ClustXRAI: Interactive Cluster Exploration and Explanation for Process Mining with Generative AI*

Christopher Lorenz Werner[1,†], Jonas Amling[2,†], Christian Dormagen[1,†] and Stephan Scheele[3,*,†]

[1] *Otto-Friedrich-Universität Bamberg, An der Weberei 5, 96049 Bamberg, Germany*
[2] *dab: Daten - Analysen & Beratung GmbH, Hans-Obser-Straße 12, 94469 Deggendorf, Germany*
[3] *Ostbayerische Technische Hochschule Regensburg (OTH Regensburg), Seybothstraße 2, 93053 Regensburg, Germany*

Abstract

Process mining enables business analysts to analyze and optimize business processes using event data. However, a key challenge remains in interpreting the complex, known as "spaghetti" models that often arise. We introduce a demonstrator ClustXRAI that addresses this challenge by integrating clustering techniques with eXplainable Artificial Intelligence (XAI) and generative AI. Clustering reduces model complexity, while XAI provides rule-based explanations that describe and differentiate the clusters. As we designed our demonstrator for business analysts — who may be unfamiliar with interpreting rule-based explanations — we integrate background knowledge with a large language model (LLM) to verbalize explanations. The demonstrator comprises two core features: an interactive dashboard for visual exploration of process mining clusters and a chatbot interface that delivers clear, natural language explanations. We show how these complementary tools enable intuitive exploration and interpretation of process mining results, making the technology more accessible to a broader audience beyond technical experts.

Keywords

eXplainable Artificial Intelligence, Process Mining, Clustering, Natural Language Explanations

1. Introduction and Background

Process mining focuses on extracting insights from real-world event data by analyzing event logs recorded in information systems, enabling the discovery, monitoring, and optimization of business processes [1]. However, the complexity of real-world data often leads to highly intricate process models, characterized by numerous transitions, overlapping paths, and frequent loops. Such "Spaghetti-Models" [2], illustrated in Fig. 2, are difficult to interpret, hindering the extraction of meaningful insights from the underlying data. Their complexity makes it challenging for analysts to identify key patterns, detect deviations, or uncover opportunities for process improvement.

By grouping similar processes and identifying sub-processes, clustering techniques can significantly reduce model complexity, thereby improving interpretability [3, 4]. However, as an unsupervised learning method, clustering lacks inherent interpretability [5], making it difficult for users – especially business analysts without a background in machine learning – to understand the rationale behind the clustering process. Yet, this group is the primary target audience for process mining.

Research on eXplainable Artificial Intelligence (XAI) addresses this challenge by developing methods to explain the outcomes of AI algorithms [5, 6]. Within process mining, the integration of XAI promises to provide clear and meaningful explanations of the clustering algorithm's outcomes, helping users

Late-breaking work, Demos and Doctoral Consortium, colocated with The 3d World Conference on eXplainable Artificial Intelligence: July 09–11, 2025, Istanbul, Turkey
* Funded by the Bavarian Ministry of Economic Affairs (BayVFP) under KIGA: DIK0313.
* Corresponding author.
† These authors contributed equally.
christopher-lorenz.werner@uni-bamberg.de (C. L. Werner); jonas.amling@dab-gmbh.de (J. Amling); christian.dormagen@uni-bamberg.de (C. Dormagen); stephan.scheele@oth-regensburg.de (S. Scheele)
0009-0002-0329-8829 (C. L. Werner); 0009-0000-1554-9950 (J. Amling); 0009-0006-7744-746X (C. Dormagen); 0000-0003-0787-3181 (S. Scheele)

comprehend the results and gain insight into the algorithmic reasoning. By leveraging the capabilities of a large language model (LLM), XAI explanations can be transformed into more accessible, interactive formats, offering human-readable narratives [7]. The integration of XAI and LLMs into process mining has only recently gained traction, with emerging applications in predictive model interpretability, concept drift handling, prescriptive monitoring, and process explanation [8, 9, 10, 11, 12].

This paper presents the demonstrator ClustXRAI that implements the XAI method for explainable clustering in process mining [13]. This work demonstrates how retrieval-augmented LLMs can operationalize rule-based explanations in process mining. The presented demonstrator enables business analysts to interactively explore, understand, and compare process clusters using visual analytics and natural language explanations. We present our contribution to interactive explainable clustering in process mining: An architecture/implementation integrating process discovery, a clustering algorithm, and XAI techniques, along with a dashboard with visualizations and an LLM-powered chatbot. Our approach enhances the process discovery workflow and fosters deep insights into business processes. Section 2 outlines the system architecture and explains the individual components and technologies used. Section 3 showcases the capabilities of our demonstrator through an example of interactive process exploration. Finally, we provide an overview of the demonstrator's features and contributions in Section 4.

2. System Overview

2.1. Architecture

Our architecture, shown in Fig. 1, integrates a comprehensive process mining workflow including data extraction, process discovery, clustering, explainability, and generative AI. Starting with raw event data extracted from an ERP system such as SAP, we generate an event log, capturing the sequence of activities for individual process instances. We apply process discovery algorithms to generate a process model of the event log and a clustering algorithm to group similar cases and variants. Then, we use our XAI explainers to generate rule-based explanations for the identified clusters. These explanations, alongside contextual background knowledge, are passed into a Retrieval-Augmented Generation (RAG) system, which leverages a Large Language Model (LLM) to generate natural language responses. The user interface consists of a dashboard that visualizes key metrics and cluster insights, complemented by a chatbot that allows users to interactively explore process models and request intuitive explanations in natural language, making process mining insights more accessible to users. The individual components are implemented and deployed as containerized microservices using docker.

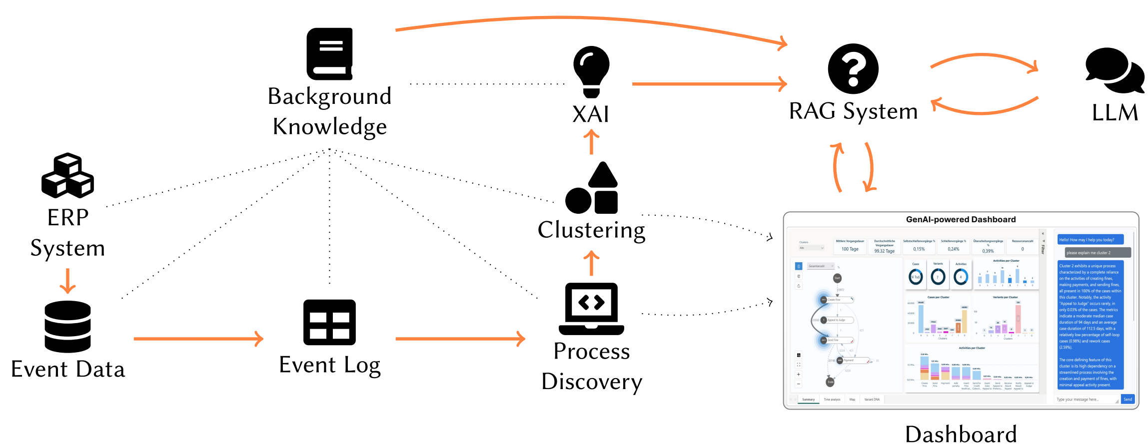

Figure 1: Architecture Diagram.

2.2. Components

2.2.1. Process Discovery

For process discovery, we employ the process mining tool of Microsoft's power automate suite[1]. When applied to the event data, the tool discovers a process model represented as a Directly-Follows Graph (DFG), reflecting the underlying structure of the event log.

2.2.2. Clustering

For clustering, we employ a state-of-the-art, density-based algorithm in combination with activity profiling and variant-based downsampling. Since clustering is performed at the variant-level, each case in the event log is first mapped to its corresponding process variant.

Variants are encoded as boolean vectors indicating whether specific constraints – such as activity presence or absence or temporal relationships – hold at the trace level. In order to obtain an accurate estimation of the underlying case-level distribution, we sample additional variants from the existing case-level distribution. Similarities between the sampled variants are captured by computing an inter-instance distance matrix. Next, the clustering algorithm is applied to this matrix, identifying meaningful clusters while filtering out noise. In the final step, the variant clusters are mapped to the corresponding cases in the original event log to enable a case-level analysis.

2.2.3. XAI for Clustering

We combine two complementary explainers: a domain-agnostic explainer for clustering based on ClusterExplainR [5], and a domain-specific explainer leveraging Declare constraints [14, 15]. Both aim to highlight the key characteristics that differentiate the clusters and run as containerized microservices. The ClusterExplainR computes local feature importance scores to derive rule-based descriptions based on the presence or absence of activities within a cluster. In contrast, the Declare-based Explainer formulates rules as Declare constraints, a temporal logic specialized for the process mining domain.

2.2.4. LLM / RAG

Our approach is designed to be agnostic regarding the specific LLM used. We choose GPT-4o-mini[2] to balance performance and resource consumption. To reduce hallucinations and ensure concise, reliable responses, we apply chain-of-thought prompting, which structures the interaction into redefined tasks for the LLM, each accompanied by explicit instructions.

Since the task definitions, contextual instructions, and core background knowledge are required for every interaction, they are provided through the system prompt. Contextual information incorporates essential background knowledge such as core concepts of process mining, explanations of key performance metrics and their importance, guidelines for structuring cluster-related information, and instructions for interpreting the XAI-generated rules associated with each cluster.

In contrast, task-specific information is retrieved using a Retrieval-Augmented Generation (RAG) approach. This information is encoded as JSON files, with their structure and keys documented in the system prompt to enable the LLM to access and navigate the data reliably. The task-specific knowledge includes precomputed data on activities within a cluster, associated metrics, and rankings. Additionally, we provide the corresponding rules from the explainers.

This approach enables the LLM to perform two core tasks: (i) Generate clear, natural language explanations of individual clusters and (ii) Compare multiple clusters to highlight key differences. It can also explain key metrics, provide a comprehensive overview of the process, and suggest meaningful cluster names. When explaining a cluster, the LLM highlights its unique characteristics through absent and present activities and their temporal relationships, supported by relevant metrics, and summarizes

[1]https://learn.microsoft.com/en-us/power-automate/process-mining-overview, 25 March 2025.
[2]https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/, accessed 2025-03-22.

the core defining features. For comparisons, the model outlines key differences in structure and performance, including activity patterns, metrics, and XAI rules, concluding with a brief summary of distinguishing factors. In the case of users requesting information beyond the LLM's scope, it is instructed to refrain from answering and instead clarify the set of tasks it is able to perform.

2.2.5. Dashboard

The dashboard shown in Fig. 2 combines an LLM-based chat interface (C) with interactive process-visualizations (A), displays key process metrics (B), and is implemented using Power BI.

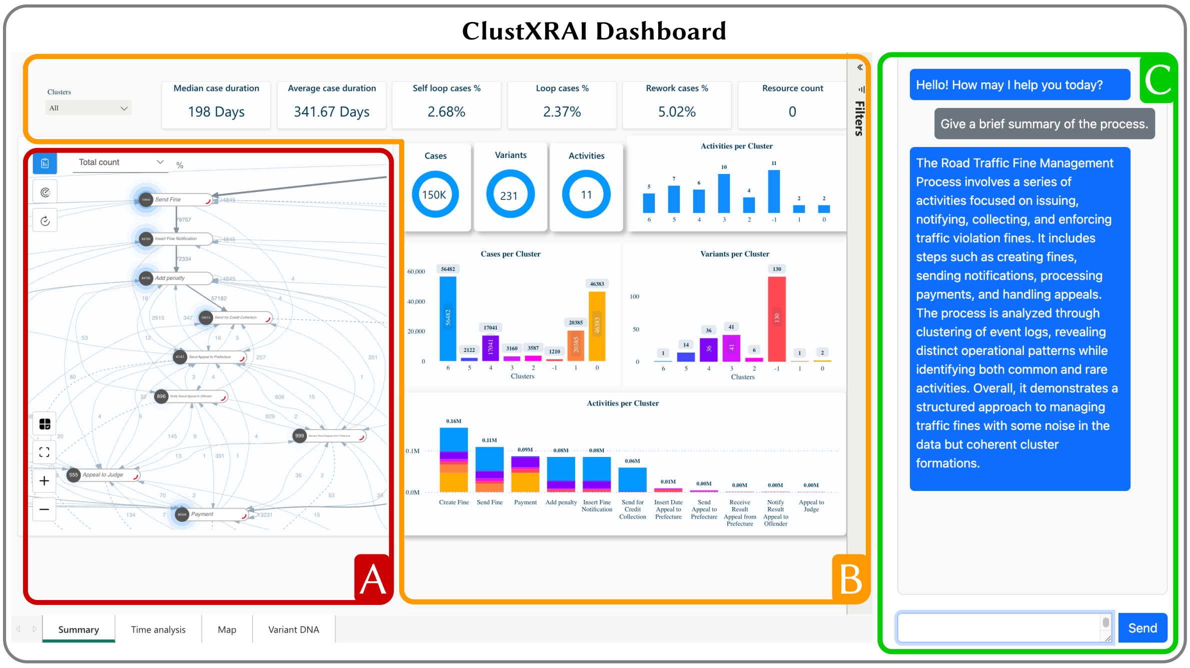

Figure 2: Power BI dashboard for Interactive Explanatory Process Mining.

At the center of the dashboard, users can see the results of the clustering (B). Six panels at the top display core metrics for both the entire dataset and the individual clusters, such as median case duration, average case duration, self-loop cases, loop cases, rework cases, and the number of resources used in the process. Below these panels, users can explore the clusters through multiple interactive bar charts, which provide visualizations of several key features. These include the number of cases per cluster (the total number of individual cases from the process log mapped to each cluster), the number of variants per cluster (the number of distinct case variants grouped within each cluster) as well as the number of activities per cluster (the number of unique activities within a given cluster). The activities are visualized in two complementary ways: In the upper-right section, a chart displays the count of distinct activities per cluster – for example, Cluster 2 includes four different activities. Meanwhile, the bar chart at the bottom visualizes the total occurrences of each activity across the entire event log. For instance, the activity "Create Fine" may appear 0.16 million times. The bars associated with each activity reveal the distribution of these occurrences across clusters, with percentages visually differentiated using color coding. Additionally, three distinct panels summarize key metrics: the total number of cases, variants, and activities.

Users can interactively explore the data by selecting a cluster by clicking on the corresponding bar in the graphs. All visualizations update dynamically to highlight information specific to the selected cluster, enabling users to quickly – and intuitively – grasp its key characteristics and metrics.

When a cluster is selected, the process model in panel (A) on the left is updated accordingly. By default, this panel displays a Directly-Follows Graph (DFG) representing the entire dataset (see Fig. 2 Panel (A)), as generated by process discovery. The nodes, representing activities, and the transitions between them are annotated with the total number of occurrences and transition frequencies, providing users with additional context to better understand the process flow. Upon selecting a specific cluster – or multiple clusters – from the chart, the model view updates to reflect the corresponding cluster(s) of the process. For example, Fig. 3 displays the combined process model for selected clusters 4 and 6. This allows users to explore the structure of individual subprocesses, identify overlapping behavior, and observe key differences in activity sequences between clusters.

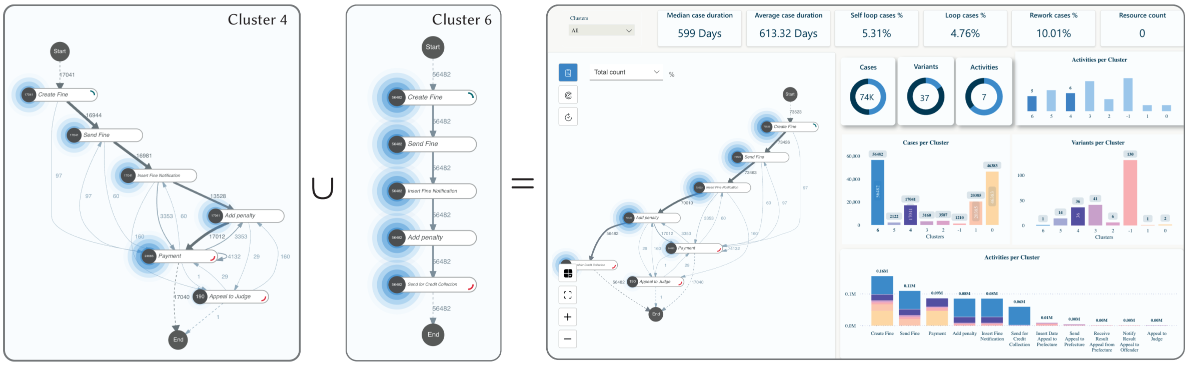

Figure 3: Combined Process Model of selected Clusters 4 and 6.

Complementing the interactive charts, the dashboard features a chat interface on the right in panel (C), which provides access to interactive LLM-generated explanations. This enables users to engage with the data through dialogue-based, natural language narratives, enhancing accessibility and interpretability.

3. Interactive Process Exploration: A Walkthrough of ClustXRAI

For the following walkthrough, we use the Road Traffic Fine Management Process (RTFMP) dataset[3], which is derived from an ERP system used by the Italian authorities. It captures the complete lifecycle of managing traffic fines: from issuing fines and notifying recipients to processing payments, handling appeals, and escalating unpaid cases to credit collection. This dataset contains a variety of distinct process paths, offering a rich yet manageable example for demonstrating interactive exploration and analysis. To highlight the most relevant aspects, we present selected extracts from screenshots that focus on key parts of the demonstrator. First, we ask for an overview of the dataset (Fig. 4).

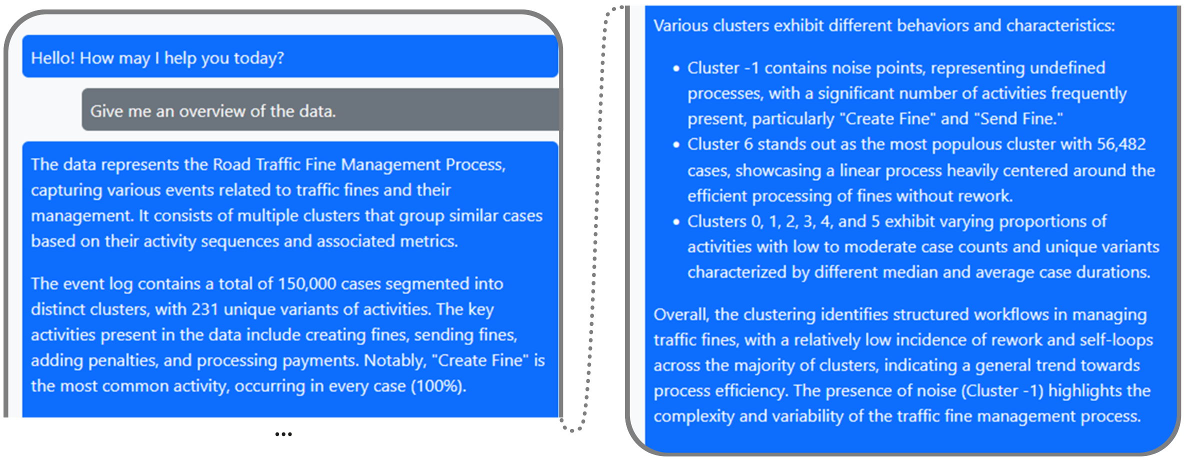

Figure 4: Overview of the data.

[3] https://data.4tu.nl/articles/dataset/Road_Traffic_Fine_Management_Process/12683249, accessed 25 March 2025.

Our next step is to identify the most important clusters. Which of the clusters warrant further investigation, and why? We can directly ask the LLM to dive into the data, shown in Fig. 5.

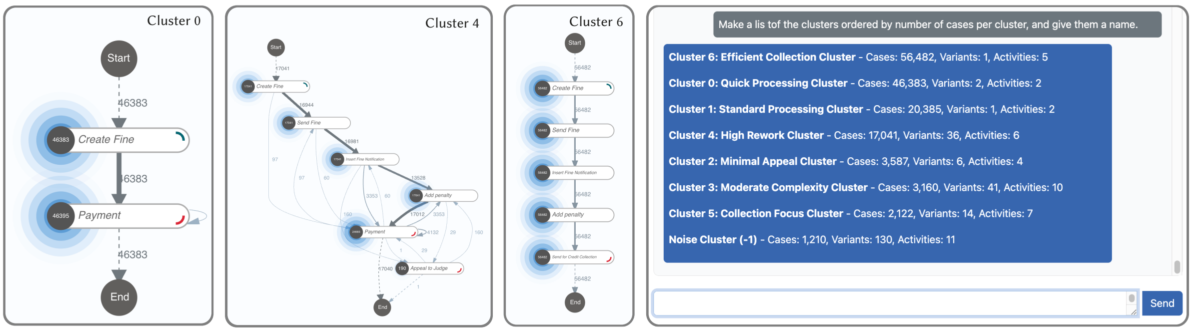

Figure 5: List of important clusters, alongside clusters 0, 4 and 6.

We receive a list of the most significant clusters, each with a short description of its key features. To explore further, we might ask: What defines a specific cluster? How does it differ from others? What are its key metrics? We continue by investigating Cluster 6, which covers a significant part of the data (Fig. 6).

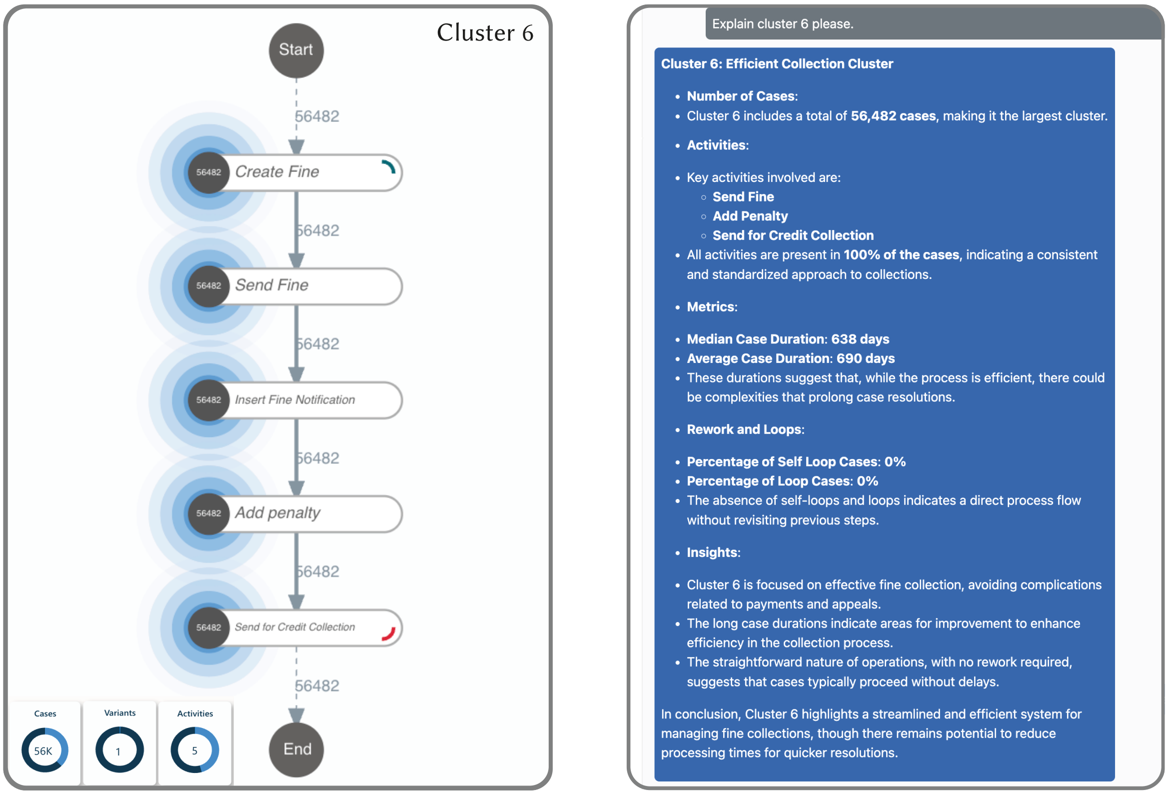

Figure 6: Directly-Follows Graph of cluster 6 and its explanation.

The responses provide detailed explanations of each cluster's activities, key metrics, and defining characteristics. Since Cluster 6 is described as streamlined, while Cluster 4 appears more complex and involves additional activities, it is useful to request a direct comparison as shown in Fig. 7.

The comparison points out that – in contrast to Cluster 6 – Cluster 4 involves different activities related to appeals and notifications and a notable share of rework activity, resulting in it being more complex and with potential room for optimization. As becomes apparent in this example, our approach enables users to interactively explore the clustered data without the need for domain-specific knowledge.

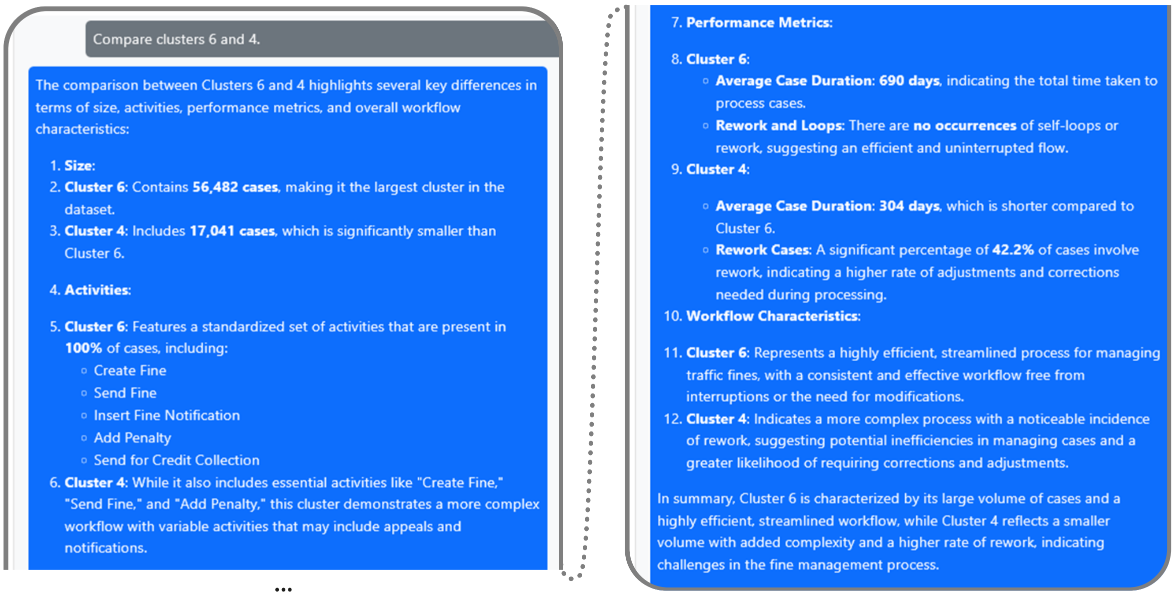

Figure 7: Comparison of cluster six and four.

Explanations and comparisons in natural language allow business analysts to intuitively learn about the clustered data and gain insights, making explainable process mining clustering accessible for end users.

4. Conclusion

Our demonstrator implements a complete workflow for interactive, explanatory clustering in process mining. Event Data is processed by a discovery algorithm to generate a baseline process model. A density-based clustering algorithm then identifies clusters and subprocesses. To explain these results, we integrate two XAI methods: ClusterExplainR, which highlights activity presence and absence, and DeclareExplainer, which derives declarative rules using Declare constraints. An interactive Power BI dashboard combines visualizations with an LLM-powered chatbot to support intuitive cluster exploration. Explanations from the XAI modules, along with contextual and task-specific background knowledge, are provided to the LLM using a RAG-based approach. This setup complements visual exploration – such as DFGs and cluster metrics – with natural language explanations of key patterns and differences. By improving interpretability, it helps non-specialists gain deeper process insights and identify opportunities for improvement.

For future work, we plan to incorporate formal ontologies to support rule-based reasoning and richer knowledge integration. We also aim to explore multi-dimensional clustering across behavioral, performance, and organizational aspects, such as object-centric perspectives. Furthermore, we plan to leverage LLM fine-tuning to improve explanation accuracy and relevance, supported by user studies.

Acknowledgments

Funded by the Bavarian Ministry of Economic Affairs (BayVFP) under KIGA: DIK0313.

Declaration on Generative AI

During the preparation of this work, the authors used Grammarly for spell checking. The chatbot responses in the Figures are generated using ChatGPT 4o-mini.

References

[1] W. van der Aalst, et al., Process mining manifesto, in: F. Daniel, K. Barkaoui, S. Dustdar (Eds.), Business Process Management Workshops, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 169–194. doi:10.1007/978-3-642-28108-2_19.

[2] G. M. Veiga, D. R. Ferreira, Understanding spaghetti models with sequence clustering for prom, in: S. Rinderle-Ma, S. Sadiq, F. Leymann (Eds.), Business Process Management Workshops, Springer Berlin Heidelberg, Berlin, Heidelberg, 2010, pp. 92–103. doi:10.1007/978-3-642-12186-9_10.

[3] T. S. Madhulatha, An overview on clustering methods, CoRR abs/1205.1117 (2012). doi:10.48550/arXiv.1205.1117.

[4] M. Song, C. W. Günther, W. M. P. van der Aalst, Trace clustering in process mining, in: D. Ardagna, M. Mecella, J. Yang (Eds.), Business Process Management Workshops, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 109–120. doi:10.1007/978-3-642-00328-8_11.

[5] J. Amling, S. Scheele, E. Slany, M. Lang, U. Schmid, Explainable ai for mixed data clustering, in: L. Longo, et al. (Eds.), Explainable Artificial Intelligence, Springer Nature Switzerland, Cham, 2024, pp. 42–62. doi:10.1007/978-3-031-63797-1_3.

[6] M. Alvarez-Garcia, R. Ibar-Alonso, M. Arenas-Parra, A comprehensive framework for explainable cluster analysis, Information Sciences 663 (2024) 120282. doi:10.1016/j.ins.2024.120282.

[7] I. Grau, L. Hernandez, A. Sierens, S. Michel, N. Sergeyssels, C. Middag, V. Froyen, A. Now'e, Talking to your data: Interactive and interpretable data mining through a conversational agent, in: 33rd Benelux Conference on Artificial Intelligence and 30th Belgian-Dutch Conference on Machine Learning, BNAIC/BeneLearn 2021, University of Luxembourg, 2021, pp. 745–747.

[8] N. Mehdiyev, P. Fettke, Explainable Artificial Intelligence for Process Mining: A General Overview and Application of a Novel Local Explanation Approach for Predictive Process Monitoring, Springer International Publishing, Cham, 2021, pp. 1–28. doi:10.1007/978-3-030-64949-4_1.

[9] A. Buliga, M. Vazifehdoostirani, L. Genga, X. Lu, R. Dijkman, C. Di Francescomarino, C. Ghidini, H. A. Reijers, Uncovering patterns for local explanations in outcome-based predictive process monitoring, in: A. Marrella, M. Resinas, M. Jans, M. Rosemann (Eds.), Business Process Management, Springer Nature Switzerland, Cham, 2024, pp. 363–380. doi:10.1007/978-3-031-70396-6_21.

[10] D. Fahland, F. Fournier, L. Limonad, I. Skarbovsky, A. J. Swevels, How well can a large language model explain business processes as perceived by users?, Data & Knowledge Engineering 157 (2025) 102416. doi:10.1016/j.datak.2025.102416.

[11] J. N. Adams, S. J. van Zelst, T. Rose, W. M. van der Aalst, Explainable concept drift in process mining, Information Systems 114 (2023) 102177. doi:10.1016/j.is.2023.102177.

[12] K. Kubrak, L. Botchorishvili, F. Milani, A. Nolte, M. Dumas, Explanatory capabilities of large language models in prescriptive process monitoring, in: A. Marrella, M. Resinas, M. Jans, M. Rosemann (Eds.), Business Process Management, Springer Nature Switzerland, Cham, 2024, pp. 403–420. doi:10.1007/978-3-031-70396-6_23.

[13] J. Amling, S. Scheele, E. Slany, M. Lang, U. Schmid, Bridging the interpretability gap in process mining: A comprehensive approach combining explainable clustering and generative ai, in: (forthcoming) 3rd World Conference on Explainable Artificial Intelligence, 09-11 July, 2025.

[14] C. Di Ciccio, M. Montali, Declarative Process Specifications: Reasoning, Discovery, Monitoring, Springer International Publishing, Cham, 2022, pp. 108–152. doi:10.1007/978-3-031-08848-3_4.

[15] M. Pesic, H. Schonenberg, W. M. van der Aalst, Declare: Full support for loosely-structured processes, in: 11th IEEE International Enterprise Distributed Object Computing Conference (EDOC 2007), 2007, pp. 287–287. doi:10.1109/EDOC.2007.14.