the same WP, and, on the other hand, investigating each sub-problem as its own WP, to ensure sufficient depth. Further, the current research takes a strongly technology-centered approach. However, in the same way our proposed changes affect BPMSs in many ways, they also affect humans involved with the processes. Encoding the knowledge in the system and using it for reasoning takes agency from the humans, who are thus far the main knowledge holders and decision makers. This can lead to phenomena such as skill decay, depersonalization, and automation bias [21]. While the scope of this project is sufficiently loaded, this important gap motivates future work.

Declaration on Generative AI

The author has not employed any Generative AI tools.

References

[1] M. Dumas, M. La Rosa, J. Mendling, H. A. Reijers, Fundamentals of Business Process Management, Springer, Berlin, Heidelberg, 2018. doi:10.1007/978-3-662-56509-4.

[2] H. Vu, H. Leopold, H. van der Aa, What is Business Process Automation Anyway?, in: 56th Hawaii International Conference on System Sciences (HICSS 2023), ScholarSpace, 2023, pp. 5462–5471.

[3] C. Di Ciccio, A. Marrella, A. Russo, Knowledge-Intensive Processes: Characteristics, Requirements and Analysis of Contemporary Approaches, Journal on Data Semantics 4 (2015) 29–57. doi:10.1007/s13740-014-0038-4.

[4] H. Kir, N. Erdogan, A knowledge-intensive adaptive business process management framework, Information Systems 95 (2021) 101639. doi:10.1016/j.is.2020.101639.

[5] M. Dumas, F. Fournier, L. Limonad, A. Marrella, M. Montali, J.-R. Rehse, R. Accorsi, D. Calvanese, G. De Giacomo, D. Fahland, A. Gal, M. La Rosa, H. Völzer, I. Weber, AI-Augmented Business Process Management Systems: A Research Manifesto, ACM Transactions on Management Information Systems 14 (2023) 1–19. doi:10.1145/3576047. arXiv:2201.12855.

[6] L. Bein, L. Pufahl, Knowledge Graphs: A Key Technology for Explainable Knowledge-Aware Process Automation?, in: Business Process Management Workshops, volume 534, Springer, Cham, 2025, pp. 18–30. doi:10.1007/978-3-031-78666-2_2.

[7] A. Beheshti, J. Yang, Q. Z. Sheng, B. Benatallah, F. Casati, S. Dustdar, H. R. M. Nezhad, X. Zhang, S. Xue, ProcessGPT: Transforming Business Process Management with Generative Artificial Intelligence, in: 2023 IEEE International Conference on Web Services (ICWS), IEEE, Chicago, IL, USA, 2023, pp. 731–739. doi:10.1109/ICWS60048.2023.00099.

[8] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, J. E. L. Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A.-C. N. Ngomo, S. M. Rashid, A. Rula, L. Schmelzeisen, J. Sequeda, S. Staab, A. Zimmermann, Knowledge Graphs, ACM Computing Surveys 54 (2022) 1–37. doi:10.1145/3447772. arXiv:2003.02320.

[9] A. Buliga, C. Di Francescomarino, C. Ghidini, F. M. Maggi, Counterfactuals and Ways to Build Them: Evaluating Approaches in Predictive Process Monitoring, in: Advanced Information Systems Engineering, Lecture Notes in Computer Science, Springer, Cham, 2023, pp. 558–574. doi:10.1007/978-3-031-34560-9_33.

[10] O. Bodenreider, The Unified Medical Language System (UMLS): Integrating biomedical terminology, Nucleic Acids Research 32 (2004) D267–D270. doi:10.1093/nar/gkh061.

[11] C. Ruiz, M. Zitnik, J. Leskovec, Identification of disease treatment mechanisms through the multiscale interactome, Nature Communications 12 (2021) 1796. doi:10.1038/s41467-021-21770-8.

[12] S. Bachhofner, E. Kiesling, K. Revoredo, P. Waibel, A. Polleres, Automated Process Knowledge Graph Construction from BPMN Models, in: Database and Expert Systems Applications, volume 13426, Springer, Cham, 2022, pp. 32–47. doi:10.1007/978-3-031-12423-5_3.

[13] A. Annane, N. Aussenac-Gilles, M. Kamel, BBO: BPMN 2.0 Based Ontology for Business Process Representation, in: 20th European Conference on Knowledge Management (ECKM 2019), volume 1, 2019, pp. 49–59. doi:10.34190/KM.19.113.

[14] M. Arias, J. Munoz-Gama, M. Sepúlveda, Towards a Taxonomy of Human Resource Allocation Criteria, in: Business Process Management Workshops, volume 308, Springer, Cham, 2018, pp. 475–483. doi:10.1007/978-3-319-74030-0_37.

[15] A. Seidel, S. Haarmann, M. Weske, Model-based decision support for knowledge-intensive processes, Journal of Intelligent Information Systems 61 (2023) 143–165. doi:10.1007/s10844-022-00770-0.

[16] C. Di Francescomarino, C. Ghidini, M. Rospocher, L. Serafini, P. Tonella, Reasoning on Semantically

Annotated Processes, in: D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, B. J. Krämer, K.-J. Lin, P. Narasimhan (Eds.), Service-Oriented Computing – ICSOC 2007, volume 4749, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 132–146. doi:10.1007/978-3-540-89652-4_13.

[17] C. Corea, P. Delfmann, Detecting Compliance with Business Rules in Ontology-Based Process Modeling, Wirtschaftsinformatik 2017 Proceedings (2017).

[18] C. Corea, M. Fellmann, P. Delfmann, Ontology-Based Process Modelling - Will We Live to See It?, in: Conceptual Modeling, Springer, 2021. doi:10.1007/978-3-030-89022-3_4.

[19] K. Peffers, T. Tuunanen, M. Rothenberger, S. Chatterjee, A design science research methodology for information systems research, Journal of Management Information Systems 24 (2007) 45–77. doi:10.2753/MIS0742-1222240302.

[20] A. R. Hevner, S. T. March, J. Park, S. Ram, Design Science in Information Systems Research, MIS Quarterly 28 (2004) 75. doi:10.2307/25148625.

[21] F. Klessascheck, L. Bein, J. Haase, L. Pufahl, A Critical Investigation of Rationalities in Automation with BPM, in: 2024 26th International Conference on Business Informatics (CBI), IEEE, 2024, pp. 30–39. doi:10.1109/CBI62504.2024.00014.