





Bridging Domain Knowledge and Process Discovery Using Large Language Models^{*}

Ali Norouzifar¹, Humam Kourani^{2,1}, Marcus Dees³, and Wil van der Aalst¹

¹ RWTH University, Aachen, Germany

{ali.norouzifar, wvdaalst}@pads.rwth-aachen.de

² Fraunhofer FIT, Sankt Augustin, Germany

humam.kourani@fit.fraunhofer.de

³ UWV Employee Insurance Agency, Amsterdam, Netherlands

marcus.dees@uwv.nl

Abstract. Discovering good process models is essential for different process analysis tasks such as conformance checking and process improvements. Automated process discovery methods often overlook valuable domain knowledge. This knowledge, including insights from domain experts and detailed process documentation, remains largely untapped during process discovery. This paper leverages Large Language Models (LLMs) to integrate such knowledge directly into process discovery. We use rules derived from LLMs to guide model construction, ensuring alignment with both domain knowledge and actual process executions. By integrating LLMs, we create a bridge between process knowledge expressed in natural language and the discovery of robust process models, advancing process discovery methodologies significantly. To showcase the usability of our framework, we conducted a case study with the UWV employee insurance agency, demonstrating its practical benefits and effectiveness.

Keywords: Process Mining · Process Discovery · Process Knowledge · Large Language Models.

1 Introduction

Recorded event data within information systems provides a rich source of information for process mining applications, enabling organizations to gain insights and improve their operational processes. In the field of process mining, various automated techniques are utilized to discover descriptive models that explain process executions. Despite the development of numerous methodologies for process discovery, the task remains inherently complex and challenging [3]. Discovering process models that do not align with domain knowledge presents significant challenges, particularly when these models are intended for conformance checking and process improvement.

^{*} This research was supported by the research training group “Dataninja” (Trustworthy AI for Seamless Problem Solving: Next Generation Intelligence Joins Robust Data Analysis) funded by the German federal state of North Rhine-Westphalia.




Fig. 1: Our proposed framework to integrate process knowledge in the IMr framework employing LLMs.

In addition to the extracted event data from information systems, we often have access to domain experts, process documentation, and other resources collectively referred to as *domain knowledge*, which cannot be directly used for process discovery. These valuable resources typically remain untapped when aiming to discover process models. Incorporating domain knowledge into the discovery of process models poses several challenges. For instance, domain experts usually have a thorough understanding of their processes, but they can only explain them in natural language. Furthermore, textual process documents, although potentially rich in detail, also pose integration difficulties. In our paper, we address these challenges by enabling the direct involvement of such information in process discovery through the use of Large Language Models (LLMs). LLMs have demonstrated the ability to handle user conversations and comprehend human reasoning effectively.

Our framework builds upon the IMr framework proposed in [12]. IMr is an inductive mining-based framework that recursively selects the process structure that best explains the actual process. Within this framework, the algorithm encounters various possibilities for constructing the process structure. To guide this selection, rules are introduced as inputs to prune the search space and eliminate potentially suboptimal process structures. Although the concept of rules is broad, the Declare rule specification language is proposed as an example [11]. Declarative rules are advantageous due to their similarity to human reasoning and logic, supported by extensive literature. They are based on logical statements and have specific semantics, though it is unrealistic to expect users to provide these rules directly.

Our proposed framework, illustrated in Fig. 1, leverages LLMs and prompt engineering to integrate domain knowledge into process discovery. Starting with an event log, it employs process knowledge in various forms. LLMs play a crucial role by translating textual inputs into declarative rules, which IMr then integrates. This framework allows for the use of textual process descriptions prior to initiating process discovery, enables domain experts to provide feedback on the discovered models, and facilitates interactive conversations with domain experts to gather information and improve the models.

2 Related Work

In traditional process discovery, event data are often used as the primary source of information to create process models [3]. However, additional information resources, such as various forms of process knowledge, can significantly enhance the quality of the discovered models [14]. When available, this supplementary knowledge can be utilized before discovery to filter the event log [5], during the discovery phase to influence the process model structure [12], or within an interactive framework [4,15]. Despite these benefits, the direct involvement of process experts is often limited due to the complexities involved in integrating their knowledge into process discovery.

In [12], declarative rules are used as an additional input for process discovery, which can be provided by the user or generated by automated methods. However, expecting users to be proficient in declarative rule specification language is not always feasible. The proposed method in [4] requires users to engage at a low level to position transitions and places based on guiding visualizations. The approach in [15] begins with an initial model discovered from a user-selected subset of variants and incrementally allows adding more variants to update the process model. Some research focuses on repairing process models after discovery, primarily to improve the correspondence between the process models and event logs, rather than incorporating process knowledge [13]. In contrast, our paper aims to minimize the effort required from domain experts by using natural language conversations to influence process discovery.

Translating the natural language to process models using natural language processing is investigated in [1]. Anomaly detection is examined in [2] by focusing on semantic inconsistencies in event labels within event logs, utilizing natural language processing to identify anomalous behavior. Recently, LLMs have been employed for various process mining tasks. The opportunities, strategies, and challenges of using LLMs for process mining and business process management are discussed in [16]. Additionally, several studies propose the extraction of process models directly from textual inputs [6,8,9]. Unlike these approaches, our method maintains the event log as the main source of information while incorporating textual process knowledge into the discovery process.

3 Background

The blue box in Fig. 1 highlights one recursion of the IMr framework [12]. Each recursion extracts a Directly Follows Graph (DFG) from the event log, representing the set of activities Σ and their direct succession. The algorithm searches for all binary cuts that divide Σ into two disjoint sets considering a structure specification type, i.e., sequence, exclusive choice, concurrent, or loop type. IMr filters out candidate cuts that may violate any rule $r \in R$, where R is the set of rules given by the user or discovered using automated methods. While [12] incorporates declarative constraints listed in Table. 1, the framework is flexible to support other rule specification languages. Cost functions evaluate the quality of candidate cuts, based on counting the number of deviating edges and estimating

Table 1: Declarative templates supported by IMr [12].

Declarative Template	Description
$at-most(a)$	a occurs at most once.
$existence(a)$	a occurs at least once.
$response(a, b)$	If a occurs, then b occurs after a .
$precedence(a, b)$	b occurs only if preceded by a .
$co-existence(a, b)$	a and b occur together.
$not-co-existence(a, b)$	a and b never occur together.
$not-succession(a, b)$	b cannot occur after a .
$responded-existence(a, b)$	If a occurs in the trace, then b occurs as well.

the number of missing edges considering parameter $sup \in [0, 1]$. In each recursion, the algorithm selects the cut with the minimum cost, splits the event log accordingly, and recursively processes each sub-log until only base cases remain.

4 Motivating Example

To motivate the research question addressed in this paper, consider the following event log extracted from a synthetic process $L = [\langle A\text{-created, A-canceled} \rangle^{300}, \langle A\text{-created, Doc-checked, Hist-checked, A-accepted} \rangle^{200}, \langle A\text{-created, Hist-checked, Doc-checked, A-accepted} \rangle^{50}, \langle A\text{-created, Doc-checked, Hist-checked, A-rejected} \rangle^{300}, \langle A\text{-created, Hist-checked, Doc-checked, A-rejected} \rangle^{80}, \langle A\text{-created, A-canceled, A-accepted} \rangle^{20}, \langle A\text{-created, A-canceled, A-rejected} \rangle^{15}, \langle A\text{-created, Doc-checked, Hist-checked, A-rejected, A-accepted} \rangle^{35}]$, where A stands for application, Doc for documents, and Hist for history.


Figure 2a illustrates the process model discovered using the IMf algorithm as a state-of-the-art process discovery technique [10]. The IMr framework with parameter $sup = 0.2$ and utilizing the Declare Miner [11] with $confidence = 1$ discovers the same process model. Consider that in addition to the provided event log, we have some additional process knowledge that helps us verify this model and pinpoint the possible unexpected behavior represented in the process model. In this paper, ChatGPT refers to ChatGPT-4o. We provided a text as feedback on this discovered model and asked ChatGPT to translate natural language feedback into understandable rules for the IMr framework. Here is our written feedback:

The discovered process does not fully adhere to our intuitions. Specifically, if a claim is canceled, the application cannot be either rejected or accepted. Furthermore, a claim cannot be both rejected and accepted for a single individual. Additionally, the history is always checked after the documents have been reviewed.

The following declarative rules, as explained in this paper, were extracted by ChatGPT:

not-co-existence(A-cancelled,A-accepted), not-co-existence(A-cancelled,A-rejected),
not-co-existence(A-accepted,A-rejected), response(Doc-checked,Hist-checked)

Figure 2b presents the process model discovered using our proposed pipeline. In this approach, we utilized ChatGPT to interpret the textual feedback and generate declarative constraints, which are then used as input for the IMr framework.



(a) Discovered model with deviations from the process knowledge. This model is discovered using IMf with $f=0.2$ and IMr with $sup=0.2$ and rules discovered employing Declare Miner [11] with $confidence=1$.



(b) The desired process model considering both the event log and process knowledge. IMr with $sup=0.2$ and rules extracted from process description employing ChatGPT discover this model.

Fig. 2: Discovered models from the motivating example event log using different techniques.

5 Domain-Enhanced Process Discovery with LLMs

In this section, we present our framework that leverages LLMs to integrate domain knowledge into the process discovery task. Figure 3 illustrates an overview of our proposed framework. The core idea is to utilize domain knowledge to generate a set of rules R which serves as input for the IMr framework. This can be done before starting the discovery by encoding process descriptions as rules, or after the process discovery by having a domain expert review the process model and provide feedback. Engaging in interactive conversations with LLMs in both scenarios helps address uncertainties and improve the quality of the extracted rules. An implementation of the framework is publicly available¹.




Fig. 3: Different components of the designed framework to bridge domain knowledge and process discovery using LLMs.

¹ <https://github.com/aliNorouzifar/IMr-LLM.git>

5.1 Task Definition

As outlined in [9], role promoting, knowledge injection, few-shot learning, and negative prompting techniques have significant potential to effectively prepare LLMs for specific process mining tasks. In our initial prompt, we define the role of the LLM as an interface between the domain expert and process discovery framework, such that LLM should encode the domain knowledge to declarative constraints as we need in IMr. Despite the similarity of declarative templates to human logic and reasoning, we observed the difficulties of LLMs in adhering to strict expectations. Therefore, we explain in our prompt the set of constraints we support, detailing both the syntax and the semantics of these constraints (cf. Table 1). We leverage the LLM’s ability to derive insights from examples by providing multiple pairs of textual process descriptions and their corresponding declarative constraints. Additionally, we include instructions to avoid common issues, such as syntactic mistakes, and extend our learning pairs to include examples of undesirable constraints. The detailed written prompt is available in our GitHub repository².

5.2 Rule Extraction

After introducing the task, the LLM is ready to receive textual input and produce output as declarative constraints. As illustrated in Fig. 3, domain experts can contribute in three distinct ways: providing business context, offering feedback after reviewing process models, and engaging in interactive conversations with the LLM. In the following sections, we explain these contributions in detail and discuss their respective roles.

Business Context The domain expert can introduce the actual business process to the LLM, providing a general overview, detailing the relationships between specific activities, or even including constraints written in natural language. This flexibility allows the domain expert to tailor the input based on their unique insights and the specifics of the process at hand. It is important to note that the LLM is unaware of specific activity labels used in the recorded event data. The list of activities can be automatically derived from the event log, ensuring that all relevant actions are accurately captured in the generated constraints. Alternatively, the domain expert can provide the list of activities and add context to guide the LLM in relating the process description with the activity labels, resulting in constraints that involve the correct activity labels.

Feedback integration After generating the initial process model, it is presented to the domain expert for review. The domain expert is expected to examine the process model for accuracy, completeness, and practical alignment with real-life scenarios. In case of finding errors in the represented model, the domain expert can provide a written feedback and explain the behaviors that do not make sense in the real process. The LLM then adjusts and refines the declarative constraints based on this feedback.

² <https://github.com/aliNorouzifar/IMr-LLM/blob/master/files/prompts.pdf>

Interactive Refinement In typical scenarios, LLMs tend to provide answers that appear confident and definitive, often without indicating any uncertainty [7]. We facilitate a more detailed understanding of the provided textual descriptions by encouraging the LLM to express uncertainty and address it by asking questions. This stage involves a dynamic dialogue between the LLM and the domain expert. Should it encounter gaps in its knowledge or find ambiguities in the process descriptions, the LLM is encouraged to formulate and pose relevant questions. These questions are directed towards the domain experts, who then provide responses. The quality and precision of the responses provided by domain experts play a significant role in enhancing the quality of the generated constraints.

5.3 Rule Validation

An essential step in the framework is checking the extracted declarative constraints from the LLM’s response. The LLM is instructed to encapsulate the constraints within specific tags in the response and to write them in a predefined language with no additional text or descriptions. Following extraction, the constraints undergo a validation process. This includes checking that the syntax of each constraint conforms to our predefined language, e.g., checking the type identifier and the number of activities specified within the constraint. Additionally, the labels of activities are verified against the activities recorded in the event log. If any errors are detected during validation, an error-handling loop is initiated. A new prompt specifies the problem and its location, prompting the LLM to adjust its output.

6 Case Study

A case study with the UWV employee insurance agency is conducted to demonstrate the usability of our approach in a real-life setting. UWV is responsible for managing unemployment and disability benefits in the Netherlands. For this case study, one of UWV’s claim-handling processes is selected. Figure 4 depicts the normative model of this process, which was developed in collaboration with process experts who have a thorough understanding of the workflow. The event log used in this study contains 144,046 cases, 16 unique activities, and 1,309,719 events. Our GitHub repository provides the full prompting history and more readable process models³.





Fig. 4: Normative model of the UWV claim handling process, extracted manually in collaboration with domain experts [12].


³ <https://github.com/aliNorouzifar/IMr-LLM/blob/master/files/prompts.pdf>

6.1 Process Discovery Without Including Process Knowledge


Our initial attempt to discover a process model using the IMf algorithm with $f = 0.2$ resulted in the model shown in Fig. 5a. When compared to the normative model, significant differences are observed, e.g., *Receive Claim* and *Start Claim* are the first mandatory steps but the process model allows for skipping them or for many other activities occurring before them. Fig. 5b illustrates the process model discovered using the IMr algorithm with $sup = 0.2$ and an empty set of input rules. Although this model shows more structural similarities to the normative model, it still contains some nonsensical differences. For instance, *Block Claim 1* should only be relevant if the claim is planned to be accepted, but this model permits it for rejected cases as well. Similarly, *Receive Objection 2* should only occur if the claim is rejected, yet the model allows it for accepted cases as well.




(a) Discovered model with using IMf with $f = 0.2$.



(b) Discovered model using IMr with $sup = 0.2$ without any rules.



(c) Discovered model using IMr with $sup = 0.2$ and the rules provided by ChatGPT.



(d) Discovered model using IMr with $sup = 0.2$ and the rules provided by ChatGPT after incorporating the domain expert feedback.

Fig. 5: Discovered models from UWV event log using different strategies.

6.2 Employing ChatGPT to Extract the Rules

We experimented with Gemini and various versions of ChatGPT to translate the process knowledge into declarative rules. ChatGPT-4o provided the best constraints and demonstrated a superior understanding of the task. By incorporating rules extracted by ChatGPT into the IMr framework, we obtained the process model shown in Fig. 5c. After iterating with feedback from domain experts, the model is further refined and improved, resulting in the process model depicted in Fig. 5d.

The process knowledge for this case study is incorporated as a set of activity labels and a process description provided by a domain expert. The blue boxes indicate the prompts given by domain experts, while the red boxes display the responses generated by ChatGPT.

Our process is designed to handle claim applications within an insurance agency, where applicants request various types of benefits, and our system processes these applications through a structured workflow. There are three types of blocking mechanisms: Block Claim 1, Block Claim 2, and Block Claim 3, each with corresponding resolutions: Unblock Claim 1, Unblock Claim 2, and Unblock Claim 3. The process begins once a claim is received. Block Claim 1 occurs when initial issues are detected, leading to necessary corrections before the claim is unblocked and accepted, resulting in the client receiving between one to three payments. Block Claim 2 is applied when the claim is deemed ineligible for receiving any payments, resulting in rejection and potentially an objection from the client. Block Claim 3 is triggered if there is missing information, such as incorrect bank account details, after a claim is accepted; the client is notified to provide the correct details, and the claim is unblocked upon receipt of this information. Additionally, clients may file an objection after receiving payments, leading to the withdrawal of the claim and repayment of the received money to the agency.

Following this process description, we encouraged ChatGPT to ask clarifying questions whenever it encountered uncertainties about the declarative rules. The box below provides an example of questions posed by ChatGPT.

Q1: Can "Block Claim 1," "Block Claim 2," and "Block Claim 3" occur in the same process instance, or are they mutually exclusive? **Q2:** Are the corresponding unblocking activities ("Unblock Claim 1," "Unblock Claim 2," and "Unblock Claim 3") strictly linked to their respective blocking activities, or can there be other sequences? **Q3:** Is there any restriction on the order of "Receive Objection 1" and "Receive Objection 2," or can they occur independently of each other? ...

The provided questions demonstrate the LLM's correct approach to understanding the declarative constraints between the activities in the process. For instance, the first question indicates that ChatGPT is attempting to discern whether there is a co-existence or non-co-existence relationship between block claim types. Question 3 is also pertinent, as the process description does not clearly distinguish between Receive Objection 1 and Receive Objection 2. The domain expert provided the following answers to these questions:

A1: Some cases may have both Block Claim 1 and Block Claim 3. But Block Claim 2 can never occur in the trace with Block Claim 1 or Block Claim 3. **A2:** Each blocking mechanism (Block Claim 1, Block Claim 2, Block Claim 3) has a corresponding unblocking mechanism (Unblock Claim 1, Unblock Claim 2, Unblock Claim 3), and each pair (Block, Unblock) occurs in a strict sequence. **A3:** Receive Objection 2 only occurs in the rejected cases and Receive Objection 1 only occurs in the accepted

cases after they receive some payments. In general, Receive Objection 1 and Receive Objection 2 cannot occur together. ...

The set of declarative constraints extracted after answering the questions is:

not-co-existence(Block Claim 2, Block Claim 1), not-co-existence(Block Claim 2, Block Claim 3), co-existence(Block Claim 1, Unblock Claim 1), co-existence(Block Claim 2, Unblock Claim 2), co-existence(Block Claim 3, Unblock Claim 3), precedence(Block Claim 1, Unblock Claim 1), precedence(Block Claim 2, Unblock Claim 2), precedence(Block Claim 3, Unblock Claim 3), not-co-existence(Receive Objection 1, Receive Objection 2), precedence(Reject Claim, Receive Objection 2), precedence(Payment Order, Receive Objection 1), at-most(Correct Claim), precedence(Block Claim 1, Correct Claim), precedence(Correct Claim, Unblock Claim 1), response(Withdraw Claim, Repayment), responded-existence(Accept Claim, Payment Order), responded-existence(Payment Order, Execute Payment)

These rules are validated using a predefined program to ensure that the activity labels are consistent with the event log and that the declarative constraints are free of syntax errors. Then, the rules are used as input for the IMr framework in addition to the event log, and the process model represented in Fig. 5c is discovered. These rules are aligned with the process description and the follow-up questions and answers. For example, the answer to the first question (A1) led to the extraction of *not-co-existence(Block Claim 2, Block Claim 1)* and *not-co-existence(Block Claim 2, Block Claim 3)*, correctly illustrating the relationship between these activities. These rules help IMr avoid the incorrect positioning of *Block Claim 1* observed in Fig. 5b. Another improvement is achieved by considering *not-co-existence(Receive Objection 1, Receive Objection 2)*, which prevents *Receive Objection 1* and *Receive Objection 2* from occurring in the same trace. Fig. 5b allows for *Block Claim 3* without the existence of *Unblock Claim 3*. The rule *co-existence(Block Claim 3, Unblock Claim 3)* guides IMr to avoid placing *Unblock Claim 3* as the re-do part of a loop. We presented this process model to domain experts for feedback. They identified some potential issues, which are then provided to ChatGPT to generate a better set of declarative templates.

The discovered process model is interesting but we observe some issues. For example, Receive Objection 1 should occur after all the payments are executed. This time of objection can occur at most one time per claim. Withdraw Claim can not be followed by another payment. Usually, after the payments are executed, the applicant has the option to send an objection, withdraw the claim, and repay the received benefits. Withdraw Claim only occurs at most once per claim.

After the consideration of the domain expert input, these constraints are added by ChatGPT to the previous set of constraints:

precedence(Execute Payment, Receive Objection 1), at-most(Receive Objection 1), not-succession(Withdraw Claim, Payment Order), not-succession(Withdraw Claim, Execute Payment), at-most(Withdraw Claim)

The discovered model using the modified set of constraints is illustrated in Fig. 5d. While this model differs from the normative model in Fig. 4, it better represents the actual process compared to the models in Fig. 5a and Fig. 5b, which were discovered without considering process knowledge. In comparison to Fig. 5c some improvements are achieved considering the provided feedback. The constraint *at-most(Receive Objection 1)* prevents *Receive Objection 1* from being included in a loop, and *precedence(Execute Payment, Receive Objection 1)* ensures it is positioned after *Execute Payment*. Additionally, *Withdraw Claim* is no longer in a loop due to the *at-most(Withdraw Claim)* constraint and is correctly positioned after *Payment Order* and *Withdraw Claim* because of the constraints *not-succession(Withdraw Claim, Payment Order)* and *not-succession(Withdraw Claim, Execute Payment)*.

7 Conclusion

The integration of process knowledge in the discovery of process models is often overlooked in the literature. In this paper, we leveraged advancements in large language models to demonstrate their capabilities in encoding textual domain knowledge into comprehensible rules for process discovery. Our proposed framework not only facilitates the integration of feedback from domain experts but also enables interactive improvement of process models. Through a comprehensive case study, we demonstrated the effectiveness of our framework in generating process models that better align with process knowledge. While the extracted set of declarative constraints from LLMs shows great promise, there is still room for improvement in precision and completeness. Future work focuses on expanding the range of declarative templates within the IMr framework and developing additional rule specification patterns. Additionally, providing more detailed examples in task definition steps helps LLMs capture a broader context, further enhancing the quality of the extracted constraints.

References

1. van der Aa, H., Ciccio, C.D., Leopold, H., Reijers, H.A.: Extracting declarative process models from natural language. In: Advanced Information Systems Engineering - 31st International Conference, CAiSE 2019, Proceedings. Lecture Notes in Computer Science, vol. 11483, pp. 365–382. Springer (2019)
2. van der Aa, H., Rebmann, A., Leopold, H.: Natural language-based detection of semantic execution anomalies in event logs. *Inf. Syst.* **102**, 101824 (2021)
3. Augusto, A., Conforti, R., Dumas, M., Rosa, M.L., Maggi, F.M., Marrella, A., Mecella, M., Soo, A.: Automated discovery of process models from event logs: Review and benchmark. *IEEE Trans. Knowl. Data Eng.* **31**(4), 686–705 (2019)

4. Dixit, P.M., Verbeek, H.M.W., Buijs, J.C.A.M., van der Aalst, W.M.P.: Interactive data-driven process model construction. In: *Conceptual Modeling - 37th International Conference, ER 2018, Proceedings*. Lecture Notes in Computer Science, vol. 11157, pp. 251–265. Springer (2018)
5. van Eck, M.L., Lu, X., Leemans, S.J.J., van der Aalst, W.M.P.: PM²: A process mining project methodology. In: *Advanced Information Systems Engineering - 27th International Conference, CAiSE 2015, Proceedings*. Lecture Notes in Computer Science, vol. 9097, pp. 297–313. Springer (2015)
6. Grohs, M., Abb, L., Elsayed, N., Rehse, J.: Large language models can accomplish business process management tasks. In: *Business Process Management Workshops - BPM 2023 International Workshops, 2023, Revised Selected Papers*. Lecture Notes in Business Information Processing, vol. 492, pp. 453–465. Springer (2023)
7. Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., Liu, T.: A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *CoRR* **abs/2311.05232** (2023)
8. Klietsova, N., Benzin, J., Kampik, T., Mangler, J., Rinderle-Ma, S.: Conversational process modelling: State of the art, applications, and implications in practice. In: *Business Process Management Forum - BPM 2023 Forum, Proceedings*. Lecture Notes in Business Information Processing, vol. 490, pp. 319–336. Springer (2023)
9. Kourani, H., Berti, A., Schuster, D., van der Aalst, W.M.P.: Process modeling with large language models. In: *Enterprise, Business-Process and Information Systems Modeling - 25th International Conference, BPMDS 2024, and 29th International Conference, EMMSAD 2024, Proceedings*. Lecture Notes in Business Information Processing, vol. 511, pp. 229–244. Springer (2024)
10. Leemans, S.J.J., Fahland, D., van der Aalst, W.M.P.: Discovering block-structured process models from event logs containing infrequent behaviour. In: *Business Process Management Workshops - BPM 2013 International Workshops , Revised Papers*. Lecture Notes in Business Information Processing, vol. 171, pp. 66–78. Springer, Cham (2013)
11. Maggi, F.M., Bose, R.P.J.C., van der Aalst, W.M.P.: Efficient discovery of understandable declarative process models from event logs. In: *Advanced Information Systems Engineering - 24th International Conference, CAiSE 2012, Proceedings*. Lecture Notes in Computer Science, vol. 7328, pp. 270–285. Springer (2012)
12. Norouzifar, A., Dees, M., van der Aalst, W.M.P.: Imposing rules in process discovery: An inductive mining approach. In: *Research Challenges in Information Science - 18th International Conference, RCIS 2024, Proceedings, Part I*. Lecture Notes in Business Information Processing, vol. 513, pp. 220–236. Springer (2024)
13. Polyvyanyy, A., van der Aalst, W.M.P., ter Hofstede, A.H.M., Wynn, M.T.: Impact-driven process model repair. *ACM Trans. Softw. Eng. Methodol.* **25**(4), 28:1–28:60 (2017)
14. Schuster, D., van Zelst, S.J., van der Aalst, W.M.P.: Utilizing domain knowledge in data-driven process discovery: A literature review. *Comput. Ind.* **137**, 103612 (2022)
15. Schuster, D., van Zelst, S.J., van der Aalst, W.M.P.: Cortado: A dedicated process mining tool for interactive process discovery. *SoftwareX* **22**, 101373 (2023)
16. Vidgof, M., Bachhofner, S., Mendling, J.: Large language models for business process management: Opportunities and challenges. In: *Business Process Management Forum - BPM 2023 Forum, 2023, Proceedings*. Lecture Notes in Business Information Processing, vol. 490, pp. 107–123. Springer (2023)