

LLM4PM: A case study on using Large Language Models for Process Modeling in Enterprise Organizations

Clara Ziche and Giovanni Apruzzese

University of Liechtenstein

Abstract. We investigate the potential of using Large Language Models (LLM) to support process model creation in organizational contexts. Specifically, we carry out a case study wherein we develop and test an LLM-based chatbot, `PRODIGY` (PROcess moDellIng Guidance for You), in a multinational company, the Hilti Group. We are particularly interested in understanding how LLM can aid (human) modellers in creating process flow diagrams. To this purpose, we first conduct a preliminary user study (n=10) with professional process modellers from Hilti, inquiring for various pain-points they encounter in their daily routines. Then, we use their responses to design and implement `PRODIGY`. Finally, we evaluate `PRODIGY` by letting our user study’s participants use `PRODIGY`, and then ask for their opinion on the pros and cons of `PRODIGY`. We coalesce our results in actionable takeaways. Through our research, we showcase the first practical application of LLM for process modelling in the real world, shedding light on how industries can leverage LLM to enhance their Business Process Management activities.

1 Introduction

Organizations perform business processes to deliver value-adding outcomes to their customers. Hence, Business Process Management (BPM) capabilities, such as *process modeling*, are a pivotal task in modern enterprises [3]. However, despite decades of efforts [7], process modeling still remains a costly activity due to, e.g., the difficulty of providing clear, up-to-date and easy-to-retrieve *documentation* [3] to those tasked to carry out such activities—the process modelers.

Inspired by recent developments in artificial intelligence (AI), such as large language models (LLM), researchers have proposed various techniques that can facilitate BPM-related tasks (e.g., [6]). Indeed, LLM can elaborate large collections of documents. Hence, by receiving an input from a given user, LLM can quickly produce an output that *(i)* accounts for existing documentation, while simultaneously *(ii)* answering the request of the user—i.e., a human. Yet, we found no evidence of *practical* applications of LLM for BPM in real contexts and, in particular, for process modeling. Hence, there is a need to investigate the effectiveness of such automation in industry [15]. Here, we tackle this challenge and showcase how a large enterprise, **Hilti**, can benefit from a LLM-powered chatbot—which we developed ad-hoc for Hilti—for BPM.

CONTRIBUTIONS. We present a (the first) real-world case study showcasing the application of LLM for Process Modeling in operational contexts. Specifically, to advance the state of the art on BPM, we:

- describe the *problems faced by the considered organization*, Hilti, providing evidence of the necessities of modern enterprises (§2);
- carry out a *requirement analysis by conducting interviews* with Hilti’s employees, shedding light on the pain-points of professional process modelers (§3.1);
- use our interviews as a scaffold to *develop an original LLM-based chatbot*, PRODIGY (Fig. 1), designed to support Hilti’s process modelers (§3.2);
- evaluate the ability of PRODIGY to *generate practical value* by (i) having Hilti’s employees use PRODIGY and (ii) collecting and analysing their feedback (§3.3). Our results (§4) show that PRODIGY is generally well-received, and identify room for improvement. We also derive lessons learned that future work can use to drive practical deployment of LLM-based technologies (§5).

Industrial Secret and Ethics: In this paper, we describe some elements pertaining to the internal processes of Hilti: we have been granted permission to share such information. Furthermore, we carry out our user studies ethically: our institutions are aware of this research, our participants have been informed of their rights and of the purpose of their contributions, and we have their consent to post them.

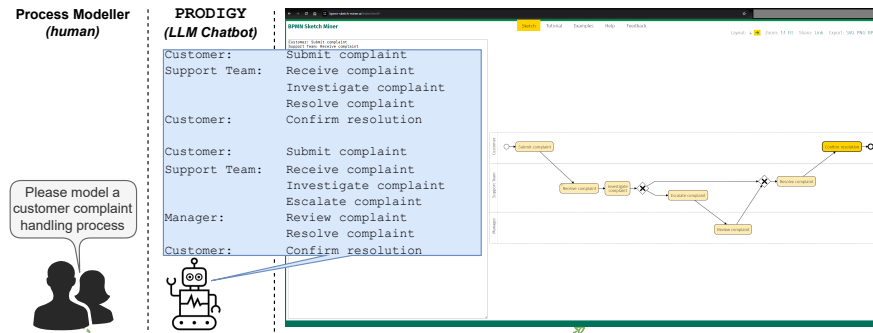


Fig. 1: An exemplary usage of PRODIGY – Our LLM-powered chatbot can fulfill various BPM-related tasks. Its most appreciated functionality is generating an output for a complete process model (if copy-pasted into the open-source tool BPMN Sketch Miner [10])

2 Organizational Context and Problem Statement

Our case organization, Hilti Group, is a multinational company that was founded in 1941 in Schaan, Liechtenstein. It is a world market leader in fastening and demolition technology for construction professionals and provides tools, technologies, software and services to the global construction industry. In 2023, Hilti’s workforce consists of about 33.000 employees in more than 120 countries, making it a highly diverse, distributed organization that operates in complex and competitive markets all over the world. The size, complexity and business model of Hilti make it an ideal use case for testing the capabilities of LLM for process modeling: BPM is essential to ensure cooperation and consistent outcomes within Hilti’s ecosystem; furthermore, it is crucial for Hilti to optimize customer-facing processes. Hence, a smooth process modelling is pivotal for Hilti.

Challenge. Hilti has an extensive and heterogeneous documentation landscape which adds to the intrinsically complex nature of process modeling. Hilti’s employees spend abundant time searching through such documentation (our interviews revealed an average of ~ 40 minutes of search *before* modelling a process). Hence, to improve the productivity of their process modellers, and to actively explore innovative technologies, Hilti is interested in novel solutions that facilitate the routines of their employees.

Technological Gap. LLM-based solutions are common (in 2024). However, existing techniques cannot be applied to Hilti’s use case. This is because of the *confidential nature* of Hilti’s documents: publicly available models (e.g., ChatGPT) should not be able to access Hilti’s data to shape their responses; furthermore, even *interacting* with certain LLM (or their APIs) from within Hilti’s networks triggers warnings, preventing a reliable usage of these solutions—which are leveraged also by renown prior work, such as [4,6,11,12,2,18].

OUR GOAL. We seek to design, develop and evaluate an LLM-based solution that facilitates the job of Hilti’s process modellers. The development of such a solution should be driven by Hilti’s distinctive organizational’s context—including its employee’s viewpoint, and its existing documentation.

3 Research and Methods

Inspired by Peffers et al. [14], we followed a Design Science Research (DSR) process consisting of four phases depicted in Fig. 2.¹ DSR is appropriate given our goal of examining LLMs for process modelling in organizations, as DSR emphasizes the creation of innovative solutions (in our case, PRODIGY) while also considering the context in which these solutions will be applied.

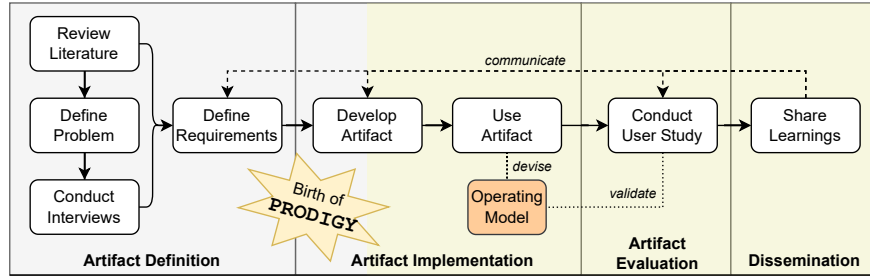


Fig. 2: **Method.** We rely on design science research to design, develop, and deploy our artifact. During the implementation of PRODIGY, we also devise an “operating model” (which we validate in the evaluation of PRODIGY) through which we explain how PRODIGY should be used in real organizations.

3.1 Artifact Definition

As a preliminary step, we carried out a systematic literature review [20] which we used as a foundation to investigate the state of the art and define the scope

¹ **Background:** DSR is a methodology that focuses on creating and evaluating artifacts to solve complex problems. Such procedure is rooted on the coming together of people, organizations and technology, with the ultimate intention of “extending the boundaries of human and organizational capabilities” [8].

of our project (see §2). Then, we carried out structured interviews [5] meant to identify pain-points and desiderata by professional process modelers² working for Hilti. We found an agreement with 10 employees, summarised in Table 1. The complete questionnaire is provided in our repository [1]. Among the most relevant questions, we ask: “what *challenges* do you experience when modelling processes?”, “how helpful is *existing documentation* when you model processes?” and “what would you like to see in a new AI artifact that supports process modelling?”; we also provide a list of functionalities for the AI artifact and ask to rate them on a 1–5 scale, as well as potential concerns. Finally, we inquire about the time spent looking for, and reviewing, existing documentation.

Table 1: **Overview of process modellers.** Our participants pertain to various geographical locations of Hilti, and have diverse backgrounds. [Demographics] Each participant has 5–25 years of experience in BPM, and they are within 26–60 years of age. Seven hold a MSc. degree. The male:female ratio is 6:4. They all have “above average” or “advanced” computer knowledge, and all have a basic understanding of LLM. Five perform process modelling activities at least weekly.

#	Job Title	Functional Area	Location
1	Business Process Excellence Manager	Corporate	Schaan, FL
2	Global Process Manager	Communications	Schaan, FL
3	Business Process Excellence Manager	Quality Management	Schaan, FL
4	Business Process Excellence Senior Manager	Corporate	Schaan, FL
5	Business Process Excellence Expert	Corporate	Schaan, FL
6	Regional Process Manager	Customer Service	Plano, US
7	Regional Process Manager	Logistics	Kaufering, GER
8	Business Process Excellence Lead	Corporate	Schaan, FL
9	Global Process Manager	Repair	Schaan, FL
10	Global Process Manager	Repair	Schaan, FL

3.2 Artifact Implementation

We use the results of our interviews alongside those of our investigation of the state of the art to define the requirements of our technical artifact, i.e., the LLM-based chatbot `PRODIGY`. To develop `PRODIGY`, we rely on `Botpress`, a platform to build custom AI chatbots powered by GPT-based LLMs; for our prototype version of `PRODIGY`, we used GPT-3.5 Turbo, which we found provided satisfactory performance while also requiring less resources to generate an output.

A crucial aspect of `PRODIGY` is its reliance on the `BPMN Sketch Miner` tool [10]. The syntax for this tool is entirely text-based, human-readable and light in terms of token consumption, making it appropriate for our case study. Therefore, we use few-shot prompting to teach `PRODIGY` to provide an output that matches the format expected by `BPMN Sketch Miner`. This output serves as the input for the model generation and transformation pipeline of `BPMN Sketch Miner` [9]. Such a design choice enables users of `PRODIGY` to directly paste the AI outputs into the online tool and get their model visualized (see Fig. 1).

Furthermore, we have leveraged retrieval-augmented generation (RAG) [13] to embed Hilti’s documentation into `PRODIGY`. Such documentation included: process descriptions from Hilti’s internal documentation repository (anonymised); and information about Hilti’s process management, and how to model processes

² This is in stark contrast with a closely related work that does not carry out any user study [12], thereby preventing to fully capture the organizational context.

at Hilti (taken verbatim from the learning platform for Hilti’s process modellers). These procedures enabled us to instill some knowledge about Hilti’s processes in PRODIGY—a functionality that was heavily endorsed by our interviewees.

Demonstration: We recorded a [video](#) [1] showing the functionalities of PRODIGY.

3.3 Artifact Evaluation

We conducted a user study with our artifact and process modellers. Our aim was to answer evaluative questions on the quality of PRODIGY for Hilti.

First, the process modellers tested all functionalities of PRODIGY by creating custom prompts, with the intention of simulating their routine tasks. Their inputs and the corresponding AI-generated outputs are [fully observable](#) in our repository [1]. Then, we carried out semi-structured interviews during which the participants answered 28 questions. Among these, we ask to give an 1–5 rating to the statement “Using PRODIGY would make it easier for me to do process modeling tasks.” The [complete questionnaire](#) is provided in our repository [1].

3.4 Dissemination and Communication of the Results

To conclude our DSR process, we formalized our learnings and made them accessible to interested parties. We documented our observations, analyzed our findings, identified lessons learned, stated limitations, and recommended directions for future work. We shared our learnings within Hilti Group and the wider BPM community in academia and practice—some companies reached out to us and expressed their interest about the development process of PRODIGY.

4 Key Findings and Lessons Learned

We first summarise the major results of our user studies, and then outline our proposed “operating model” for our developed LLM-based chatbot, PRODIGY.

Confidentiality Statement: To protect the privacy of the participants to our user studies, we cannot reveal the full transcript of their interviews. However, we are able to answer questions about their generic viewpoint on some specific issues.

4.1 Preliminary Interviews: what do Hilti process modellers want?

These open interviews lasted for 60 minutes, and the results shed light on the pain-points and desiderata of our participants. We found that, before modelling a process, 60% spend between 15–60 minutes *searching* for documentation; and also 60% spend between 5–60 minutes to *review* such documentation. As a matter of fact, 90% state that it is “extremely important” that an LLM-based chatbot has access to Hilti’s documentation; however, we also found that, on a 1–10 rating (low to high) scale, the average usefulness of current Hilti’s documentation is 6.7—indicating helpfulness, but with huge margins for improvement. Nonetheless, with respect to AI-related concerns, some stated that “humans may misinterpret the AI’s outputs” or “AI may negatively impact collaboration with colleagues” or even about accountability (“the mindset that [the machine] does everything and we no longer have to worry about it is dangerous”).

TAKEAWAYS. After analysing all our responses, we identified two design objectives which we used as basis to develop our LLM-based chatbot, **PRODIGY**.

- The chatbot should *support process modellers* in creating BPMN models. In doing so, the chatbot should hint at the larger picture, i.e., emphasize and guide in purpose, usage, and value creation of the resulting models.
- The chatbot should be able to access and utilize existing documentation, and hence *be aware of organizational specifics*. Such knowledge should drive the formulation of the output, which will be tailored to the organization.

The name **PRODIGY** stands for “**PRO**cess **mo**Delling **GU**idance for **Y**ou”.

4.2 Evaluation: what do Hilti process modellers say about **PRODIGY**?

After letting our process modellers use **PRODIGY**, we collected their feedback via 90-minutes long semi-structured interviews; one participant to the preliminary interviews did not provide feedback since they were not available, so we obtained responses from nine employees. **The general opinion was positive.** Five of our participants asserted that they would use **PRODIGY** on a daily or weekly basis (i.e., whenever they have to carry out process-modeling duties). Moreover, six participants asserted that **PRODIGY would speed-up** their tasks (three remained neutral), and eight believe that **PRODIGY makes their tasks easier** (one remained neutral). Finally, we report in Fig. 3 the participants’ perception on the functionalities we integrated in **PRODIGY**, showing great appreciation.




Fig. 3: **Helpfulness of **PRODIGY**'s functionalities.** According to our participants, most of our implemented features are helpful—especially for *creating process models* and *supporting human requests*.

4.3 Operating model: how should **PRODIGY** be used in practice?

During our implementation, we devised an *operating model* that describes how **PRODIGY** should be leveraged by real organizations; we have further refined our model (shown in Fig. 4) after receiving the feedback by our interviewees.

At a high-level, our model emphasizes continuous improvement through regular evaluations, feedback, and updates. To this end, our model delineates the interaction between a *Governance Team* (i.e., the set of employees within a company that oversee the development and maintenance of PRODIGY) and a *Process Modeller* (i.e., the end-users of PRODIGY). These two actors work collaboratively to ensure that the system performs well over time. For instance, the Process Modeller should be familiar with existing documentation and with the specific process, and scrutinize the response of PRODIGY accordingly; they should also be willing to provide feedback (collected in a dedicated repository) and receive guidance from the Governance Team—who must, in turn, define clear performance indicators for PRODIGY and periodically review the performance of prodigy (e.g., by analysing logs [19]) and apply updates if needed; as well as ensure that existing documentation is properly embedded in PRODIGY (in a timely manner).

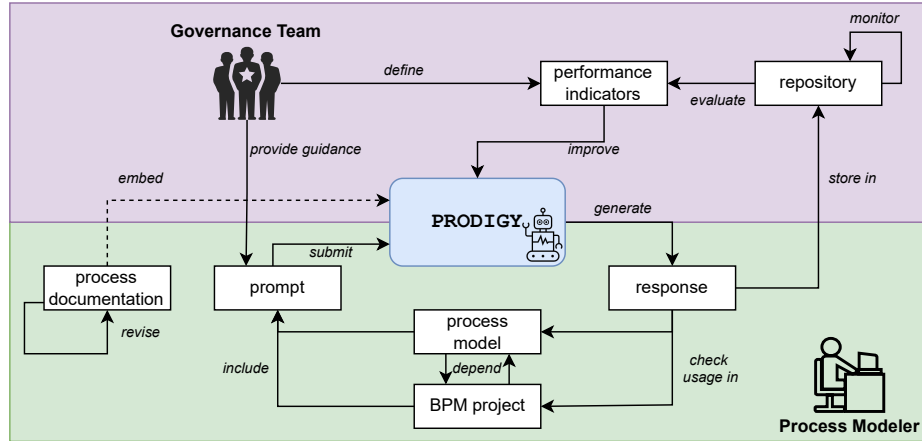


Fig. 4: **Operating model of PRODIGY** – We visualize the interactions between the *Governance Team* (e.g., developers and managers) of a given organization with the *Process Modeller* (i.e., the end-users of PRODIGY) that ensure a smooth operation of PRODIGY for real-world deployments.

5 Significance and Relevance in Research and Practice

Besides our key findings we underscore three orthogonal aspects of our research.

The perspective of process modellers in organizations. We coalesce the responses—not pertaining to AI—of our preliminary interviews, and derive an *original framework* representing dynamics of process modellers’ issues at Hilti. This is instructive because, during our literature analysis, we found some works mentioning pitfalls of process modeling (e.g., [16]) but without accounting for context. Our framework (displayed in Fig. 5, and described in the caption of Fig. 5) attempts to rectify this shortcoming, providing guidance for future work.

Evaluating LLM-/AI-based solutions. Upon further analysing the results of our evaluation interview, we have found that the reception of PRODIGY by our process modellers was highly dependant on their expectations and overall attitude towards AI and IT innovation. Indeed, some participants had “lower expectations” and provided prompts that were “more aligned” to the expected

input of PRODIGY—and these participants rated PRODIGY more positively. In contrast, participants who were expecting that PRODIGY would “do their work for them” by issuing a single (and typically poorly phrased and/or unclear) prompt were more skeptical of PRODIGY’s helpfulness. These results underscore the importance of (i) accounting for each end-user’s expectations while evaluating the performance of operational AI-based solutions; as well as (ii) educating end-users on the potential (and limitations) of AI-based tools.

The role played by higher education. With this paper, we (also) seek to bridge three domains: industrial practice, scholarly literature, and higher-education institutions [17]. This research has been predominantly carried by Clara Ziche for her MSc. thesis, during which she was working part-time at Hilti. The development of PRODIGY was driven by following the guidelines of prior academic literature, and the resulting artifact was appreciated by Hilti as well as by other companies that witnessed its capabilities. On this note, we find it instructive to trace the timeline of this research by outlining the path followed by Clara Ziche to bring our findings to light. In Sept–Dec 2023, after attending the BPM’23 conference, Clara investigated the state of the art and designed the interviews for the requirement analysis. In Jan 2024, Clara carried out the interviews, and began familiarizing with current LLM technologies. In Feb 2024, Clara developed PRODIGY and designed the questionnaire for its evaluation—which took place in March 2024. Disseminations occurred in Apr–May 2024.

Technical feasibility. After having collected the input from experts, the development of PRODIGY took only two weeks from a MSc. student in Information Systems.

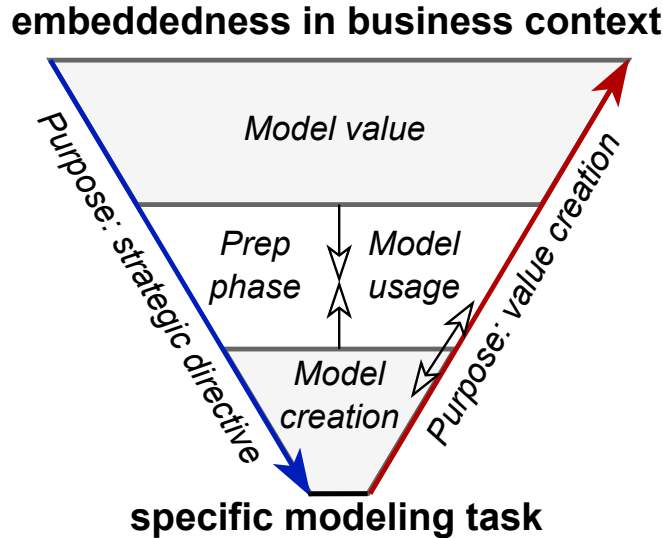


Fig. 5: **Issues of process modellers.** Our framework has three levels with two-way transitions between each level—each having its own set of issues. *Model value*→*Model creation*: the organization must communicate a clear strategic directive for process model creation, balancing the cost/benefit of model creation and usage [issue: during the “Prep phase”, process modellers find existing communication to lack clarity, leading to time waste and unproductive discussions among various stakeholders]. *Model creation*→*Model value*: While creating the model, process modelers should have a clear vision of “who and how” is going to use the model (which is what leads to the model becoming valuable) [issue: lack of clarity and/or poor documentation may lead to process models representing “standalone exercises” which do not bring any value to the company.]

6 Discussion: Scope and Limitations

We showcased an exemplary application of an LLM-based chatbot that can assist process modellers in a large organization, Hilti. In doing so, we have carried out a twofold user study with 10 employees of Hilti, and developed an original artifact, *PRODIGY*. Our research has a number of limitations. For instance, we do not claim that our findings can apply to other organizations—irrespective of their similarity to Hilti. Moreover, we cannot claim that even our own findings can apply to the entirety of Hilti: The participants of our user study are mostly based in Liechtenstein, and therefore cover the global headquarters perspective rather than regional and local perspectives. Furthermore, *PRODIGY* uses GPT-3.5 Turbo (which is not privacy-compliant), and it relies on BPMN Sketch Miner: if such a tool is taken down, the output of *PRODIGY* may lose its immediate usefulness. Finally, even our own participants have pointed out some shortcomings of *PRODIGY*, such as a poor “knowledge” of Hilti’s documentation. Such a result, however, was expected: the documents that *PRODIGY* has access to (with RAG) are just a drop in the deluge of files and logs included in Hilti’s databases (and we, as researchers, do not have complete access to such data).

7 Conclusions

We have presented the first case study showcasing how LLM can be used for process modeling in large enterprises—specifically, Hilti Group. We follow DSR guidelines and develop an original LLM-based chatbot, *PRODIGY*, which we test with professional process modellers from Hilti. Our findings revealed that end-users appreciate the functionalities of *PRODIGY*. However, concerns were raised about the poor alignment of *PRODIGY*’s output with Hilti’s specifications. Such a shortcoming underscores the importance of integrating LLM-based solutions with the organization’s documentation—which is a task outside the responsibilities of process modellers. Hence, deployment of similar solutions in real contexts should be done with the support of the organization’s governance team: it is unrealistic to expect that “off the shelf” solutions work properly to drive the process modeling routines of complex and large organizations (see Fig. 6).

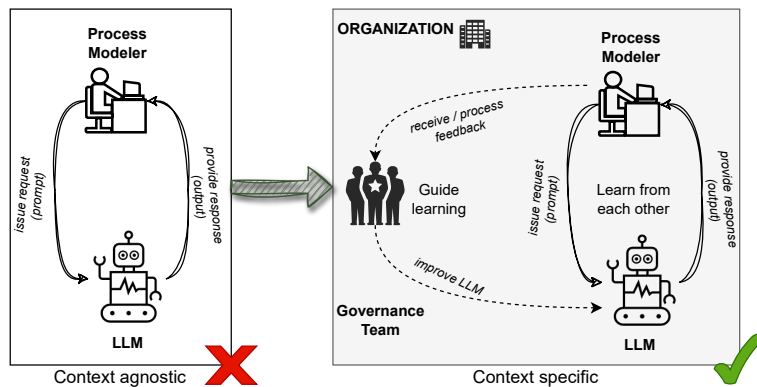


Fig. 6: **Takeaway.** LLMs are hardly usable for process modeling in a context-agnostic setting [left]. Deployment of LLM in organizations for process modeling should follow a context-specific approach, in which the governance team ensures that LLM and end-users “learn from each other” [right].

Acknowledgements. We thank Hilti for enabling and funding this research; and the participants to our user study for their contributions and feedback.

References

1. Our repository. https://github.com/Nouronihar/BPM24_LLM4PM
2. Bellan, P., Dragoni, M., Ghidini, C.: Extracting business process entities and relations from text using pre-trained language models and in-context learning. In: International Conference on Enterprise Design, Operations, and Computing (2022)
3. Dumas, M., La Rosa, M., Mendling, J., Reijers, H.A., et al.: Fundamentals of business process management. Springer (2018)
4. Fill, H.G., Fettke, P., Köpke, J.: Conceptual modeling and large language models: impressions from first experiments with chatgpt. EMISAJ (2023)
5. Franch, X., Palomares, C., Quer, C., Chatzipetrou, P., Gorschek, T.: The state-of-practice in requirements specification: an extended interview study at 12 companies. Requirements Engineering (2023)
6. Grohs, M., Abb, L., Elsayed, N., Rehse, J.R.: Large language models can accomplish business process management tasks. In: Int. Conf. BPM (2023)
7. Herbst, J., Karagiannis, D.: An inductive approach to the acquisition and adaptation of workflow models. In: IJCAI (1999)
8. Hevner, A.R., March, S.T., Park, J., Ram, S.: Design science in information systems research. MIS quarterly (2004)
9. Ivanchikj, A., Serbout, S., Pautasso, C.: From text to visual bpmn process models: Design and evaluation. In: ACM/IEEE MoDELS (2020)
10. Ivanchikj, A., Serbout, S., Pautasso, C.: Live process modeling with the bpmn sketch miner. Software and systems modeling (2022)
11. Klievtsova, N., Benzin, J.V., Kampik, T., Mangler, J., Rinderle-Ma, S.: Conversational process modelling: state of the art, applications, and implications in practice. In: International Conference on Business Process Management (2023)
12. Kourani, H., Berti, A., Schuster, D., van der Aalst, W.M.: Process modeling with large language models. arXiv:2403.07541 (2024)
13. Lewis, P., et al.: Retrieval-augmented generation for knowledge-intensive nlp tasks. Advances in Neural Information Processing Systems (2020)
14. Peffers, K., Tuunanen, T., Rothenberger, M.A., Chatterjee, S.: A design science research methodology for information systems research. JMIS (2007)
15. Plattfaut, R., Borghoff, V., Godefroid, M., Koch, J., Trampler, M., Coners, A.: The critical success factors for robotic process automation. Comp. Ind. (2022)
16. Rosemann, M.: Potential pitfalls of process modeling: part a. BPM Journal (2006)
17. Senkus, P., Berniak-Wozny, J., Gabryelczyk, R., Napieraj, A., Podobińska-Staniec, M., Sliż, P., Szelkagowski, M.: Bridging the gap: An evaluation of business process management education and industry expectations—the case of poland. In: International Conference on Business Process Management (2023)
18. Sola, D., van der Aa, H., Meilicke, C., Stuckenschmidt, H.: Activity recommendation for business process modeling with pre-trained language models. In: European Semantic Web Conference (2023)
19. Stein Dani, V., Leopold, H., van der Werf, J.M.E., Lu, X., Beerepoot, I., Koorn, J.J., Reijers, H.A.: Towards understanding the role of the human in event log extraction. In: International Conference on Business Process Management (2021)
20. Vom Brocke, J., Simons, A., Riemer, K., Niehaves, B., Plattfaut, R., Clevn, A.: Standing on the shoulders of giants: Challenges and recommendations of literature search in information systems research. CAIS (2015)