17 Hans-Jörg Schulz and Chris Weaver. Transient visual analytics. In Mennatallah El-Assady and Hans-Jörg Schulz, editors, *Proceedings of the 15th International EuroVis Workshop on Visual Analytics (EuroVA'24)*. Eurographics Association, 2024.

18 Wil M. P. van der Aalst. *Process Mining – Data Science in Action, Second Edition*. Springer, 2016.

19 Jarke J Van Wijk. Views on visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):421–432, 2006.

4.3 Interactivity: Visual Feedback and Feedforward for Process Exploration

Francesca Zerbato (TU Eindhoven, NL), Irit Hadar (University of Haifa, IL), Philipp Koytek (Celonis Labs GmbH – München, DE), Christian Tominski (Universität Rostock, DE), and Maria-Cruz Villa-Uriol (University of Sheffield, GB)

This working group focused on the role that interactivity plays in supporting the exploration of processes in process mining (PM). The group first analyzed the current PM practice and its limitations. Then existing works related to interactive visual data exploration were collected from the Visual Analytics (VA) literature. Based on that, preliminary formalizations were sythesized and initial design ideas sketched. In particular, the focus was on enhancing PM with informative visual feedback and feedforward techniques from the VA realm.

4.3.1 The Problem of (Un)Informed Process Exploration

Like interactive visual data analysis in general [19], process mining in particular is exploratory and human-driven. Process analysts typically have to engage in iterative exploration cycles of process visualizations (see Figure 17) to build an understanding of the process, examine different scenarios, and generate hypotheses [26]. Hypotheses are then tested, refined, or discarded based on intermediate insights, which, in turn, guide analysts in choosing what to explore next [18] and lead to the crystallization of knowledge [21].

Many process mining tools support this *interactive exploration* through components such as filter masks or sliders that allow analysts to create views and isolate data subsets of interest. However, these interactions often lack transparency and context, making it difficult for users to anticipate the effects of their actions before executing them (*What should I usefully do?*) and understand the effect once an action has been executed (*Have I achieved the desired outcome?*).

As an example, consider a process mining analyst, let's name him Bob, who is in charge of analyzing a Road Fine Management process [7] visualized as a directly-follows graphs (DFG). His goal is to investigate cases where offenders do not pay their fines. Bob's analysis steps are sketched in Figure 18.

As a first step, Bob loads the raw data provided by the police (L_0) into a process mining tool. His goal is to get an initial understanding of the structure of the process. To achieve this goal, he intents to focus on the most common behavior using a *variant filter* to remove infrequent cases. However, to find a suitable abstraction level for his analysis, Bob has to go through the costly procedure of applying the filter three times (o_1, o_2, o_3). He selects different filtering thresholds of 75%, 85%, and 90% of the cases in the log, and each time he

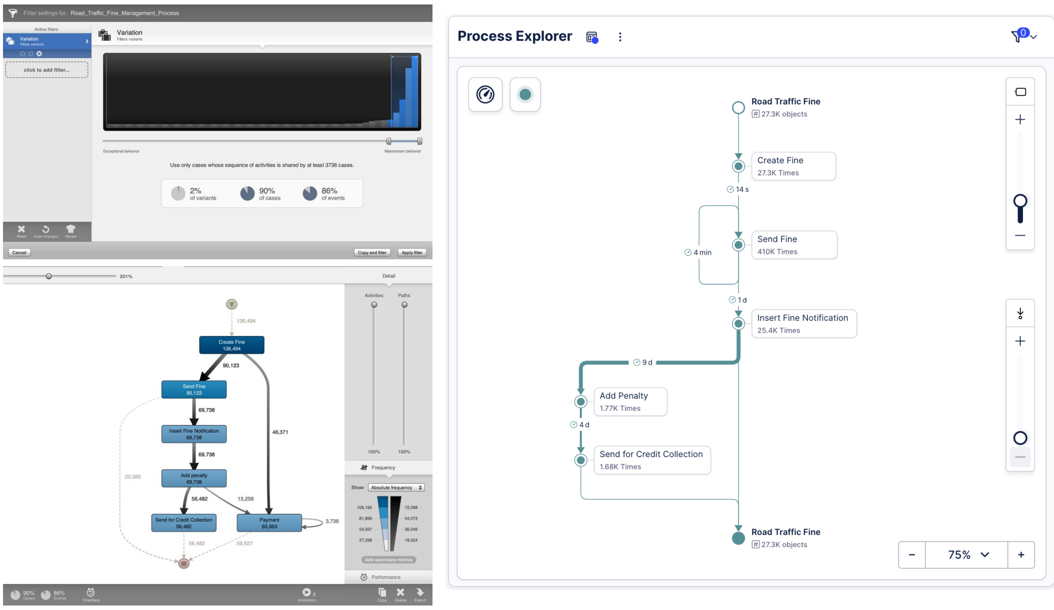

Figure 17 Most process mining tools provide visualizations of DFGs that can be filtered using sliders. However these sliders typically do not provide immediately visual feedback, which can make it difficult to understand their effects on the data and the visualizations.

	Id	Operation	I/O		Timestamp	User Intent
DFG Exploration	o_1	variantFilter(cases, keep, 75%)	L_0	L_1	07/10/22 10:01:18	Focus on most frequent behavior in DFG
	v_1	showDFG()	L_1	DFG_1	07/10/22 10:01:50	
	o_2	variantFilter(cases, keep, 85%)	L_0	L_2	07/10/22 10:02:03	
	v_2	showDFG()	L_2	DFG_2	07/10/22 10:02:32	
	o_3	variantFilter(cases, keep, 90%)	L_0	L_3	07/10/22 10:03:11	
	v_3	showDFG()	L_3	DFG_3	07/10/22 10:03:29	

Figure 18 Bob's first analysis steps adapted from [25] represented as **Operations**, the input and output **I/O**, the **Timestamp** at which each operation occurred and the higher level **User Intent**.

inspects the resulting DFG (v_1, v_2, v_3) to assess visually the effects of the filters. While the first two filter configurations remove too many cases, he settles for o_3.

This interactive selection of most frequent cases is a common first step of many process mining analyses. However, during this process, Bob runs into several limitations:

- There is no preview mechanism to support the decision for a suitable filter threshold: Bob must fully apply each filter to see what the result looks like;
- Comparisons across multiple visualizations resulting from the filtering are lost: Since each filter operation completely replaces the DFG with a new one, Bob is forced to take screenshots or compare from short-term memory;
- The DFG shows the impact of his filter on the control-flow only: Bob would need to create different views of the data to see how his filtering impacts other facets of the process.

The difficulties Bob faces demonstrate some of the challenges that stem from limitations of interaction functionalities currently being used in PM practice. These limitations introduce constraints to the process of process mining (PPM) [18] and potentially hinder performance when making sense of event log data. The objective of our working group is to investigate these challenges and outline the opportunity of enhancing the process exploration based on established concepts, methods, and techniques from the realm of VA.

With our working group, we aim to improve the overall process of making sense during the PPM. To this end, we compiled relevant previous work from the PM and VA communities, both to understand some of the cognitive challenges during process exploration and to bring together models and approaches that can inform the design of advanced process exploration techniques to overcome these challenges.

4.3.2 Relevant Works Related to Interactivity in Process Exploration

During the group discussion, we considered several works from PM and VA. These works are concerned with the cognitive, processing, and interactive mechanisms that are relevant during visual process exploration.

Related work on Process Mining

The process of process mining (PPM) is a recent stream of research studying the sequence of activities – both from behavioral and cognitive points of view [18, 26, 27]. Such studies are important for informing efforts toward providing support to the cognitive processes underlying the PPM. The PPM starts based on a general goal (e.g., identifying obstacles in the process), building on available event datasets, and continues to additional operations, such as filtering the data to explore it from different angles, interpreting the data, and trying to make sense of them in order to find insights relevant to the goal at hand.

Looking at the PPM from a cognitive perspective, during process mining, the data serve as input signals coming from the 'external world'. The sense-making process entails an iterative cycle, where attention is focused according to a set goal, leading to the generation of hypotheses about the process, which are then tested and reconsidered against the data for minimizing the prediction error [18]. This process not only aligns with general knowledge generation processes in VA [21, 15], but also with the post-cognitivism principle of prediction error minimization (PEM) [6, 11] in particular.

PEM conceives the brain as a probabilistic inference system, which attempts to predict the input it receives by constructing models of the possible causes of this input. While aiming to minimize the prediction error (i.e., the gap between the predicted and the actual input), it either introduces small refinements to the model or substantial revisions (or even a complete replacement of the model), depending on the size of the error. This process is iteratively performed until the prediction error is satisfactorily small.

Figure 19 illustrates the adapted model proposed by Sorokina et al. [18] named PEM4PPM. The model captures the sequence of PM steps and their corresponding cognitive operations. It begins with high-level business goals that can be decomposed or refined into more specific ones as needed. The refinement process iterates until the goals are concrete enough to be achieved through available PM operations. To focus attention on studied aspects of the input data, a relevant subset of the data is filtered and organized, enabling subsequent exploration of the data to identify behavioral patterns that are of interest. Data exploration is conducted to uncover behavior patterns that may be relevant to the set goals. Based on the exploration results, concrete predictions are generated, in the form of hypotheses or artifacts (e.g., process models) and then tested. The results obtained from these steps are assessed against the original goal or hypothesis to evaluate prediction errors and take actions for their minimization. This assessment serves as a basis for determining whether the goal has been achieved, thus leading to a conclusion, or if further refinement is needed, in which case the process continues in another iteration [18].

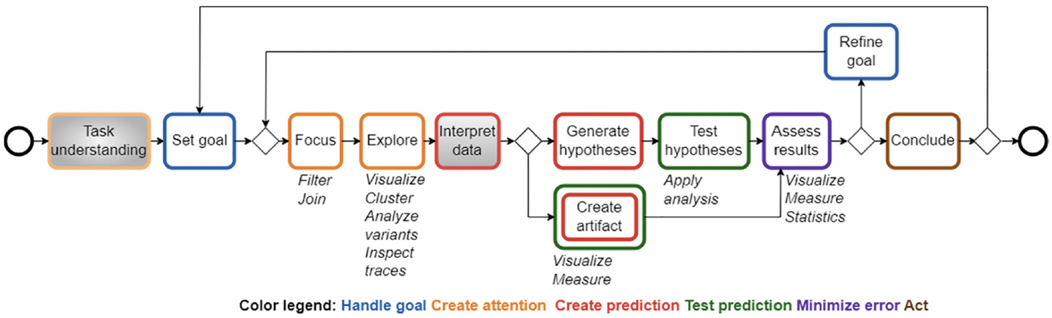

Figure 19 The PEM4PPM model illustrating the cognitive steps of the Process of Process Mining. Image from [18].

Related Works from Visual Analytics

The VA community has worked extensively on interactive visual data exploration. There are several important related works that are relevant in this regard, underpinning effective user interaction and data exploration:

- On a more abstract level than the PEM4PPM, **Norman's action cycle** [14] provides a crucial framework for understanding the stages users go through when interacting with a system, emphasizing the gulfs of execution and evaluation that interactive visualization interfaces should aim to bridge in order to minimize interaction costs [13].
- **Shneiderman's visual information seeking mantra** [17] characterizes the general process of interactive data exploration as: *Overview first, zoom and filter, then details on demand.* This mantra has been expanded to the **Visual Analytics Mantra** [12]: *Analyze first, show the important, zoom and filter, and analyze further, details on demand.*
- **Brushing & linking** [2] and **dynamic queries** [16] provide techniques that enable simultaneous highlighting and filtering of related data in different views. They offer users immediate and continuous visual feedback as they manipulate query parameters, fostering an iterative and exploratory analysis process across multiple perspectives.
- **Fluid interaction** [8] has been conceived to create seamless and responsive visualization interfaces that minimize the cognitive load of interaction, allowing users to stay in the analysis flow and focus on data insights.
- **Guidance mechanisms** [4, 3] aim to actively support users in their analytical process, ranging from subtle visual cues to more explicit recommendations, helping them navigate complex datasets and analysis tasks effectively.
- **Visual Feedback and Feedforward** [22] are essential principles for designing intuitive interactive systems. Commonly, visual feedback informs users about the results of their actions. However, only rarely is visual feedforward applied to provide users with cues suggesting available options and potential interaction outcomes.

These interconnected concepts collectively contribute to the design of powerful and user-friendly VA tools. This working group particularly focused on visual feedback and feedforward as promising, yet so far under-explored mechanisms to enhance process exploration in PM. By dynamically adding information to existing visual process representations, they can support the understanding of interaction effects and the decisions of the user about their next activity.

4.3.3 Conceptualization of Feedback and Feedforward

In an attempt to pinpoint the fundamental conceptual aspects of the desired process exploration support, we came up with the following (incomplete) list of notations inspired by the section on interactive selection and accentuation in [19]:

- D, the data to be visualized, explored, and understood
- $D_+ \subseteq D$, the currently relevant focus data, subject to change frequently during process exploration
- $D_- = D \setminus D_+$, the data currently not being of relevance for the process exploration
- S, a state capturing the data underlying the visualization
- S_{cur}, the "current" state
- S_{old}, the "old" state
- $\{S_{a_1}, S_{a_2}, \ldots, S_{a_n}\}$, a set of possible (useful) "alternative" states that can be entered through alternative interactions
- $\delta(S_i, S_j)$, the explicit difference(s) between two states
- ...

Based on these notations, we defined exploration as the repeated refinement and change of D^+ (and D_- respectively), which usually involves numerous state changes (e.g., the three different filtering states in Bob's exploration example). Moreover, it seems that understanding state changes is crucial for effective exploration. Possible options for supporting the understanding of state changes can be based on Gleicher et al.'s [10] strategies for visual comparison:

- **Juxtaposition:** Visualize S_i and S_j side by side
- **Superposition:** Superimpose the visualization of S_i over the visualization of S_j
- **Explicit encodingg:** Visualize $\delta(S_i, S_j)$ directly

So far, these strategies are not sufficiently integrated into existing process exploration practice!

To understand the users' needs better, we further conceptualized a cycle of interaction for a seamless analysis that is based on feedback and feedforward. Generally, in interactive exploration/analysis, the analyst interacts with the visualizations for dicing, slicing and relating different parts of the data. So, the exploration cycle starts with the analyst expressing an **intent to change** the visual representation. Yi et al. [24] identified several different categories of interaction intents called "Show me...", of which we focus on the intents related to changing the focused subset of the data exploration:

- **Show me something else:** A different subset of the data (e.g., navigate in time) will be visualized; $D_+ \rightarrow D'_+$.
- **Show me more/less:** A subset of different size or level of aggregation (e.g., reduce number of nodes in DFG) will be visualized; $|D_+| \neq |D'_+|$.
- **Show me something conditionally:** A subset that fulfills certain (filter) condition(s) (e.g., filter for frequent variants) will be visualized; $C(D_+) = \mathit{true}$.
- **Show me related things:** A subset that is (in some way) related to the currently shown subset (e.g., brushing and linking across multiple faceted views) will be visualized; $R(D_+, D'_+)$, typically $D_+ \sim D'_+$.

These intents lead to interactions to which the system responds by providing visual feedback, and what we would like to emphasize, also visual feedforward. We envisioned the following scenario in which a user executes an interaction and receives the corresponding visual feedback and feedforward.

1. The user's intention is typically communicated to the system through different ways of interaction (e.g., hovering a visual mark in the visualization or clicking a button or slider in the user interface).
2. The system interprets the user's action and then provides *relevant* context and suggests possible next steps with previews of their impact. Here we can explore a large design space of different useful visual feedback and feedforward, which generally are dependent on the semantics of the interaction and will incorporate different facets of the data. The additional context and possible next steps help the analyst to decide if the intended interaction outcome has been obtained, and if not, how to execute their alternative more fruitful interactions.
3. Based on the visual feedback and feedforward, the analyst can now better understand interaction effects and can more easily decide what to do next. The cycle starts again as the analyst continues the exploration and expresses their new intents through new interactions.

With this general scenario now being clear, the question that remains is how to design the visual feedforward and feedback concretely. However, given the huge design space, this is quite a challenging task.

4.3.4 Preliminary Design Examples

For our design sketches, we drew inspiration from previous work on enhancing interaction with visual feedback and feedforward. In particular, we considered:

- Scented widgets [23, 5] embed small miniature visualizations (e.g., histograms) directly into graphical control elements such as sliders.
- Small multiples and large singles [20] is a concept to preview thumbnails of alternative parameterizations of visual representations.
- Guidance visual cues [9] can be embedded into visualization views to indicate potentially interesting next navigation targets.
- Octopocus [1] is an interaction technique that provides feedforward as an interactive gesture is performed to indicate possible interaction outcomes.

The sketches used these inspirational techniques to outline possible solutions for informative visual feedback and feedforward for process exploration. Figure 20 shows a selection of our sketches, including conceptualization of interaction intents, general interface ideas, scented slider widgets, and preview thumbnails. From these and further similar sketches, we abstracted the following design dimensions that can play a role when implementing enhanced process exploration mechanisms:

Interaction control: What type of control is used to carry out the interaction? Slider, button, hover area, spoken command, etc.

Interaction integration: Where is the interaction control located? Integrated in the visualization vs. external to the visualization in a separate user interface.

Visual feedback/feedforward integration: Where is the visual feedback/feedforward shown? Visualization enhancement integrated into the visualization vs. Interface enhancement integrated into the interaction control.

Summary

In summary, the working group made some preliminary first steps toward overcoming the current PM limitations by integrating VA approaches. It became clear that completely solving the problem remains a formidable challenge for future work. Not only would it be

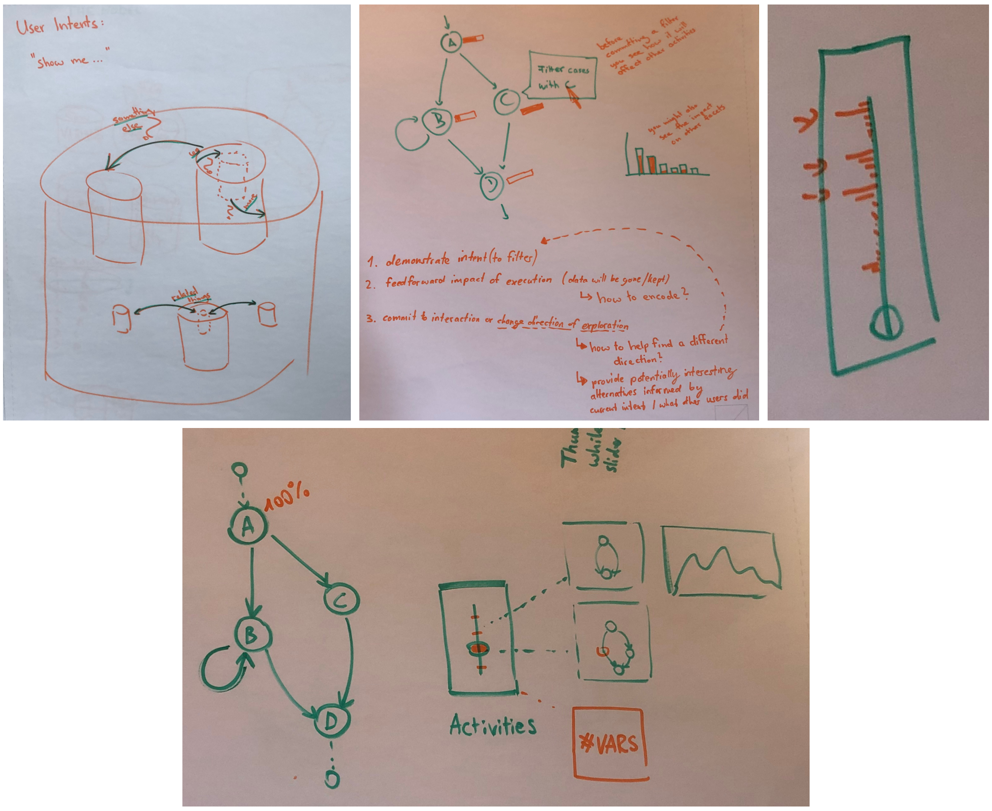

Figure 20 Selected sketches showing different "Show me..." interactions (top left), general action loop for DFG filtering (top center), scented slider widget (top right), and preview thumbnails attached to a slider (bottom).

necessary to more comprehensively map the design space of visual feedback and feedforward, but one would also need to implement the new designs into PM tools, which may just not be ready for handling the multiple states, visual feedbacks and feedforwards in their underlying architecture. Therefore, we suggest putting more research and development efforts into interactivity to support process exploration.

References

1 Olivier Bau and Wendy E. Mackay. Octopocus: a dynamic guide for learning gesture-based command sets. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology*, UIST '08, page 37–46, New York, NY, USA, 2008. Association for Computing Machinery.

2 Richard A. Becker and William S. Cleveland. Brushing Scatterplots. *Technometrics*, 29(2):127–142, 1987.

3 Davide Ceneda, Natalia Andrienko, Gennady Andrienko, Theresia Gschwandtner, Silvia Miksch, Nikolaus Piccolotto, Tobias Schreck, Marc Streit, Josef Suschnigg, and Christian Tominski. Guide me in analysis: A framework for guidance designers. *Computer Graphics Forum*, 39(6):269–288, 2020.

4 Davide Ceneda, Theresia Gschwandtner, Thorsten May, Silvia Miksch, Hans-Jörg Schulz, Marc Streit, and Christian Tominski. Characterizing guidance in visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):111–120, 2017.

5 Vasile Ciorna, Nicolas Médoc, Guy Melançon, Frank Petry, and Mohammad Ghoniem. A tour through the zoo of scented widgets. In Christina Gillmann, Michael Krone, and Simone Lenti, editors, *25th Eurographics Conference on Visualization, EuroVis 2023 – Posters, Leipzig, Germany, June 12-16, 2023*, pages 9–11. Eurographics Association, 2023.

6 Andy Clark. Busting out: Predictive brains, embodied minds, and the puzzle of the evidentiary veil. *Noûs*, 51(4):727–753, 2017.

7 Massimiliano De Leoni and Felix Mannhardt. Road traffic fine management process. *Eindhoven University of Technology, Dataset*, 2015.

8 Niklas Elmqvist, Andrew Vande Moere, Hans-Christian Jetter, Daniel Cernea, Harald Reiterer, and TJ Jankun-Kelly. Fluid interaction for information visualization. *Information Visualization*, 10(4):327–340, 2011.

9 Stefan Gladisch, Heidrun Schumann, and Christian Tominski. Navigation Recommendations for Exploring Hierarchical Graphs. In George Bebis, Richard Boyle, Bahram Parvin, Darko Koracin, Baoxin Li, Fatih Porikli, Victor Zordan, James Klosowski, Sabine Coquillart, Xun Luo, Min Chen, and David Gotz, editors, *Advances in Visual Computing: Proceedings of the International Symposium on Visual Computing (ISVC)*, volume 8034 of *Lecture Notes in Computer Science*, pages 36–47. Springer, 2013.

10 Michael Gleicher, Danielle Albers, Rick Walker, Ilir Jusufi, Charles D. Hansen, and Jonathan C. Roberts. Visual Comparison for Information Visualization. *Information Visualization*, 10(4):289–309, 2011.

11 Jakob Hohwy. The self-evidencing brain. *Noûs*, 50(2):259–285, 2016.

12 Daniel A. Keim, Florian Mansmann, Jörn Schneidewind, and Hartmut Ziegler. Challenges in Visual Data Analysis. In *International Conference Information Visualisation (IV)*, pages 9–16. IEEE Computer Society, 2006.

13 Heidi Lam. A Framework of Interaction Costs in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1149–1156, 2008.

14 Donald A. Norman. *The Design of Everyday Things, Revised and Expanded Edition.* Basic Books, 2013.

15 Dominik Sacha, Andreas Stoffel, Bum Chul Kwon, Geoffrey Ellis, and Daniel A. Keim. Knowledge Generation Model for Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1604–1613, 2014.

16 Ben Shneiderman. Dynamic Queries for Visual Information Seeking. *IEEE Software*, 11(6):70–77, 1994.

17 Ben Shneiderman. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *IEEE Symposium on Visual Languages (VL)*, pages 336–343. IEEE Computer Society, 1996.

18 Elizaveta Sorokina, Pnina Soffer, Irit Hadar, Uri Leron, Francesca Zerbato, and Barbara Weber. PEM4PPM: A cognitive perspective on the process of process mining. In *International Conference on Business Process Management*, pages 465–481. Springer, 2023.

19 Christian Tominski and Heidrun Schumann. *Interactive Visual Data Analysis*. AK Peters Visualization Series. CRC Press, 2020.

20 Stef van den Elzen and Jarke J. van Wijk. Small Multiples, Large Singles: A New Approach for Visual Data Exploration. *Computer Graphics Forum*, 32(3):191–200, 2013.

21 Jarke J. van Wijk. Views on Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):421–433, 2006.

22 Jo Vermeulen, Kris Luyten, Elise van den Hoven, and Karin Coninx. Crossing the bridge over norman's gulf of execution: revealing feedforward's true identity. In *Conference on Human Factors in Computing Systems*, pages 1931–1940. ACM, 2013.

23 Wesley Willett, Jeffrey Heer, and Maneesh Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Trans. Vis. Comput. Graph.*, 13(6):1129–1136, 2007.

24 Ji Soo Yi, Youn ah Kang, John T. Stasko, and Julie A. Jacko. Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, 2007.

25 Francesca Zerbato, Andrea Burattin, Hagen Völzer, Paul Nelson Becker, Elia Boscaini, and Barbara Weber. Supporting provenance and data awareness in exploratory process mining. In *International Conference on Advanced Information Systems Engineering*, pages 454–470. Springer, 2023.

26 Francesca Zerbato, Pnina Soffer, and Barbara Weber. Initial insights into exploratory process mining practices. In *Business Process Management Forum: BPM Forum 2021, Rome, Italy, September 06–10, 2021, Proceedings 19*, pages 145–161. Springer, 2021.

27 Francesca Zerbato, Pnina Soffer, and Barbara Weber. Process mining practices: Evidence from interviews. In *Int. Conf. on Business Process Management (BPM)*, LNCS, pages 268–285. Springer, 2022.

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

Barbara Weber (Universität St. Gallen, CH), Gennady Andrienko (Fraunhofer IAIS – Sankt Augustin, DE), Andreas Kerren (Linköping University, SE), Manuel Resinas (University of Sevilla, ES), Stef Van den Elzen (TU Eindhoven, NL), and Peilin Yu (Linköping University, SE)

Working group D (see Figure 21) focused on new visual representations for multi-faceted process exploration.