Milana: Bridging process mining and visual analytics through task classification

Lisa Zimmermann[1][0000−0002−6149−7060], Katerina Vrotsou[2][0000−0003−4761−8601], Pnina Soffer[3][0000−0003−4659−883X], Philipp Koytek[4], Barbara Weber[1][0000−0002−6004−4860], and Shazia Sadiq[5][0000−0001−6739−4145]

[1] University of St. Gallen, Switzerland
[2] Linköping University, Campus Norrköping, Sweden
[3] University of Haifa, Israel
[4] Celonis Labs GmbH, Germany
[5] The University of Queensland, Australia

Abstract. Process mining is a powerful approach for analyzing event data, benefiting greatly from human-in-the-loop methods due to its reliance on human interpretation and decision-making. However, current tools do not fully exploit the potential of integrating process mining with interactive visual support. To address this gap, we adopt a design science research approach to systematically connect task concepts from both domains. As a result, we introduce Milana, a method that links process mining tasks, expressed as analysis questions, to established visual analytics requirements. Milana fosters a shared vocabulary, improves communication between the communities, and offers practical guidance for designing effective visualizations tailored to process mining.

Keywords: process mining · visual analytics · analysis tasks · visualization design

1 Introduction

Effective visualization is critical for empowering process analysts to derive meaningful insights from complex process data [3]. Over the last two decades, process mining (PM) has developed into an established research field and offers a collection of techniques and tools that provide immense value for the analysis of processes, enabling diverse analyses such as process discovery, conformance checking, performance analysis, or prediction [4]. While techniques in these areas are constantly advancing, human-centered aspects of PM, including the use of existing tools and the perception of artifacts produced by these techniques, have received less attention [10,19]. As analysts often face complex decision-making tasks, providing appropriate and intuitive visual support is essential for facilitating insight generation and ensuring that conclusions drawn from PM analyses lead to meaningful impact in practice [17].

Concurrently, the field of visual analytics (VA) [15] offers an extensive body of knowledge about how visualizations can support analytical reasoning and decision making [22]. Established VA taxonomies describe tasks, intents, and visualization techniques across diverse domains [7,9,24] and can help guide visualization design. However, the application of these taxonomies to domain-specific needs of PM remains underexplored [19,28]. One key obstacle is a missing shared vocabulary. For example, what VA defines as comparison or correlation analysis may manifest as conformance checking or performance benchmarking in PM. So far, there is a lack of systematic mapping, and without a common language, it remains difficult to translate PM needs into visual solutions and vice versa.

In this paper, we take a step toward bridging this gap. Rather than defining a new taxonomy of PM tasks, we build on an established VA task typology [9] and existing categorization schemes in PM [8,31]. We consider PM tasks as analytical activities that are often driven by concrete questions, whose answers are sought by PM analysts based on process data [12]. Accordingly, our approach starts from a set of PM analysis questions, which we use as a practical entry point to abstract task types. This in turn provides the basis for abstracting the set of operations and data types required for making decisions on appropriate visual encodings [21]. The result is a structured mapping that helps to describe PM tasks in terms that are meaningful within the VA community. Our goal is to support the selection, design, and evaluation of visualizations tailored to the actual analytical needs of PM users.

In particular, we introduce *Milana*[6], a method developed through a design science approach as a conceptual bridge between these two domains. Milana reuses and connects well-established categories from VA and PM and enables analysts, visualization designers, and researchers to communicate and systematically reflect on what kind of visual support is suitable for specific PM tasks.

2 Related Work

This section reviews existing classification systems for analytical tasks and visualization needs. We focus on established schemes from VA and PM, which form the conceptual basis for our work.

2.1 Tasks and Task Definitions in Visual Analytics

In the existing VA literature, the concept of "task" is used in varying ways, with taxonomies reflecting different levels of abstraction and granularity. For example, they might be reflected as low-level operations of analytic activity [6], users' interactions intents [29] or visualization usage [9,13]. Munzner [21] especially highlighted that a task can be described at different levels, i.e., from domain-specific problems to abstract operations.

Rind et al. [24] extended this idea by formalizing task definitions within a conceptual space called TaskCube, consisting of three orthogonal dimensions:

[6] Milana is an Urdu word that means "compound" or "to unite" (see `https://www.urdupoint.com/dictionary/urdu-to-english/milana-meaning-in-english/19252.html`, accessed 12.06.2025).

abstraction, composition, and perspective. *Abstraction* spans from generic analytical tasks to those specific to a data type (e.g., network or temporal data), a domain, or a specific tool. *Composition* refers to the granularity, ranging from concise low-level to broader high-level tasks. *Perspective* distinguishes the user's objective (why a task is done) from the action taken (how it is done). An *objective* is "a question on data [...] to solve a problem [...]" while an *action* is "a discrete step towards addressing an objective".

Moving to concrete classification schemes for tasks, one of the most recognized frameworks is the typology by Brehmer and Munzner [9]. It structures tasks along three dimensions — *why*, *how*, and *what* — capturing the user's intent, the means of execution, and the data involved. The *why* dimension is especially relevant to our work, as it clarifies the intentions of users rather than describing the execution methods, which tend to be more specific to the domain and tools used [9]. Depending on the user's knowledge of the tasks' target and its location, the *why* dimension describes high-level goals (e.g. consume vs. produce), mid-level search behaviors (e.g. look up or explore) and low-level queries (e.g. identify or compare). This typology serves as a strong foundation for translating domain-specific problems into abstract tasks [24] and can inform visualization design. We therefore use it as a basis for abstracting relevant aspects of PM tasks in order to clarify their visualization requirements.

2.2 Tasks in Process Mining

While task taxonomies in VA are well-established, PM lacks a shared understanding of tasks. As Klinkmüller et al. [16] note, little is known about the types of questions practitioners address in practice. Still, some studies have categorized PM tasks from various perspectives. We review key developments below.

The Process Mining Manifesto [1] outlines a widely cited classification of three core PM techniques: process discovery, conformance checking, and process enhancement. While useful, these categories cover diverse use cases. To provide more detail, van der Aalst [4] later introduced six types: process discovery, conformance checking, performance analysis, comparative PM, predictive PM, and action-oriented PM, each reflecting distinct purposes and tool capabilities.

A growing body of research complements this classification of PM techniques with task categorizations derived from case studies. For example, Klinkmüller et al. [16] revealed a list of common domain problems and Milani et al. [20] examined business questions driving PM projects. The latter identified twelve use cases grouped under five business objectives: transparency, efficiency, quality, compliance, and agility. To our knowledge, these works offer the most comprehensive overviews of PM use cases and specific tasks derived from case studies.

Additionally, Barbieri et al. [8] proposed a multi-dimensional taxonomy to categorize questions answerable via their natural language interface. Though focused on interface evaluation rather than task modeling, the taxonomy describes PM tasks along multiple dimensions (e.g., task perspective, filtering, context, composition). Similarly, Zimmermann [31] highlighted the need to classify PM analysis questions along a well-developed taxonomy or classification scheme to structure the analysis phase and ensure a common understanding of the analysis goals. Their tool is built on a six-dimensional classification scheme, including the main use case, process perspective, and data level needed to answer a question.

This work does not propose a new task taxonomy or definition of PM tasks. Instead, we build on existing literature and, following recent work advocating a question or goal-driven analysis approach [30,31,26], treat analysis questions as the starting point for PM tasks. Our aim is to derive and frame visualization requirements from these questions.

3 Design Science Approach

To develop *Milana*, we adopt a Design Science Research (DSR)[23] approach as outlined in Fig. 1. In particular, in this paper, we focus on the problem motivation, the objectives of our solution, the design and development steps, and the demonstration of Milana. While our demonstration already reveals its value, a complete evaluation remains for future work.

3.1 Identify Problem and Motivate

PM relies on human-centered analysis processes in which analysts engage with process data iteratively, formulate questions, interpret results, and derive insights given their domain and context [26]. However, during an analysis, individual analysts struggle with diverse aspects, including the interpretations of visualizations produced by PM algorithms [32].

Overall, the application of VA principles in PM remains underdeveloped. As noted in Sect.1, a key reason is the lack of a shared conceptual foundation. PM and VA rely on different task models, terminologies, and design rationales, making it difficult to translate requirements across domains [28]. Current process analysis tools offer little structured support to align human analytical tasks with appropriate visual interfaces, essential for understanding, exploring, and decision-making. This gap is especially problematic in collaborative settings or for tool development, where process analysts define high-level goals while visualization experts must decide how to support them visually. Without a shared

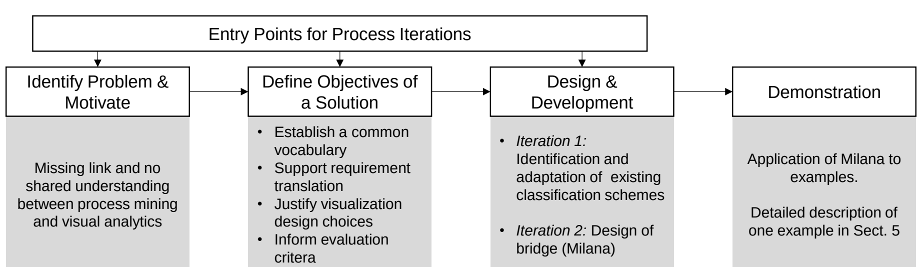

Fig. 1. Overview of our method, adapted from the DSR Process [23].

vocabulary or task-to-visualization mapping, communication suffers and design potential is lost. Therefore, our goal is to enable a structured translation of requirements from PM tasks into visualization solutions that are cognitively appropriate and analytically meaningful to respective VA experts.

3.2 Define Objectives of a Solution

Having identified the gap between PM and VA, we define the following objectives for the proposed solution.

- **Establish an aligned vocabulary across domains.** Our first objective is to provide a conceptual foundation that aligns key concepts, tasks, and terminology across PM and VA. To achieve this, we abstract from specific tasks, typically expressed in PM terms, to ensure a broader, shared understanding.
- **Translate PM requirements into the Visual Design Space.** We aim to support the structured translation of analytical goals and user needs in PM into actionable visualization requirements. This includes mapping abstract tasks (e.g., discovering deviations, profiling behaviors) to suitable operations and visual encodings, enabling VA experts to design effective solutions without deep PM expertise.
- **Support the justification of visualization design choices.** By connecting process tasks to visual design options, our mapping should serve as a reference framework to justify why specific visualizations are appropriate for given tasks. This might facilitate design rationale documentation and enhance interpretability of VA components in PM tools.
- **Inform evaluation criteria for process visualizations.** Our mapping should also form a basis for the development of evaluation protocols by making explicit the expected support for a given PM task. This enables researchers and practitioners to assess not only correctness but also usability and task fit of visualizations based on explicit requirements.

3.3 Design and Development

After identifying the problem and defining clear objectives, we initiated the design and development phase of Milana. Following the principles of DSR, our target artifact can be understood as a method [14] that offers a conceptual and structured approach to clarify tasks and guide visualization design. Its development was iterative, with ongoing refinement and validation. We grounded our work in empirical observations, analyzing PM analysis questions from [8] and [31], which stem from real-world and educational use cases. These questions helped us understand analysts' intentions and cognitive demands. We sought to characterize what information is essential for visualization experts to meaningfully support these tasks. This empirically grounded basis helped to ensure that Milana remains rooted in realistic analysis settings and not only theoretical assumptions. The result of our iterative refinement is presented in Sect. 4 and can be summarized in the following main development components:

1. **Identification and adaption of existing typologies and classification schemes.** We identified foundational classification schemes to describe PM tasks in terms of visualization and data requirements, and developed guiding questions to clarify additional key aspects in a less formal structure.
2. **Design of Milana.** We combined the elements of the first component into meaningful statements to finalize our artifact.

3.4 Demonstration

After its development, we performed a demonstration of how Milana can be used to guide VA experts in making informed choices for effective visualizations relative to the PM task (cf. Sect. 5). The demonstration indicates the capability of Milana to establish an aligned vocabulary across the two domains and enable the translation of PM requirements into the visual design space. It also informs the evaluation criteria for process visualizations which we aim to further develop towards a more detailed and empirical evaluation in future work.

4 Result: Milana

In this section, we present the results of the design and development phases of our DSR approach. Milana involves four steps as outlined in Fig. 2 and described below: classification of task(s) according to adapted visualization and PM task typologies, clarification of key aspects, and formulation of the Milana bridge.

4.1 Identification and Adaptation of Existing Typologies

Identification of base typology. As pointed out in Sect. 2, in comparison to PM, VA provides a more cohesive understanding of visualization and analysis tasks. We chose the well-known and widely accepted typology of Brehmer and Munzner [9] as a starting point for our method due to its flexibility and expressiveness. It defines tasks in three dimensions, indicating *why* it is performed, *how* it is performed, and *what* kind of input and output (data) is produced. Focusing on the *why* dimension provides us with a path to abstract a given PM question into a visualization task objective which can then inform the design of a visualization solution supporting the underlying intention of the task.

Applying the *why* dimension to PM tasks, i.e., a set of analysis questions, revealed certain interpretative challenges. In particular, the notions of *target* and *location*, which are central to the mid-level (search) part of the dimension, required further contextualization. We, therefore, developed refinements tailored to the PM domain:

- **Target**, in PM, refers to the condition an analyst is trying to isolate or investigate. We consider it *known* when this condition can be directly specified and evaluated on the data (e.g., "identify cases longer than 5 days"). The target is *unknown* when solving the task requires first evaluating a comparative or abstract criterion (e.g., "cases that take the most time"), where the benchmark or reference emerges from the data.

- **Location** defines where in the process data the target is expected to be found. In PM, this includes the levels of log, trace, and event. The location is *known* when the task explicitly constrains the data subset (e.g., "cases processed by department B and lasting longer than 5 days" - which narrows the search to a specific subset of the cases (traces) in the log). It is *unknown* when no such reference is made, requiring broader exploration across all traces of the log (e.g., "which cases take longer than 5 days?").

The resulting PM tailored version of the typology which we adopt for Milana is presented in Fig. 2A.

Classification of PM tasks. Next, we turned to identify further relevant categories of PM tasks to specify the nature of the required PM analysis in more detail. To this end, we built on existing question classification approaches described in Sect. 2. After applying different combinations of classification schemes to our set of questions and assessing their relevance, we identified three dimensions inspired by the classification scheme of Zimmermann [31] as both analytically relevant and practically useful for building our translation bridge (cf. Fig. 2B):

- **Use Case** refers to the underlying purpose or motivation for analysis. Rooted in common PM objectives, we distinguish four primary use cases: transparency, performance, compliance, and automation. These are aligned with established PM categories [20,4] and tool functionalities.
- **Perspective** denotes the angle from which the process is analyzed. Following the definitions by van der Aalst [2], we include perspectives such as *control-flow*, *time*, *resources*, *data*, and identified a need for an *"other"* category, allowing to describe tasks that require the consideration of multiple perspectives or perspectives not part of the previous list. The perspective of a task describes what aspect of the process is foregrounded in an analysis.
- **Data Level**, inspired by the XES standard [5], reflects the granularity of the data involved: *log*, *trace*, *event*, or *other*. It indicates the scope of data that is required to answer a question, helping visualization experts assess the necessary resolution and abstraction level.

Clarification of Key Aspects. Throughout the development iterations of Milana, it became evident that even questions with similar surface structure could differ substantially in their analytic implications. Therefore, beyond classification, we introduced two clarifications of finer-grained aspects of task formulation (cf. Fig. 2C):

- **C1**: What is expected from the answer, i.e., is a definite value (binary/numeric) expected or a qualitative description/explanation? If latter applies, what perspective constitutes a relevant context for the qualitative analysis?
- **C2**: How are the specific concepts referred to by the questions defined (e.g., bottleneck, "working as agreed upon,..")? In case these concepts require knowledge/information (e.g., a threshold, a normative model), specify it!

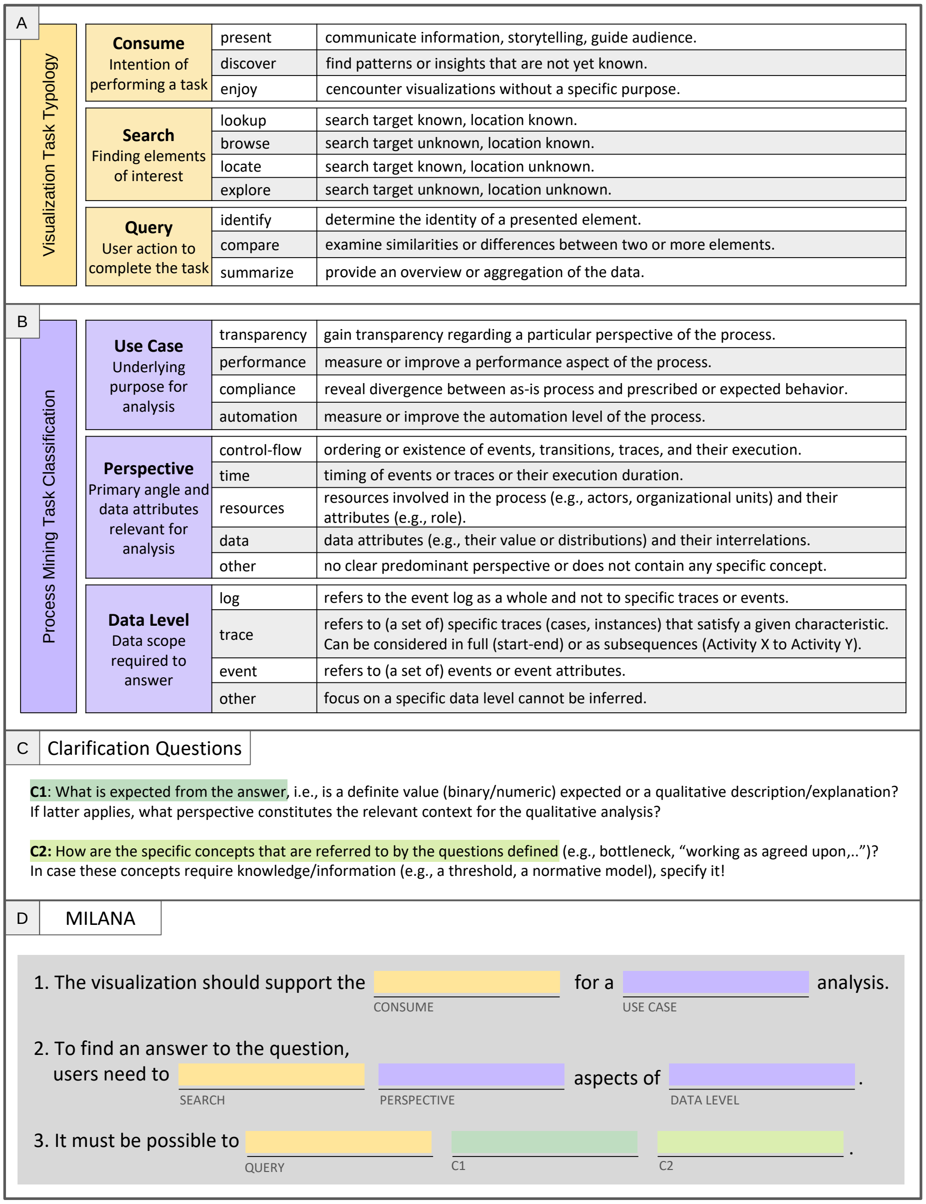

Fig. 2. Overview of the different parts of the Milana method. (A) The "why" perspective of Brehmer and Munzner's visualization task typology [9]. (B) Classification of PM Tasks inspired from [31]. (C) Additional clarifications of key aspects. (D) Milana bridge, which makes use of all three parts to translate PM analysis tasks into actionable design guidance for VA experts.

4.2 Design of Milana

After completion of the iterations on the identification and adaptation of existing typologies, the second main component of the development phase was to operationalize the connection between PM tasks and VA methods. Therefore, we formulated a structured bridge that translates PM analysis tasks into actionable design guidance for VA experts. After testing several versions of it on our set of empirical analysis questions and discussing among the authors, Milana emerged as the composition of the following three sentences. They act as the final translation layer informing a task-oriented design of visualizations in PM (cf. Fig. 2D):

1. **"The visualization should support the *[Consume]* for a *[Use Case]* analysis."**
   This sentence sets the analytical objective by expressing the user's intention ("why" dimension) and grounding it in a concrete PM use case (e.g., compliance, performance). This framing aligns the visualization with the underlying motivation for analysis.
2. **"To find an answer to the question, users need to *[Search]* *[Perspective]* aspects of *[Data Level]*.**
   This sentence links the analytical objective to the required user actions and the relevant process characteristics. It draws from the "how" and "what" dimensions of the visualization task typology, connecting them to PM perspectives (e.g., time, control-flow) and the data granularity (log, trace, event).
3. **"It must be possible to *[Query]* *[C1]* *[C2]*".**
   This sentence provides technical design implications by specifying how users must be able to interact with the data. The placeholders [C1] and [C2] denote task-specific contextual constraints, such as comparison operators, filtering techniques, or temporal relationships that must be supported to accomplish the task (cf. Sect. 4.1).

5 Demonstration

To illustrate Milana's use, consider the following scenario: a PM expert prepares the bridge specifications for a PM task. A VA expert then takes these specifications and proposes a suitable VA interface design. Our overall vision is that in a realistic setting applying Milana to a collection of relevant PM tasks for a given (business) process will help identify the common analytical needs and functionality requirements to guide the design of a tailored (interactive) visual interface for analysis.

For this demonstration, we choose a common and recurring task in PM projects: ***"Where are the bottlenecks in the process?"***. Identifying bottlenecks is relevant for uncovering concrete inefficiencies and has been identified as a commonly raised question in projects [18]. Several techniques, which can be attributed to the broader category of performance analysis [4], have been proposed and applied to support the analysis of bottlenecks. The relevance of

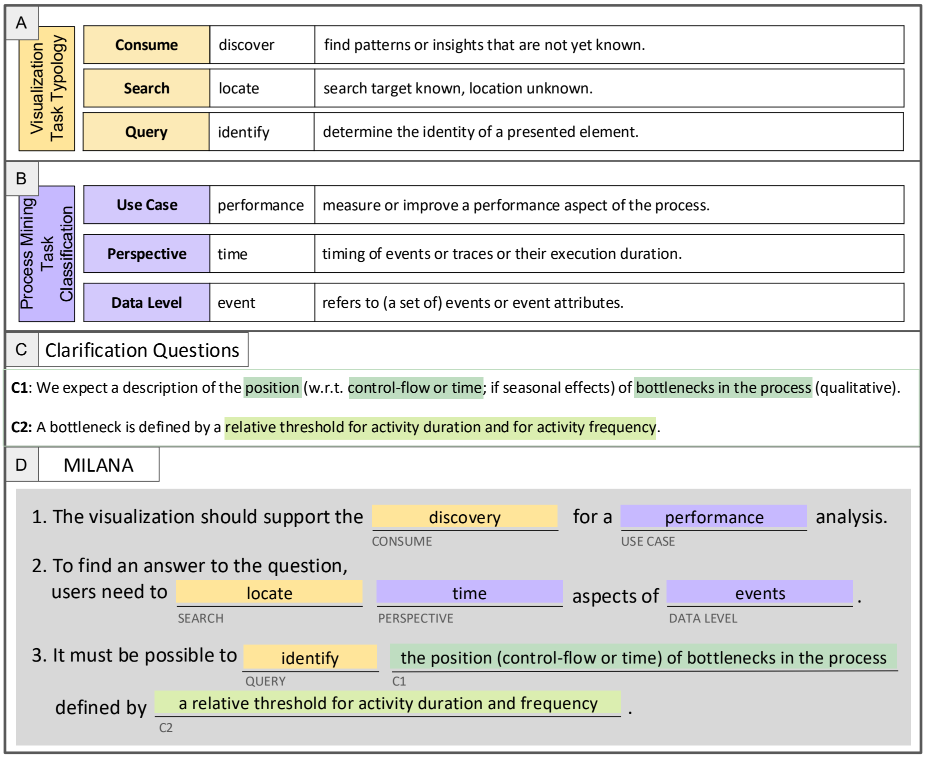

Fig. 3. (A)-(C) Classifications made by the PM expert as input to Milana. (D) Application of Mialna to the demonstration example.

this question has also been highlighted in the context of the Business Process Intelligence Challenge (BPIC) 2020[7], where it was among the questions posed by the process owners.

The classifications that were established by the PM expert for the task at hand are provided in Fig. 3A-C and form the input for Milana, which is provided in Fig. 3D. Based on the Milana bridge, the following visualization-related interpretations are made by the VA expert. The visualization interface should support *discovery* for performance analysis; this implies that the interface should enable *interactive exploration* of the process and its relevant performance aspects (in this case *temporal aspects*). Specifically, the interface needs to enable a user to *identify* where in the process (w.r.t. the control flow or time) bottleneck *events* occur. Bottleneck events are defined by a threshold for frequency and duration (*time aspects*) which is set interactively by the analyst.

This interpretation of the Milana bridge outlines the requirements for the visual interface enabling the VA expert to sketch initial mockup views, as provided

[7] https://data.4tu.nl/collections/BPI$_{C}$*hallenge*$_{2}$020/5065541/1

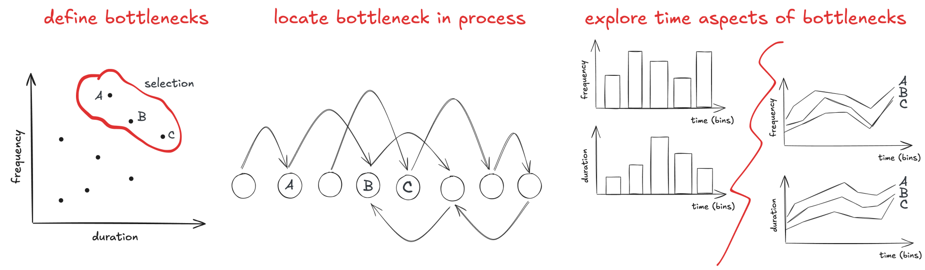

Fig. 4. Concept mockups of potential complementary views for bottleneck analysis

in Fig. 4. To set the bottleneck threshold a representation is needed that shows an overview of the duration and frequency of events in which a user can make selections (qualitative identification). A representation for displaying continuous 2-dimensional data is the scatter plot (Fig. 4 left). A scatter plot mapping the average frequency vs average duration of events would reveal common events with long durations which could be selected as (potential) bottlenecks for further exploration. To locate these events in the context of the process and explore their temporal characteristics, a set of temporal views are needed. There are several alternatives that can be equivalently valid and the selection could be guided by surveys of visualization techniques for different data types (e.g., the survey of visualization techniques for time-oriented data [27]) as presented in Fig. 5). For displaying the process in a simplified manner a directional arc diagram was sketched (Fig. 4 middle) displaying the transitions between events with the bottlenecks highlighted. Temporal aspects such as average frequency and duration of events over time can with advantage be explored in well-established temporal representations such as histograms or line plots (Fig. 4 right).

Based on the sketched representation alternatives, a mockup of a prototype interface was created to experiment with the proposed views (cf. Fig. 6) based on the BPIC dataset 2020 [11]. The proposed interface comprises a scatter plot representation (Fig. 6, top left) showing the relationship between frequency and average duration, enabling a user to define the conditions (thresholds) for the bottleneck exploration by interactively selecting suspicious event types. Selecting events will highlight them by rendering the remaining semi-transparent.

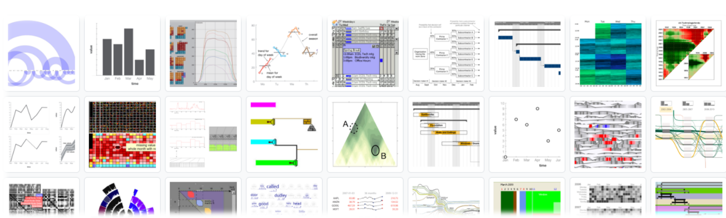

Fig. 5. Visualization techniques for time-oriented data from the TimeViz Browser [27].

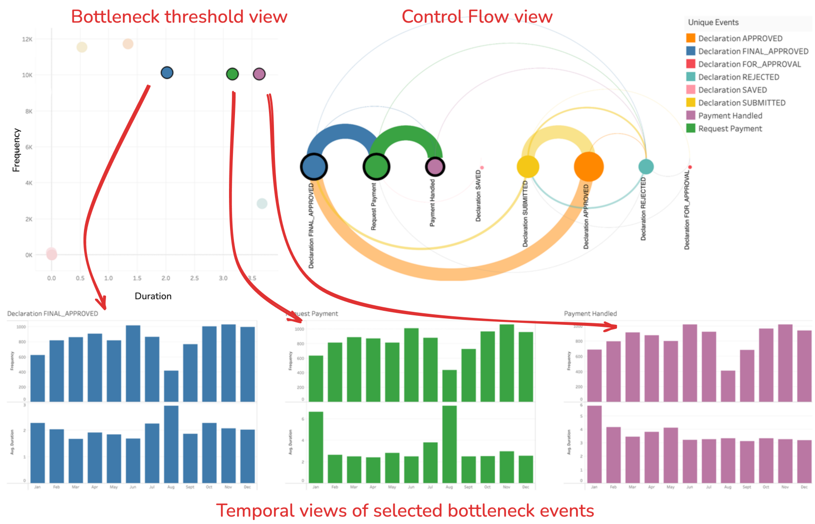

Fig. 6. Visualization interface mockup for bottleneck exploration elicited from Milana.

The control-flow is captured through an arc diagram showing the process overview (Fig. 6, top right). Nodes are drawn along a horizontal axis with connecting arcs marking transitions between events, enabling quick identification of the order and positioning of events of interest. Nodes are drawn proportional to their frequency and arcs are weighted by number of transitions. "Forward" transitions are drawn above the nodes and "backward" ones below. The nodes are by default sorted to reduce edge crossings but can be sorted according to different criteria such as a benchmark event order. Upon selection of bottleneck events in the scatter plot, these are highlighted in the arc diagram. In addition, corresponding histogram views are drawn for each bottleneck at the bottom of the interface (Fig. 6, bottom) revealing temporal aspects (duration and frequency over time) and providing additional context for analysis (e.g., the identification of seasonal bottlenecks). As such, the arc diagram provides an ordered view of the events whose timings can be explored in more detail in the histograms.

Exploring the domestic travel declarations dataset from the BPIC 2020[8] in the proposed interface reveals that events with a high average duration and relevant frequencies are: `Payment Handled`, `Request Payment`, and `Declaration FINAL_APPROVED`. The arc diagram shows how these are connected to the other process events. The frequency distribution of all three is similar showing a notable drop in August. The distribution of durations varies between the selected events. `Declaration FINAL_APPROVED` peeks in August and shows a notable increase in January, July, October (> 2 days). `Request Payment` displays large peeks in January and August while `Payment Handled` takes longest in January.

[8] Before analysis, we preprocessed the events and removed role names from the activity names

6 Discussion and Future Work

In this paper, we proposed a novel method for bridging PM tasks to visualization requirements. In particular, Milana serves as a conceptual bridge between domain-specific questions and VA design considerations, offering a shared foundation for PM analysts and visualization experts. To develop Milana, we reviewed existing task classifications from both VA and PM, and applied them to practical PM tasks until we converged to a version that could be applied to our selected subset of PM questions and successfully informed visualization design decisions, as demonstrated in Sec. 5. Thereby, our objectives when developing Milana were fourfold: (1) to establish an aligned vocabulary; (2) to translate PM-specific requirements into the visual design space; (3) to justify design decisions for visualization interfaces; and (4) to inform evaluation criteria. Milana addresses each of these goals by offering a structured format for articulating analysis needs by abstracting from concrete tasks, mapping them to actionable design elements, and fostering traceable rationales for interface decisions.

While Milana contributes a decisive step towards closer cooperation between the communities, there is substantial room for further research. For example, we observed that the final translation of the Milana sentences by VA experts is not deterministic. Users may arrive at different, yet equally valid, design solutions for the same analytical question. This flexibility is needed since it reflects the creative nature of design and the diversity of visualization strategies. However, future work could investigate whether specific combinations of task characteristics tend to correlate with certain visualization patterns that are more effective, thereby allowing more prescriptive guidance.

Additionally, while we applied our approach to multiple PM tasks in the process of developing it, a comprehensive evaluation needs be conducted to determine if Milana leads to *good* visualizations as well as to assess the overall utility of the bridge mechanism. The evaluation should follow an experimental design involving two groups: (1) a control group tasked with designing visualizations without structured guidance; and (2) a treatment group receiving guidance through Milana. This comparative approach will help isolate the impact of the bridge on users' ability to identify, interpret and utilize suitable visualizations effectively.

To ensure the successful application of Milana, we plan to develop structured guidelines of use. At the moment, it is essential to involve a domain or PM expert to accurately categorize the analytical question. Their expertise is crucial in interpreting the context and ensuring that the question is framed correctly. However, we expect additional guidance will be needed to support the use of Milana, especially when not exclusively applied by experts. This may include common classification patterns, examples of common question types and corresponding visualizations, or navigation support through existing surveys of visualization techniques, such as provided by [27,25], to guide users in selecting the right visualization based on Milana.

Acknowledgements. This paper resulted from discussions at Dagstuhl Seminar 23271.

The authors thank the input of all participants on this topic. Further, this work was supported by the SNF (Grant No. 200021_197032), ISF (Grant No. 2005/21), and ARC Training Centre for Information Resilience IC200100022.

References

1. van der Aalst, W.M.P., Adriansyah, A., de Medeiros, A.K.A., et al.: Process mining manifesto. In: Daniel, F., Barkaoui, K., Dustdar, S. (eds.) Business Process Management Workshops (BPM 2011). pp. 169–194. Springer (2012). https://doi.org/10.1007/978-3-642-28108-2_19
2. van der Aalst, W.M.: Process discovery from event data: Relating models and logs through abstractions. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery **8**(3) (2018)
3. van der Aalst, W.M., de Leoni, M., ter Hofstede, A.: Process mining and visual analytics: Breathing life into business process models. BPM Center Report BPM-11-15, BPMcenter. org **17**, 699–730 (2011)
4. van der Aalst, W.: Process mining: A 360 degree overview. In: van der Aalst, W., Carmona, J. (eds.) Process Mining Handbook, pp. 3–34. Springer (2022). https://doi.org/10.1007/978-3-031-08848-3_1
5. Acampora, G., Vitiello, A., Di Stefano, B., van der Aalst, W., Günther, C., Verbeek, E.: Ieee 1849: The xes standard. IEEE Comp. Intelligence Magazine pp. 4–8 (2017)
6. Amar, R., Eagan, J., Stasko, J.: Low-level components of analytic activity in information visualization. In: IEEE Symp on Information Visualization. pp. 111–117. IEEE (2005)
7. Andrienko, N., Andrienko, G.: Exploratory Analysis of Spatial and Temporal Data. Springer Science & Business Media (2006)
8. Barbieri, L., Madeira, E., Stroeh, K., van der Aalst, W.: A natural language querying interface for process mining. J. Intell. Inf. Syst **61**(1), 113–142 (2023)
9. Brehmer, M., Munzner, T.: A multi-level typology of abstract visualization tasks. IEEE Trans on Vis. and Comp. Graphics **19**(12), 2376–2385 (2013)
10. vom Brocke, J., Jans, M., Mendling, J., Reijers, H.A.: A Five-Level Framework for Research on Process Mining. Business and Information Systems Engineering **63**(5), 483–490 (2021). https://doi.org/10.1007/s12599-021-00718-8
11. van Dongen, B.: BPI Challenge 2020. https://doi.org/10.4121/uuid:52fb97d4-4588-43c9-9d04-3604d4613b51 (2020), 4TU.ResearchData collection
12. van Eck, M., Lu, X., Leemans, S., van der Aalst, W.: Pm2: A process mining project methodology. In: Int. Conf. on Advanced Information Systems Engineering. Springer (2015). https://doi.org/10.1007/978-3-319-19069-3_19
13. Heer, J., Shneiderman, B.: Interactive dynamics for visual analysis: A taxonomy of tools that support the fluent and flexible use of visualizations. Queue **10**(2), 30–55 (2012)
14. Hevner, A.R., March, S.T., Park, J., Ram, S.: Design science in information systems research. MIS quarterly pp. 75–105 (2004)
15. Keim, D., Andrienko, G., Fekete, J.D., Görg, C., Kohlhammer, J., Melançon, G.: Visual analytics: Definition, process, and challenges. In: Information visualization: Human-centered issues and perspectives, pp. 154–175. Springer (2008)
16. Klinkmüller, C., Müller, R., Weber, I.: Mining process mining practices: An exploratory characterization of information needs in process analytics. In: Int. Conf. on Business Process Management. Springer (2019). https://doi.org/10.1007/978-3-030-26619-6_21

17. Kubrak, K., Milani, F., Nolte, A.: A visual approach to support process analysts in working with process improvement opportunities. Bus. Process Manag. J. **29**(8), 101–132 (2023)
18. Mans, R.S., van der Aalst, W.M.P., Vanwersch, R.J.B., Moleman, A.J.: Process mining in healthcare: Data challenges when answering frequently posed questions. In: Process Support and Knowledge Representation in Health Care. pp. 140–153. Springer (2013). https://doi.org/10.1007/978-3-642-36438-9_10
19. Miksch, S., Di Ciccio, C., Soffer, P., Weber, B.: Visual analytics meets process mining: Challenges and opportunities. IEEE Computer Graphics and Applications **44**(6), 132–141 (2024). https://doi.org/10.1109/MCG.2024.3456916
20. Milani, F., Lashkevich, K., Maggi, F., Di Francescomarino, C.: Process Mining: A Guide for Practitioners. In: Research Challenges in Information Science. Springer (2022). https://doi.org/10.1007/978-3-031-05760-1_16
21. Munzner, T.: A nested model for visualization design and validation. IEEE Trans on Vis. and Comp. Graphics **15**(6), 921–928 (2009)
22. Oral, E., Chawla, R., Wijkstra, M., Mahyar, N., Dimara, E.: From Information to Choice: A Critical Inquiry Into Visualization Tools for Decision Making . IEEE Trans on Vis. & Comput. Graphics **30**(01), 359–369 (Jan 2024). https://doi.org/10.1109/TVCG.2023.3326593
23. Peffers, K., Tuunanen, T., Rothenberger, M.A., Chatterjee, S.: A design science research methodology for information systems research. J. Manag. Inf. Syst. **24**(3), 45–77 (2007)
24. Rind, A., Aigner, W., Wagner, M., Miksch, S., Lammarsch, T.: Task Cube: A three-dimensional conceptual space of user tasks in visualization design and evaluation. Information Visualization **15**(4), 288–300 (2016). https://doi.org/10.1177/1473871615621602
25. Schulz, H.J.: Treevis.net: A tree visualization reference. IEEE Computer Graphics and Applications **31**(6), 11–15 (2011). https://doi.org/10.1109/MCG.2011.103
26. Sorokina, E., Soffer, P., Hadar, I., Leron, U., Zerbato, F., Weber, B.: Pem4ppm: A cognitive perspective on the process of process mining. In: Intl Conf on Business Process Management. pp. 465–481. Springer (2023)
27. Tominski, C., Aigner, W.: The TimeViz Browser – A Visual Survey of Visualization Techniques for Time-Oriented Data (2023), https://browser.timeviz.net
28. Yeshchenko, A., Mendling, J.: A survey of approaches for event sequence analysis and visualization. Information Systems **120**, 102283 (2024)
29. Yi, J.S., ah Kang, Y., Stasko, J., Jacko, J.A.: Toward a deeper understanding of the role of interaction in information visualization. IEEE Trans on Vis. and Comp. Graphics **13**(6), 1224–1231 (2007)
30. Zerbato, F., Koorn, J., Beerepoot, I., Weber, B., Reijers, H.: On the origin of questions in process mining projects. In: Int. Conf. on Enterprise Design, Operations, and Computing. Springer (2022). https://doi.org/10.1007/978-3-031-17604-3_10
31. Zimmermann, L.: The process mining question forge. In: Proc of the Best Dissertation Award, Doctoral Consortium, and Demonstration and Resources Forum at BPM 2024 (2024)
32. Zimmermann, L., Zerbato, F., Weber, B.: What makes life for process mining analysts difficult? a reflection of challenges. Softw. Syst. Model. **23**(6), 1345–1373 (2024). https://doi.org/10.1007/s10270-023-01134-0