A Direct Manipulation Interface for LLM-based Process Modeling

Sergio Pardo Gutierrez
Politecnico di Milano
Milan, Italy
sergio.pardo@mail.polimi.it

Lina Alrabie
Politecnico di Milano
Milan, Italy
lina.alrabie@polimi.it

Giulia Di Fede
Politecnico di Milano
Milan, Italy
giulia.difede@polimi.it

Monica Vitali
Politecnico di Milano
Milan, Italy
monica.vitali@polimi.it

Salvatore Andolina
Politecnico di Milano
Milan, Italy
salvatore.andolina@polimi.it

Abstract

This paper presents HyperMod, a human-centered system that leverages Large Language Models (LLMs) to automate complex process modeling tasks. The tool generates process models from natural language and incorporates a direct manipulation interface that enables users to interact with specific elements of the model by asking questions, requesting changes, or exploring alternatives. By combining generative AI with interactive control, the system aims to reduce technical barriers and support flexible human-AI interaction in digital process automation.

CCS Concepts

• **Human-centered computing → Human computer interaction (HCI)**.

Keywords

Business Process Management, BPMN, Direct Manipulation, Large Language Model, Digital Process Automation

ACM Reference Format:
Sergio Pardo Gutierrez, Lina Alrabie, Giulia Di Fede, Monica Vitali, and Salvatore Andolina. 2025. A Direct Manipulation Interface for LLM-based Process Modeling. In *CHItaly 2025: 16th Biannual Conference of the Italian SIGCHI Chapter (CHItaly 2025), October 06–10, 2025, Salerno, Italy*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3750069.3755960

1 Introduction and Related Work

LLMs have opened up new possibilities for natural language interfaces across diverse domains, from brainstorming [3] to code editing [16] and business process management [14]. Automating the transformation of natural language into process models has traditionally relied on rule-based NLP and computational linguistics techniques, such as part-of-speech tagging and syntactic parsing, to extract keywords, concepts, and structural patterns from the text for constructing Business Process Model and Notation (BPMN) diagrams [4, 7, 11]. While effective, these approaches often required highly structured input and offered limited usability for non-experts.

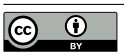

CHItaly 2025, Salerno, Italy

ACM ISBN 979-8-4007-2102-1/25/10
https://doi.org/10.1145/3750069.3755960

The advent of LLMs showed the potential to overcome these limitations by providing more flexible ways to interact with process modeling tools. Recent systems, such as PRODIGY [17], introduced a fully integrated conversational interface using Botpress and GPT-3.5 Turbo, allowing users to iteratively build models through guided dialogue, template generation, and next element recommendations. In contrast, ProMoAI [9] focused on improving the structural soundness of models through hierarchical representations using POWL (a subclass of Petri nets), integrating few-shot prompting, error resolution mechanisms, and a feedback loop for higher-quality outputs. BPMN-Chatbot [8] takes a different approach, providing a cost-effective solution featuring a JSON-to-BPMN module and support for versioning and voice input.

While existing BPMN systems support process modeling through natural language input, they often lack the fine control and immediate feedback needed for effective visual model refinement. From a Human-Centered AI perspective [5, 13], this limits user agency, particularly in complex tasks that involve frequent revisions and the exploration of alternatives. To address this, the concept of direct manipulation has become central in interactive system design, allowing users to act on visible representations through intuitive actions such as selecting, clicking, and dragging [6, 12]. These principles remain relevant in modern AI interfaces [10, 15]. Systems like DirectGPT [10] show that combining language models with direct manipulation improves efficiency and usability. Using "physical actions through prompt-object interactions" can help convey intent unambiguously and reduce the mental load needed to refer to specific visual elements through natural language [10].

This work proposes **HyperMod**, a direct manipulation interface for LLM-based process modeling, grounded in the principles of Human-Centered AI [5]. Unlike traditional chatbot-based modeling systems, our approach leverages the generative and interpretive strengths of LLMs while enabling users to ground prompts in selected visual elements of the process model. Users can ask questions, refine logic, or request suggestions in the context of specific model components. We expect that this will make the modeling experience more fluid and intuitive and provide users with a stronger feeling of control over the LLM-powered system [2].

2 HyperMod

HyperMod is composed of two main parts: the chat (Figure 1a, left) and the BPMN workspace (Figure 1a, right). Inside the chat panel, the user can type in prompts and get textual responses that typically involve the generation of a BPMN model rendered on the

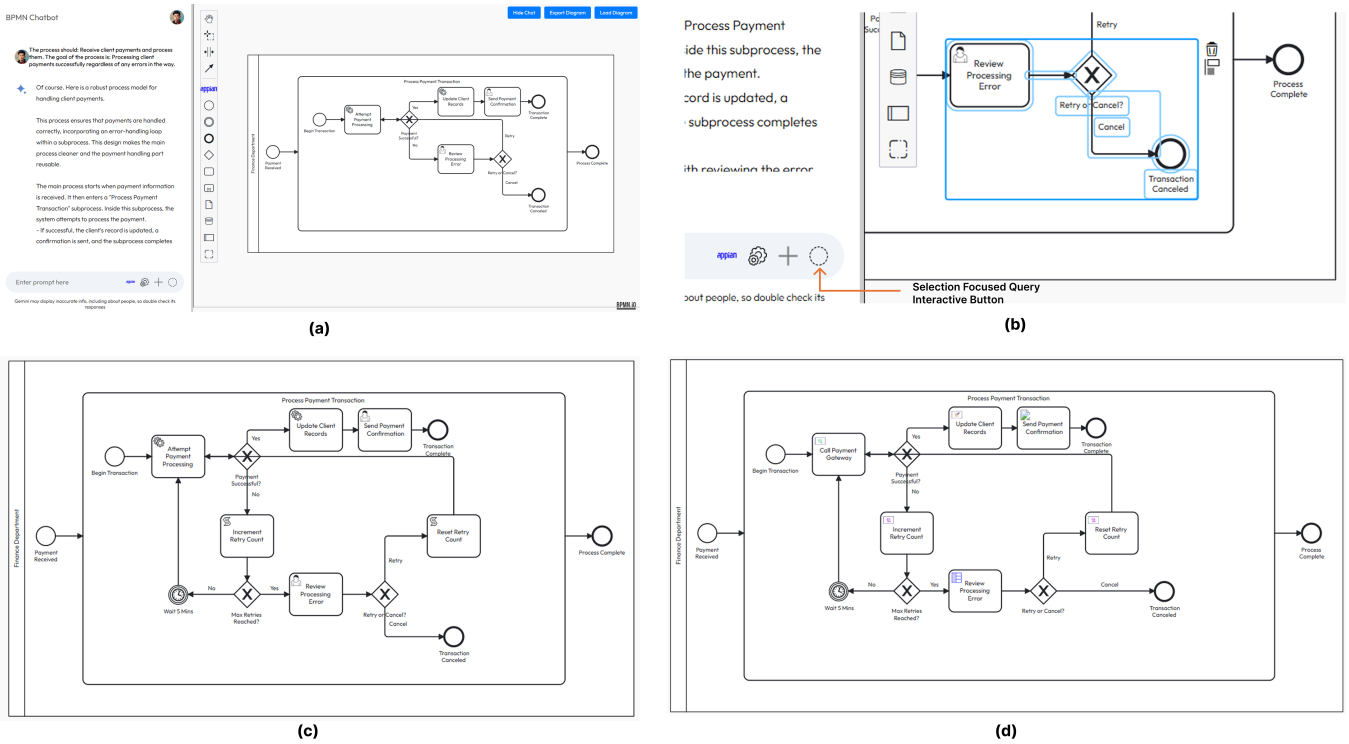

Figure 1: HyperMod's UI and walkthrough example. (a) The UI is composed of two main parts: the chat (left) and the BPMN workspace (right). In this example, the user requests the generation of a process model for handling client payments and HyperMod provides a first model. (b) Using the selection tool, relevant sub-processes are highlighted to be refined with extended error-handling logic. (c) The resulting process model from the selection focused query. (d) The refined model is adapted into Appian's BPMN notation.

BPMN workspace (Figure 1a, right). In the workspace, users can edit and export process model diagrams. The user can also activate the selection-focused query mode, which allows performing actions on specific elements via a combination of prompts written in natural language and direct manipulation. (Figure 1b). To render, modify, and extend the process model, BPMN.io was used as the primary modeling tool. Drawing on insights from early pilot studies, HyperMod was extended to support non-standard notations. In its current version, the system integrates modeling elements from Appian, a widely adopted commercial low-code automation platform for enterprise process applications [1]. To enable this integration, Appian's official documentation was provided to the model via its URL. This approach enables the specification of platform-specific modules (see Figure 1d), empowering even non-expert users to model processes using domain-specific or proprietary notations.

To illustrate how this works in practice, consider an e-commerce platform that needs to implement a robust customer payment workflow. To achieve this, a manager queries HyperMod: *"Generate a process model that ensures client payments are successfully received and processed"*. After inspecting the resulting process model (Figure 1a), the manager decides to enhance the error handling mechanism; therefore, he makes use of the selection tool to highlight the relevant sub-processes (Figure 1b) and queries the model: *"Could you extend the error-handling logic for the selected sub-processes?"* To which, HyperMod generates a more reliable process model (Figure 1c). As the next step, the manager requests that the current process model be adapted to Appian's BPMN notation. HyperMod adapts the process model and delivers the new implementation (Figure 1d). Finally, the manager performs a set of manual refinements to the diagram, ensuring it effectively meets internal standards.

3 Conclusion

In this work, we presented HyperMod, an LLM-based process modeling tool inspired by the principles of Human-Centered AI. By combining natural language interaction with the ability to perform actions through direct manipulation, it enables users to flexibly generate, explore, and refine BPMN diagrams. Prompts can be grounded in selected visual elements, making instructions more contextualized and responses more relevant. The system can help non-expert users generate initial process models and get acquainted with domain-specific or proprietary notations.

Acknowledgments

This work was partially supported by the European Union - Next Generation EU under the Italian NRRP, M4C2, Investment 1.3, D53C22002380006, PE00000013, FAIR.

References

[1] Appian Corporation. 2025. Process Modeling with Appian. Appian Documentation, version 25.2. https://docs.appian.com/suite/help/25.2/process_modeling.html Accessed on 31 July 2025.

[2] Giulia Di Fede, Lina Alrabie, and Salvatore Andolina. 2025. *Human-centered LLM.* Springer Nature Singapore, Singapore.

[3] Giulia Di Fede, Davide Rocchesso, Steven P. Dow, and Salvatore Andolina. 2022. The Idea Machine: LLM-based Expansion, Rewriting, Combination, and Suggestion of Ideas. In *Proceedings of the 14th Conference on Creativity and Cognition* (Venice, Italy) *(C&C '22)*. Association for Computing Machinery, New York, NY, USA, 623–627. doi:10.1145/3527927.3535197

[4] Fabian Friedrich, Jan Mendling, and Frank Puhlmann. 2011. Process model generation from natural language text. In *Advanced Information Systems Engineering: 23rd International Conference, CAiSE 2011, London, UK, June 20-24, 2011. Proceedings 23*. Springer, 482–496.

[5] Ozlem Ozmen Garibay, Brent Winslow, Salvatore Andolina, Margherita Antona, Anja Bodenschatz, Constantinos Coursaris, Gregory Falco, Stephen M. Fiore, Ivan Garibay, Keri Grieman, John C. Havens, Marina Jirotka, Hernisa Kacorri, Waldemar Karwowski, Joe Kider, Joseph Konstan, Sean Koon, Monica Lopez-Gonzalez, Iliana Maifeld-Carucci, Sean McGregor, Gavriel Salvendy, Ben Shneiderman, Constantine Stephanidis, Christina Strobel, Carolyn Ten Holter, and Wei Xu. 2023. Six Human-Centered Artificial Intelligence Grand Challenges. *International Journal of Human–Computer Interaction* 39, 3 (2023), 391–437. doi:10.1080/10447318.2022.2153320

[6] Edwin L Hutchins, James D Hollan, and Donald A Norman. 1986. Direct manipulation interfaces. In *User centered system design*. CRC Press, 87–124.

[7] Ana Ivanchikj, Souhaila Serbout, and Cesare Pautasso. 2020. From text to visual BPMN process models: Design and evaluation. In *Proceedings of the 23rd ACM/IEEE international conference on model driven engineering languages and systems*. 229–239.

[8] Julius Köpke and Aya Safan. 2024. Introducing the BPMN-Chatbot for Efficient LLM-Based Process Modeling. (2024).

[9] Humam Kourani, Alessandro Berti, Daniel Schuster, and Wil MP van der Aalst. 2024. Process modeling with large language models. In *International Conference on Business Process Modeling, Development and Support*. Springer, 229–244.

[10] Damien Masson, Sylvain Malacria, Géry Casiez, and Daniel Vogel. 2024. Directgpt: A direct manipulation interface to interact with large language models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–16.

[11] Maximilian Riefer, Simon Felix Ternis, and Tom Thaler. 2016. Mining process models from natural language text: A state-of-the-art analysis. *Multikonferenz Wirtschaftsinformatik (MKWI-16), March* (2016), 9–11.

[12] Ben Shneiderman. 1983. Direct manipulation: A step beyond programming languages. *Computer* 16, 8 (1983), 57–69.

[13] Ben Shneiderman. 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human–Computer Interaction* 36, 6 (2020), 495–504.

[14] Maxim Vidgof, Stefan Bachhofner, and Jan Mendling. 2023. Large language models for business process management: Opportunities and challenges. In *International conference on business process management*. Springer, 107–123.

[15] Liwenhan Xie, Chengbo Zheng, Haijun Xia, Huamin Qu, and Chen Zhu-Tian. 2024. WaitGPT: Monitoring and Steering Conversational LLM Agent in Data Analysis with On-the-Fly Code Visualization. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) *(UIST '24)*. Association for Computing Machinery, New York, NY, USA, Article 119, 14 pages. doi:10.1145/3654777.3676374

[16] J.D. Zamfirescu-Pereira, Eunice Jun, Michael Terry, Qian Yang, and Bjoern Hartmann. 2025. Beyond Code Generation: LLM-supported Exploration of the Program Design Space. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 153, 17 pages. doi:10.1145/3706598.3714154

[17] Clara Ziche and Giovanni Apruzzese. 2024. LLM4PM: A case study on using Large Language Models for Process Modeling in Enterprise Organizations. In *International Conference on Business Process Management*. Springer, 472–483.