Coordinated Projections: A New Approach to Multi-Faceted Process Exploration

Stef van den Elzen[1], Gennady Andrienko[2,3], Andreas Kerren[4,5], Manuel Resinas[6], Barbara Weber[7], and Peilin Yu[8]

[1] Eindhoven University of Technology, The Netherlands s.j.v.d.elzen@tue.nl
[2] Fraunhofer Institute IAIS, Germany gennady.andrienko@iais.fraunhofer.de
[3] City St.George's University of London, UK gennady.andrienko.1@city.ac.uk
[4] Linköping University, Sweden andreas.kerren@liu.se
[5] Linnaeus University, Sweden andreas.kerren@lnu.se
[6] Universidad de Sevilla, Spain resinas@us.es
[7] University of St. Gallen, Switzerland barbara.weber@unisg.ch
[8] Linköping University, Sweden peilin.yu@liu.se

Abstract. Process exploration is a critical task in process mining, enabling analysts to uncover insights and generate hypotheses about process behavior from event log data. Traditional approaches often rely on static, single-facet (control-flow) visualizations, such as Directly-Follows Graphs, that limit flexibility and obscure multi-perspective dependencies. In this paper, we introduce a novel visual analytics approach that leverages coordinated projections to support multi-faceted process exploration. Through dimensionality reduction techniques (e.g., UMAP, t-SNE) and topic modeling, our method generates coordinated views that dynamically link different facets of process data, enabling cross-perspective exploration. Interaction techniques, such as zooming-panning, brushing-and-linking, and attribute color, shape, and size encodings further enhance the analyst's ability to correlate patterns across dimensions. We demonstrate our approach using a real-world traffic fines event log.

1 Motivation

Process exploration constitutes a fundamental task in process mining, primarily aimed at facilitating the exploration and generation of hypotheses about the underlying process behavior captured in event log data. In particular, during the initial phases of process mining projects, analysts engage in various exploratory activities—they dedicate time to familiarize themselves with the data, develop a preliminary understanding of the process, formulate or refine analytical questions, and uncover unexpected patterns or insights [35]. This iterative exploration plays a crucial role in shaping the direction of subsequent hypothesis testing [24].

The way process exploration is currently performed typically involves the use of discovered *Process Maps*, encoded through techniques such as the Directly-Follows Graphs (DFGs) [36,5,16], where nodes represent activities and edges indicate direct successions based on event log timestamps. Using the DFG as an encoding and visual representation to enable exploration generally implies that the primary facet (i.e., the first-class citizen of our analytical workflow) is the temporal order of activities [29,27].

More details are often represented with additional visual channels, such as projecting performance metrics (e.g., activity duration and waiting time) or frequency information (e.g., how often transitions occur) on top of the DFG structure. These additional visual cues enrich the visualization but do not alter the primary facet and encoding. This enables the analyst to gain insight into different aspects of the process, such as identifying activities (e.g., activities with high duration, infrequent activities, endpoint activities), fragments (e.g., the most frequent fragment), transitions (e.g., transitions with high durations), or bottlenecks [21]. Less commonly, the DFG can be restructured entirely by using different encodings that bring forward resources as the primary facet [23]. In resource-centric DFGs, nodes represent individuals, roles, or organizational units, and the edges reflect handovers or collaborative interactions. This shift enables analysis of organizational dynamics, revealing patterns such as teamwork structures, silos, or handover inefficiencies.

Despite the value of these perspectives, a major limitation is that the primary facet is typically fixed, constraining the flexibility of exploration. Real-world analytical tasks often require users to continuously shift between facets [16]—from case-centric to activity-centric, from control flow to resource flow, or from process variants to attribute analysis to discover inter- and intra-dependencies [30]. Fixing the facet can obscure relevant patterns, reduce user agency, and lead to a cognitive mismatch when visualizations do not align with the user's current task or mental model [14].

To truly support hypothesis generation and sense-making, process exploration tools must support flexible encodings and transitions between facets, allowing analysts to dynamically reframe and encode the data perspective depending on their line of inquiry [16,15]. Typically, users are interested in exploring *similarities* and *differences* between (combinations of) facets, i.e., the visual analysis approach should support more complex comparison tasks. To support this, dimensionality reduction might help here. Dimensionality reduction (DR) is a class of techniques aimed at transforming high-dimensional data into a lower-dimensional representation while preserving its structure and relationships as much as possible. DR enables the exploration of complex datasets by revealing latent structures that are otherwise difficult to discern in their original high-dimensional form; it is flexible in what facets are considered and helps to reveal similarities (local neighborhoods) in the encoded high-dimensional space. As a first step, we explore combining multiple encodings for the same trace, for example, one based on control-flow and another on case-level features or data-level features. This aligns with the principles of multi-perspective process mining, enabling analysts to explore not only what happens in a process (via control-flow) but also how and why it happens (via data attributes, timing, or resources). This approach facilitates deeper insights compared to relying on a single view [27].

2 Background and Related Work

Our work is closely related to prior research on trace encodings in the field of process mining. Trace encoding refers to the transformation of process traces (i.e., sequences of events or activities) into structured representations that can be used for a variety of analytical tasks. In process mining, such encodings are essential for enabling techniques such as clustering [2], anomaly detection [31], predictive process monitoring [18], classification, [20], and exploratory visualization [22,35]. While the development of trace encodings has been extensively studied, their application in the context of advanced, coordinated visualizations of process data has received comparatively limited attention.

Broadly, existing encodings can be categorized into three main types: 1) control-flow encodings, 2) data-aware encodings, and 3) embedding-based encodings. *Control-flow encodings* capture the order and structure of activities, often using n-grams, bags-of-activities, or transition abstractions (e.g., [4]). *Data-aware encodings* enrich the trace representation by incorporating contextual information such as case attributes, event-level data, or resource involvement (e.g., [12]). Time performance information is also typically included by computing derived attributes such as elapsed time since the previous event or elapsed time since the beginning of the case [20]. *Embedding-based encodings* leverage deep learning models—such as LSTMs, Transformers, or autoencoders—to learn dense vector representations of traces (e.g., [6,18]). A comprehensive benchmarking of trace encodings with a focus on control-flow can be found in [25], while multi-perspective encodings for classification, clustering, and anomaly detection tasks are evaluated in [20].

To support the exploration of these multiple perspectives, we employ UMAP and t-SNE, two widely-used non-linear DR techniques for visualizing complex high-dimensional data [8]. In our approach, these methods are used to construct coordinated views, where clusters identified in one projection (e.g., based on control-flow encodings) can be interactively examined in another projection (e.g., based on data attributes). This setup enables linked, cross-perspective exploration, allowing users to discover and analyze patterns that may not be visible when perspectives are viewed in isolation.

This interactive, user-centered approach aligns with recent developments in visual analytics, which aim to augment algorithms with exploratory visualization and human reasoning. Integrating visual analytics with process mining has emerged as a powerful approach to address the challenges posed by complex event data. The potential of visual analytics, combining visualizations, interactions, and mining techniques, to enhance process mining has been emphasized as a key opportunity to advance the field [28]. The synergy between human perception and computational power facilitates deeper insights into process data [10]. Unlike traditional process mining approaches that focus predominantly on algorithmic output, visual analytics emphasizes interactive visualizations that enable users to iteratively explore, refine, and interpret complex processes.

From a visualization research perspective, combining multiple encodings of the same trace aligns with established principles in visual analytics, which em-

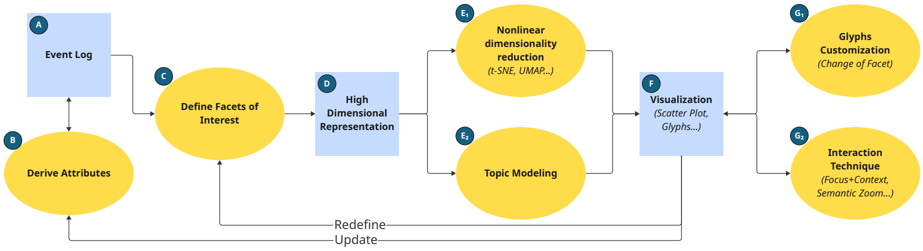

Fig. 1: Multi-Faceted Process Exploration approach going from A) event log data with B) derived attributes, to user-defined facets of interest C). These facets are then represented in high-dimensional space through different encodings D). From the high-dimensional representations, we apply DR E_1) or topic modeling E_2) for visualization purposes F). Interaction techniques and the use of glyphs G) enable exploration and analysis to close the sense-making loop.

phasize the value of coordinated views and multi-perspective representations to enhance analytical reasoning [19]. By integrating control-flow with case- and data-level features, such approaches support multi-perspective analysis, more nuanced exploration and interpretation of process behavior, fostering deeper insights through visual comparison and contextualization [1,27].

3 Multi-Faceted Process Exploration Approach

To enable detailed analysis of event logs through more interpretable subsets, we adopt a multi-faceted decomposition approach [14]. Each resulting subset can be further examined using compact and comprehensible representations, such as directly-follows graphs (DFGs). Our methodology integrates both topic modeling [3] and embedding-based techniques [11][9], by transforming event logs into either textual documents or structured feature vectors. We employ representations that capture event logs both as collections of activities and as sequences of activity transitions. These representations are designed to reflect the temporal structure of events, including activity order, durations, and other contextual attributes. Additionally, we leverage multiple coordinated views to support interactive exploration and to facilitate a more comprehensive understanding of the decomposed event log subsets.

The result of these investigations is summarized in the workflow presented in Fig. 1. In the following, we describe each of the steps involved in the workflow.

A. Original Data Set The publicly available Road Traffic Fine event log is used as an example [7]. The Road Traffic Fines event log documents the handling of traffic fines by a local police force in Italy. It contains approximately 561,470 events across 150,370 cases, recorded between January 2000 and June 2013. The process involves 11 activities and 12 data attributes.

---

[9] `https://va-embeddings-browser.ivis.itn.liu.se`

Each case starts with a *Create Fine* event, which includes the fine amount next to other attributes. The offender can pay the fine at any time via a *Payment* event. The amount paid is recorded in the attribute *paymentAmount*. If not paid, a *Send Fine* action sends a letter, possibly incurring additional charges (*expenses*). This is followed by *Insert Fine Notification* and, if necessary, *Add Penalty*, which increases the amount due. If the fine remains unpaid, an event *Send for Credit Collection* indicates an escalation to a collection agency. Offenders may also appeal to the prefecture or a judge, triggering events such as *Insert Appeal* and *Notify Result Appeal to Offender*. If an appeal is successful, the fine is marked as dismissed via the *dismissal* attribute.

This dataset was selected because it supports exploration across multiple meaningful facets. It includes a rich combination of control-flow, temporal, and data perspectives. Events are timestamped, enabling the discovery of relevant temporal constraints (e.g., the fine notification must be sent within 90 days of fine creation; otherwise, there is no obligation to pay). In addition, the event log captures various attributes, e.g., *amount*, *expense*, *totalPaymentAmount*, and appeal results (attribute *dismissal* containing a flag whether and by whom the fine is dismissed). These attributes enable the identification of patterns within specific subpopulations of cases and help correlate behavioral differences with underlying data characteristics. This multidimensionality makes the dataset particularly well-suited for studying the need for flexible faceting and supporting dynamic transitions between different analytical perspectives during process exploration.

B. From Event Log to Enriched Case Log Next, the event log is transformed into a case log and enriched using case predicates as suggested in [32]. More specifically, the case log is enriched with process outcomes. For this, each case was assigned a specific outcome of the process (that is, fully paid, dismissed, credit collected, and unresolved) according to [32]. Fully paid cases are cases where the last outstanding balance is ≤ 0. Dismissed cases are cases with the dismissal code $\in \{\#, G\}$. Cases are assigned the label credit collected if the activity *Send for Credit Collection* is present in the trace. The remaining cases were then classified as unresolved. Moreover, the value of the dismissal attribute of the last activity of the case was added to the enriched case log, as well as a derived attribute *outstandingBalance*, which is calculated as the sum of *amounts* plus the sum of expenses minus *totalPaymentAmount*.

C. Define Facets of Interest To explore the facets of interest within the process data, we employ DR and topic modeling techniques to project high-dimensional attribute representations into a two-dimensional space suitable for visualization (step E). This enables us to generate scatterplots in which each point corresponds to a specific process element, such as a case, activity, variant, or resource, depending on the analytical perspective adopted (e.g., similar to earlier work on dynamic network exploration [26]). In this step, users determine what process elements and associated attributes are relevant and how to encode these to serve as an input to the DR. The goal of this step is therefore to determine which process elements will be represented as individual points and which attributes will define their position in the projection space (i.e., what encoding is selected for the DR). These attributes are selected based on the analytical goal and may capture various process dimensions. These include control-flow characteristics (e.g., activity sequences, frequency patterns), outcome-related indicators (e.g., outstanding balance, dismissal codes), contextual attributes (e.g., vehicle class, notification type), temporal aspects (e.g., activity duration, case throughput time), or process variants (e.g., distinct execution paths). The outcome of this step is an encoding of the facets of interest that serves as an input vector to the DR. Each input vector is represented as a point in the visualization (step F). In this work, as a first exploration, we focus exclusively on representing cases as points in the scatterplot, using attributes related to control flow and outcome indicators to construct the projections.

D. High Dimensional Representation—Multi-faceted Trace Encoding Based on the selected facets of interest, the original and enriched case log are used to implement different alternative encodings in a high-dimensional space. These encodings can next be used as input for the topic modeling and DR step.

For topic modeling, representing event logs as sets of activities and as sets of direct transitions between the activities are considered. For example, having an event log of consecutive activities A, B, and C, we represent either as a string A B C or as a string A_B B_C, reflecting activities and their transitions accordingly. The third variant combines both representations, uniting two alternative viewpoints A B C A_B B_C.

For DR, two different encodings are constructed that cover different facets of the dataset. The first encoding focuses on attributes deemed relevant to the outcome of each case [32]. Specifically, two attributes are included to represent whether the *outstandingBalance* and the *totalPaymentAmount* are greater than zero. Furthermore, the *dismissal* code is included from the enriched case log and the last activity recorded in each case, since it helps to determine the outcome of the case if it goes to credit collection. Since both the dismissal code and the last activity are categorical attributes, one-hot encodings are applied to enable the application of a DR algorithm in the next step.

The second encoding is designed to capture the sequence of activities performed within each case. For this purpose, the sequence of activities for every case is extracted while preserving their order of execution. Since cases can vary in the number of activities, shorter sequences are padded with a special padding token to ensure all sequences are of equal length, matching the case with the highest number of activities. This pre-processing step results in an $n \times m$ matrix, where n is the total number of cases, and m represents the length of the longest case. Finally, one-hot encoding is applied to each column of the matrix, expanding the $n \times m$ matrix into a higher-dimensional format where categorical activity labels are represented in a numerical form.

E. Non-Linear Dimensionality Reduction & Topic Modeling Non-linear DR (i.e., UMAP and t-SNE [8]) and topic modeling techniques are applied to the multi-faceted trace encodings described above. DR and topic modeling complement each other by capturing different aspects. Topic modeling reveals the underlying *semantic themes*, while DR projects high-dimensional represen-

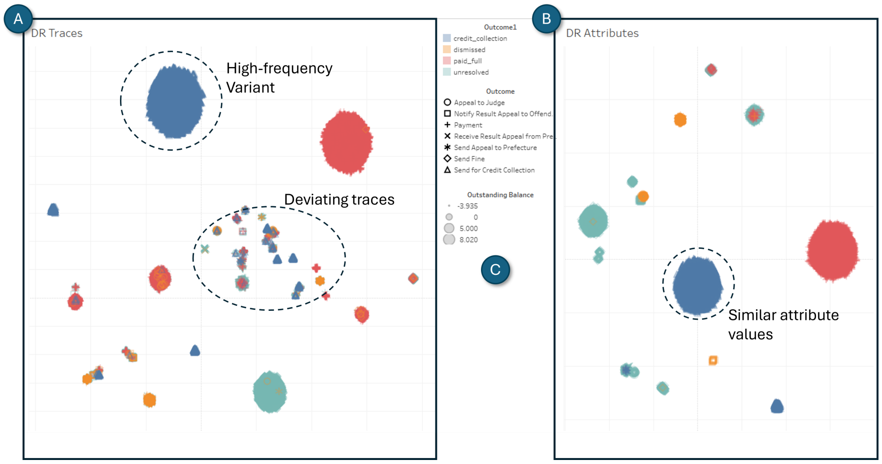

Fig. 2: Multiple coordinated view of two scatterplots showing the DR results of A) traces and B) attributes (outcome, outstanding balance, payment amount, dismissal). C) An interactive legend enables filtering. Dots in the scatterplot represent cases (i.e., traces). Jittering is used to prevent overplotting. If two dots are close, it means the traces are similar; in A) clusters of similar traces representing variants emerge (e.g., high-frequency variant) and smaller clusters of deviating traces. In B) clusters of similar attribute values emerge. Interaction enables the exploration and discovery of interactions between both representations.

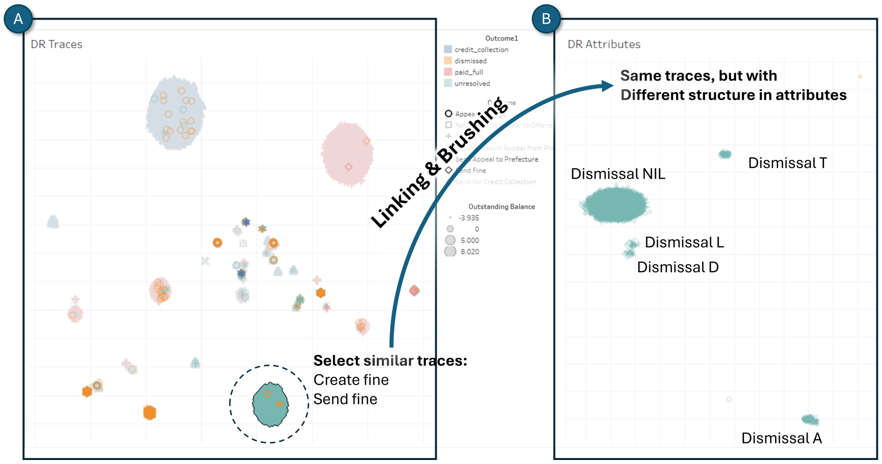

Fig. 3: Two-way brushing and linking enables multi-faceted exploration of both in the context of one another. For example, selecting a cluster of similar traces {create fine, send fine} in A) shows the corresponding structure of attribute clusters in B) formed by dismissal values.

tations into a compact space that preserves *similarity.* Combined, they provide both semantic structure and visual exploration.

F. Visualization and Interaction To support the exploration of the DR results, a coordinated multiple view approach [19] is used in which two or more visualizations are connected through brushing and linking (see Fig. 2) (i.e., if items in one visualization are highlighted or selected, the corresponding items in the other visualization are also highlighted or selected). This enables users to explore the different data facets in the context of each other, see Fig. 3 (e.g., explore the correlation between temporal order and data attributes). For the visualization of the DR result, a proven workhorse is a scatterplot-like representation. This has the advantage that the main facet of interest is encoded with the highest-ranked visual channel of position [17]. To encode the additional facets, additional visual channels can be used, such as color, shape, and size. However, as these channels are limited, we propose using glyphs (cf. the next paragraph) to encode the additional facets. Furthermore, the scatterplot also enables a scalable solution with respect to the number of items, as *jittering* and various interaction techniques can help to alleviate overplotting issues. For example, *focus+context* allows users to closely inspect a subset of the data while still maintaining an overview of the entire dataset; and *semantic zooming* provides different levels of detail depending on the zoom level.

Additionally, topic modeling (Fig. 4) is applied based on the combined representation of traces as sets of activities and transitions between activities. To validate the results, three UMAP projections (Table 1) are created of all traces based on the weights of topics, orders, and attributes of the logs. These projections have been colored according to outcomes of the traces (3rd column), main topic (fourth column), and also propagated continuous 2D color schemes (shown in the 2nd column) across the three projections (columns 5-7). By applying topic modeling, we obtain results that are consistent with those produced by alternative projection methods based on the ordering of activities and trace attributes. Notably, the use of topic modeling enables a substantial reduction in data dimensionality while also enhancing the interpretability of the results, as the identified topics can be characterized by their distinctive patterns of frequently occurring activities and transitions.

G. Glyphs Currently, through the DR, only the local distances between points are meaningful. However, there is no global axis explanation. To improve the interpretation of the additional facets, we propose replacing the dots with glyphs. Prior research demonstrates that glyph-based representations can be employed across diverse application domains and serve various analytical and communicative purposes. Various design alternatives of the most commonly used glyph types have been examined and discussed in numerous prior studies [33,9,13].

Multiple data attributes can be encoded within a single glyph to represent multiple properties of a single entity or aggregate information across multiple entities. Additionally, certain glyph types (e.g., face-based and icon-based representations) typically exhibit a one-to-one mapping between glyph and data entity. Different glyph types can be combined to encode more complex data facets. Note that in our prototype experiments, we used additional visual channels such as color, shape, and size to enhance interpretation. We leave further exploration and design of effective multi-facet glyphs for future work.

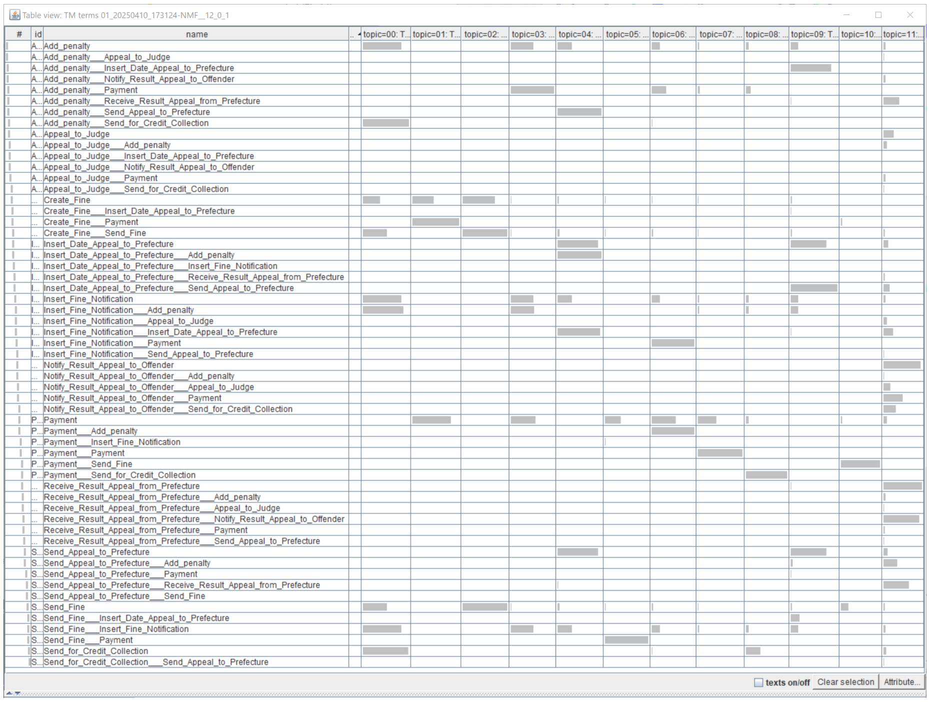

#	id	name
	A...	Add_penalty
	A...	Add_penalty___Appeal_to_Judge
	A...	Add_penalty___Insert_Date_Appeal_to_Prefecture
	A...	Add_penalty___Notify_Result_Appeal_to_Offender
	A...	Add_penalty___Payment
	A...	Add_penalty___Receive_Result_Appeal_from_Prefecture
	A...	Add_penalty___Send_Appeal_to_Prefecture
	A...	Add_penalty___Send_for_Credit_Collection
	A...	Appeal_to_Judge
	A...	Appeal_to_Judge___Add_penalty
	A...	Appeal_to_Judge___Insert_Date_Appeal_to_Prefecture
	A...	Appeal_to_Judge___Notify_Result_Appeal_to_Offender
	A...	Appeal_to_Judge___Payment
	A...	Appeal_to_Judge___Send_for_Credit_Collection
	...	Create_Fine
	...	Create_Fine___Insert_Date_Appeal_to_Prefecture
	...	Create_Fine___Payment
	...	Create_Fine___Send_Fine
	I...	Insert_Date_Appeal_to_Prefecture
	I...	Insert_Date_Appeal_to_Prefecture___Add_penalty
	I...	Insert_Date_Appeal_to_Prefecture___Insert_Fine_Notification
	I...	Insert_Date_Appeal_to_Prefecture___Receive_Result_Appeal_from_Prefecture
	I...	Insert_Date_Appeal_to_Prefecture___Send_Appeal_to_Prefecture
	I...	Insert_Fine_Notification
	I...	Insert_Fine_Notification___Add_penalty
	I...	Insert_Fine_Notification___Appeal_to_Judge
	I...	Insert_Fine_Notification___Insert_Date_Appeal_to_Prefecture
	I...	Insert_Fine_Notification___Payment
	I...	Insert_Fine_Notification___Send_Appeal_to_Prefecture
	...	Notify_Result_Appeal_to_Offender
	...	Notify_Result_Appeal_to_Offender___Add_penalty
	...	Notify_Result_Appeal_to_Offender___Appeal_to_Judge
	...	Notify_Result_Appeal_to_Offender___Payment
	...	Notify_Result_Appeal_to_Offender___Send_for_Credit_Collection
	P...	Payment
	P...	Payment___Add_penalty
	P...	Payment___Insert_Fine_Notification
	P...	Payment___Payment
	P...	Payment___Send_Fine
	P...	Payment___Send_for_Credit_Collection
	...	Receive_Result_Appeal_from_Prefecture
	...	Receive_Result_Appeal_from_Prefecture___Add_penalty
	...	Receive_Result_Appeal_from_Prefecture___Appeal_to_Judge
	...	Receive_Result_Appeal_from_Prefecture___Notify_Result_Appeal_to_Offender
	...	Receive_Result_Appeal_from_Prefecture___Payment
	...	Receive_Result_Appeal_from_Prefecture___Send_Appeal_to_Prefecture
	S...	Send_Appeal_to_Prefecture
	S...	Send_Appeal_to_Prefecture___Add_penalty
	S...	Send_Appeal_to_Prefecture___Payment
	S...	Send_Appeal_to_Prefecture___Receive_Result_Appeal_from_Prefecture
	S...	Send_Appeal_to_Prefecture___Send_Fine
	S...	Send_Fine
	S...	Send_Fine___Insert_Date_Appeal_to_Prefecture
	S...	Send_Fine___Insert_Fine_Notification
	S...	Send_Fine___Payment
	S...	Send_for_Credit_Collection
	S...	Send_for_Credit_Collection___Send_Appeal_to_Prefecture

Fig. 4: Table with bar charts demonstrates the composition of topics over terms, with bar charts representing term weights in the topics.

Base File	-	-outcome	-mainTopic	-topics	-order	-attrs
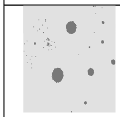	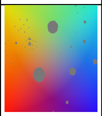	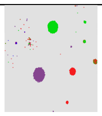	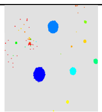		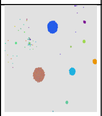	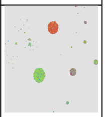
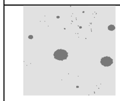	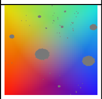	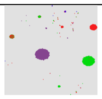	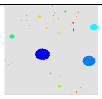	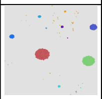		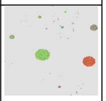
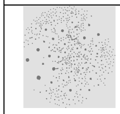	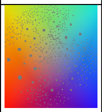	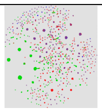	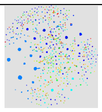	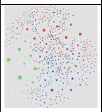	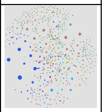	

Table 1: Results of UMAP embedding based on topics (top row), order (middle row), and attributes (bottom row).

4 Summary and Outlook

In summary, our work contributes a novel visual analytics methodology for multi-faceted process exploration that moves beyond the constraints of fixed-facet representations such as traditional DFGs. By integrating DR and topic modeling techniques, we generate coordinated views that represent different facets of event log data, enabling analysts to explore structural, behavioral, and contextual patterns in a unified framework. Our approach supports dynamic transitions between perspectives through interaction techniques like brushing-and-linking and glyph-based visual encodings, empowering users to detect similarities, differences, and dependencies across process facets. Through this combination of flexible representation, coordinated visualization, and interactive exploration, we provide a foundation for richer, more adaptive hypothesis generation and sense-making in process mining.

As part of future work, we plan to implement a comprehensive visual analytics system that builds on our current approach. In the current system design, each dot in the visualization represents a case positioned using DR techniques. Moving forward, we aim to explore alternative semantic representations for dots, such as activities or process variants. We would also like to systematically investigate how to encode activities, variants, and traces in a multi-faceted manner to best support various exploratory goals. To further enrich the visual expressiveness, we plan to integrate glyph-based representations that convey additional attributes or contextual cues (e.g., for selected subpopulations). Moreover, we intend to integrate our approach with existing process discovery algorithms to support the creation of DFGs for selected subpopulations. Additionally, we intend to expand the visual space to support additional facets, including relationships, control-flow, and resources. Another interesting research direction is to study ensemble methods for embeddings where various embedding approaches might be combined (either conceptually different state-of-the-art embedding technologies or the same embedding algorithm with various hyperparameter settings) to provide better performance [34]. A key objective of our future work will be to evaluate the effectiveness and scalability of the proposed system in supporting exploratory process mining, demonstrating its ability to accommodate a wide range of analytical tasks and user needs.

References

1. W. Aigner, S. Miksch, H. Schumann, and C. Tominski. *Visualization of time-oriented data*, volume 4. Springer, 2011.
2. Y. Bertrand, J. De Weerdt, and E. Serral. A novel multi-perspective trace clustering technique for iot-enhanced processes: A case study in smart manufacturing. In C. Di Francescomarino, A. Burattin, C. Janiesch, and S. Sadiq, editors, *Business Process Management*, pages 395–412. Springer Nature Switzerland, 2023.
3. D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *J. Machine Learning Research*, 3:993–1022, 2003.

4. R. P. J. C. Bose and W. M. P. van der Aalst. Context aware trace clustering: Towards improving process mining results. In *Proc. Int. Conf. Data Mining, SDM*, pages 401–412. SIAM, 2009.
5. C. Capitán-Agudo, M. Salas-Urbano, C. Cabanillas, and M. Resinas. Analyzing how process mining reports answer time performance questions. In *Business Process Management*, LNCS, pages 234–250. Springer Int. Publishing, 2022.
6. P. De Koninck, S. vanden Broucke, and J. De Weerdt. act2vec, trace2vec, log2vec, and model2vec: Representation Learning for Business Processes. In M. Weske, M. Montali, I. Weber, and J. vom Brocke, editors, *Business Process Management*, pages 305–321, Cham, 2018. Springer Int. Publishing.
7. M. de Leoni and F. Mannhardt. Road traffic fine management process. Eindhoven University of Technology, Dataset, 2015. Dataset ID 284.
8. M. Espadoto, R. M. Martins, A. Kerren, N. S. T. Hirata, and A. C. Telea. Toward a quantitative survey of dimension reduction techniques. *IEEE Trans. Vis. Comp. Graph.*, 27(3):2153–2173, 2021.
9. J. Fuchs, P. Isenberg, A. Bezerianos, and D. Keim. A systematic review of experimental studies on data glyphs. *IEEE Trans. Vis. Comp. Graph.*, 23(7):1863–1879, 2016.
10. T. Gschwandtner. Visual analytics meets process mining: Challenges and opportunities. In *Data-Driven Process Discovery and Analysis*, pages 142–154. Springer Int. Publishing, 2017.
11. Z. Huang, D. Witschard, K. Kucher, and A. Kerren. VA + Embeddings STAR: A State-of-the-Art Report on the Use of Embeddings in Visual Analytics. *CGF*, 42(3):539–571, 2023.
12. S. Jablonski, M. Röglinger, S. Schönig, and K. M. Wyrtki. Multi-perspective clustering of process execution traces. *Enterp. Model. Inf. Syst. Archit. Int. J. Concept. Model.*, 14:2:1–2:22, 2019.
13. M. Keck, D. Kammer, T. Gründer, T. Thom, M. Kleinsteuber, A. Maasch, and R. Groh. Towards glyph-based visualizations for big data clustering. In *Proc. 10th Int. Symp. Visual Information Comm. & Interaction*, pages 129–136, 2017.
14. J. Kehrer and H. Hauser. Visualization and visual analysis of multifaceted scientific data: A survey. *IEEE Trans. Vis. Comp. Graph.*, 19(3):495–513, 2013.
15. S. v. d. Linden, V. A. Filipov, L. Pufahl, S. Miksch, and S. v. d. Elzen. Towards Integrating Visual Analytics in Multi-Perspective Conformance Checking: A Call to Action. In H.-J. Schulz and A. Villanova, editors, *EuroVis Workshop on Visual Analytics (EuroVA)*. The Eurographics Association, 2025.
16. S. Miksch, C. D. Ciccio, P. Soffer, and B. Weber. Visual analytics meets process mining: Challenges and opportunities. *IEEE Comp. Graph. Appl.*, 44(6):132–141, 2024.
17. T. Munzner. *Visualization Analysis and Design.* AK Peters Visualization Series, CRC Press, 2015.
18. E. Rama-Maneiro, J. C. Vidal, and M. Lama. Deep Learning for Predictive Business Process Monitoring: Review and Benchmark. *IEEE Trans. Services Computing*, 16(01):739–756, Jan. 2023.
19. J. C. Roberts. State of the art: Coordinated & multiple views in exploratory visualization. In *Fifth Int. Conf. Coordinated and Multiple Views in Exploratory Visualization*, pages 61–71, 2007.
20. A. Rullo, F. Alam, and E. Serra. Trace encoding techniques for multi-perspective process mining: A comparative study. *WIREs Data. Mining. Knowl. Discov.*, 15(1), 2025.

21. M. Salas-Urbano, C. Capitán-Agudo, C. Cabanillas, and M. Resinas. Designing a user interface to explore collections of directly-follows graphs for process mining analysis. In *25th Int. Conf. Enterprise, Business-Process and Information Systems Modeling*, volume 511 of *LNBIP*, pages 35–47. Springer, 2024.
22. D. Schuster, L. Schade, S. J. van Zelst, and W. M. P. van der Aalst. Visualizing trace variants from partially ordered event data. In J. Munoz-Gama and X. Lu, editors, *Process Mining Workshops*, pages 34–46. Springer Int. Publishing, 2022.
23. M. Song and W. M. P. van der Aalst. Towards comprehensive support for organizational mining. *Decision Support Systems*, 46(1):300–317, 2008.
24. E. Sorokina, P. Soffer, I. Hadar, U. Leron, F. Zerbato, and B. Weber. PEM4PPM: A Cognitive Perspective on the Process of Process Mining. In C. D. Francescomarino, A. Burattin, C. Janiesch, and S. Sadiq, editors, *21st Int. Conf. Business Process Management*, volume 14159 of *LNCS*, pages 465–481. Springer, 2023.
25. G. M. Tavares, R. S. Oyamada, S. Barbon, and P. Ceravolo. Trace encoding in process mining: A survey and benchmarking. *Eng. Appl. Artificial Intelligence*, 126:107028, 2023.
26. S. van den Elzen, D. Holten, J. Blaas, and J. J. van Wijk. Reducing snapshots to points: A visual analytics approach to dynamic network exploration. *IEEE Trans. Vis. Comp. Graph.*, 22(1):1–10, 2016.
27. S. van den Elzen, M. Jans, N. Martin, F. Pieters, C. Tominski, M.-C. Villa-Uriol, and S. J. van Zelst. Towards multi-faceted visual process analytics. *Information Systems*, 133:102560, 2025.
28. W. van der Aalst, A. Adriansyah, A. K. A. de Medeiros, F. Arcieri, et al. Process mining manifesto. In *Business Process Management Workshops*, pages 169–194, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
29. W. M. P. van der Aalst. *Process Mining – Data Science in Action, Second Edition*. Springer, 2016.
30. S. van der Linden, B. Wulterkens, M. van Gilst, S. Overeem, C. van Pul, A. Vilanova, and S. van den Elzen. Flexevent: going beyond case-centric exploration and analysis of multivariate event sequences. *Computer Graphics Forum*, 42(3):161–172, 2023.
31. F. Vitale, M. Pegoraro, W. M. van der Aalst, and N. Mazzocca. Control-flow anomaly detection by process mining-based feature extraction and dimensionality reduction. *Knowledge-Based Systems*, 310:112970, 2025.
32. H. Völzer, F. Zerbato, T. Sulzer, and B. Weber. A fresh approach to analyze process outcomes. In *5th Int. Conf. Process Mining*, pages 97–104. IEEE, 2023.
33. M. O. Ward. Multivariate data glyphs: principles and practice. In *Handbook of Data Visualization*, pages 179–198. Springer, 2008.
34. D. Witschard, I. Jusufi, R. M. Martins, and A. Kerren. A statement report on the use of multiple embeddings for visual analytics of multivariate networks. In *Proc. 16th Int. Joint Conf. Computer Vision, Imaging and Computer Graphics Theory and Applications*, pages 219–223, 2021.
35. F. Zerbato, P. Soffer, and B. Weber. Initial insights into exploratory process mining practices. In A. Polyvyanyy, M. T. Wynn, A. V. Looy, and M. Reichert, editors, *Business Process Management Forum*, volume 427 of *LNBIP*, pages 145–161. Springer, 2021.
36. L. Zimmermann, F. Zerbato, K. Vrotsou, and B. Weber. A process-oriented approach to analyze analysts' use of visualizations: Revealing insights into the what, when, and how. *Computer Graphics Forum*, 2025.