

# ProMoAI: Process Modeling with Generative AI

Humam Kourani<sup>1,2</sup>, Alessandro Berti<sup>1,2</sup>, Daniel Schuster<sup>1,2</sup>, Wil M. P. van der Aalst<sup>1,2</sup>

<sup>1</sup>Fraunhofer Institute for Applied Information Technology FIT, Sankt Augustin, Germany

<sup>2</sup>RWTH Aachen University, Aachen, Germany

{humam.kourani, alessandro.berti, daniel.schuster, wil.van.der.aalst}@fit.fraunhofer.de

## Abstract

ProMoAI is a novel tool that leverages Large Language Models (LLMs) to automatically generate process models from textual descriptions, incorporating advanced prompt engineering, error handling, and code generation techniques. Beyond automating the generation of complex process models, ProMoAI also supports process model optimization. Users can interact with the tool by providing feedback on the generated model, which is then used for refining the process model. ProMoAI utilizes the capabilities LLMs to offer a novel, AI-driven approach to process modeling, significantly reducing the barrier to entry for users without deep technical knowledge in process modeling.

## 1 Introduction

Process modeling is a fundamental task in analyzing and optimizing organizational workflows. Traditional process modeling approaches require significant effort to capture workflow details and deep expertise in process modeling languages [Indulska *et al.*, 2009]. Such approaches often require specialized knowledge, making them inaccessible to non-experts.

Large Language Models (LLMs) [Feuerriegel *et al.*, 2024] have shown advanced capabilities in a wide array of applications, ranging from content generation to complex problem-solving. By leveraging vast amounts of data, LLMs have demonstrated remarkable advancements in understanding and generating human-like text [Li *et al.*, 2021; Zhou *et al.*, 2023]. Their ability to interpret natural language and generate coherent, contextually relevant responses makes them particularly suited for tasks that involve understanding complex textual descriptions and translating them into structured outputs.

ProMoAI introduces an innovative approach to process modeling by utilizing the capabilities of state-of-the-art LLMs, such as GPT-4 [OpenAI, 2023], and employing advanced prompt engineering, error handling, and code generation techniques. The primary contribution of our tool is the automation of process model generation. Furthermore, a feedback loop allows users to interactively refine the generated process models, ensuring they accurately reflect the intended processes. ProMoAI is available under



## Process Model Generator with Generative AI

Enter the OpenAI model name: gpt-4

Enter your OpenAI API key: .....

Enter the process description:

Consider a process for purchasing items from an online shop. The user starts an order by logging in to their account. Then, the user simultaneously selects the items to purchase and sets a payment method. Afterward, the user either pays or completes an installment agreement. After selecting the items, the user chooses between multiple options for a free

Run

Model generated successfully!

Feedback:

the user may skip the reward selection

Update Model

Export Model

Download BPMN

Download PNML

Feedback History

[1] model the item selection using an activity Add Items that can be repeated

[2] the user may skip the reward selection

Select a view:

Petri net

View Image

Figure 1: A screenshot of ProMoAI.

<https://promoai.streamlit.app/>, and a demonstration video is available under <https://youtu.be/t8fkx9rJmBE>.

Fig. 1 shows a screenshot of ProMoAI. The user initially sets the LLM<sup>1</sup>, enters the corresponding API key, and pro-

<sup>1</sup>Currently, only OpenAI LLMs (<https://openai.com/>) are supported. We are planning to support more LLMs in the future as our system's architecture is independent of the used LLM (cf. Fig. 2).

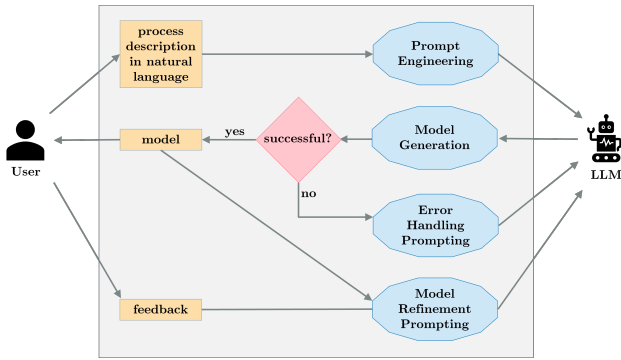


Figure 2: System architecture.

vides a textual description of the process. After the initial process model generation, the user can provide feedback, and the process model is updated accordingly. The generated process model can be viewed and exported in the standard modeling notation used in the fields of business process management and process mining (BPMN [von Rosing *et al.*, 2015] and PNML [Hillah *et al.*, 2010]).

## 2 Application Domain

ProMoAI allows for a more intuitive and efficient creation and optimization of process models, marking a novel application of LLMs in the field of process modeling. The tool can be employed in a wide range of domains including business process management, workflow automation, and systems engineering. It is particularly beneficial in environments where rapid prototyping of process models is required to visualize and optimize workflows, making it a valuable asset for business analysts, systems engineers, and process managers.

## 3 Related Work

Previous works have explored utilizing natural language processing and text mining for the generation of process models (e.g., [Friedrich *et al.*, 2011] and [de A. R. Gonçalves *et al.*, 2011]). An overview of different techniques for the extraction of processes from text is provided in [Bellan *et al.*, 2020].

Recently, LLMs have been evaluated on business process management (BPM) tasks. In [Busch *et al.*, 2023] and [Vidgof *et al.*, 2023], the authors discuss the potentials and challenges of using LLMs for BPM tasks. In [Berti *et al.*, 2023], LLMs are evaluated on various process mining tasks. In [Klietsova *et al.*, 2023], the authors discuss the potential of LLMs in generating process models through conversational modeling. In [Grohs *et al.*, 2023], the authors illustrated LLMs capabilities in translating textual statements into procedural and declarative process model constraints. In [Fill *et al.*, 2023], the authors explored potential applications of LLMs for conceptual modeling.

## 4 System Overview

This section provides an overview of our methodologies.

ProMoAI implements our framework for LLM-based process modeling [Kourani *et al.*, 2024]. The high-level architecture of the underlying system is illustrated in Fig. 2. The

user submits a textual description of a process in natural language. Our system creates a prompt that guides the LLM in generating Python code that utilizes a predefined set of functions we designed for the generation of process models (cf. Section 4.2). This prompt is then sent to the LLM, and based on the LLM’s answer, the system attempts to generate a process model within a secure execution environment, designed to safely execute the generated code while protecting against malicious inputs or unintended operations (cf. Section 4.3). If errors are encountered during this process, an error handling prompt is sent back to the LLM, utilizing self-refinement techniques to fix the code (cf. Section 4.4). Finally, the generated model is presented to the user using standard process modeling notations (BPMN and Petri nets), with an option for exporting the process model in these formats. Additionally, the user can provide feedback, which is integrated into the system to optimize and refine the process model in an iterative refinement loop (cf. Section 4.5).

### 4.1 Utilization of POWL

ProMoAI employs the Partially Ordered Workflow Language (POWL) [Kourani and van Zelst, 2023] for the generation of process models. POWL and similar hierarchical modeling languages are employed in the automated discovery of process models (e.g., in [Kourani *et al.*, 2023] and [Leemans, 2022]) due to the quality guarantees they inherently provide. For example, it is not possible to generate a POWL model with dead parts that can never be reached. POWL allows for modeling complex, non-hierarchical process constructs while preserving the quality guarantees of hierarchical process modeling languages [Kourani and van Zelst, 2023].

### 4.2 Prompt Engineering Techniques

ProMoAI uses sophisticated prompt engineering techniques to guide the LLM to understand the process descriptions accurately and generate the desired process models, without the extensive resource requirements and computational costs of fine-tuning [Li *et al.*, 2022; Thangarasa *et al.*, 2023]. First, we employ *role prompting* [Xu *et al.*, 2023] by establishing a clear role for the LLM as an expert in process modeling and as a process owner who is familiar with the process context at the same time. Second, we provide a comprehensive knowledge base about POWL, offering detailed insights into its hierarchical structure, the semantics of the different POWL components, and the way POWL models are generated. Recently, LLMs have demonstrated strong capabilities in solving programming tasks and generating executable code [Vidan and Fiedler, 2023]. ProMoAI leverages these capabilities by instructing the LLM to transform natural language descriptions into Python code that utilizes a predefined set of functions we designed for the safe generation of POWL models. Moreover, we implement *negative prompting* [Miyake *et al.*, 2023] by instructing the LLM to avoid common errors and mistakes. Finally, we apply *few-shot learning* [Brown *et al.*, 2020] by providing the LLM with concrete examples that illustrate how to generate POWL models starting from process descriptions (including descriptions from the PET data set [Bellan *et al.*, 2022]).




Figure 3: BPMN generated by ProMoAI for the order process using GPT-4. This model conforms to the reference model of the process, and it additionally captures the feedback requests for modeling the item selection as a loop and allowing for skipping the reward selection.

The prompt created for our example process in Section 6 can be found under [https://github.com/humam-kourani/ProMoAI/blob/main/example\\_online\\_shop/prompt\\_1.txt](https://github.com/humam-kourani/ProMoAI/blob/main/example_online_shop/prompt_1.txt).

### 4.3 Process Model Generation

To ensure a high level of security and reliability, we restrict the execution of the generated Python code to a controlled environment. We implement a validation process to confirm that the generated code strictly adheres to our defined validation rules and syntax, expressly prohibiting the use of any external libraries or constructs that could introduce security risks. This methodology ensures a secure execution environment for the generation of valid process models.

ProMoAI transforms the generated POWL models into Petri nets and BPMN models [Kourani and van Zelst, 2023]. The tool provides options for viewing and exporting the generated models in these notations, which are widely recognized in the field of business process management.

### 4.4 Error Handling

ProMoAI features an error-handling mechanism that utilizes the self-evaluation and refinement capabilities of LLMs through *iterative prompting* [Jha *et al.*, 2023]. If an error occurs during the execution of the validated Python code (e.g., due to a violation of our validation rules), the system produces an informative feedback message detailing the nature of the error. This message is then incorporated into the conversation history with the LLM, enabling the LLM to refine the generated code in response to the identified error.

### 4.5 Process Model Refinement

A key feature of ProMoAI is its model optimization mechanism, realized through an interactive refinement loop. Users can provide feedback on the generated process models, which the system uses to generate a refinement prompt. This prompt is integrated into the conversation history with the LLM, allowing for the iterative improvement of the process model to ensure it accurately meets the user’s requirements.

## 5 Future Work and Extensions

Recognizing the rapid advancement in LLM technology, ProMoAI is designed to be forward-compatible, enabling users to utilize newer OpenAI models as they become available. This flexibility ensures that ProMoAI remains at the cutting edge of AI-driven process modeling, capable of delivering improved results by leveraging the latest advancements. We also plan to support further LLMs beyond those offered by OpenAI (e.g., Google Gemini [Anil *et al.*, 2023]).

We are exploring the direct generation of BPMN models from textual process descriptions without the intermediate POWL representation. This will offer more flexibility in

representing complex process structures and will allow us to enrich the process models with informative, context-rich annotations. However, moving away from POWL’s structured guarantees requires the development of more advanced process model generation and validation techniques to maintain the quality of the generated process models.

## 6 Example Application

We applied ProMoAI using GPT-4 on the process used in [Kourani and van Zelst, 2023] to describe the user’s journey to order items through an online shop. This process contains typical process constructs (choice, loop, concurrency, sequence) as well as more complex non-hierarchical dependencies. After generating the initial process model, we applied two feedback comments to refine the model (cf. Fig. 1). The initial process description, the feedback comments, all generated prompts, and GPT-4’s responses are available under [https://github.com/humam-kourani/ProMoAI/tree/main/example\\_online\\_shop](https://github.com/humam-kourani/ProMoAI/tree/main/example_online_shop).

The BPMN view of the generated process model is shown in Fig. 3. The generated model conforms to the reference process model, and it additionally integrates the feedback requests correctly into the model. This process contains complex non-hierarchical dependencies between selecting the items, setting a payment method, the reward selection, and the payment choice. As discussed in [Kourani and van Zelst, 2023], traditional hierarchical process modeling notations such as process trees [Leemans, 2022] fail to model such structures. The use of POWL enables ProMoAI to detect such structures, offering a level of flexibility and precision that traditional hierarchical process modeling languages lack.

Rerunning the application on the same input may lead to different results, e.g., GPT-4 might place the reward selection after setting the payment method. With the continuous development and advances added to LLMs, we expect future LLMs to perform better in understanding and modeling complex dependencies, leading to more consistency in the results.

## 7 Conclusion

ProMoAI represents a significant advancement in the field of process modeling, offering a user-friendly, efficient, and interactive approach to generating and refining process models. The tool allows users to describe workflows in natural language, which are then automatically translated into formal process models. By utilizing the capabilities of large language models, ProMoAI simplifies the creation of complex process models, opening up new possibilities for optimizing organizational workflows.

## References

- [Anil *et al.*, 2023] Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Slav Petrov, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy P. Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul Ronald Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, Ryan Doherty, Eli Collins, Clemens Meyer, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, George Tucker, Enrique Piqueras, Maxim Krikun, Iain Barr, Nikolay Savinov, Ivo Danihelka, Becca Roelofs, Anaïs White, Anders Andreassen, Tamara von Glehn, Lakshman Yagati, Mehran Kazemi, Lucas Gonzalez, Misha Khalman, Jakub Sygnowski, and et al. Gemini: A family of highly capable multimodal models. *CoRR*, abs/2312.11805, 2023.
- [Bellan *et al.*, 2020] Patrizio Bellan, Mauro Dragoni, and Chiara Ghidini. A qualitative analysis of the state of the art in process extraction from text. In Giuseppe Vizzari, Matteo Palmonari, and Andrea Orlandini, editors, *Proceedings of the AIXIA 2020 Discussion Papers Workshop co-located with the 19th International Conference of the Italian Association for Artificial Intelligence (AIXIA2020), Anywhere, November 27th, 2020*, volume 2776 of *CEUR Workshop Proceedings*, pages 19–30. CEUR-WS.org, 2020.
- [Bellan *et al.*, 2022] Patrizio Bellan, Han van der Aa, Mauro Dragoni, Chiara Ghidini, and Simone Paolo Ponzetto. PET: an annotated dataset for process extraction from natural language text tasks. In Cristina Cabanillas, Niels Frederik Garmann-Johnsen, and Agnes Koschmider, editors, *Business Process Management Workshops - BPM 2022 International Workshops, Münster, Germany, September 11-16, 2022, Revised Selected Papers*, volume 460 of *Lecture Notes in Business Information Processing*, pages 315–321. Springer, 2022.
- [Berti *et al.*, 2023] Alessandro Berti, Daniel Schuster, and Wil M. P. van der Aalst. Abstractions, scenarios, and prompt definitions for process mining with llms: A case study. In Jochen De Weerd and Luise Pufahl, editors, *Business Process Management Workshops - BPM 2023 International Workshops, Utrecht, The Netherlands, September 11-15, 2023, Revised Selected Papers*, volume 492 of *Lecture Notes in Business Information Processing*, pages 427–439. Springer, 2023.
- [Brown *et al.*, 2020] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [Busch *et al.*, 2023] Kiran Busch, Alexander Rochlitzer, Diana Sola, and Henrik Leopold. Just tell me: Prompt engineering in business process management. In Han van der Aa, Dominik Bork, Henderik A. Proper, and Rainer Schmidt, editors, *Enterprise, Business-Process and Information Systems Modeling - 24th International Conference, BPMDS 2023, and 28th International Conference, EMMSAD 2023, Zaragoza, Spain, June 12-13, 2023, Proceedings*, volume 479 of *Lecture Notes in Business Information Processing*, pages 3–11. Springer, 2023.
- [de A. R. Gonçalves *et al.*, 2011] João Carlos de A. R. Gonçalves, Flávia Maria Santoro, and Fernanda Araújo Baião. Let me tell you a story - on how to build process models. *J. Univers. Comput. Sci.*, 17(2):276–295, 2011.
- [Feuerriegel *et al.*, 2024] Stefan Feuerriegel, Jochen Hartmann, Christian Janiesch, and Patrick Zschech. Generative AI. *Bus. Inf. Syst. Eng.*, 66(1):111–126, 2024.
- [Fill *et al.*, 2023] Hans-Georg Fill, Peter Fettke, and Julius Köpke. Conceptual modeling and large language models: Impressions from first experiments with chatgpt. *Enterp. Model. Inf. Syst. Archit. Int. J. Concept. Model.*, 18:3, 2023.
- [Friedrich *et al.*, 2011] Fabian Friedrich, Jan Mendling, and Frank Puhlmann. Process model generation from natural language text. In Haralambos Mouratidis and Colette Rolland, editors, *Advanced Information Systems Engineering - 23rd International Conference, CAiSE 2011, London, UK, June 20-24, 2011. Proceedings*, volume 6741 of *Lecture Notes in Computer Science*, pages 482–496. Springer, 2011.
- [Grohs *et al.*, 2023] Michael Grohs, Luka Abb, Nourhan Elsayed, and Jana-Rebecca Rehse. Large language models can accomplish business process management tasks. In Jochen De Weerd and Luise Pufahl, editors, *Business Process Management Workshops - BPM 2023 International Workshops, Utrecht, The Netherlands, September 11-15, 2023, Revised Selected Papers*, volume 492 of *Lecture Notes in Business Information Processing*, pages 453–465. Springer, 2023.
- [Hillah *et al.*, 2010] Lom-Messan Hillah, Fabrice Kordon, Laure Petrucci, and Nicolas Trèves. PNML framework: An extendable reference implementation of the petri net markup language. In Johan Lilius and Wojciech Penczek, editors, *Applications and Theory of Petri Nets, 31st International Conference, PETRI NETS 2010, Braga, Portugal, June 21-25, 2010. Proceedings*, volume 6128 of *Lecture Notes in Computer Science*, pages 318–327. Springer, 2010.

- [Indulska *et al.*, 2009] Marta Indulska, Jan Recker, Michael Rosemann, and Peter F. Green. Business process modeling: Current issues and future challenges. In Pascal van Eck, Jaap Gordijn, and Roel J. Wieringa, editors, *Advanced Information Systems Engineering, 21st International Conference, CAiSE 2009, Amsterdam, The Netherlands, June 8-12, 2009. Proceedings*, volume 5565 of *Lecture Notes in Computer Science*, pages 501–514. Springer, 2009.
- [Jha *et al.*, 2023] Susmit Jha, Sumit Kumar Jha, Patrick Lincoln, Nathaniel D. Bastian, Alvaro Velasquez, and Sandeep Neema. Dehallucinating large language models using formal methods guided iterative prompting. In *IEEE International Conference on Assured Autonomy, ICAA 2023, Laurel, MD, USA, June 6-8, 2023*, pages 149–152. IEEE, 2023.
- [Kliedtsova *et al.*, 2023] Nataliia Kliedtsova, Janik-Vasily Benzin, Timotheus Kampik, Juergen Mangler, and Stefanie Rinderle-Ma. Conversational process modelling: State of the art, applications, and implications in practice. In Chiara Di Francescomarino, Andrea Burattin, Christian Janiesch, and Shazia W. Sadiq, editors, *Business Process Management Forum - BPM 2023 Forum, Utrecht, The Netherlands, September 11-15, 2023, Proceedings*, volume 490 of *Lecture Notes in Business Information Processing*, pages 319–336. Springer, 2023.
- [Kourani and van Zelst, 2023] Humam Kourani and Sebastian J. van Zelst. POWL: partially ordered workflow language. In Chiara Di Francescomarino, Andrea Burattin, Christian Janiesch, and Shazia Sadiq, editors, *Business Process Management - 21st International Conference, BPM 2023, Utrecht, The Netherlands, September 11-15, 2023, Proceedings*, volume 14159 of *Lecture Notes in Computer Science*, pages 92–108. Springer, 2023.
- [Kourani *et al.*, 2023] Humam Kourani, Daniel Schuster, and Wil M. P. van der Aalst. Scalable discovery of partially ordered workflow models with formal guarantees. In *5th International Conference on Process Mining, ICPM 2023, Rome, Italy, October 23-27, 2023*, pages 89–96. IEEE, 2023.
- [Kourani *et al.*, 2024] Humam Kourani, Alessandro Berti, Daniel Schuster, and Wil M. P. van der Aalst. Process Modeling With Large Language Models. In *Proceedings of the 29th International Working Conference on Exploring Modeling Methods for Systems Analysis and Development (EMMSAD 2024)*. Springer, 2024. In press.
- [Leemans, 2022] Sander J. J. Leemans. *Robust Process Mining with Guarantees - Process Discovery, Conformance Checking and Enhancement*, volume 440 of *Lecture Notes in Business Information Processing*. Springer, 2022.
- [Li *et al.*, 2021] Junyi Li, Tianyi Tang, Wayne Xin Zhao, and Ji-Rong Wen. Pretrained language model for text generation: A survey. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4492–4499. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Survey Track.
- [Li *et al.*, 2022] Yuchao Li, Fuli Luo, Chuanqi Tan, Mengdi Wang, Songfang Huang, Shen Li, and Junjie Bai. Parameter-efficient sparsity for large language models fine-tuning. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 4223–4229. ijcai.org, 2022.
- [Miyake *et al.*, 2023] Daiki Miyake, Akihiro Iohara, Yu Saito, and Toshiyuki Tanaka. Negative-prompt inversion: Fast image inversion for editing with text-guided diffusion models. *CoRR*, abs/2305.16807, 2023.
- [OpenAI, 2023] OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023.
- [Thangarasa *et al.*, 2023] Vithursan Thangarasa, Abhay Gupta, William Marshall, Tianda Li, Kevin Leong, Dennis DeCoste, Sean Lie, and Shreyas Saxena. SPDF: sparse pre-training and dense fine-tuning for large language models. In Robin J. Evans and Ilya Shpitser, editors, *Uncertainty in Artificial Intelligence, UAI 2023, July 31 - 4 August 2023, Pittsburgh, PA, USA*, volume 216 of *Proceedings of Machine Learning Research*, pages 2134–2146. PMLR, 2023.
- [Vidan and Fiedler, 2023] Andy Vidan and Lars H. Fiedler. A composable just-in-time programming framework with llms and FBP. In *IEEE High Performance Extreme Computing Conference, HPEC 2023, Boston, MA, USA, September 25-29, 2023*, pages 1–8. IEEE, 2023.
- [Vidgof *et al.*, 2023] Maxim Vidgof, Stefan Bachhofner, and Jan Mendling. Large language models for business process management: Opportunities and challenges. In Chiara Di Francescomarino, Andrea Burattin, Christian Janiesch, and Shazia W. Sadiq, editors, *Business Process Management Forum - BPM 2023 Forum, Utrecht, The Netherlands, September 11-15, 2023, Proceedings*, volume 490 of *Lecture Notes in Business Information Processing*, pages 107–123. Springer, 2023.
- [von Rosing *et al.*, 2015] Mark von Rosing, Stephen White, Fred Cummins, and Henk de Man. Business process model and notation - BPMN. In Mark von Rosing, Henrik von Scheel, and August-Wilhelm Scheer, editors, *The Complete Business Process Handbook: Body of Knowledge from Process Modeling to BPM, Volume 1*, pages 429–453. Morgan Kaufmann/Elsevier, 2015.
- [Xu *et al.*, 2023] Benfeng Xu, An Yang, Junyang Lin, Quan Wang, Chang Zhou, Yongdong Zhang, and Zhendong Mao. Expertprompting: Instructing large language models to be distinguished experts. *CoRR*, abs/2305.14688, 2023.
- [Zhou *et al.*, 2023] Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.