

Process Copilots: Scaling Process Mining Through Conversational Interfaces

Carolin Ullrich¹, Teodora Lata¹ and Alexander Seeliger¹

¹Celonis, Munich, Germany

Abstract

Large Language Models (LLM) continue to increase their adoption and have reached the Business Process Management (BPM) domain. Traditionally, process insights have been consumed via tailored dashboards built by specialized analysts. Progressions in conversational Artificial Intelligence (AI) promise a paradigm shift reducing cost and enhancing speed to process insight. We introduce Celonis Process Copilots, an LLM-powered conversational interface readily integrated into the Celonis platform. With its access to event data and business annotations, it empowers a broader user base to directly query and understand complex process dynamics. Applied in various processes and domains, the tool fundamentally democratizes access to process mining insights. This demo paper showcases the tool's capabilities and current maturity in delivering actionable and on-demand process intelligence.

Keywords

Process Mining, Large Language Models, LLM, Conversational Interface

1. Introduction

With over 50% of U.S. citizens using large language models (LLMs)¹ and OpenAI adding 1 million new users within one hour², the application and usage of LLMs is rapidly growing, also reaching the Business Process Management (BPM) field. BPM is an interdisciplinary field concerned with the management of processes to realize improvement opportunities within companies which is not only researched by academia but also actively practiced [1]. Vidgof et al. [2] identified promising LLM use cases along the BPM life cycle as well as outlined research directions on this intersection. Building on this, LLMs can be used for creating process definitions based on already available documentation. The BPM community has seen many case studies on using LLMs for process modeling such as [3, 4] and Franzoi et al. [5] described its real-world application and potential through studying a multinational company. Conversely, applications for generating textual descriptions derived from existing process models have been presented in the past [6].

However, the BPM life cycle consists of more than producing process models. In this broader context, process mining refers to techniques extracting process insights based on event logs in order to discover, monitor and enhance processes, and can be applied during multiple steps of the BPM life cycle, e.g., during process analysis or process monitoring. The use of LLM applications in the field of process mining is still in an early stage with first experiments [7] and proof-of-concepts [8], however, these early attempts already show promising results. Currently, analytical and technical skills [9] coupled with tool knowledge [10], present barriers to successful process mining initiatives. However, accessing process mining insights and identifying process improvement opportunities through a conversational interface might alleviate these challenges. In addition, the hope is that LLMs can be used for automated process discovery [11], a studied success factor in process mining initiatives.

Proceedings of the Best BPM Dissertation Award, Doctoral Consortium, and Demonstrations & Resources Forum co-located with 23rd International Conference on Business Process Management (BPM 2025), Seville, Spain, August 31st to September 5th, 2025.

✉ c.ullrich@celonis.com (C. Ullrich); t.lata@celonis.com (T. Lata); a.seeliger@celonis.com (A. Seeliger)

ORCID 0000-0002-0649-6225 (A. Seeliger)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://imaginingthefuture.org/reports-and-publications/close-encounters-of-the-ai-kind/close-encounters-of-the-ai-kind-main-report/>

²<https://x.com/sama/status/1906771292390666325>

In this paper, we introduce Celonis Process Copilots, a task-specific conversational LLM grounded in companies' object and event data. This tool has been applied in practice across a variety of processes and domains, and in this demo, we demonstrate how the use of LLMs can advance the practical application of BPM.

2. Process Copilots

Process Copilot is a conversational interface within the Celonis platform that allows users to explore and analyze process data connected to the system. It allows users asking questions, building visual components, looking for insights, or viewing process flows. Users can analyze data by asking natural language questions or using the predefined prompts to get a response in a variety of formats. For example, it can respond in text format but also visual elements like KPI cards, tables or charts.

We built highly customizable tools that equip the LLM with additional capabilities to access data, context, and knowledge from the Celonis platform. These tools have been built to answer process mining specific questions but also enrich the context of the LLM by injecting specific instructions for reliability and consistency. We explored several tools and instructions with different capabilities and granularity (e.g., splitting and merging tools, or fine-tuning descriptions and arguments). The current set of tools in Process Copilots showed most fit for many use cases and customers as well as offer a wide range of customization options.

2.1. Setting up Process Copilots

Process Copilots must be tailored to specific use cases to ensure reliable performance. Unlike consumer-facing tools like ChatGPT, where there is high tolerance for mistakes, enterprise environments demand much greater accuracy. Narrowing a Copilot's scope to a focused task better grounds the AI and reduces the risk of costly errors in business operations. Figure 1 shows the step-by-step setup interface of Process Copilots.

First, an analyst user defines the **knowledge inputs** for the Process Copilot. For that, Process Copilots need to be connected to a *Knowledge Model*, which is a representation of raw process data transformed into standardized business objects and events and enhanced with semantic meaning. Then relevant KPIs, record attributes, event logs and filters are selected from the Knowledge Model to define the scope of the Copilot. The response accuracy of Process Copilots relies on the quality of the Knowledge Model, e.g., the use of clean descriptions of business objects or defining formats, units, and desired directions of KPIs. The system automatically suggests improvement opportunities for the underlying Knowledge Model definitions to increase output quality. In a second step, the analyst selects **tools**, which equip Process Copilots with additional capabilities to complete tasks. There are eleven tools for a wide range of capabilities including calculation, visualization, exploratory process analysis, recommended insights, clarification, automation, and data retrieval. In a third step, the user defines a **system prompt**, consisting of a task description, explaining specific terminology in the context at hand, and providing example questions and answers. Next, the user selects the underlying **LLM** that is used to generate responses. Analysts can choose between OpenAI GPT-4o, GPT-4o mini, Anthropic Claude Sonnet 3.5, Claude Sonnet 3.7, or they can bring their own model. Optionally, the analyst defines the start screen of the Process Copilot to ease end-user on-boarding and expectation management, for example by adding frequently asked questions or pinning important KPIs. Once testing and troubleshooting are completed, Process Copilots can be published so they become available to a broader audience in the company.

In a monitoring section, the analyst can see how users are interacting with the Process Copilot and use this information to further improve its configuration.

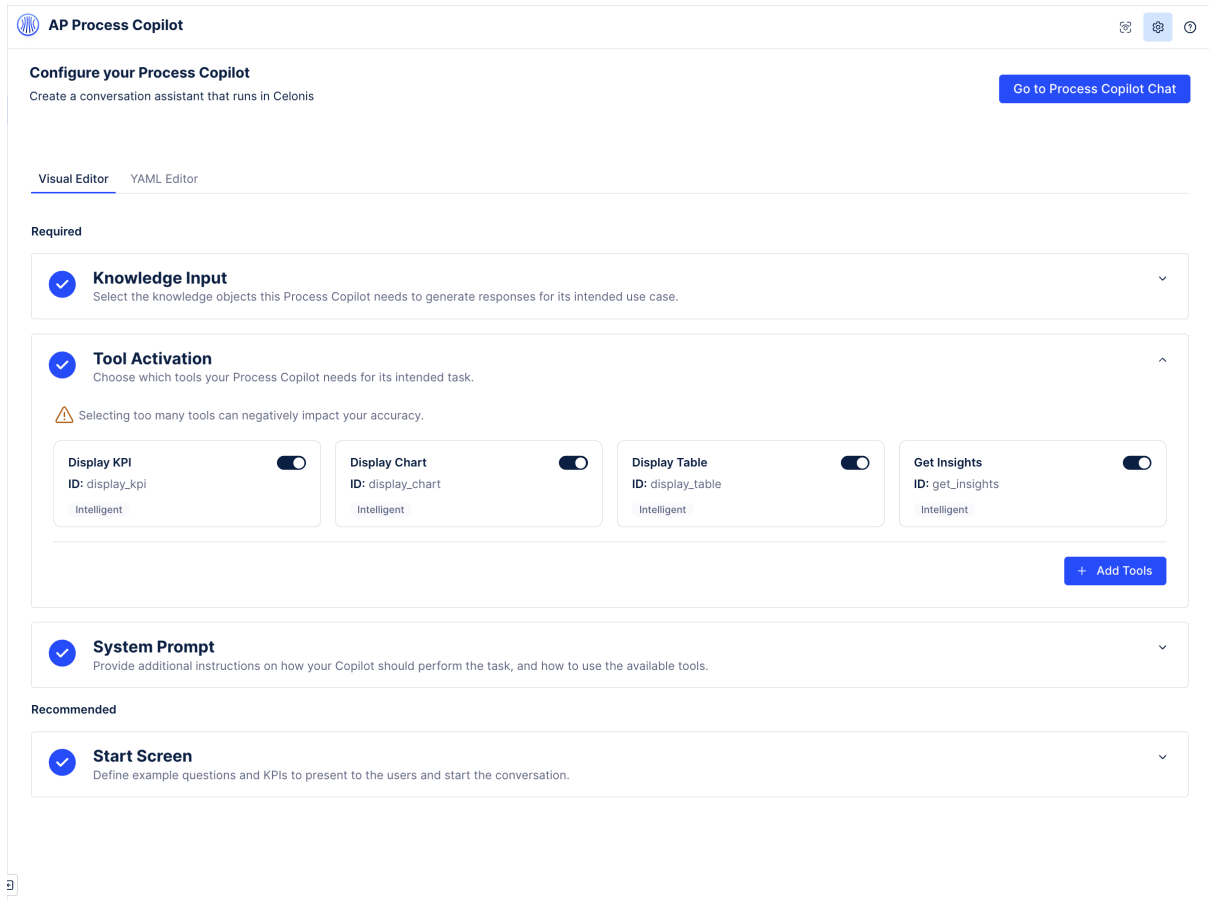


Figure 1: Setup interface of Process Copilot, showing the tools section to equip the LLM with certain capabilities from Celonis.

2.2. Using Process Copilots

Any user of a Celonis instance can interact with published Process Copilots and, thus, engage in conversations with underlying process data (see Figure 2). Users can choose between four methods for interacting with Process Copilots. When a conversation is initialized, **quick start questions** appear on the screen as useful conversation starters, if configured by an analyst. Starter questions are helpful to kick off a conversation, indicating which types of questions can be answered by this particular instance as well as make users familiar what to ask when using Process Copilots the first time.

In addition, Process Copilots offer **template prompts** which aim to both generate high-quality LLM responses and educate the users in crafting successful prompts. The template prompts are grouped by use case into analyzing metrics, finding opportunities, exploring the process and building a visualization.

At the bottom of each Process Copilot is a **free text field** that can be used to craft own questions. By clicking in the free text field, a modal with additionally **suggested questions** appears that the user can choose from. These are generated based on previous user interaction with the Copilot, available data, and commonly asked questions.

After the user has decided to enter a question or task, the tool generates an answer based on the knowledge input and scope that was previously defined. It can answer requests with text responses, but also with rich component responses such as charts, tables, and other visualizations. These visualizations can be reused in new and existing dashboards via a YAML copy functionality, bridging traditional and new ways of consuming process data.

The user can provide feedback by rating a response as helpful or not helpful as well as re-generate it. Users can also validate responses by inspecting the actual computation steps that the Process Copilot took, as well as the underlying data it used to generate the response. In addition, users can navigate

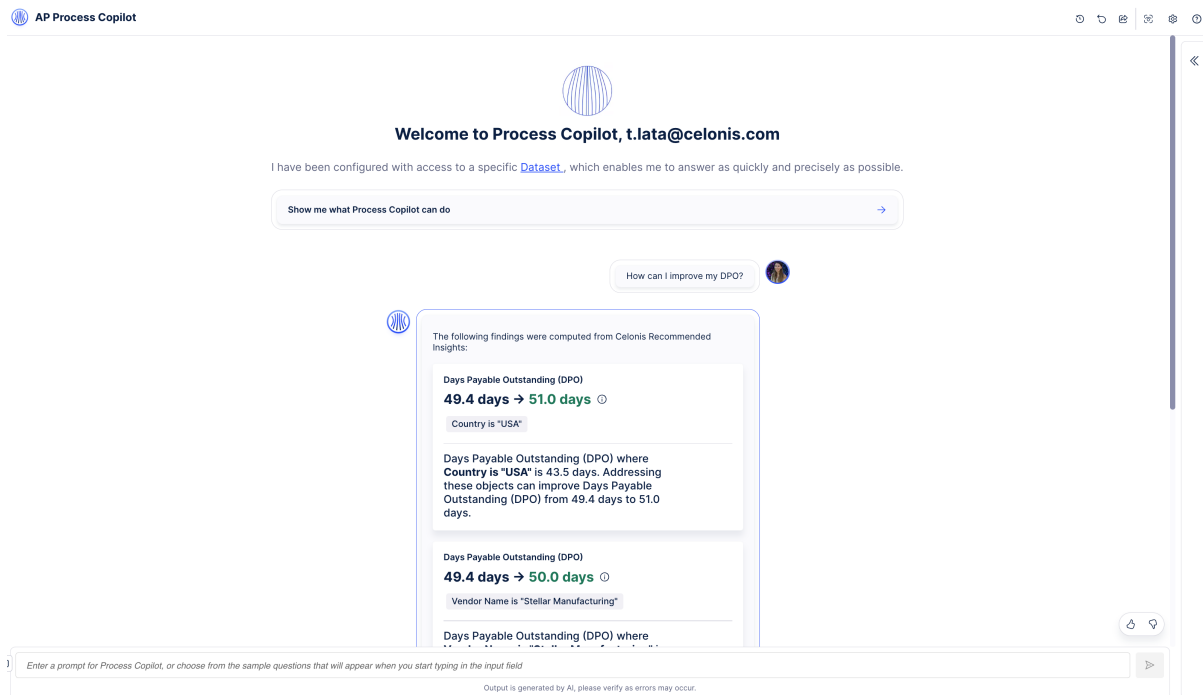


Figure 2: Process Copilot conversational interface, showing the recommended insights for the “Days Payable Outstanding (DPO)” KPI.

their own chat history if needed or share their conversation with other users.

3. Maturity of the Tool

Since 2025, Process Copilot is generally available but requires a dedicated license, available through standard sales channels. It has been deployed across various processes and domains, including Accounts Payable, Accounts Receivable, Procurement, Inventory Management, Order Management, IT Service Management, and Claims Management. It has been tested with hundreds of users, encompassing a wide range of roles: from executives who need transparency in process health to improve decision making, to shop floor workers who use the tool to find available spare parts faster or manage truck loads more efficiently. Process Copilots are accessible directly within the Celonis instance or can be embedded in chat tools such as Microsoft Teams. A short demo of the tool³ is available as well as real-world use case demonstrations featuring a Copilot user⁴.

Process Copilot’s design emphasizes task-specificity rather than being a general-purpose LLM, which necessitates an effective setup and scope definition. Providing high-quality event data, selecting meaningful business objects, choosing the right tools and writing effective prompts are prerequisites for the usage and significantly influence the end-user experience and accuracy of results. Thus, we invest in extensive user documentation⁵ and online training materials⁶ to support users in prompt engineering and to shorten the setup time for the different task-specific Copilots. Process Copilots rely on third-party LLMs such as provided by OpenAI or Anthropic to generate responses, thus, the quality of its results is influenced by the capabilities of these models. As a mitigation, Process Copilot is model-agnostic, and users can bring their own models, also fine-tuned and domain-specific variants.

³<https://videos.celonis.com/watch/A2pknwmiMLgbWS8C5ZoWqq>

⁴<https://www.celonis.com/celosphere/2024/recordings/watch/?search=copilot&modalId=WwnWyma5IoArt5tPfGRzA#celosphere-24-use-gen-ai>

⁵<https://docs.celonis.com/en/process-copilot.html>

⁶<https://academy.celonis.com/courses/configure-process-copilot>

4. Conclusion and Future Work

Process Copilots are task-specific, conversational interfaces within the Celonis platform that allow a diverse range of users with different backgrounds and skill-levels to interact with process data. It offers a guided setup flow to tailor the scope of the Process Copilot and it offers quick start questions and example prompts to users interacting with the Copilot.

In future work, Process Copilot will ingest and search through process documentation as well as process models to enrich user queries with contextual information. Moreover, we work on making the existing tools Process Copilot can use available to external agents such as Microsoft Studio Copilot.

Declaration on Generative AI

During the preparation of this work, the authors used Google Gemini 2.5 Pro in order to: Improve readability and language. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] M. Dumas, M. La Rosa, J. Mendling, H. A. Reijers, Introduction to business process management, in: *Fundamentals of Business Process Management*, Springer Berlin Heidelberg, 2018, p. 1–33. doi:10.1007/978-3-662-56509-4_1.
- [2] M. Vidgof, S. Bachhofner, J. Mendling, Large language models for business process management: Opportunities and challenges, in: *Business Process Management Forum*, Springer Nature Switzerland, 2023, p. 107–123. doi:10.1007/978-3-031-41623-1_7.
- [3] J. Köpke, A. Safan, Introducing the bpmn-chatbot for efficient llm-based process modeling, in: *Proceedings of the Best BPM Dissertation Award*, CEUR Workshop Proceedings, 2024.
- [4] M. Grohs, L. Abb, N. Elsayed, J.-R. Rehse, Large language models can accomplish business process management tasks, in: *Business Process Management Workshops*, Springer Nature Switzerland, 2024, p. 453–465. doi:10.1007/978-3-031-50974-2_34.
- [5] S. Franzoi, M. Delwaulle, J. Dyong, J. Schaffner, M. Burger, J. vom Brocke, Using large language models to generate process knowledge from enterprise content, in: *Business Process Management Workshops*, Springer Nature Switzerland, 2025, p. 247–258. doi:10.1007/978-3-031-78666-2_19.
- [6] H. Leopold, J. Mendling, A. Polyvyanny, Generating natural language texts from business process models, in: *Active Flow and Combustion Control 2018*, Springer International Publishing, 2012, p. 64–79. doi:10.1007/978-3-642-31095-9_5.
- [7] L. Barbieri, E. Madeira, K. Stroeh, W. van der Aalst, A natural language querying interface for process mining, in: *Journal of Intelligent Information Systems*, volume 61, Springer Science and Business Media LLC, 2022, p. 113–142. doi:10.1007/s10844-022-00759-9.
- [8] K. Lashkevich, F. Milani, M. Avramenko, M. Dumas, Redesigning business processes to reduce waiting times using large language models, in: *Proceedings of the Best BPM Dissertation Award*, CEUR Workshop Proceedings, 2024.
- [9] N. Martin, D. A. Fischer, G. D. Kerpedzhiev, K. Goel, S. J. J. Leemans, M. Röglinger, W. M. P. van der Aalst, M. Dumas, M. La Rosa, M. T. Wynn, Opportunities and challenges for process mining in organizations: Results of a delphi study, in: *Business & Information Systems Engineering*, volume 63, Springer Science and Business Media LLC, 2021, p. 511–527. doi:10.1007/s12599-021-00720-0.
- [10] L. Zimmermann, F. Zerbató, B. Weber, What makes life for process mining analysts difficult? a reflection of challenges, in: *Software and Systems Modeling*, volume 23, Springer Science and Business Media LLC, 2023, p. 1345–1373. doi:10.1007/s10270-023-01134-0.
- [11] A. Mamudu, W. Bandara, M. T. Wynn, S. J. J. Leemans, A process mining success factors model, in: *Business Process Management*, Springer International Publishing, 2022, p. 143–160. doi:10.1007/978-3-031-16103-2_12.