Towards Informed Process Exploration: Better Previews, Reviews, and Views Beyond

Christian Tominski[1*], Irit Hadar[2], Philipp Koytek[3], Maria-Cruz Villa-Uriol[4], Francesca Zerbato[5]

[1*]Institute of Visual and Analytic Computing, University of Rostock, A.-Einstein-Str. 22, 18059 Rostock, Germany.
[2]Department of Information Systems, University of Haifa, Hanamal St. 65, Haifa, Israel.
[3]Celonis Labs GmbH, Theresienstr. 6, 80333 Munich, Germany.
[4]School of Computer Science, University of Sheffield, Regent Court, 211 Portobello St, Sheffield, S1 4DP, United Kingdom.
[5]Department of Mathematics and Computer Science, Eindhoven University of Technology, Groene Loeper 3, Eindhoven, 5612 AE, The Netherlands.

*Corresponding author(s). E-mail(s): christian.tominski@uni-rostock.de; Contributing authors: hadari@is.haifa.ac.il; p.koytek@celonis.com; m.villa-uriol@sheffield.ac.uk; f.zerbato@tue.nl;

Abstract

The exploration of large event log data in search of process insights and bottlenecks is at the heart of process mining (PM). This search is inherently based on combining the analytical power of automatic computations with the advanced reasoning skills of human domain experts. Yet, PM research tends to focus on algorithmic components, paying less attention to the human side of interactive process exploration facilitated through visual interfaces. Our paper is a call to action to strengthen the aspects of interactivity in process exploration. We propose to enhance process exploration with visual analytics (VA) methods that have so far not been considered in the PM community. Our work establishes conceptual links between PM and VA based on relevant interaction intents and user questions. Reviewing the level of support for these intents and questions in current PM tools reveals uninformed exploratory interactions as a key challenge. We outline how VA methods can help tackle this challenge and propose corresponding research directions towards better informed process exploration.

Keywords: Process mining, visual analytics, human-computer interaction.

1 Introduction

Process mining (PM) deals with extracting, validating, and understanding processes based on multivariate event log data [1]. While PM employs various algorithmic methods, e.g., to discover process models from event data [5], many PM tasks require human expertise to steer and complement the automatic mining part [82].

Analysts cannot rely solely on automated processes. Instead, they have to engage in iterative exploration cycles, interacting with visualizations to build an understanding of the process, examining different scenarios, and generating hypotheses that ultimately inform their choice of what techniques to apply on which subsets of the log [67, 83]. This iterative exploration is achieved through several discrete steps, such as choosing expressive visual representations, finding suitable levels of abstraction to represent activities in the process, setting effective filter thresholds, and selecting relevant data attributes.

To support analysts in their exploration, many PM tools have increasingly integrated interactive visualizations [62, 80], thanks also to the teaming up with the discipline of visual analytics (VA) [16, 49]. VA's core mission is to combine automatic methods with interactive visual interfaces to support the analysis of large, complex data [39, 72, 74]. In the context of PM, this means providing interactive functionalities, such as filter masks, sliders, and abstraction controls, that allow analysts to isolate data subsets and simplify the process models obtained.

However, these interactions often fail to bridge the gap between user intent and system response, resulting in what we refer to as *uninformed exploration*. While tools provide the functional means to filter and manipulate data, they lack the transparency and context necessary for the analyst to make informed decisions and generate insights [71]. This leads to trial-and-error workflows where users cannot anticipate the effects of an action before executing it and also may struggle to evaluate the effects of an action taken.

This state of uninformed exploration represents a critical gap in both tool design and research. Despite the importance of the human-in-the-loop, the question of how to move from uninformed to informed exploration in PM remains under-researched.

This work is a call to action for more research on strengthening the human side of PM using VA technology. Our goal is to address the challenge of uninformed exploration by grounding PM interaction in established VA principles. After introducing the reader to background concepts in Section 2, we present a systematic view on process exploration by applying Yi et al.'s interaction intents [81] and Spence's user questions [68] as reference frameworks in Section 3. In Section 4, we look at how PM tools commonly used in practice support these interaction intents and user questions, focusing on the example of directly-follows graphs. From this tool review, we identify three key concerns regarding process exploration: limited previews, limited reviews, and limited views beyond in Section 5. Addressing these concerns, we suggest that established VA methods are integrated into PM tools for enhanced process exploration with better previews, reviews, and views beyond. To guide such an integration of VA and PM in future work, we outline a research vision and corresponding research directions in Section 6.

Overall, we make the following contributions that, taken together, inform and form our call to action:

- Conceptual grounding of process exploration based on fundamental interaction intents and user questions
- Review of how common PM tools support interaction intents and user questions, using directly-follows graphs as a case study
- Identification and discussion of interactive exploration limitations in PM practice
- Illustration of the potential of VA technology to address existing limitations
- Development of a research vision towards informed process exploration

2 Background

Next, we provide relevant background information on exploratory analysis in PM (see Section 2.1), on directly-follows graphs, which we use as a running example (see Section 2.2), and on the field of VA and the role of interactivity in VA (see Section 2.3).

2.1 Exploration in process mining

PM supports many tasks, including the discovery of process models, the analysis of process performance, and the checking of conformance between actual process executions and desired process models [1]. In practice, many projects start with an exploratory analysis [40, 59, 82]. An iterative exploration allows process analysts to better understand the process and its domain, identify potential issues, and formulate hypotheses about the process [67]. These hypotheses are then tested, refined, or discarded based on intermediate insights, which, in turn, guide the analysts in choosing what to explore next [77, 85].

To support process analysts in their work, there are many commercial and open-source tools available. These tools are typically interactive and offer various ways to explore and analyze process data through visual representations such as process models, bar charts, or density plots [22, 80]. Next to existing tools, PM research has advanced on creating new visualizations to support analysts in their work. For example, the *Time-Order Map* [58] enables seamless zooming between abstract process models and concrete process instances. For fine-grained performance insights, the *Performance Spectrum* [15] provides an approach to visualize the dynamics of the process over time. Furthermore, new methods have been developed to define and visualize trace variants, even when dealing with partially ordered event data [61].

During exploration, many visualizations and analysis operations can be applied to event logs, depending on the tool used and the analysis task [22, 40]. However, it is up to the analyst to perform the *right* tasks and iterate through the different phases of the analysis [82]. Therefore, supporting analysts in determining and carrying out the right task remains essential [9].

2.2 Directly-follows graphs as starting point

A common starting point for process exploration is a directly-follows graph (DFG), where nodes represent the activities of a process, and edges represent directly-follows relationships between them [1]. DFGs are typically visualized as node-link diagrams. The example DFG in Fig. 1 shows the process of managing road traffic fines [45]. The process starts with 150K cases of a "Create Fine" activity followed by "Send Fine" in 100K cases and "Payment" in 47K cases (see arrow labels). The links are visually weighted to show how frequently the activities follow one another, providing a visual overview of the process flow, including common paths and potential deviations.

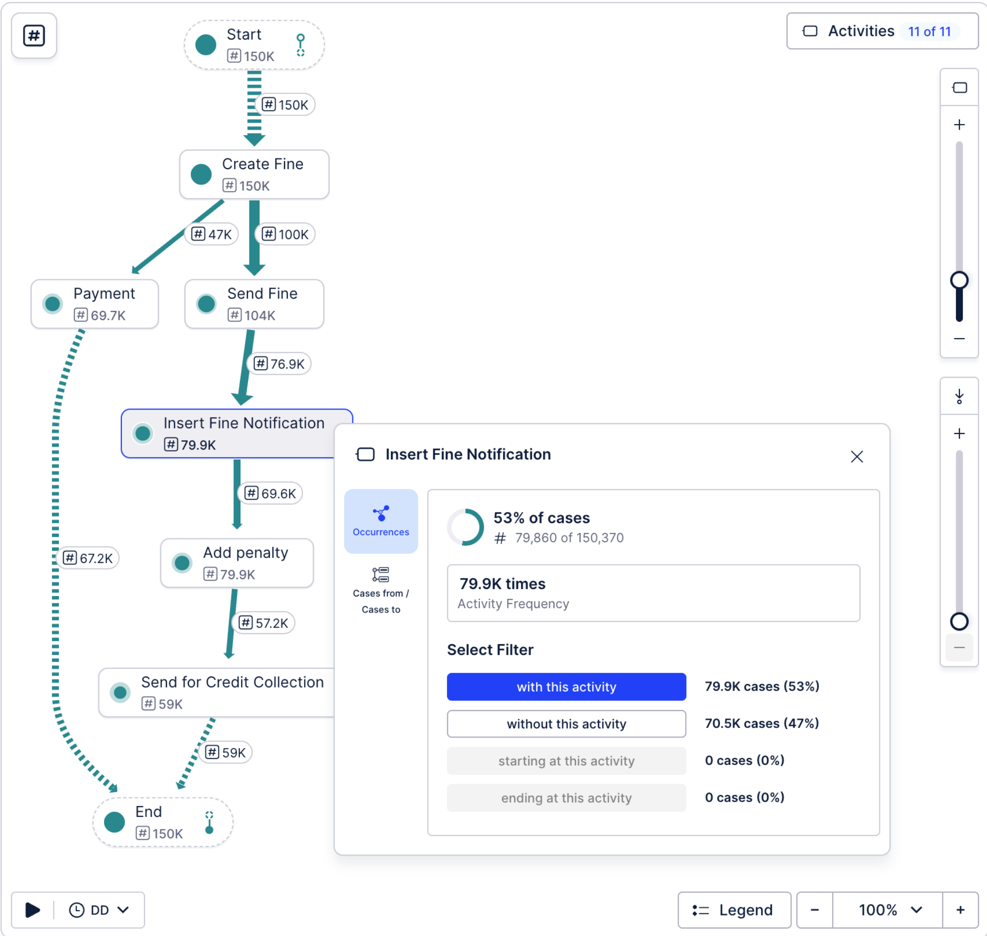

Fig. 1 Celonis Process Explorer is an example of an interactive DFG to analyze a process. Here, the DFG shows the process of road traffic fine management. The popup shows details and filter options for the selected activity "Insert Fine Notification".

Real-life processes often result in a DFG with hundreds or thousands of nodes and edges. Therefore, most PM tools start with a simplified DFG showing only the most common process behavior and filtering out less frequent behavior. To find a suitable balance between simplicity and details, analysts have to adjust filter thresholds (see sliders on the right in Fig. 1) to gradually add or remove nodes and edges based on their frequency. In addition, it is common for PM tools to allow analysts to filter on process cases with certain behavior, for example, to filter only on process cases that include a certain activity such as "Insert Fine Notification" (see popup in Fig. 1).

Due to their simple notation, DFGs are widely adopted [22, 44] in PM tools, such as Apromore, Celonis, Disco, and ProM. Therefore, we use DFGs as our running example to study the challenges of process exploration. Previous research has already identified limitations and challenges of DFGs in terms of their accuracy and interpretability [2, 50, 66], and different algorithms have been developed to balance the accuracy of the event log representation with the perceived complexity of the resulting graph [11]. What remains an open question is how process exploration can be improved through better informed and more transparent interaction. Focusing on DFGs first to answer this question seems reasonable given their popularity in PM.

2.3 Visual analytics and interactivity

VA technology has the potential to enhance process exploration. VA has been defined as "the science of analytical reasoning facilitated by interactive visual interfaces" [72]. It essentially combines advanced visualization, interaction, and computational analysis methods to facilitate data-driven analytical work [74]. This work typically involves exploratory journeys during which a data analyst creates various visualizations of the data and employs different computational methods for data abstraction (e.g., clustering, feature detection, trend analysis).

Because the human analyst is at the center, driving the analysis, interaction techniques play a key role in VA [18, 34, 56]. Various interaction intents [81] and interaction patterns [63] have been studied and are employed in VA solutions, including selection, navigation, zoom, alignment, and visual comparisons. These operations rely on the concept of direct manipulation [48, 64], which is understood as the ability to directly interact with and manipulate the visual elements that constitute a data visualization.

In general, VA aims to satisfy three quality criteria: expressiveness, effectiveness and efficiency [74]. An interactive visual data analysis solution is *expressive* if it communicates the relevant information, and users can carry out the actions required to obtain that information. *Effectiveness* measures the degree to which users can achieve their task. And last, *efficiency* relates to a VA solution's benefits outweighing its costs, typically consisting of computational resources plus human effort. Along these lines, previous works have looked into developing heuristics to measure how well interactive visual solutions work [24, 53, 90] and assess VA solutions based on value [76, 77].

To obtain expressive, effective, and efficient VA solutions, a variety of factors need to be considered: *what* needs to be analyzed (data characteristics) and *why* (tasks), *who* will conduct the analyses (human factors), *where* analyses will be conducted (device properties), and *when* the analytical activities take place (workflow) [74]. Focusing on interactivity, there are helpful guidelines to inform the design of VA solutions. Shneiderman's Golden Rules [65] set several principles to guide how interfaces should be designed, including *striving for consistency*, *offering informative feedback*, *permitting easy reversal of actions*, and *keeping users in control*. Norman [54] provides complementary guidelines from the perspective of how humans think, perceive and learn. For example, facilitating *discoverability* of interactions is key to allow users to build a mental model of how to interact with a VA system. Ideally, interactions are fluid and smooth to allow users to fully immerse in the data analysis [20].

While the aforementioned criteria and guidelines are well-known in the VA community, they have not found widespread consideration for process exploration. Existing research at the intersection of PM and VA has primarily focused on visual encodings and analytic tasks, for example, [22, 30, 46, 80, 89]. This leaves interactivity for process exploration to be investigated in more depth.

3 A systematic view on process exploration

The topic of interactivity is not easy to grasp. It can mean and imply many different things and can be approached from various perspectives [18, 36, 56, 73]. Therefore, we start with looking at some conceptual foundations and applying them to interactive process exploration using DFGs as a running example. First, we review Yi et al.'s fundamental *interaction intents* [81], a seminal work that categorized VA interactions systematically. Then, we discuss the challenges that users face during interactive exploration and structure them according to Spence's *user questions* [68], which provide a basis for user-oriented thinking about interactivity. Establishing a connection to PM, we will link the otherwise application-agnostic interaction intents and user questions to concrete examples in the context of process exploration with DFGs.

3.1 Why to interact: Interaction intents

For our systematic view, we resort to the established categorization of interactions according to seven interaction intents by Yi et al. [81]. These intents capture *why* data analysts need to or want to interact while exploring visual data representations. Note that the intents do *not* specify how they can be satisfied through specific interaction techniques for visual data representations. The seven intents are:

I1: Mark something as interesting. As a fundamental operation, users need to be able to mark parts of the visual representation as relevant for the task at hand, either temporarily (hover) or permanently (select). Marking is often required for follow-up interactions. In a DFG, users may want to mark a path corresponding to a group of process instances deviating from the happy path to filter them out.
I2: Show me something else. Given limited display space, it is typically only possible to show parts of the data or the processes under analysis. Therefore, users need to move from one partial view to the other and progressively create an overall picture. In a DFG, users may want to pan or scroll to look at different parts of the graph.
I3: Show me a different arrangement. A specific visual arrangement typically serves a limited set of analytical tasks. Rearranging the graphical primitives of a visual

representation can facilitate diverse tasks and help make specific parts of the data easier to follow and understand. An example would be to rearrange the nodes in a DFG to centrally align a path of interest.

I4: Show me a different representation. Multi-faceted analyses involve different visual representations each offering a specific perspective on processes. Integrating the different perspectives leads to comprehensive understanding. For example, once users identified an interesting DFG path, they may want to switch to a time-line visualization to quantify temporal distances between events.

I5: Show me more or less detail. Interesting patterns or behaviors can surface at different levels of granularity. Therefore, interactive exploration typically involves visual and data abstractions that can be adjusted to the task at hand. A DFG, for example, could visualize activity clusters as abstracted meta-nodes rather than showing individual events.

I6: Show me something conditionally. If users have some hypotheses about processes, they want to interactively specify them as conditions to dynamically query or filter the data and work on a reduced subset that meets the conditions. In a DFG, users may filter based on frequency such that infrequent behavior is not shown.

I7: Show me related things. Upon seeing some interesting patterns or behaviors in the data, users often need to answer the question whether related findings can be made elsewhere in the data space. For example, users may spot an interesting path in a DFG and may be interested to see additional paths that share a common sub-path.

While these seven interaction intents address different information needs of users during interactive visual data exploration, they all share a common characteristic: By interacting, users navigate a space of possibilities. When switching to a different visual representation, users make a choice among the available visual representations. When setting a filter threshold, users make a choice from the filter's value range. When marking data items, users select a subset among all possible subsets of the data.

Abstractly, all these interactions can be understood as setting a visualization parameter, that is, selecting a specific parameter value (or subset of values) from the parameter's value range [37]. The key challenge of interactive exploration is to make *informed* decisions on what parameter values to set next to ensure analytic progress. *Uninformedness* when setting parameter values, on the other hand, is a critical hindering factor, which could degrade data exploration to a search for a needle in a haystack.

3.2 What users face: Questions during data exploration

Uninformedness relates to users having difficulties answering relevant questions towards making decisions for exploration steps. In this regard, Spence [68] identified the following questions that users need to be able to answer quickly and effortlessly for effective and efficient interactive exploration:

Q1: Where am I? At any time during the exploration, users need to know where they are. This pertains to where the current partial view is located in the data space and also awareness of the currently set parameter values. For example, when zooming into specific details of a DFG, an overview should constantly indicate where in the data space the zoomed-in detail is located.

Q2: Am I where I want to be? Users should be able to validate whether their current view and parameter settings match the task at hand. For example, when filtering a DFG, it should be clear whether the currently filtered view contains the relevant information to make analytical progress. If this is not the case or if the task changes, further exploration steps are necessary.

Q3: Where can I go? Users must be able to get an overview of the space of possible next steps for the data exploration. Therefore, the available options should be clearly visible in the system. For example, the value range of possible thresholds for filtering operations on a DFG could be indicated by means of a slider control.

Q4: How do I get there? In addition to knowing possible next steps, users also need to be able to infer how to actually take the next steps. For standard interface controls accompanying a DFG, such as buttons or sliders this is typically granted. However, interactions performed directly within the DFG visualization should also be easy to discover, learn, and execute.

Q5: What lies beyond? When users see only a partial view of the data, they need to be able to estimate or make informed guesses on what relevant information lies beyond the current view. For a filtered DFG, for example, one may wonder whether relevant information might be found beyond the current filtering threshold.

Q6: Where can I most usefully go? Among the usually very many possible next steps for further data exploration, users need to decide on a good, ideally the best next step. However, this can be challenging for users to decide and for systems to provide valuable suggestions because what constitutes the most useful exploration step is generally hard to define.

Q7: Where have I been? As exploration involves surveying a space of possibilities and making discoveries, users need to have awareness of the parts of the space that have already been visited. This facilitates linking findings to higher-level insights and helps to avoid redundant exploration. A system could for example track which parts of a DFG have been shown at which level of detail.

These seven user questions provide a useful framework for the design of visual data exploration tools. Designers can critically reflect whether a developed approach offers adequate support for each of these questions. If adequate support is lacking, interaction costs increase and a smooth analytical discourse is hampered [43].

What is relevant: A moving target

It has to be noted that adequate support also requires a reasonable definition of what is relevant for a user during process exploration. Only with the help of such a definition can we say that an exploratory step is promising (if it leads to something relevant) or has a low prospect of success (if it leads to something less relevant). However, a concrete definition of what is relevant is very much dependent on the goal that a user aims to achieve and the tasks that have to be carried out along the way [8, 60]. And because goals and tasks typically change during exploration, such a definition is a moving target.

Conceptually, the relevancy of the data items being explored can be modeled via degree-of-interest (DoI) functions, originally introduced by Furnas [27] and later revised and expanded in other works [4, 29, 33]. Such DoI functions assign an interest value between 0 and 1 to data items and may consist of several weighted components that capture different aspects (e.g., a priori relevance, distance to focus, user interest, data item already visited). Later in Section 5, we will give a brief example of how such DoI functions can be used to guide the exploratory navigation of graphs.

Yet, before looking at this and other examples, we qualitatively examine how Yi et al.'s interaction intents and Spence's user questions are supported in interactive DFG interfaces of selected PM tools.

4 Tool support for process exploration

To contextualize the notion of *informed* process exploration, we review representative examples of interactive process exploration interfaces available in current process mining tools. For this purpose, we focus on visual interfaces featuring directly-follows graphs (DFGs) as they are widely used in practice for exploratory analysis [22, 32]. Beyond their prevalence, DFGs are particularly suitable for our review because they tightly couple visual representation and interaction: Analysts can directly adjust the DFG representations through interactive filtering and navigation operations. Examples of such interfaces in commercial tools are shown in Fig. 1 and Fig. 2. In PM tools, DFGs (also called "process maps") are often the entry point for exploration [87, 88].

Note that our goal is not to systematically evaluate existing tools, but to illustrate what interactions are available in selected tools (see Section 4.1) and discuss shared limitations that may hinder informed process exploration (see Section 4.2).

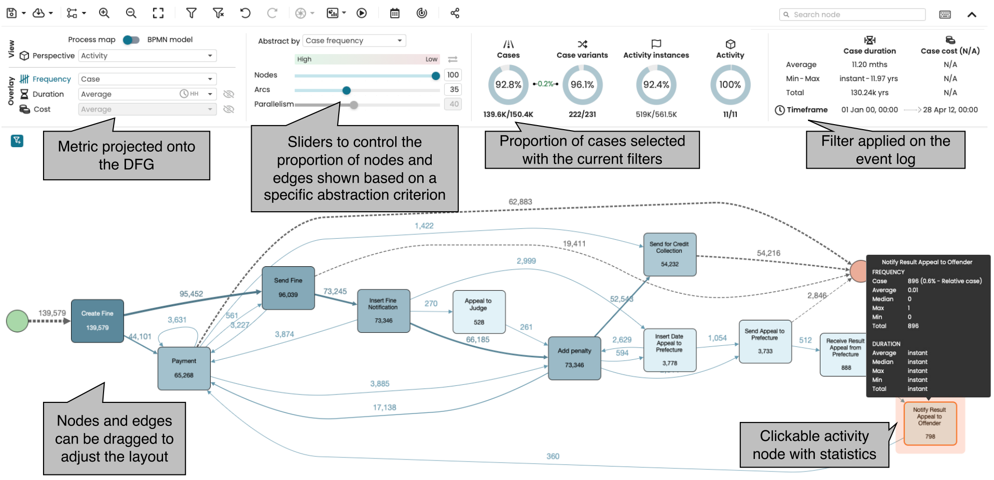

Fig. 2 Example of an interactive DFG, available in the *Process Discoverer* of Apromore. Some of the interactions available, such as node selection, node and edge dragging, and filtering are described in the call-outs. 🅭🅯

4.1 Support for DFG exploration in existing tools

We considered four widely used PM tools that feature interactive DFG-based exploration interfaces: the commercial tools Apromore[1], Celonis[2], and Disco by Fluxicon[3], as well as the open-source framework ProM[4].

The observations presented in the following emerged from discussions initiated within the author group during a Dagstuhl Seminar [17]. As a starting point, we jointly reflected on process exploration and established a shared understanding of typical interaction patterns and challenges in PM tools, supported by screen recordings and interaction logs collected in previous empirical studies [67, 83]. Then, we reviewed the selected tools regarding their support for the interaction intents and user questions introduced in Section 3.

The review was conducted by two authors, one from academia and one from industry, both with multiple years of experience at the intersection of PM and VA. Before conducting the review independently, the authors aligned on how the interaction concepts apply to DFG-based process exploration interfaces to ensure that the interaction intents and user questions were used consistently. Both authors replayed the same exploratory process mining scenario in different tools, focusing on steps involving the DFG. The scenario captures the realistic analysis of road traffic fines [45] as outlined in Section 2.2, in which an analyst iteratively applies filters, tests hypotheses, compares results, and revisits earlier exploration steps to investigate why certain fines remain unpaid. Individual observations were then consolidated and discussed in a joint session with a third author, an expert in VA, resulting in the qualitative support scores for each tool, interaction intent, and user question as summarized in Table 1. Finally, all co-authors condensed the review material to the observations presented in Section 4.2.

Details of our review of the four tools Apromore, Celonis, and Disco, and ProM are described in the following.

Apromore

Apromore, as shown in Fig. 2, provides an interactive DFG where nodes and edges can be explored through scrolling (*I2*) and by re-arranging nodes and edges to fit a specific layout, such as a user-defined happy path (*I3*). The user can switch between a DFG and a BPMN model (*I4*) and select which nodes are shown based on a specific abstraction chosen by the user among those available, using sliders (*I5*). Through the widgets showing case and variant statistics, the user can filter the event log for specific conditions (*I6*) and open further interfaces to explore individual cases.

Apromore's interactive DFG provides clear feedback to the user on their current focus (*Q1: Where am I?*) and what can be done next (*Q3: Where can I go?*). For example, clicking on a node (activity) or an edge (transition) highlights it and displays further statistical details. The interactions available with the DFG are also clear. Questions related to user intent and guidance, such as *Q6: Where can I most usefully go?* and *Q4: How do I get there?*, are only partially supported by color highlights that

[1] https://apromore.com
[2] https://www.celonis.com
[3] https://fluxicon.com/disco/
[4] https://promtools.org

indicate what stands out. Filtering, for instance, is handled in a separate interface rather than through direct manipulation of paths on the graph. The DFG dynamically updates to reflect any filters applied to the event log. While the user is made aware of the size of the data subset being analyzed, the excluded subsets (*Q5: What lies beyond?*) are not shown. Similarly, the tool provides no support for tracking a user's analytical journey (*Q7: Where have I been?*), for example, to highlight or pin previously explored elements or slider settings.

Celonis

Celonis provides a pan and zoom interface for navigating through different areas of the DFG (*I2*), which is shown in Fig. 1. Each graph element is interactive and can be clicked for more details (*I5*). Using slider and filter controls, users can conditionally filter on specific process behaviors (*I6*). However, it is not possible to mark something as interesting, change the arrangement and representation of the information, or identify similar patterns (*I1, I3, I4, I7*).

What data is in and out of scope of an analysis can be understood from the slider controls and the filter bar, which displays the active filter criteria (*Q1: Where am I?, Q5: What lies beyond?*). The DFG supports direct interaction with nodes and edges to retrieve detailed information and filter on specific process behaviors (*Q3: Where can I go?, Q4: How do I get there?*). A few visual cues like color highlight and edge thickness hint at potential areas of interest in the DFG, providing partial support for the question *Q6: Where can I most usefully go?*. While the filter bar provides an indication of what filter criteria are active, the question *Q7: Where have I been?* cannot be answered.

Disco

The DFG in Disco is the first view that is shown to the user after loading an event log. Disco allows users to scroll and zoom to see different areas of the DFG (*I2*) and to show more or less details by applying node and edge filters (*I5*). By clicking on specific nodes and edges, users can visualize statistics about that node (edge) and access a filter panel that allows them to filter the event log based on specific conditions (*I6*).

Disco allows an analyst to effectively answer *Q1: Where am I?* and *Q3: Where can I go?*, providing clear hints about what can be selected or clicked. Questions like *Q4: How do I get there?* and *Q6: Where can I most usefully go?* are partially supported by visual cues like color-coded highlights to indicate potential areas of interest rather than explicit guidance. Regarding *Q5: What lies beyond?*, Disco's DFG visualization is also re-rendered whenever filters are applied, and while the tool indicates the scope of the filtered log, it does not display the active filter criteria alongside the DFG or what data was left out. Lastly, Disco offers no functionality to address question *Q7: Where have I been?*, as it does not log or visualize the exploration history of the DFG.

ProM

In ProM, we reviewed the Directly Follows visual Miner (DFvM) [44]. Users can pan and zoom to navigate different parts of the graph (*I2*), and apply frequency-based filters using sliders to show or hide edges and nodes based on occurrence thresholds (*I5*). The tool also supports changing the event classifier to allow the user to control what attribute of the event log is shown as "activity" in the graph (*I2*). By clicking on nodes or edges, users can see detailed statistics about selected activities and paths, and use them to directly filter the event log (*I6*).

The DFvM provides consistent feedback on the current view through hover details and filter indicators (*Q1: Where am I?*), making available actions clear via sliders and buttons (*Q3: Where can I go?*). The controls are made available to the user via a menu that offers access to the functionality to manipulate the visualization (*Q4: How do I get there?*). The DFG also displays faded, filtered-out paths, providing a sense of the behavior that exists outside the current selection for one filtering step. However, the filter conditions are not shown explicitly (*Q5: What lies beyond?*). The integration of deviations and performance metrics projected directly onto nodes and edges can suggest where the user might go next (*Q6: Where can I most usefully go?*). However, the tool does not currently maintain a history of the DFG exploration, leaving *Q7: Where have I been?* unsupported.

The summarized support scores in Table 1 shows a consistent pattern of how interactive exploration of DFGs is supported across the reviewed tools. All of them enable basic navigation through pan and zoom, allow the inspection of nodes and edges, and provide mechanisms to filter the event log and view updated statistics. Concerning interaction intents, Apromore allows user to interact with the DFG to adjust the layout and change the modeling notation, something that is not available in the other tools. However, intents such as *I1: Mark something as interesting* and *I7: Show me related things* remain unsupported.

Table 1 Tool support (● full, ◐ partial ○ none) for interaction intents and user questions. For user questions, ◐ indicates that users have to face some trial-and-error exploration to answer a question.

Interaction Intents [81]	**Apromore**	**Celonis**	**Disco**	**ProM (DFvM)**
I1: Mark something as interesting	○	○	○	○
I2: Show me something else	●	●	●	●
I3: Show me a different arrangement	●	○	○	○
I4: Show me a different representation	●	○	○	○
I5: Show me more or less detail	●	●	●	●
I6: Show me something conditionally	●	●	●	●
I7: Show me related things	○	○	○	○
User Question [68]				
Q1: Where am I?	●	●	●	●
Q2: Am I where I want to be?	○	○	○	○
Q3: Where can I go?	●	●	●	●
Q4: How do I get there?	◐	◐	◐	◐
Q5: What lies beyond?	○	●	○	◐
Q6: Where can I most usefully go?	◐	◐	◐	◐
Q7: Where have I been?	○	○	○	○

Across the four tools, questions such as *Q1: Where am I?* and *Q3: Where can I go?* are consistently well supported, while *Q2: Am I where I want to be?*, *Q6: Where can I most usefully go?*, and *Q7: Where have I been?* remain largely unaddressed. This reveals a common limitation: Although the tools provide basic navigation and inspection capabilities, they do not explicitly take user intents into account. As a consequence, interactions remain mostly reactive, and users are required to navigate their analysis without adaptive guidance linked to their objectives.

While not exhaustive, our review illustrates the core interactions that many PM platforms provide for DFGs. Although the broader landscape of PM software offers a wide range of capabilities [47], the DFG remains the primary entry point for exploratory analysis. Given their importance and the significant effort that analysts invest in their study, it is reasonable to assume that our review of DFGs is relevant for understanding a large part of process exploration and how current tools support different interaction intents and user questions. Still, we acknowledge that our review focuses on a specific interface component and cannot capture every nuance of the broader functionality that PM systems offer.

4.2 Observations

In in-depth discussions based on our tool review, we identified several shared limitations in supporting informed process exploration, particularly when it comes to guiding the user in making decisions about the future effects of an interaction based on the current context, which is shaped by past interactions. Next we discuss these limitations along the lines of limited previews, limited reviews, and limited views beyond.

Limited previews

Overall, there is little support for forward-looking guidance (*Q4: How do I get there?* and *Q6: Where can I most usefully go?*). While tools use visual cues like color gradients or edge thickness to highlight aspects that stand out, such as less frequent or slow paths, they do not proactively guide the user towards what is most useful to explore, e.g., based on their implicit intent or a defined analytical goal. In particular, due to the multivariate nature of the process data [22], there are many possible ways to look at a process (many *views* that can be created on different process perspectives), as well as many alternative next steps that can be chosen at any point in the analysis [69]. This availability of many possible exploration paths makes it difficult for analyst to keep the exploration focused on selected aspects and to know what would be worthwhile to explore next.

Limited reviews

A significant concern is the general lack of support for backward-looking analysis (*Q7: Where have I been?*), a concept central to analytic provenance [55]. While some tools like RapidProM [3] and bupaRflow [70], or, simply, scripting tools allow analysts to construct reusable pipelines or code, these artifacts serve more as a script for re-enacting an analysis than for reviewing it. The resulting trail documents the operations performed but provides no insight into the analyst's journey nor cognitive support; it does not show what data was explored, what hypotheses were considered, or why certain paths were taken. This becomes even more problematic in interactive tools where exploration is fluid and (past) user interactions are not recorded. Without a mechanism to look back on this journey, users struggle to trace their reasoning, compare findings, and understand the context behind their conclusions [84].

Limited views beyond

The tools also do not offer explicit support to answer *Q5: what lies beyond*, for example, when a filter is applied to the event log. In most commercial tools, once a filter is applied, the visualization is entirely replaced with a new one that is derived on the subset of the original event log that met the filtering conditions. This interaction paradigm does not provide a persistent view of the original scope of the data and, in some cases, the criteria used for selection. The analyst is shown only the filtered result but loses sight of how this was obtained and what was left out. This can lead to misinterpretations, such as the analyst mistaking a pattern present only in a narrow subset for a general characteristic of the overall process. The DFvM in ProM is an exception to this, as all the activities and edges that are filtered out are not removed but faded. However, this works only for one filter applied: As soon as another filter is applied, information on the previous data selections is lost.

The points above, particularly the lack of support for backward and forward looking in the analysis, might explain the challenge of getting "lost" in the analysis [86]. Next, we outline some possible solutions to the identified problems.

5 Improving process exploration

A great amount of inspiration for solving the current challenges in process exploration can be drawn from the field of VA, which covers and combines research in visualization, interaction, and automatic analysis methods [74]. Given the limited previews, limited reviews, and limited views beyond identified earlier, the goal of this section is to suggest selected VA solutions for better previews, better reviews, and better views beyond to be integrated into PM practice.

A valuable theoretical foundation for such improved views can be found in Furnas' conceptual considerations on efficient view navigation [28]. Translated to process exploration, efficient view navigation corresponds to taking exploratory steps in a large space of options (e.g., filter thresholds, subset selections, types of visual representation). Furnas argues that exploring such spaces efficiently requires (1) traversibility, that is, the space of options can be traversed with reasonably small steps, and (2) navigability, that is, individual options are discoverable in order to actually navigate to them. Traversibility typically requires a suitable structuring of the space of options (e.g., through binning, sorting, or hierarchies). Navigability is dependent on suitable visual indicators of (selected) options. Interestingly, Furnas' work provides theoretical evidence how these aspects lead to effective navigation. Yet, it is beyond the scope of this work to go into the theoretical details. Instead, we next focus more on practical examples that illustrate the potential of considering VA technology, typically grounded in or inspired by Furnas' work, for informed process exploration.

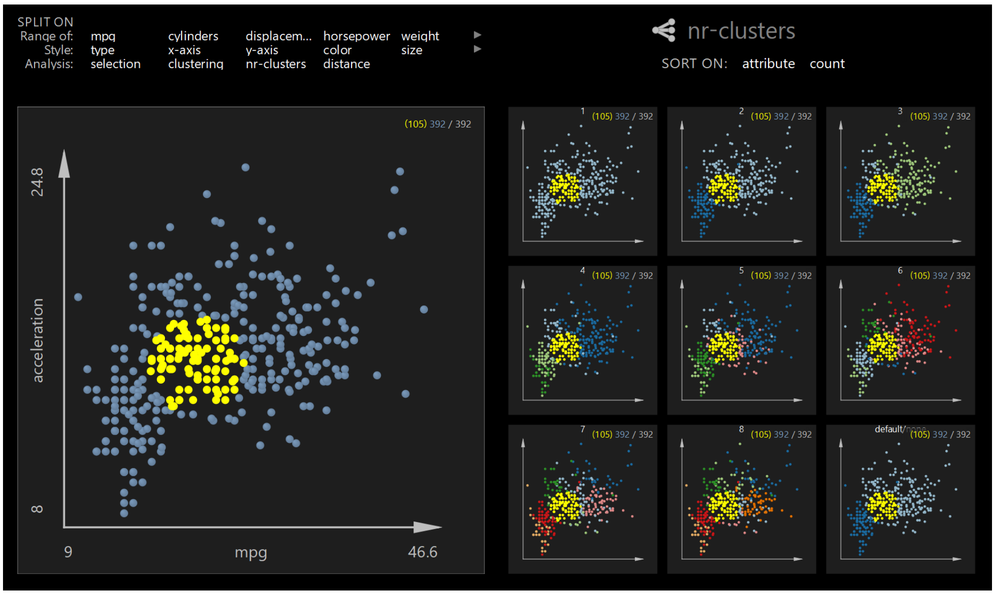

Fig. 3 Small previews (right) next to a large main visualization (left). Generated with [21]. 🅭🅯

5.1 Better previews

The purpose of previews is to indicate interaction results before the interaction has actually finished. In the literature, previews also relate to *foreshadowing* and *feed-forward*, where some scholars differentiate foreshadowing as previews that indicate possible interaction results from feedforward as previews that indicate possible ways of completing an interaction [19, 75]. In any case, previews inform about potential interaction outcomes and allow users to steer the exploration toward promising directions. As such, informative previews can reduce the amount of time spent on fruitless exploration paths.

Yet, previews are not easy to design. They share the same display space already occupied by the actual process visualization and need to be integrated within existing layouts of graphical elements and visual mapping strategies. In other words, previews require visual resources but should not conflict with the existing visual representation of the processes and data. Moreover, there is a computational overhead, simply because many different futures of an ongoing interaction are possible and may be worthwhile previewing. A non-trivial decision has to be made, which of these futures should be checked with regard to suitability and usefulness as a preview, and which of them should actually be presented to the user.

How foreshadowing can support the visual data exploration has been nicely demonstrated in the work by van Elzen et al. [21]. They describe a system where a large single view shows the main visualization (see Fig. 3). Next to it, several smaller previews show possible future visualizations that could be generated by varying a visualization parameter (note traversibility and navigability [28]). Instead of blindly using some user interface to adjust the visualization parameter, users can simply choose from the previews the one that they deem most useful to continue the exploration. Once a preview has been selected, it is fetched to the large main view and new previews can be generated for the next exploration step. The overall process implements a filmstrip metaphor to guide the exploratory navigation. Similar previews could also be shown alongside DFG visualizations, for example, to preview DFGs for different filter thresholds.

A nice example of a feedforward approach is Octopocus [6]. This technique helps users perform stroke gestures for triggering different actions. While a user is performing a stroke gesture, the system monitors the partial gesture and indicates the most likely valid gestures and the actions potentially being triggered. Users can immediately see if they perform gestures correctly toward the desired action and correct the gesture stroke should they be on a wrong path. Such interface assistance could be useful, for example, for collaborative process exploration on touch table devices. For the classic interface controls prevalent in PM tools, there are also suitable feedforward solutions [14].

In general, foreshadowing and feedforward are means for guiding users with well-known advantages [9, 10]. However, they are still rarely found in (commercial) visual exploration solutions. To a large degree this is due to the involved design challenges and the increased implementation and computation costs. Future research can address these challenges and reduce costs so that enhanced process exploration with better previews can be established in the long run.

5.2 Better reviews

Keeping track of performed exploration steps, especially their corresponding outcomes as intermediate results during process exploration can be mentally demanding for PM analysts. With better reviews, we refer to approaches that automatically keep track of the exploration history, enable undo/redo of exploration steps, and assist in maintaining analytical provenance in general. Being able to easily review or return to previous exploration states helps PM analysts integrate partial views into a big picture of the process. Reviews also help in understanding how analytical results were obtained, should it be necessary to replicate them at a later time.

The challenge with better reviews is that exploratory activities can result in complex, branching trace structures of states [25, 35, 42], where each state potentially has to hold a full replication of entire data (sub)sets or appropriate analytical abstractions of them. Therefore, effective strategies are needed to determine which exploration states should be stored at which level of abstraction. Supporting *I1: Mark something as interesting* can play a key role in this regard.

Given the size and complexity of trace structures, dedicated visual interfaces need to be provided to enable access to the exploration history. Such interfaces typically show the branching exploration history [42], and where possible, miniature visualizations indicate the different states of the exploration [35]. Once again, there is a conflict over display resources between the interface required for better reviews and the main process visualization.

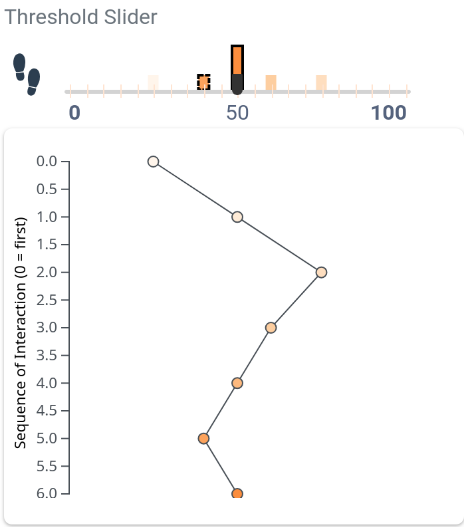

Fig. 4 A threshold slider enhanced with ProvenanceWidgets [52] shows aggregated provenance information as colored bars above the slider track, here the number of times a certain threshold was set. On demand, a separate view can be unfolded below the slider to see the sequence of past threshold changes leading to the currently set value of 50.

Therefore, detailed reviews are typically shown only on demand like, for example, in the ProvenanceWidgets by Narechania et al. [52]. They enhance classic control widgets with provenance information. For example, the slider in Fig. 4 is enriched with colored bars that indicate aggregated provenance information, in this case, the number of times a certain slider value (e.g., filter threshold) has been set. For a more detailed analysis, a dedicated view can be expanded below the slider to disclose how thresholds were chosen in the past.

Such enhanced widgets could also be beneficial in process exploration, especially for the filter sliders of DFG visualizations. An analyst could see which filter thresholds have already been explored, where frequently set thresholds may indicate useful ones that lead to insightful DFG representations. The positive effects on the exploration can be even stronger in collaborative settings, where the collective provenance information of multiple process analysts can enhance the process exploration in teams.

For a deeper discussion of provenance in VA, the interested reader is referred to a comprehensive survey on the analysis of user interactions and visualization provenance [79].

5.3 Better views beyond

Finally, process exploration can benefit from better views beyond. By that, we mean views or visual cues that indicate information that is not currently visible on the screen. Views beyond, or so-called off-screen visualizations [26, 31], allow users to discover items of the information space being explored even when it is represented only partially. Seeing beyond can help in understanding the current view on processes or data and planning the next exploration steps. As such, seeing beyond plays a key role in terms of Furnas' navigability [28]. Still, it is often considered optional support in VA solutions, and hardly present in PM tools.

Similar to the challenges regarding better previews, better views beyond also require additional computational resources to determine the information (or information residue) to be shown and visual resources to actually represent these additional pieces of information. To obtain scalable solutions, off-screen information can be indicated at different levels of abstraction [74]: (1) simply the fact whether task-relevant off-screen information is present, (2) direction, location, and/or distance where the off-screen information can be found, and (3) properties of the off-screen information that make it interesting or relevant in the current situation. While (1) requires rather little additional resources, (2) and (3) are more challenging to implement.

A classic approach that allows users to see beyond the current view are *scented widgets* [78], which is a precursor to the provenance widgets mentioned earlier. The idea of scented widgets is to enhance interface controls with additional information about the data. For example, a slider can be combined with a histogram that indicates how many data items exhibit certain values. During the exploration, an analyst can use the depicted value distribution to make informed decisions to explore either dense or sparse parts of the data. For a brief overview of related works on scented widgets, we refer the reader to [13]. Given the wide use of sliders in DFG interfaces in PM, it seems sensible to introduce these solutions to process exploration.

In addition to enhancing widgets, one can also enrich the visualization itself to facilitate seeing beyond the current view. An example is Gladisch et al.'s work [29] on supporting the exploration of hierarchical graphs (which are quite similar to the different levels of abstraction of DFGs). Based on a user-adjustable degree-of-interest (DoI) function, potential navigation targets are determined and ranked. For the top-ranked targets, visual cues in the form of wedges [31] are embedded at the view borders as shown in Fig. 5. A wedge indicates where and at what distance a relevant (wrt. to the DoI function) target can be found. To increase informativeness, small bar charts are embedded into the wedges to visualize the components of the DoI function

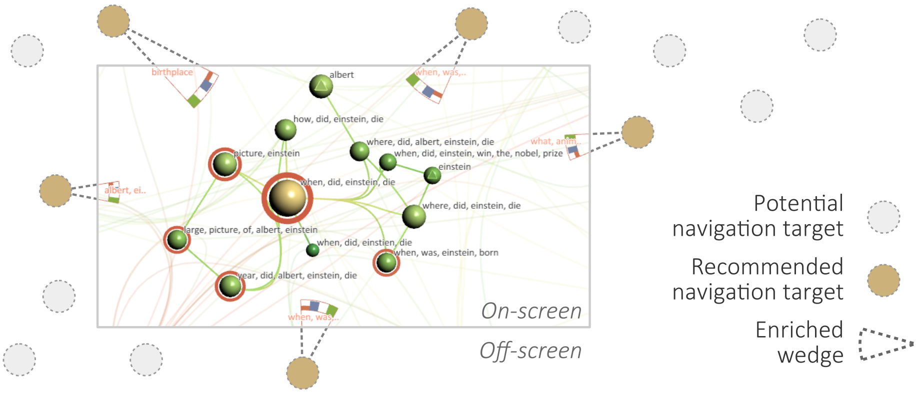

Fig. 5 Wedges at the border of the view can indicate possible next navigation steps to off-screen targets to support the exploration of hierarchical graphs. [74]

and indicate why the system thinks that a target might be of interest to the user. It seems natural that similar ideas could also enhance the exploration of DFGs or other graph-based process models.

Last but not least, there is the concept of "brushing & linking" [7, 12] that is universally applicable for better previews, reviews, and views beyond. Brushing enables *I1: Mark something as interesting*, to which the system reacts immediately with a suitable highlighting of the marked information. Linking then propagates the highlighting of related information to all other views, irrespective of whether they are alternative visualizations, previews, overviews, provenance views, or off-screen visualizations. This way, brushing & linking helps users understand information across different views. For a DFG, it can, for example, support before/after comparison for a given filter. When filtering on a subset of process cases, the system can overlay the filtered DFG onto the original DFG to help users understand the impact of the filter. While brushing & linking is a popular approach available in different variations in most VA solutions [41], current PM tools can still benefit from integrating it.

An example of how such an integration might look can be seen in Figure 6. It illustrates how the impact of applying a filter can be foreshadowed. One view is used to "brush" larger data values for filtering, while all "linked" views apply visual cues to indicate how the filter will change the visual representation. Note that this example is from a research prototype and not yet publicly released.

In summary, this section shed some light on potential solutions that could help improve process exploration. Yet, while these (and many more) conceptual solutions exist, more work is necessary to adjust them to the specific requirements in PM and integrate them into PM tools. The next section will outline corresponding research directions.

6 Towards informed process exploration

Our work identified the problem of uninformed process exploration as a key challenge in current PM practice. The lack of suitable, supportive information forces analysts into cognitively demanding iterations of *guess-then-evaluate-based-on-memory*, hindering their ability to build a coherent understanding of the process under investigation. Addressing this challenge, our vision is to enhance PM by leveraging VA theories and techniques, such as those described in Section 5. Next, we outline implications and advantages of this vision and discuss research directions for realizing it. By doing so, we call on the community to contribute to the systematic development of this body of knowledge, not only limited to DFGs but for the broader spectrum of PM tools.

6.1 Research vision

Our vision aspires to enhance *the process* of process exploration by integrating VA principles into PM tools. We see this as an interesting direction for developing effective means of identifying and bridging the gap between an analyst's intent and the insights derived from the visually represented data. In fact, our vision advocates a paradigm shift, moving from a one-way, command-and-inspect sequence to two-way informed

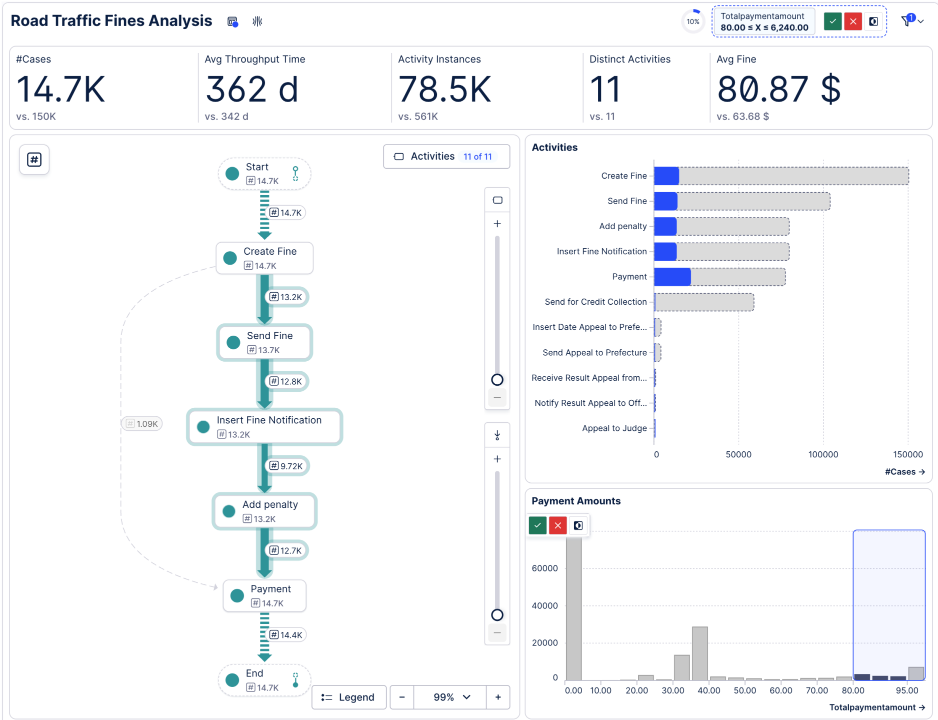

Fig. 6 Prototype of brushing & linking in a dashboard including a DFG. After filtering for cases with "Payment Amounts" larger than $80, a before and after comparison is shown across all views. In the DFG, nodes and edges that disappear after applying the filter are dimmed in gray, while those that appear are highlighted with a halo. This way, the impact of applying the filter is made clear to the user: While the original happy path, which consisted of a direct connection between "Create Fine" and "Payment", becomes gray, focusing on high payment amounts reveals additional events "Send Fine", "Insert Fine Notification" and "Add Penalty" (halo highlight).

analytical dialogue. In such a dialogue, the analyst would express their intents through interactions, and the system would respond not with a single, static result but with rich and dynamically visible context, promising alternatives, and potential consequences of future actions.

This proposed type of dialogue is facilitated by the three key improvements discussed above: better previews, reviews, and views beyond. Implementing these improvements would result in reduced cognitive load of analysts, by removing the burden of remembering past actions, mentally simulating the outcomes of future operations, and comparing different analytical states based on the analyst's bounded short-term memory and cognitive capacity. As discussed earlier, current PM tools tend to force process analysts to blindly apply different filter thresholds and compare the results from memory, leading to a cognitive strain. We advocate offloading these cognitive tasks (i.e., remembering, predicting, and comparing) onto the system, which can then present the required information visually as needed. By offloading the cognitive overhead, PM analysts are freed to focus on higher-level cognitive activities such as pattern detection, data interpretation, hypothesis generation and testing, sense-making, and conclusions [67]. Ultimately, the interactive environment is aligned more closely with the goals of process exploration, fostering analytical flow [20], making the process more effective, and leading to successful outcomes.

More concretely, the objective is to enable analysts to (a) initiate an action with a clear, preemptive understanding of its consequences, (b) reflect on the history of their analysis to recall and compare findings from previous paths, and (c) be informed about targets of interests beyond the current visualization displayed on the screen. Achieving these objectives would yield the following benefits. First, it would provide effective means for managing the analytical process, by making the exploration more transparent, guided, and reflective. Second, less time and effort will be invested in fruitless exploration paths that have a low prospect of success. Finally, the visual mechanisms will serve for identifying high-yield and low-risk analytical paths, enabling analysts to make informed decisions as they choose the direction of their exploration.

6.2 Future research directions

The main challenges of our research vision are (i) implementing appropriate visual analytic mechanisms for process exploration and (ii) evaluating if and how they enhance a more informed dialogue between the analyst and the PM tool. To address these challenges, future efforts must draw on knowledge from the domains of empirical PM research and VA theory.

A foundational step would be a systematic identification of the specific interaction challenges and cognitive bottlenecks that lead to uninformed exploration in current PM practice beyond just DFGs. Because most research efforts on PM have thus far focused on technical aspects rather than cognitive aspects, complementary empirical case studies and practitioner surveys are needed to establish a comprehensive understanding of the cognitive load and challenges faced by process analysts.

Building upon this understanding, future work should focus on designing and integrating advanced visual, interactive, and computational means to facilitate an enhanced analytic dialogue. Specifically, developing effective visual previews requires efficient algorithms capable of performing instantaneous dry runs on large-scale event logs. Methods from the area of *progressive VA* [23, 57] can generate increasingly mature results incrementally, which can help in tackling the computational challenges. Enhancing visual reviews necessitates provenance-aware interaction components that can capture the user's analytical journey without competing for primary screen real estate. Existing concepts for interaction histories [51] can be used as starting points. Providing views beyond the immediate screen requires the creation of context-aware visual cues that summarize off-screen process characteristics. Here, the literature on off-screen visualization [38] can provide valuable inspiration. Integrating all these VA concepts into the specific context of PM is a non-trivial endeavor that will require close collaboration among VA researchers, PM researchers, and tool providers. Open-source frameworks such as ProM[5] and KNIME [6], to only name two, can provide the infrastructure and technical basis for such collaborations.

[5] https://promtools.org/
[6] https://www.knime.com/

Finally, the impact of these enhanced tools must be rigorously evaluated. Future studies should assess the degree to which integrated VA techniques – encompassing previews, reviews, and views beyond – can mitigate cognitive load, improve parameter selection accuracy, and increase the generation of actionable insights. Importantly, because the value of these interactive techniques for exploratory purposes emerges naturally over time, testing their efficacy cannot rely solely on controlled lab experiments. Field experiments within real-world industry contexts and constraints will be essential to provide realistic, generalizable conclusions regarding analyst confidence, performance, and satisfaction.

7 Conclusion

In summary, this paper shifted the often more algorithmically-oriented research in process mining (PM) towards a more human-oriented perspective informed by visual analytics (VA). Our multi-disciplinary research at the intersection of PM and VA reported on the conceptual identification and discussion of uninformed process exploration. By assessing commonly used PM tools, in particular focusing on DFGs as the primary entry point for process exploration, we identified the currently limited previews, limited reviews, and limited views beyond, forcing analysts into a cognitively straining *guess-then-evaluate-based-on-memory* cycle. To address this, we propose a shift in the interface of PM tools toward an informed dialogue with the analyst based on VA principles, including feedforward mechanisms, analytic provenance, and off-screen visualization. These principles are to offload the heavy cognitive burdens of memory and prediction from the human analyst to the system.

Despite the potential of these proposed solutions, several limitations are to be acknowledged. The integration of advanced interactions and corresponding visual enhancements like foreshadowing and off-screen cues may introduce additional design challenges for visual PM interfaces and computational overhead. As process data grow in complexity, there is a risk that these guidance mechanisms would contribute to visual clutter if not carefully balanced. The anticipated effectiveness of these techniques in PM contexts, currently rests on VA knowledge alone. The true value of *informed* exploration in PM remains to be empirically evaluated through observational field studies in diverse industrial settings.

Still, the prospect for this research is highly promising. By following the proposed research directions from empirical challenge identification to VA-driven design to rigorous evaluation, we can move towards a new generation of PM tools that foster a cognitively effective analytical flow. The expected contribution of realizing this vision is the empowerment of the analyst: By offloading low-level cognitive tasks to the interface, we free the human analyst to focus on higher-level activities of sensemaking and knowledge generation. With this accomplished, we envision a future where informed process exploration enables a transparent, guided, and reflective experience that consistently leads to more accurate and actionable business insights.

Acknowledgements. We gratefully acknowledge that this work has been initiated during Dagstuhl-Seminar 25152 “Multi-Faceted Visual Process Mining and Analytics”.

Declarations

Author contribution

All authors conceptualized, wrote, revised, and approved the manuscript.

Funding

No funding was received to assist with the preparation of this manuscript.

Ethics

Not applicable.

References

[1] van der Aalst W (2016) Process Mining. Springer, https://doi.org/10.1007/978-3-662-49851-4

[2] van der Aalst WMP (2019) A Practitioner's Guide to Process Mining: Limitations of the Directly-Follows Graph. Procedia Computer Science 164. https://doi.org/10.1016/j.procs.2019.12.189

[3] van der Aalst WMP, Bolt A, van Zelst SJ (2017) RapidProM: Mine Your Processes and Not Just Your Data. CoRR abs/1703.03740. https://doi.org/10.48550/arXiv.1703.03740

[4] Abello J, Hadlak S, Schumann H, et al (2014) A Modular Degree-of-Interest Specification for the Visual Analysis of Large Dynamic Networks. IEEE Transactions on Visualization and Computer Graphics 20(3). https://doi.org/10.1109/TVCG.2013.109

[5] Augusto A, Conforti R, Dumas M, et al (2018) Automated Discovery of Process Models from Event Logs: Review and Benchmark. IEEE Transactions on Knowledge and Data Engineering 31(4). https://doi.org/10.1109/TKDE.2018.2841877

[6] Bau O, Mackay WE (2008) OctoPocus: A Dynamic Guide for Learning Gesture-Based Command Sets. In: Proceedings of the ACM Symposium on User Interface Software and Technology (UIST). ACM, https://doi.org/10.1145/1449715.1449724

[7] Becker RA, Cleveland WS (1987) Brushing Scatterplots. Technometrics 29(2). https://doi.org/10.2307/1269768

[8] Brehmer M, Munzner T (2013) A Multi-Level Typology of Abstract Visualization Tasks. IEEE Transactions on Visualization and Computer Graphics 19(12). https://doi.org/10.1109/TVCG.2013.124

[9] Ceneda D, Gschwandtner T, May T, et al (2017) Characterizing Guidance in Visual Analytics. IEEE Transactions on Visualization and Computer Graphics 23(1). https://doi.org/10.1109/TVCG.2016.2598468

[10] Ceneda D, Andrienko N, Andrienko G, et al (2020) Guide Me in Analysis: A Framework for Guidance Designers. Computer Graphics Forum 39(6). https://doi.org/10.1111/cgf.14017

[11] Chapela-Campa D, Dumas M, Mucientes M, et al (2022) Efficient Edge Filtering of Directly-Follows Graphs for Process Mining. Information Sciences 610. https://doi.org/10.1016/j.ins.2022.07.170

[12] Chen H (2004) Compound Brushing Explained. Information Visualization 3(2). https://doi.org/10.1057/palgrave.ivs.9500068

[13] Ciorna V, Médoc N, Melançon G, et al (2023) A Tour Through the Zoo of Scented Widgets. In: Eurographics Conference on Visualization - Posters. Eurographics Association, https://doi.org/10.2312/EVP.20231055

[14] Coppers S, Luyten K, Vanacken D, et al (2019) Fortunettes: Feedforward about the Future State of GUI Widgets. Proceedings of the ACM on Human-Computer Interaction 3(EICS). https://doi.org/10.1145/3331162

[15] Denisov V, Fahland D, van der Aalst WM (2018) Unbiased, Fine-Grained Description of Processes Performance from Event Data. In: Proceedings of the International Conference on Business Process Management (BPM). Springer, https://doi.org/10.1007/978-3-319-98648-7_9

[16] Di Ciccio C, Miksch S, Soffer P, et al (2024) Human in the (Process) Mines (Dagstuhl Seminar 23271). Dagstuhl Reports 13(7). https://doi.org/10.4230/DagRep.13.7.1

[17] Di Ciccio C, Soffer P, Tominski C, et al (2025) Multi-Faceted Visual Process Mining and Analytics (Dagstuhl Seminar 25152). Dagstuhl Reports 15(4). https://doi.org/10.4230/DagRep.15.4.28

[18] Dimara E, Perin C (2020) What is Interaction for Data Visualization? IEEE Transactions on Visualization and Computer Graphics 26(1). https://doi.org/10.1109/TVCG.2019.2934283

[19] Djajadiningrat T, Overbeeke K, Wensveen S (2002) But how, Donald, tell us how? On the creation of meaning in interaction design through feedforward and inherent feedback. In: Proceedings of the Conference on Designing Interactive Systems (DIS). ACM, https://doi.org/10.1145/778712.778752

[20] Elmqvist N, Moere AV, Jetter HC, et al (2011) Fluid Interaction for Information Visualization. Information Visualization 10(4). https://doi.org/10.1177/

1473871611413180

[21] van den Elzen S, van Wijk JJ (2013) Small Multiples, Large Singles: A New Approach for Visual Data Exploration. Computer Graphics Forum 32(3). https://doi.org/10.1111/CGF.12106

[22] van den Elzen S, Jans M, Martin N, et al (2025) Towards Multi-Faceted Visual Process Analytics. Information Systems 133. https://doi.org/https://doi.org/10.1016/j.is.2025.102560

[23] Fekete JD, Fisher D, Sedlmair M (eds) (2024) Progressive Data Analysis. Eurographics Association, https://doi.org/10.2312/pda.20242707

[24] Forsell C, Johansson J (2010) An Heuristic Set for Evaluation in Information Visualization. In: Santucci G (ed) Proceedings of the Conference on Advanced Visual Interfaces (AVI). ACM, https://doi.org/10.1145/1842993.1843029

[25] Freire J, Silva CT (2012) Making Computations and Publications Reproducible with VisTrails. Computing in Science & Engineering 14(4). https://doi.org/10.1109/MCSE.2012.76

[26] Frisch M, Dachselt R (2013) Visualizing Offscreen Elements of Node-Link Diagrams. Information Visalization 12(2). https://doi.org/10.1177/1473871612473589

[27] Furnas GW (1986) Generalized Fisheye Views. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). ACM, https://doi.org/10.1145/22339.22342

[28] Furnas GW (1997) Effective View Navigation. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). ACM, https://doi.org/10.1145/258549.258800

[29] Gladisch S, Schumann H, Tominski C (2013) Navigation Recommendations for Exploring Hierarchical Graphs. In: Advances in Visual Computing: Proceedings of the International Symposium on Visual Computing (ISVC). Springer, https://doi.org/10.1007/978-3-642-41939-3_4

[30] Guo Y, Guo S, Jin Z, et al (2022) Survey on Visual Analysis of Event Sequence Data. IEEE Transactions on Visualization and Computer Graphics 28(12). https://doi.org/10.1109/TVCG.2021.3100413

[31] Gustafson S, Baudisch P, Gutwin C, et al (2008) Wedge: Clutter-Free Visualization of Off-Screen Locations. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). ACM, https://doi.org/10.1145/1357054.1357179

[32] Häge MC, Rehse JR (2026) Mental Maps in Process Mining: Does Stabilizing DFGs Improve Process Analysis Performance? In: Proceedings of the International Conference on Business Process Management (BPM). Springer, https://doi.org/10.1007/978-3-032-02867-9_31

[33] van Ham F, Perer A (2009) Search, Show Context, Expand on Demand: Supporting Large Graph Exploration with Degree-of-Interest . IEEE Transactions on Visualization and Computer Graphics 15(6). https://doi.org/10.1109/TVCG.2009.108

[34] Heer J, Shneiderman B (2012) Interactive Dynamics for Visual Analysis. Communications of the ACM 55(4). https://doi.org/10.1145/2133806.2133821

[35] Heer J, Mackinlay JD, Stolte C, et al (2008) Graphical Histories for Visualization: Supporting Analysis, Communication, and Evaluation. IEEE Transactions on Visualization and Computer Graphics 14(6). https://doi.org/10.1109/TVCG.2008.137

[36] Hornbæk K, Oulasvirta A (2017) What Is Interaction? In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). ACM, https://doi.org/10.1145/3025453.3025765

[37] Jankun-Kelly T, Ma KL, Gertz M (2007) A Model and Framework for Visualization Exploration. IEEE Transactions on Visualization and Computer Graphics 13(2). https://doi.org/10.1109/TVCG.2007.28

[38] Jäckle D, Kwon BC, Keim DA (2015) Off-Screen Visualization Perspectives: Tasks and Challenges. In: Proceedings of the Symposium on Visualization in Data Science (VDS) at IEEE VIS, URL https://scibib.dbvis.de/publications/view/632

[39] Keim D, Kohlhammer J, Ellis G, et al (eds) (2010) Mastering the Information Age – Solving Problems with Visual Analytics. Eurographics Association, https://doi.org/10.2312/14803

[40] Klinkmüller C, Müller R, Weber I (2019) Mining Process Mining Practices: An Exploratory Characterization of Information Needs in Process Analytics. In: Proceedings of the International Conference on Business Process Management (BPM), Springer, https://doi.org/10.1007/978-3-030-26619-6_21

[41] Koytek P, Perin C, Vermeulen J, et al (2017) MyBrush: Brushing and Linking with Personal Agency. IEEE Transactions on Visualization and Computer Graphics 24(1). https://doi.org/10.1109/TVCG.2017.2743859

[42] Kreuseler M, Nocke T, Schumann H (2004) A History Mechanism for Visual Data Mining. In: Symposium on Information Visualization (InfoVis). IEEE Computer Society, https://doi.org/10.1109/INFVIS.2004.2

[43] Lam H (2008) A Framework of Interaction Costs in Information Visualization. IEEE Transactions on Visualization and Computer Graphics 14(6). https://doi.org/10.1109/TVCG.2008.109

[44] Leemans SJ, Poppe E, Wynn MT (2019) Directly Follows-Based Process Mining: Exploration & a Case Study. In: Proceedings of the International Conference on Process Mining (ICPM), IEEE Computer Society, https://doi.org/10.1109/icpm.2019.00015

[45] de Leoni M, Mannhardt F (2015) Road Traffic Fine Management Process. Dataset, Eindhoven University of Technology, https://doi.org/10.4121/uuid:270fd440-1057-4fb9-89a9-b699b47990f5

[46] van der Linden S, de Fouw E, van den Elzen S, et al (2023) A Survey of Visualization Techniques for Comparing Event Sequences. Computers & Graphics 115. https://doi.org/10.1016/j.cag.2023.05.016

[47] Loyola-González O (2023) Process Mining: Software Comparison, Trends, and Challenges. International Journal of Data Science and Analytics 15(4). https://doi.org/10.1007/s41060-022-00379-0

[48] McGuffin MJ, Fuhrman CP (2020) Categories and Completeness of Visual Programming and Direct Manipulation. In: Proceedings of the Conference on Advanced Visual Interfaces (AVI). ACM, https://doi.org/10.1145/3399715.3399821

[49] Miksch S, Ciccio CD, Soffer P, et al (2024) Visual Analytics Meets Process Mining: Challenges and Opportunities. IEEE Computer Graphics and Applications 44(6). https://doi.org/10.1109/MCG.2024.3456916

[50] Montana L, Resinas M, Villa-Uriol MC (2025) Addressing Directly-Follows Graphs Limitations with Visualization of Event Sequences. In: Proceedings of the Visual Process Analytics Workshop (VIPRA). Springer, https://doi.org/10.1007/978-3-032-13426-4_42

[51] Nancel M, Cockburn A (2014) Causality: A Conceptual Model of Interaction History. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). ACM, pp 1777–1786, https://doi.org/10.1145/2556288.2556990

[52] Narechania A, Odak K, El-Assady M, et al (2025) ProvenanceWidgets: A Library of UI Control Elements to Track and Dynamically Overlay Analytic Provenance. IEEE Transactions on Visualization and Computer Graphics 31(1). https://doi.org/10.1109/TVCG.2024.3456144

[53] Nielsen J (1994) Enhancing the Explanatory Power of Usability Heuristics. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). ACM, https://doi.org/10.1145/191666.191729

[54] Norman DA (2013) The Design of Everyday Things, revised and expanded edition edn. Basic Books, https://doi.org/10.15358/9783800648108

[55] North C, Chang R, Endert A, et al (2011) Analytic Provenance: Process+Interaction+Insight. In: CHI Extended Abstracts on Human Factors in Computing Systems. ACM, https://doi.org/10.1145/1979742.1979570

[56] Pike WA, Stasko JT, Chang R, et al (2009) The Science of Interaction. Information Visualization 8(4). https://doi.org/10.1057/IVS.2009.22

[57] Pérez-Messina I, Angelini M, Ceneda D, et al (2025) Coupling Guidance and Progressiveness in Visual Analytics. Computer Graphics Forum 44(3):e70115. https://doi.org/10.1111/cgf.70115, URL https://doi.org/10.1111/cgf.70115

[58] Rubensson C, Mendling J (2025) Time-Order Map for Seamless Zooming between Process Models and Process Instances. In: Proceedings of the International Conference on Process Mining (ICPM). IEEE Computer Society, https://doi.org/10.1109/ICPM66919.2025.11220736

[59] Salas-Urbano M, Capitán-Agudo C, Cabanillas C, et al (2026) Interactive Creation, Visualization, and Exploration of Process Model Collections. Software and Systems Modeling https://doi.org/10.1007/s10270-025-01346-6

[60] Schulz HJ, Nocke T, Heitzler M, et al (2013) A Design Space of Visualization Tasks. IEEE Transactions on Visualization and Computer Graphics 19(12). https://doi.org/10.1109/TVCG.2013.120

[61] Schuster D, Zerbato F, van Zelst SJ, et al (2024) Defining and Visualizing Process Execution Variants from Partially Ordered Event Data. Information Sciences 657. https://doi.org/10.1016/j.ins.2023.119958

[62] Schuster D, Aigner W, Di Francescomarino C, et al (2026) KAVA-PM: Knowledge-Assisted Visual Process Mining. Information Systems 137. https://doi.org/10.1016/j.is.2025.102638

[63] Sedig K, Parsons P (2016) Design of Visualizations for Human-Information Interaction: A Pattern-Based Framework. No. 4 in Synthesis Lectures on Visualization, Morgan & Claypool, https://doi.org/10.2200/S00685ED1V01Y201512VIS005

[64] Shneiderman B (1983) Direct Manipulation: A Step Beyond Programming Languages. IEEE Computer 16(8). https://doi.org/10.1109/mc.1983.1654471

[65] Shneiderman B, Plaisant C, Cohen M, et al (2017) Designing the User Interface: Strategies for Effective Human-Computer Interaction, 6th edn. Pearson

[66] Shurkhovetskyy G, Schumacher P, Koytek P (2024) Interacting with Large Process Data: Challenges in Visual Exploration. In: Proceedings of the Visual Process

Analytics Workshop (VIPRA). The Eurographics Association, https://doi.org/10.2312/vipra.20241102

[67] Sorokina E, Soffer P, Hadar I, et al (2023) PEM4PPM: A Cognitive Perspective on the Process of Process Mining. In: Proceedings of the International Conference on Process Mining (ICPM), Springer, https://doi.org/10.1007/978-3-031-41620-0_27

[68] Spence R (2014) Information Visualization, 3rd edn. Springer, https://doi.org/10.1007/978-3-319-07341-5

[69] Steegen S, Tuerlinckx F, Gelman A, et al (2016) Increasing Transparency Through a Multiverse Analysis. Perspectives on Psychological Science 11(5). https://doi.org/10.1177/1745691616658637

[70] Steukers B, Janssenswillen G, van Hulzen GAWM, et al (2022) bupaRflow: A Workflow Interface for bupaR. In: Proceedings of the International Conference on Business Process Management (BPM) – Demonstration & Resources Track, URL https://ceur-ws.org/Vol-3216/paper_245.pdf

[71] Tentina I, Zerbato F, Mannhardt F, et al (2025) Why Do Users Struggle to Get Insights from Process Mining? In: Proceedings of the International Conference on Process Mining (ICPM). IEEE Computer Society, https://doi.org/10.1109/ICPM66919.2025.11220617

[72] Thomas JJ, Cook KA (2005) Illuminating the Path: The Research and Development Agenda for Visual Analytics. IEEE Computer Society

[73] Tominski C (2015) Interaction for Visualization. No. 3 in Synthesis Lectures on Visualization, Morgan & Claypool, https://doi.org/10.2200/S00651ED1V01Y201506VIS003

[74] Tominski C, Schumann H (2020) Interactive Visual Data Analysis. AK Peters Visualization Series, CRC Press, https://doi.org/10.1201/9781315152707, URL https://ivda-book.de

[75] Vermeulen J, Luyten K, van den Hoven E, et al (2013) Crossing the Bridge Over Norman's Gulf of Execution: Revealing Feedforward's True Identity. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI). ACM, https://doi.org/10.1145/2470654.2466255

[76] Wall E, Agnihotri M, Matzen LE, et al (2019) A Heuristic Approach to Value-Driven Evaluation of Visualizations. IEEE Transactions on Visualization and Computer Graphics 25(1). https://doi.org/10.1109/TVCG.2018.2865146

[77] van Wijk JJ (2006) Views on Visualization. IEEE Transactions on Visualization and Computer Graphics 12(4). https://doi.org/10.1109/TVCG.2006.80

[78] Willett W, Heer J, Agrawala M (2007) Scented Widgets: Improving Navigation Cues with Embedded Visualizations. IEEE Transactions on Visualization and Computer Graphics 13(6). https://doi.org/10.1109/TVCG.2007.70589

[79] Xu K, Ottley A, Walchshofer C, et al (2020) Survey on the Analysis of User Interactions and Visualization Provenance. Computer Graphics Forum 39(3). https://doi.org/10.1111/cgf.14035

[80] Yeshchenko A, Mendling J (2024) A Survey of Approaches for Event Sequence Analysis and Visualization. Information Systems 120. https://doi.org/10.1016/j.is.2023.102283

[81] Yi JS, ah Kang Y, Stasko JT, et al (2007) Toward a Deeper Understanding of the Role of Interaction in Information Visualization. IEEE Transactions on Visualization and Computer Graphics 13(6). https://doi.org/10.1109/TVCG.2007.70515

[82] Zerbato F, Soffer P, Weber B (2021) Initial Insights into Exploratory Process Mining Practices. In: Proceedings of the Business Process Management Forum, Springer, https://doi.org/10.1007/978-3-030-85440-9_9

[83] Zerbato F, Soffer P, Weber B (2022) Process Mining Practices: Evidence from Interviews. In: Proceedings of the International Conference on Business Process Management (BPM), Springer, https://doi.org/10.1007/978-3-031-16103-2_19

[84] Zerbato F, Burattin A, Völzer H, et al (2023) Supporting Provenance and Data awareness in Exploratory Process Mining. In: Proceedings of the International Conference on Advanced Information Systems Engineering, Springer, https://doi.org/10.1007/978-3-031-34560-9_27

[85] Zerbato F, Franceschetti M, Weber B (2024) A Framework to Support the Validation of Process Mining Inquiries. In: Proceedings of the Business Process Management Forum. Springer, https://doi.org/10.1007/978-3-031-70418-5_15

[86] Zimmermann L, Zerbato F, Weber B (2023) What Makes Life for Process Mining Analysts Difficult? A Reflection of Challenges. Software and Systems Modeling 23. https://doi.org/10.1007/s10270-023-01134-0

[87] Zimmermann L, Zerbato F, Vrotsou K, et al (2025) A Process-Oriented Approach to Analyze Analysts' Use of Visualizations: Revealing Insights into the What, When, and How. Computer Graphics Forum 44(3). https://doi.org/10.1111/cgf.70104

[88] Zimmermann L, Zerbato F, Weber B (2026) Visualization Use in Process Mining Analysis. In: Proceedings of the Conference Enterprise, Business-Process and Information Systems Modeling (BPMDS). Springer, https://doi.org/10.1007/978-3-032-28274-3_4

[89] Zinat KT, Sakhamuri SN, Chen AS, et al (2025) A Multi-Level Task Framework for Event Sequence Analysis. IEEE Transactions on Visualization and Computer Graphics 31(1). https://doi.org/10.1109/TVCG.2024.3456510

[90] Zuk T, Schlesier L, Neumann P, et al (2006) Heuristics for Information Visualization Evaluation. In: Proceedings of the Workshop on Beyond time and errors: novel evaluation methods for information visualization (BELIV). ACM, https://doi.org/10.1145/1168149.1168162