

Building Process Mining Dashboards Manually or AI-assisted with Celonis Views

Carolin Ullrich^{1,*}, Teodora Lata^{2,†} and Gregory Benton^{2,†}

¹Celonis, Theresienstraße 6, Munich, BY 80333, Germany

²Celonis, 1 World Trade Center, New York, NY 10006, United States

Abstract

Transforming raw process data into actionable insights for process improvements requires defining the right data questions and answering them effectively with the right visualizations bundled in dashboards. This is not a trivial task for builders. To successfully implement a process mining project, domain expertise and technical skills are needed - a combination that can be rarely found within a single user. This paper presents Celonis Views, a simplified dashboard building experience as part of a process mining platform. Users get support through guided configuration flows, best practices being built into the component configurations, smart defaults and LLM assistants. For this reason the entry barrier to build or configure one's own process mining analyses is lowered and access to process mining insights is further democratized.

Keywords

Process Mining, LLM, Dashboards

Metadata description	Value
Tool name	Celonis Views
Current version	N/A
Legal code license	proprietary
Languages, tools and services used	Python, Typescript, Angular, OpenAI
Supported operating environment	Web browser
Download/Demo URL	https://tinyurl.com/icpm2024
Documentation URL	https://docs.celonis.com/en/creating-views.html
Source code repository	N/A
Screencast video	https://celonis-academy.wistia.com/medias/5qn5qmr7uh

1. Introduction

Process Mining enables organizations to see, monitor and improve their processes given event log data captured from information systems in use. Process Mining is not only an established research field but has been successfully applied in real-world settings across industries and process domains. To make process mining effective for an organization, insights need to be

ICPM 2024 Tool Demonstration Track, October 14-18, 2024, Kongens Lyngby, Denmark

*Corresponding author.

†These authors contributed equally.

✉ c.ullrich@celonis.com (C. Ullrich); t.lata@celonis.com (T. Lata); g.benton@celonis.com (G. Benton)

🆔 0009-0008-9892-1807 (C. Ullrich); 0009-0004-0722-2527 (T. Lata); 0009-0002-7882-9559 (G. Benton)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

made available to the user. One common way for the user to consume information provided by process mining is via *dashboards*.

Usually multiple user roles are involved in process mining projects, including process analysts that provide process mining applications and business users or process experts that consume these process mining applications to eventually act upon the findings [1]. Previous research has shown that it is important for analysts to understand the domain of the process mining project. However, it was also noted that analysts do not always possess the needed domain knowledge and that sometimes the needed collaboration between user groups is missing [2]. Others even report on unwillingness to share domain knowledge indicating missing trust between the involved roles [3]. Tools like Celonis Business Miner are aimed to automatically generate process mining insights for business users [4] to reduce the need of multiple roles working together simultaneously to deploy process mining. User feedback has shown shortcomings of this approach relying on auto-generated content without analyst involvement, such as lack of customization and missing journeys to move from a one-off analysis to process mining applications that are used continuously and built for scale. Dashboard development remains essential because it enables a deep understanding of the specific nuances of a process, acting as the translation layer between raw process data and the people who can take actionable steps to improve that process.

Today, building custom process mining analyses poses high requirements on the skill set of the user [3] and requires dedicated training [5] thus creating high entry barriers for less technical roles that might be well equipped with domain expertise. We aim to address this problem through the following innovations which ease the building process so that less technical roles can build custom process mining content.

2. Core Features

Celonis Studio [6] is an environment to build analytical or operational process mining dashboards, so-called Views. In this section we outline the core enhancements of the View editing experience as well as the latest Large Language Model (LLM) assistants for configuring View visualizations.

2.1. Enhanced View Building Experience

The enhanced View building experience consists of some major improvements along different steps of the dashboard building process. These improvements include creating a layout, choosing, and configuring the right visualizations for the data question at hand, and defining and reusing knowledge fueling these visualizations.

Layout. The layout grid offers guiding lines and magnetism to support the builder in aligning components relative to each other and to ensure a tidy and easy to read dashboard. With drag-and-drop interactions the user can add components to their View and choose between 30 components that are pre-sorted into categories to ease their selection. The separation between category charts, time series charts and distribution charts supports the builder in choosing the right visualization type for the data they want to display. As a result, good practices on

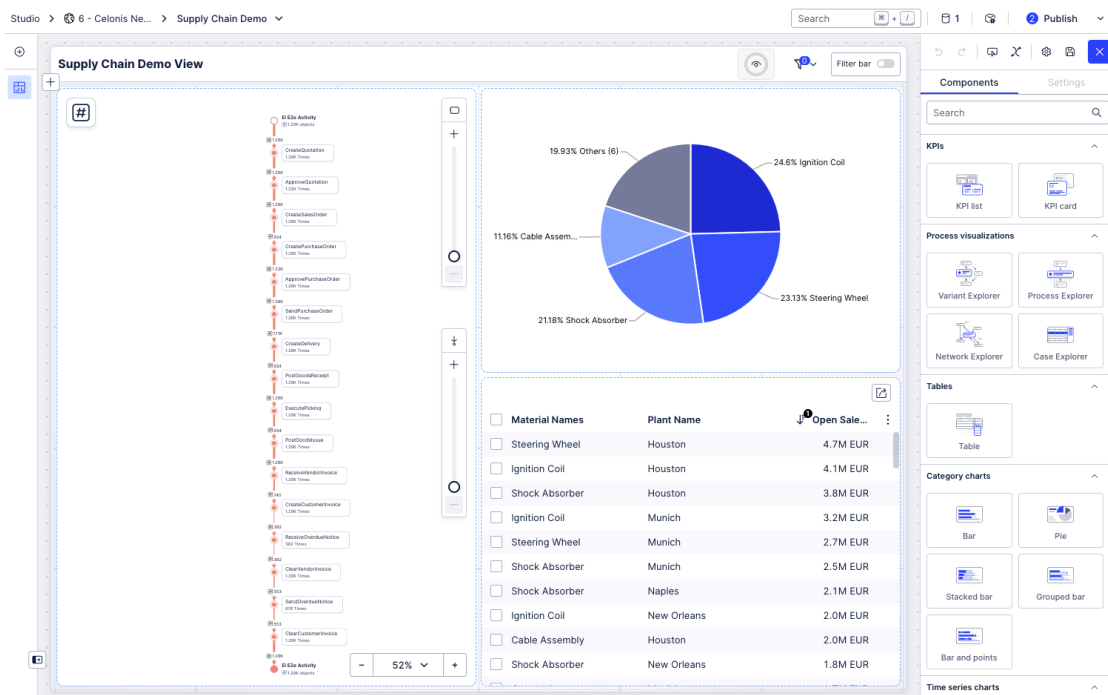


Figure 1: Screenshot of the improved View building experience with sorted components

dashboard design are built into the product now. By that we aim to tap into the opportunity 'Generating intuitive visualizations for business users' which has been rated as extremely relevant by a recent Delphi study in the process mining field [3].

Component editing. Once the component is dropped on the grid, a guided component configuration is opened guiding the user through the set-up. When applicable, a conversational style is used to make it easier for less technical users to configure a component. For instance, to set up a bar chart, the user is asked to complete the sentence "Each bar is a ..." with the dimension they are interested in. Similarly, the metric of a bar chart is defined by completing the sentence "Length of bar shows ...". The component editor is offering alternative display options for the selected data. This way the product consults the user on possible alternatives while educating the user about non-fitting visualization types. As a result, the users improve their data visualization skills over time.

Some components work out of the box after being added by choosing smart defaults for their configuration and data source. An example is the Process Explorer, a directly-follows-graph, that makes use of the default activity table set in the data layer. The default is then also shared between components, which increases consistency.

Data querying. When picking the correct data for the configuration of components, the builder has access to all tables of the data model as well as central definitions of metrics, filters

and variables they can reference. This way they can reuse business definitions of previous projects or from colleagues which leads to consistency across process mining implementations and allows non-technical users to build upon existing work. In addition, users can on the fly translate their process questions into queries, which are then executed by the Celonis query engine. The Process Query Language (PQL) [7] editor supports the data analyst with an in-built PQL reference library, inline PQL suggestions and guided error and success messages. Less-experienced users thus get trained without the need to leave their task at hand and switch to an external documentation or training environment. The entry barrier to work with PQL is now reduced.

2.2. LLM-assisted Component Configuration

In addition to the aforementioned improvements to make the dashboard building experience easier, Celonis Views also offers a LLM assistant for setting up commonly used components to further reduce the entry barrier to dashboard building. Users can communicate with the assistant via text and configure components via prompts using their natural language. The LLM assistant is designed to be user-friendly and aims to accelerate the component creation without replacing the critical, creative aspects of assembling a dashboard.

The LLM assistant can be accessed within the component configuration bar. For each component, the LLM-assistant will search the underlying data source of the dashboard for relevant items like KPIs and attributes. If the content is clear based on the user request, the LLM will configure the content into the visualization requested by the user. If not, we provide the LLM with a disambiguation tool, which will present users with options to choose from. In Figure 2 we can see an example user interaction with the LLM assistant to build a table component. In addition, specific to each component, the assistant follows predefined instructions provided by Celonis, such as understanding what a "breakdown" means in the context of a table component.

The initial version of the LLM assistant focuses on helping users with data selection and component configuration, while later versions can be enhanced to cover advanced component configurations such as conditional coloring.

3. Availability

The enhanced View building experience has been available to all Celonis customers since May 2024. It has more than 1,300 monthly active users, and users have publicly shared how the enhancements of the building journey positively impacted the success of their process mining initiatives¹. Academics and practitioners can use its full set of capabilities through the Celonis Free Plan² and access a screencast and documentation. The LLM-assisted component configurator is currently in limited availability, undergoing testing with end-users to identify gaps and necessary enhancements before a broader launch.

¹<https://videos.celonis.com/watch/iMKWfsky8Ve5r2jQLTnmrT>

²<https://www.celonis.com/solutions/free-plan/>

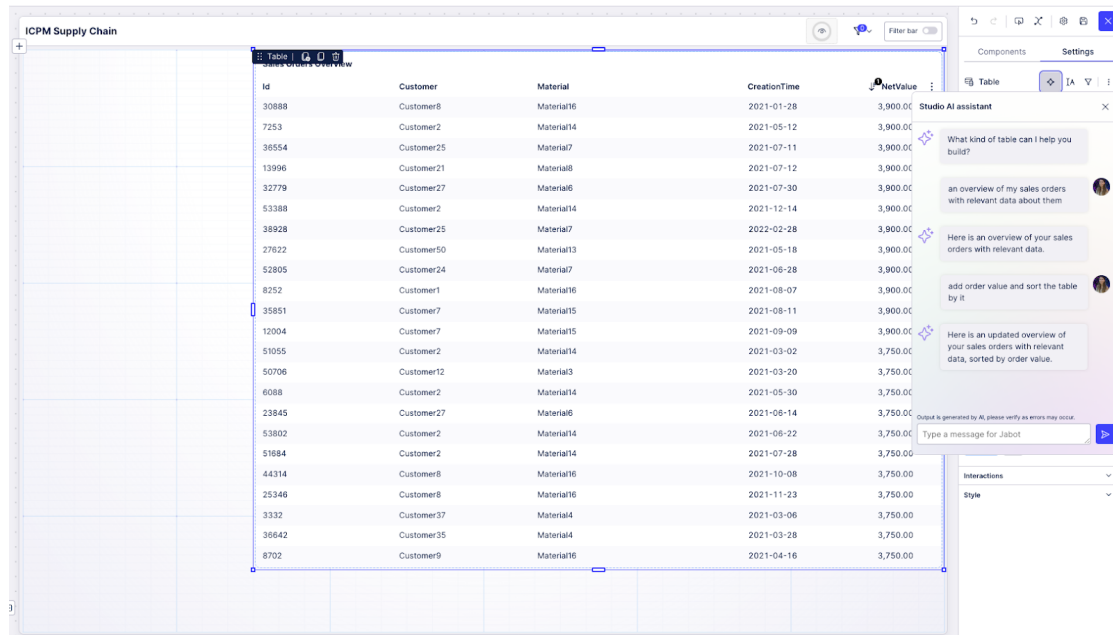


Figure 2: Screenshot of user interaction with the LLM assistant to build a table component

4. Future Work

In an ongoing user research study we examine under which conditions the LLM-assisted component configuration is most helpful. Based on the findings we plan to adjust the assistant to be tailored to specific personas (experienced vs. inexperienced analysts), tasks (building from scratch vs. adjusting an existing configuration) or user problems (choosing the right visualization vs. speeding up the set-up process). Other interesting areas of research with LLM assistants used in process mining include assistance in process insight generation, process improvements and dashboard integration.

5. Conclusion

The enhanced View building experience introduced a differentiated set of improvements on layouting, setting up process visualizations and defining process queries. It is complemented by an LLM-assisted component configuration. As a result, the entry barrier to build analytical and operational dashboards to surface process mining insights is reduced. Ultimately, more users with diverse skill sets and backgrounds can build their own process mining applications which contributes to the adoption of process mining as a discipline.

References

- [1] M. L. van Eck, X. Lu, S. J. J. Leemans, W. M. P. van der Aalst, Pm²: A process mining project methodology, in: J. Zdravkovic, M. Kirikova, P. Johannesson (Eds.), *Advanced Information Systems Engineering*, Springer International Publishing, Cham, 2015, pp. 297–313. doi:10.1007/978-3-319-19069-3_19.
- [2] F. Zerbato, P. Soffer, B. Weber, Process mining practices: Evidence from interviews, in: C. Di Ciccio, R. Dijkman, A. del Río Ortega, S. Rinderle-Ma (Eds.), *Business Process Management*, Springer International Publishing, Cham, 2022, pp. 268–285. doi:10.1007/978-3-031-16103-2_19.
- [3] 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, *Business & Information Systems Engineering* 63 (2021) 511–527. URL: <http://dx.doi.org/10.1007/s12599-021-00720-0>. doi:10.1007/s12599-021-00720-0.
- [4] C. Ullrich, T. Lata, Business miner: Process mining insights for business users, in: *2023 5th International Conference on Process Mining (ICPM), Doctoral Consortium/Demo, 2023*.
- [5] A. Mamudu, W. Bandara, M. T. Wynn, S. J. J. Leemans, A process mining success factors model, in: C. Di Ciccio, R. Dijkman, A. del Río Ortega, S. Rinderle-Ma (Eds.), *Business Process Management*, Springer International Publishing, Cham, 2022, pp. 143–160. doi:10.1007/978-3-031-16103-2_12.
- [6] C. Ullrich, T. Lata, J. Geyer-Klingenberg, Celonis studio-a low-code development platform for citizen developers, in: *International Conference on Business Process Management (BPM), PhD/Demos, 2021*, pp. 102–105.
- [7] T. Vogelgesang, J. Ambrosy, D. Becher, R. Seilbeck, J. Geyer-Klingenberg, M. Klenk, *Celonis PQL: A Query Language for Process Mining*, Springer International Publishing, Cham, 2022, pp. 377–408. URL: https://doi.org/10.1007/978-3-030-92875-9_13. doi:10.1007/978-3-030-92875-9_13.