

A Process-Oriented Approach to Analyze Analysts' Use of Visualizations: Revealing Insights into the What, When, and How

L. Zimmermann¹ , F. Zerbato² , K. Vrotsou³ , and B. Weber¹ 

¹University of St.Gallen, St.Gallen, Switzerland

²Eindhoven University of Technology, Eindhoven, The Netherlands

³Linköping University, Norrköping, Sweden

Abstract

Despite Visual Analytics (VA) tools being essential for supporting data analysis, evaluating their use in real-world analytical processes remains challenging. Traditional evaluation methods often overlook the nuanced and evolving nature of analysis processes and are not always suitable for investigating scenarios in which analysts combine multiple tools and visualization types. In this paper, we propose a flexible analysis approach for studying analysts' use of visualizations within and across VA tools. Our qualitative method allows researchers to extract user behavior and cognitive steps from screen recordings and think-aloud data and generate event sequences that capture analytic processes. This enables the analysis of usage patterns from multiple perspectives and levels of granularity and allows for the evaluation of effectiveness measures, such as efficiency and accuracy. We demonstrate our approach in the domain of process mining, where our findings provide insights into the use of existing visualizations, and we reflect on lessons learned from this application.

CCS Concepts

• **General and reference** → Empirical studies; Evaluation; • **Human-centered computing** → Visualization design and evaluation methods; Visualization systems and tools;

1. Introduction

Over the past years, a vast number of *Visual Analytics* (VA) tools have emerged, designed for diverse analytical needs and aimed to support (data) analysts in their daily tasks of analyzing and communicating information [WGK10]. Such tools are highly relevant to uncover patterns, assess trends, and construct compelling narratives in various analytical domains. Yet, evaluating VA tools remains challenging, especially when the goal is to understand how these tools and the visualizations they provide are used in diverse, real-world analytical processes [Car08, Mun09, LBI*11].

Traditional evaluation methods, such as measuring task completion times or accuracy, provide useful measures but fail to capture the nuanced ways in which VA tools support complex analysis tasks in practice [SND05, SP06]. Observational studies address this gap by capturing rich qualitative data on user behavior [IZCC08, LBI*11], often through video recordings and think-aloud protocols [Nor06]. Analyzing these recordings can reveal detailed, qualitative insights into user behavior, strategies, and potential difficulties [GK12]. However, despite its potential, the analysis of video recordings is far from straightforward, as the data is unstructured, multidimensional, and context-dependent. Existing studies often limit their scope to predefined aspects [SKFZ21] and

overlook the dynamic nature of analysis processes where multiple tools and visualizations are used.

To address these limitations, we propose a flexible, qualitative approach that utilizes screen recordings as a primary data source to study the behavior of users engaging in analytic processes. Our approach allows researchers to capture how users engage with multiple visualization types within or across VA tools as they are applied to solve real-world analytical tasks. It enables the analysis of authentic behavior patterns across analytical domains without requiring modifications of the VA tools being studied. Instead of relying solely on interaction logs, which primarily capture low-level actions such as clicks or zooming [XOW*20], our method encodes video-recorded interactions and cognitive processes into *multi-perspective event sequences*. These event sequences integrate different analytical perspectives, allowing us to study how users interact with visualizations and how specific visualization types are used over time to identify areas for improvement.

We demonstrate our approach with data collected from process analysts. This application example allows us to put our method into practice while learning about the use and effectiveness of visualizations in process mining. Indeed, although process mining analysis relies heavily on visual representations to assess and communicate process flows and trends, it is unclear how process mining analysts

actually work with these visualizations in practice and how they can be evaluated [KMW19, YM24]. Moreover, since exploratory process mining tasks require analysts to use different tools and visualizations [ZSW21], this makes it an interesting domain to apply our approach to. In summary, our contributions are as follows:

- We propose a **structured process-oriented analysis method** to convert video recordings, including think-aloud and behavior observations, into event sequences that capture analytic processes from multiple perspectives and levels of granularity.
- We present the **empirical results from a process mining use case** in which we demonstrate how the proposed approach can be applied to explore and evaluate analysts' use of visualizations for performing analytical tasks.
- We **critically reflect on the learnings** from applying our method and identify opportunities for future improvements.

By taking a process-oriented view on analysis tasks, our method complements traditional evaluation techniques, prioritizing ecological validity and bridging the gap between quantitative metrics and rich, qualitative insights. In doing so, we contribute both an approach that enables a deeper understanding of how VA tools are applied in realistic settings and insights from our use case.

2. Background and Related Work

2.1. Evaluating User Interaction and Visualization Use

Understanding and evaluating how users interact with computational systems and perceive interfaces has been a long-standing research focus. Since our work explicitly evaluates visualization use, we primarily consider literature on visualization evaluation, but also integrate relevant work on general interface use. Based on our review, existing evaluation methods can be broadly categorized into two main categories: (1) analytical methods and (2) empirical studies, including user observations and mixed-method approaches.

Analytical methods assess interfaces through theoretical assessments, expert reviews, or predefined heuristics. A prominent example from the field of Human-Computer-Interaction (HCI) is the GOMS (Goals, Operators, Methods, and Selection Rules) framework, which models user interactions to evaluate interface efficiency [JK96]. For example, Gong and Kieras [GK94] demonstrated its utility in identifying improvement opportunities for an existing interface. Other analytical methods include heuristic testing [Nie92] or cognitive walkthroughs [LW97].

Isenberg et al. [IIC*13] found that even in evaluation scenarios explicitly addressing user interactions with visualizations, such as *evaluating environments and work practice, understanding work practices, and evaluating visual data analysis and reasoning* [LBI*11], studies are conducted rather analytically, involving self-assessments by tool developers or reflections of selected domain experts. In general, analytical methods do not analyze actual user behavior and are therefore less related to our work.

Empirical studies, as we propose in our work, derive insights from real user interactions. Traditional quantitative and qualitative approaches, such as measuring task accuracy, speed, or user satisfaction, provide means to evaluate specific visualization characteristics but do not capture nuanced real-world behavior pat-

terns [Car08, vW13]. Instead, our work takes an observational approach, analyzing user behavior and its context to reveal what visualizations are used, and why and how they are used to uncover cognitive workflows and evaluate the efficiency of different visualization types for diverse analytical tasks.

Therefore, relevant to our work are especially methods that leverage actual **user observations**. One example is the *insight evaluation method* [SND05, NSD11], which assesses VA tools in realistic analysis settings by quantifying and characterizing user-generated insights. Grammel et al. [GTS10] additionally demonstrate that observational settings allow for evaluating insights and understanding generic work practices. These approaches rely on video recordings as primary data source and it has been shown that video data in general is a valuable source for capturing user behavior, as it enables the analysis of verbal exchanges, embodied actions, and interactions with the environment [Mon08]. However, neither of the mentioned works explicitly considers the temporal sequence of user actions and reasoning steps. In our approach, we maintain a temporal perspective, i.e., the order and duration of relevant *events*, enabling us to observe the evolution of analytical workflows.

A **temporal perspective** on analytic processes is central in *analysis provenance* research, where studies focus on workflow reproducibility and behavior analysis. For example, Heer et al. [HMSA08] demonstrated that maintaining a graphical history of analysis operations allows researchers to study how analysts work in practice. Typically, analysis provenance studies rely on interaction logs derived from low-level user actions [XOW*20]. However, instead of relying *only* on user interaction logs, some of the existing studies adopt mixed-method approaches, combining interaction logs with qualitative data to be able to put user interactions into context. For example, Duo et al. [DJS*09] point to difficulties in inferring analysis strategies of users solely from interaction logs and suggest validating findings with qualitative data from video recordings. Similarly, Guo et al. [GGZL16] enrich their analysis with video data to identify usability issues and behavior leading to insights. The authors point to difficulties in summarizing the high-level video observations—a gap our method addresses by providing a structured framework for coding screen recordings across multiple perspectives and levels of granularity. In addition to video recordings, Blascheck et al. [BJK*15] combine interaction logs with eye-tracking data and think-aloud protocols and propose an interactive analysis tool to identify behavior patterns. While related to our work, their approach focuses on where users look during visualization use, whereas we examine how and when visualizations are used based on screen recordings, eliminating the need for an experimental setup with eye-tracking equipment.

Overall, we found no prior work that relies on video recordings and think-aloud data to infer user actions and cognitive steps while assessing both visualization use and effectiveness. Although future enhancements could incorporate interaction logs or eye-tracking for better scalability (cf. Sect. 6), our current approach requires minimal setup, enabling the capture of realistic visualization use.

2.2. Process Mining

In process mining data records generated during the execution of processes are (visually) analyzed to derive insights [vdA22]. These

records must be transformed into *event logs*, where each activity (event) is linked to a unique case or object and a point or interval (start and end timestamp) in time. Apart from a mandatory name and a start and/or end timestamp, activities can include additional attributes (e.g., resource identifier, domain-specific information such as payment amounts). Analysts employ process mining to retrieve end-to-end visualizations and performance indicators from event logs, make sense of process-related information, make data-driven decisions, and implement interventions [BWG*22].

To support these use cases, process mining algorithms generate different types of process models (e.g., Petri nets, directly-follows graphs (DFGs), process trees) that visually represent the underlying process and are essential for exploring process structures but may be challenging to interpret [ZZW24, VDA19]. In addition to process models, process mining tools include general-purpose data operations, such as filtering, aggregating, and zooming, as well as visualizations, such as bar charts, scatterplots, or line graphs that support data exploration alongside process-specific techniques.

Despite the broad availability of visualizations, their role during analysis has only been explored to a limited extent. Previous studies rather characterized process mining analyses through analytical operations or strategies [CASUCR22, SSH*23, ZSW21]. Focusing on visualizations use, Klinkmüller et al. [KMW19] assessed analysis practices and found that mainly tables are used to communicate results. For predictions, analysts refer to general-purpose charts (e.g., line charts) and for process discovery, process models are used. Our study expands the existing body of knowledge and reveals more insights into the use and effectiveness of visualizations.

3. A Method for Creating Multi-perspective and Multi-granular Event Sequences

This section begins with an example process mining analysis (cf. Sect. 3.1) to illustrate our method. Starting from video recordings of analysis sessions (cf. Sect. 3.2), we identify events (cf. Sect. 3.3) and transform them into event sequences (cf. Sect. 3.4) for analyzing the use and effectiveness of visualizations.

3.1. An Example Process Mining Analysis

We observed and recorded an analyst, let us call her Alice, as she analyzed the Road Traffic Fine Management (RTFM) event log, a publicly available log that captures the process of traffic fine handling by the Italian police [dLM15]. Alice is asked to identify situations in which fines are not paid by offenders and explain them. Example steps made by Alice are shown in step 2 of Figure 1.

To “get acquainted with the process”, Alice opens the process mining tool ProM [VBVDvdA10] and navigates to the *Trace Variant Explorer* where she explores control-flow variants in a *Sequence Chart* visualization (cf. A in Figure 1). Alice observes that the most common variant does not involve any payment and hypothesizes that this may be a data quality issue. After exploring the control-flow of multiple variants, she decides to dive into further details using the *Log Inspector* (cf. B in Figure 1). The displayed *Trace Sequence* enables her to assess attributes linked to events for one fine. Alice hypothesizes that the fine amount might play a role

in payment completion: “If the fine is high, then there could be the need for sending it to credit collection”. Returning to the *Trace Variant Explorer*, she compares examples and rejects her hypothesis as she observes similar amounts across other variants that also involve payment: “The amount is 36, but this one was paid. So, I think the amount before didn’t tell me anything”. During her further analysis, Alice switches tools and visualizations multiple times. By exploring the process model displayed as DFG (cf. C in Figure 1) in the process mining tool Celonis [Cel24] and focusing on single activities, she discovers that a portion of unpaid fines involve appeals. Alice derives that offenders have successfully contested their fines.

Above, we focused on segments of Alice’s analysis and exemplified her intents, hypotheses, and observations. In the remainder, we outline our method, illustrating how Alice’s analysis, including both analytical and cognitive steps, can be captured and represented, allowing us to analyze and evaluate her use of visualizations throughout the entire analysis. We present our approach in generic terms to allow for its transferability to other analytical settings but contextualize them with examples from our own experience in process mining [SSH*23, ZSW21, ZSW22, ZZW24].

3.2. Video Recordings as Input Data

As shown in Figure 1, our approach starts from *video recordings* of analysis sessions, capturing both their screen (behavior observations) and the accompanying audio (think-aloud protocol). Thereby, the screen recordings show the interactions with visualizations and tools (analytical steps) and the think-aloud—which we transcribe while maintaining a timestamp associated with each utterance—reveals thought processes, reasoning, and decision-making strategies (cognitive steps). Since these two data modalities originate from the same video, they are synchronized in time which allows us to disambiguate the observable behavior and to understand the user intents driving specific actions [WAAF*24].

Using video recordings for our analysis comes with several advantages. They are inexpensive to collect, as most laptops already have recording tools. They also eliminate the need for instrumenting software, which is time-consuming, challenging with proprietary tools, and impractical in exploratory settings where relevant data points are unclear. Additionally, video recordings capture realistic interactions without restricting analysis to specific tools or visualizations. They can also be complemented by other data sources, such as recorded answers to specific questions to identify insight emergence or eye-tracking data to automate focus and cognitive load analysis [AALW23].

3.3. Event Identification

To explore analysis processes, we identify events that describe the user’s analytical and cognitive steps (cf. Figure 1, step 2). To this end, there are several key decisions to make that can lead to identifying different events based on the *perspective* and the *granularity* level at which we want to describe the analytical process. Since these decisions are highly dependent on the specific use case and research question, our goal is to highlight the importance of these choices rather than propose a one-size-fits-all approach.




Figure 1: Overview of our proposed method. Based on the input data (1), relevant segments of interest (events) are identified (2). These segments should be captured for different perspectives (P) at a level of granularity that is relevant for the analysis. The screenshots highlighted with A, B, and C are of lower quality, labels are not intended to be legible. Subsequently, identified events can be combined into event sequences (3) and diverse aspects regarding the use and effectiveness of visualizations can be analyzed (4).

In the database literature, an event is “a temporal happening” [HZZ08] with a single timestamp. In this work, we refer to an *event* as a fact of interest for the analysis that can be represented as an interval (known as *activity* in process mining). An event has a name (or *event type* [HZZ08]) and a start and end timestamp. To identify events from video recordings, we suggest following a *qualitative coding* approach [Sal15] that allows us to mark segments of varying duration in the videos, capturing observable user behavior like actions carried out in the VA system or cognitive steps verbalized in the think-aloud [SSS*14].

Perspectives. As Figure 1 shows, screen recordings and think-aloud transcripts can guide the identification of relevant events depending on the *perspective* from which we choose to look at the process. Existing taxonomies of user actions are a natural starting point to identify events that abstract analytical steps [vLFB*14, YaKSJ07, BM13] and thus can guide the identification of suitable perspectives to start from. An example is the typology by Brehmner and Munzner [BM13], which describes sequences of visualization tasks based on WHY (reason), HOW (method) and WHAT (input/output). The HOW perspective is the easiest to observe in screen recordings, as it refers to the technique, UI element, or software function used (e.g., Vis1: “Sequence Chart”; Figure 1). The WHAT perspective refers to the target object of the interaction, which can also be observed in screen recordings but often requires triangulation with the corresponding think-aloud utterance to disambiguate the actual focus of the user. In the context of an analysis process, primary objects of interest are the *data* being analyzed and its *representations* (e.g., Focus 1: “Variants”; Figure 1). The WHY perspective instead refers to the reasons or intents behind a task. Therefore, it requires coding the think-aloud transcripts. In our context, the WHY perspective can refer to different kinds of cognitive steps, including user intents. A single event can reflect multiple perspectives depending on the analysis goal. For example, using a DFG (e.g., C in Figure 1) to assess process flow could be coded as HOW (tool use) or WHAT (visualization focus) based

on the research question. In our case, the goal of the analysis is to understand how analysts use specific visualizations to, for example, understand the process flow. Therefore, the DFG belongs to the HOW perspective. Instead, if we aimed to quantify how often the user focuses on a specific visualization type, it would be categorized as WHAT. In this case, the tool in which the process model is visualized (e.g., Celonis) is a suitable event in the HOW perspective. With this coding approach, we can identify events that reflect specific perspectives of the analysis process.

It is important to note that the different perspectives can lead to multiple segmentations in the video recordings. In some cases, a single video segment may be labeled with multiple codes, as different perspectives can capture the same fact of interest from unique angles. For instance, in Figure 1, two events from two different perspectives (i.e., P1 and P2) describe the same video segment. Such overlaps are advantageous for the creation of event sequences, as we discuss later. However, in most cases, different perspectives lead to distinct segmentations, with each perspective isolating unique events in the same video.

Levels of Granularity. Each perspective of the analysis process can be described at various levels of granularity, depending on the level of abstraction we choose for its event types. For instance, the Process Model as a generic group of visualizations can be described at a fine-granular level, through its individual format (e.g., DFG), or even more fine-grained through its state of visible activities and transitions. Common starting points for defining granularity levels for events include groupings of user interface (UI) elements for the HOW perspective [AR24], or visualization types or data structures for the WHAT perspective. For the event identification and coding step, it is crucial to determine the finest level of granularity for the events and ensure that they are meaningful for the analysis. A guideline is to select a level of granularity for which there are no overlapping events from the same perspective in a single video segment. This means that a user cannot use two tool functions or formulate two hypotheses simultaneously. Our approach allows for

adjustments of the event granularity at any point in the analysis, depending on the patterns the data reveals [ZSDF*21].

3.4. Event Sequence Generation

In the third phase of our approach, we prepare one or multiple event sequences for analysis, each representing a list of events ordered by their occurrence in time [GGJ*22]. In this phase, we need to make decisions concerning the perspectives we want to analyze and the granularity of the events. Depending on these choices, different event sequences can be generated (cf. Figure 1, step 3).

Single Perspective. At the finest level of granularity, each generated sequence will be totally ordered, potentially including some temporal gaps between events. Depending on the analysis techniques we choose, we may need to code these gaps explicitly. For instance, most process mining techniques explicitly model waiting times between events, while many sequential pattern mining techniques do not account for inter-event delays. Abstraction to higher granularity levels may involve merging consecutive events (e.g., combining “Variant Sequence” and “Trace Sequence” into “Sequence Chart” in Figure 1). Whether or not to merge similar, consecutive events depend on the relevancy of repetitive patterns and temporal proximity of the events, as merging distant events may not accurately reflect the user’s actual behavior.

Multiple Perspectives. To leverage the richness of our multi-perspective and multi-granular dataset we can combine event sequences from different perspectives. For perspectives with the same segmentation (e.g., P1 and P2 in Figure 1), sequences can be merged by creating new composite events reflecting events combined from multiple perspectives. However, this step requires care, as it introduces variability in the event types. Adjusting the granularity level of the perspectives can lead to the need to merge consecutive events, as discussed above, but in this case will not lead to overlapping events. For perspectives with different segmentations, overlaps must be handled differently. To assess them, Allen’s interval algebra [All83] provides a framework for analyzing temporal relationships, such as contain, overlap, and meet. For example:

- If “Sequence Chart”(P1) overlaps with both “Variants” (P2) and “Attributes”(P2), we must decide whether to split the event of P1 or create a composite event with a label like “Sequence Chart + Variants + Attributes” (cf. Figure 1; step 3).
- If the event “Process Model” (P1) is fully contained within the event “Activities” (P2), the shorter event could form a filtered sequence aligned with the duration of events in P1.

The key idea behind our approach is to preserve the richness of the data, capturing multiple perspectives and granularity levels, rather than limiting ourselves to a single view. At the same time, each perspective is encoded into one sequence and can be analyzed independently from the others. During analysis, it is always possible to revisit the event generation phase to create relevant event sequences for new analytical aspects.

4. Application: Process Mining Analysis Sessions

This section outlines the observational study we conducted with process mining analysts, detailing data collection and identification of relevant events, with the analysis presented in Sect. 5.

4.1. Input Data

To investigate the behavior of process mining analysts, we conducted an observational study between May and July 2021 as part of a broader research project [ZZVW23].

Study Design. The observational study comprised analysis sessions with think-aloud. During the sessions, participants were asked to explore the Road Traffic Fine Management event log [dLM15], the same as used by Alice (cf. Sect. 3.1). The participants were provided with documentation about the log that included details about the process and the domain (e.g., fines being void if issued after 90 days). We asked them to identify three circumstances where traffic fines were not paid, and, if possible from the data, provide explanations for these circumstances.

Procedure. For their analysis, the participants accessed a virtual machine (VM) environment from their local browser. This ensured that everyone worked with the same environment. In the VM, we provided the event log, the documentation (pdf files), and three process mining tools: ProM [VBVDvdA10], a well-known open-source framework, Celonis Process Mining [Cel24], a leading commercial tool by Celonis, and Disco [GR12] a commercial tool by Fluxicon. This selection of tools captured a wide range of standard tools commonly used in practice at the time of the study. All three allow for an interactive and visual process mining analysis. All analysis sessions were conducted virtually, and participants shared their screens while thinking aloud.

Participants. Participants were recruited through our personal network and referrals from colleagues. To be eligible for the study, individuals had to: (1) have analyzed at least one real-life event log within the two years preceding the study; (2) be familiar with at least one of the available process mining tools (i.e., ProM, Celonis, Disco); and (3) have been involved in at least one process mining project where third-party data was analyzed. These criteria ensured that participants had relevant and practical experience in process mining, allowing us to observe a range of analytical behaviors across varying levels of expertise. A total of 33 individuals met these criteria and participated in the study. In the following, we will refer to them as “analysts”. Further details on the participants are available in the supplementary material. The recorded analysis sessions yield a total of 21.32 hours of video material, with an average analysis session duration of 38.7 minutes per analyst.

4.2. Event Identification

To create the event sequences, we followed the approach outlined in Sec. 3.3. We imported all videos and transcribed think-aloud protocols into MaxQDA [VER24], a qualitative data analysis tool that enables the annotation of events with start and end timestamps.

We developed the coding scheme for each perspective iteratively, defining granularity levels and grouping fine-grained events into abstract categories as needed, following grounded theory coding principles [Sal15]. Event identification involved collaborative coding by two authors and two external assistants, with discrepancies being resolved through team reviews and discussions. One author standardized the coding across all videos to ensure consistency.

Details on the encoded perspectives as well as the resulting event

sequences can be found in the supplementary material. The Visualization Perspective (HOW) captures events when a specific visualization was actively used, while the Focus Perspective (WHAT) describes the data aspects of the considered process [BM13]. Additionally, we identified events in three further perspectives. The Intent Perspective captures the purpose behind analytical activities, the Interpretation Error Perspective captures moments of incorrect interpretations, and the Observation Perspective captures relevant observations made by analysts.

While these perspectives were deemed relevant for the analysis of visualization use in process mining, additional perspectives could be extracted, for example inspired by other established taxonomies [AS05, GZ08, HS12, YaKSJ07]. However, for our purposes, the selected five perspectives provide a sufficient basis for a holistic view of visualization use, efficiency, and accuracy.

5. Analysis Results

In this section, we present the results of our analysis of process mining analysis sessions, focusing on the use of visualizations. We first analyze the visualization perspective, examining how a single-perspective event sequence reveals insights into the types, timing, and patterns of its events (cf. Sect. 5.1). Next, we integrate additional perspectives (cf. Sect. 5.2) and finally evaluate the efficiency and accuracy of visualization for the analysis task (cf. Sect. 5.3).

5.1. Single-Perspective Analysis: Visualization

Initially, we focus on the event sequences capturing the Visualization Perspective, shown in Figure 2. Events in this perspective capture when a specific visualization type is actively used for analysis.

5.1.1. Statistical Measurement: What visualizations are used?

We aggregated information from the 33 single-perspective event sequences displayed in Figure 2 to understand what visualization types are used during a process mining analysis. Table 1 presents the results.

The event sequences describe the use of 14 different visualizations at the most granular level (L3) and can be grouped into 10 more abstract visualization types (L2) categorized in three main groups (L1). The DFG, with its two subtypes, is the most commonly used visualization and accounts for 69% of all events in the sequences. All other visualization types are used less frequently. Other Chart Visualizations are overall used less often, with the Bar Chart accounting for less than 1% of the visualization use. Additionally, Social Network and BPMN process models are used for only 2.66 and 1.45 minutes and by 1 and 3 analysts, respectively.

On average, a visualization event has a duration of 44.1 seconds, with the longest event lasting seven minutes (see *Scatterplot* for A28 in Figure 2), potentially indicating a deep analysis. Overall, the *Scatterplot/Dotted Chart* has the longest average duration, followed by the Social Network, and Inductive Visual Miner. The BPMN model, the *Scatterplot/Scatterplot*, and Bar Charts (Bar Chart and Histogram) are used for rather shorter periods of time, which could indicate that their use was quickly discarded, as it was not perceived as beneficial, or the required information were identified fast.

Visualization Types						
L1	L2	L3	#A	min	%	$\bar{D}$
Process Models (PM)			32	371	79.5	44.2
	Directly-Follows Graph (DFG)		32	322	69	43.4
	—"	Process Explorer	31	253	54.2	46.2
	—"	Variant Explorer	11	69.2	14.8	35.5
	Petri net	Petri net	4	22.6	4.84	61.6
	Inductive Visual Miner (IvM)	IvM	4	15.6	3.34	71.8
	Causal net	Causal net	3	6.63	1.42	33.2
	Social Network	Social Network	1	2.66	0.57	79.7
	BPMN	BPMN	3	1.45	0.31	9.69
Other Process Representation			12	49.9	10.7	35.7
	Sequence Chart	Trace Sequence	12	33.8	7.24	35
	—"	Variant Sequence	5	16.1	3.46	37.2
Other Chart Visualizations			16	45.7	9.8	58.3
	Scatterplot		6	34.6	7.42	83.1
	—"	Dotted Chart	5	34.2	7.33	85.4
	—"	Scatterplot	1	< 1	0.1	26.7
	Line Graph	Line Graph	8	6.89	1.48	31.8
	Bar Chart		6	4.2	0.9	28
	—"	Bar Chart	4	3.31	0.71	28.4
	—"	Histogram	2	< 1	0.19	26.5

Table 1: Visualization Use. Statistics of the number of analysts (#A), the cumulative usage time in minutes (min), the relative duration (%), and the average event duration in seconds ($\bar{D}$). L1, L2, and L3 represent the three granularity levels of event description.

For the remaining analysis, we focus on the visualization events at granularity level L2 to allow for relevant insights but aggregating less frequent types into more meaningful groups.

5.1.2. Use over Time: When are visualizations used?

To analyze the use of visualizations over time, we look at how visualization events are distributed over an event sequence. This allows us to, for example, identify the first and last visualization events per sequence. Notably, many analysts begin using visualizations early during their analysis. For approximately 60% of the sequences (20/33), events occur already in the first 10% of the analysis time (e.g., A12, A21, A26; Figure 2). The early use suggests that analysts perceive visualizations as useful for exploratory analysis or as guiding mechanisms for their investigations. Only two analysts, A10 and A7, show a late use of visualizations.

Next, we examine the event frequency over time for each visualization type. Results are shown in Figure 3. As expected, the DFG exhibits the highest frequency of use over the complete analysis duration. While the plot does not indicate a clear trend, a slight decrease in DFG usage can be observed towards the end of an analysis. Instead, other visualization types show rather homogeneous patterns (e.g., Bar Chart, Causal Net). BPMN models and the Social Network are used later, which could indicate that these types can serve as alternative tools for analysts who have not yet achieved their goals by other means. The Sequence Chart and the IvM model also show a slight increase in usage towards the second half of the sequences, which allows for a similar explanation as before.

5.1.3. Global & Local Patterns: How are visualizations used?

Process mining techniques were applied using the tool Disco to analyze global patterns. No significant patterns were revealed in




Figure 2: Analysis sequences showing the use of different types of visualizations by analysts throughout their analysis duration.




Figure 3: Frequency of visualization type use over time.




Figure 4: DFG showing control-flow patterns in visualization use, with annotations indicating usage frequency and duration. For readability, only the top 40% most frequent transitions are shown.

terms of control flow, with only two variants being observed multiple times. Specifically, A11 and A10 use the DFG 6 times and A23 and A14 15 times without employing other visualizations during their analysis sessions. Figure 4 provides an overview of the identified control-flow aspects. The lack of consistent global patterns suggests that analysts adopt highly individualized strategies, adapting visualization choices based on experience and preferences.

While the process mining analysis of the event sequences did not uncover any significant global patterns, Figure 4 does reveal common directly-follows relationships. The most common pattern is the DFG event being followed by another DFG event. This can be observed in 31 sequences for a total of 316 transitions. They are repeated at an average interval of 86.6 seconds (median = 33 seconds). In 9 sequences, the DFG is instead followed by the Sequence Chart, in 5 by the Bar Chart, and in 4 by a Line Graph or Scatterplot. Transitions between different visualization types likely reflect shifts in analytical focus. For example, a switch from the DFG to the Scatterplot (e.g., Dotted Chart) could indicate the validation of an observation or finding. Other explanations include a switch to a more detailed analysis of specific aspects, or a change of direction as the previous approach was not successful or is completed.

5.2. Multiple-Perspectives: Visualization, Focus, and Intent

To enrich the results obtained in Sect. 5.1, we analyze the Visualization Perspective in relation to two other perspectives, namely the Analysis Focus and Analytical Intent, with a focus on when these perspectives interact, i.e. focusing on events that overlap in time.

5.2.1. Visualization and Analysis Focus: What is the analysis focus per visualization type?

We generated the event sequences combining the Visualization Perspective with the highest granularity level of the Analysis Focus Perspective (cf. Sect. 4.2). To this end, we created the multiperspective event sequences by combining each visualization event




Figure 5: Relative duration of overlaps for each visualization type with focus events. Values lower than 0.01 are not indicated.

with the focus event(s) it overlapped with and the respective start and end timestamps of the overlap. Based on the resulting event sequences, we can analyze the duration of each focus-visualization combination. Figure 5 presents these durations, normalized for each visualization type. We can observe that process models, such as Petri nets, IvM Models, DFGs, Causal Nets, and BPMNs, are predominantly used for examining a process's *Flow*. This aligns with their utility for understanding sequential patterns and transitions in a process. Interestingly, the IvM Model and the DFG are also frequently used for exploring *Statistics*, likely because they annotate events and transitions with metrics such as case and event frequency. For analyzing *Performance* aspects, i.e., aspects related to the execution time of the process, the Scatterplot and Line Graph show a higher relative utilization than the DFG, even though the DFG can also depict mean or median throughput times between events. This suggests that the Dotted Chart and Line Graph may offer a more intuitive approach for analyzing performance aspects. Sequence Charts mainly support a focus on *Variants* and *Attributes*, as demonstrated in our example from Alice (cf. Sect. 3.1).




Figure 6: Relative duration of overlaps with Statistics focus events for each visualization type.

Based on these initial observations, we decided to dive deeper into the use of visualizations for *statistical* analysis, as Figure 5 shows that a significant amount of visualizations are used for this purpose. To do so, we created new event sequences combining the Visualization and Focus Perspectives at a finer granularity, distinguishing four sub-categories of *Statistics*: “Number of cases”,

“Number of events”, “User-derived Metrics”, and “Others”. Figure 6 reveals that *User-derived Metrics* were exclusively analyzed with Scatterplots. A deeper analysis of this specific aspect based on the Focus Perspective reveals that while analysts spent a total of 53.6 minutes focusing on *User-derived Metrics*, only 27 seconds overlap with a visualization event (i.e., Scatterplot). This indicates that analysts prefer textual representations (e.g., numbers, tables) of such statistics and highlights potential limitations in integrating *User-derived Metrics* in current visualizations.

Further, Figure 6 indicates that most visualization types, when used for the analysis of statistics, focus primarily on the *Number of Cases*. In contrast, the Bar Chart tends to be utilized for analyzing the *Number of Events*, and for the Causal Net, the focus is on *Other* statistics. These are, for example, calculated by process mining algorithms and inform analysts about confidence intervals or the ratio of behavior that the model represents from the event log.

5.2.2. Visualization and Analytical Intent: What is the intent behind using each visualization type?




Figure 7: Relative duration of overlaps for each visualization type with analytical intent events.

To determine what visualization types are most frequently used to accomplish specific analytical intents, we analyze the interactions between the Visualization and Analytical Intent Perspective (cf. Sect. 4.2) following the same mapping strategy as for the overlap of visualization events with events of the Focus Perspective.

Our analysis shows that all visualization events are fully contained in a single intent event, indicating that the analysis intent is not switched during the use of a visualization type. However, multiple visualization events might be contained in the same intent event, presumably as multiple analysis steps are combined to work on one intent. Figure 7 shows the relative duration of overlaps for each visualization-intent combination (relative for each visualization type). We observe that several intents, such as Data Wrangling, Defining Needs, and Verification, are not conducted with the help of visualizations. While this is expected for the first two, Verifications could indeed be performed by referring to a visual representation of the data. One possible explanation could be a lack of trust in existing visualization types overall, leading analysts to rely on raw data for verifying their results. For other intents, Figure 7




Figure 8: Distribution of intent events, overlapping with the visualization type “DFG”.

reveals more insights. For example, the Sequence Chart and IvM Model are primarily used for Exploration and Descriptive Modeling, aligning with moments where analysts seek to discover or summarize process features. Instead, Petri nets, while still mainly used for Exploration, seem to also support Diagnostic Modeling, during which analysts aim to identify root causes for observations.

When do different intents influence the use of a visualization?

The Stacked Area Chart in Figure 8 reveals the temporal distribution of intent events overlapping with the use of the DFG – the most widely used visualization type in our study. It shows that during the first quarter of an analysis, the DFG mainly supports Profiling and Exploration. While Exploration remains a relevant intent for using the DFG, Descriptive and Diagnostic Modeling gain relevancy in the second half of the analysis. Interestingly, analysts seem to also refer to it to (re-)define the needs of the specific analysis question that they had to answer. Likely, the DFG provides the required overview of all process aspects to recall the goal of the analysis.

5.3. Assessment of Effectiveness

Zhu [Zhu07] describes three dimensions of effectiveness, covering efficiency, utility, and accuracy. Utility is partially covered in Sect. 5.2.2, where we describe the analytical intents for which visualization types are utilized and presumably perceived as useful. Below, we focus on efficiency and accuracy. Efficiency refers to the interaction between a visualization and its user, typically measured by task completion time for benchmark tasks. Accuracy reflects the correctness and clarity of the representation, linking the visualization to the data, and can be assessed through domain-specific mappings or by counting interpretation errors.

5.3.1. Efficiency: How efficiently do different visualization types allow analysts to derive observations?

To assess the efficiency of the different visualization types we use the Observation Perspective event sequences. These allow us to reveal to what extent each visualization type supports the analysts' ability to derive observations.

We identified observations overlapping with events of eight of the ten visualization types. Table 2 provides an overview of the number of observations and observations per minute for each visualization type. Notably, no observation events overlap with the use of Social Network and BPMN diagrams. The Causal Net displays

Visualization Type	Observations	Obs. per min.
DFG	196	0.61
Petri net	8	0.35
IvM	7	0.45
Causal Net	6	0.91
Sequence Chart	36	0.72
Scatterplot	24	0.69
Line Graph	4	0.58
Bar Chart	2	0.48

Table 2: Observations per different visualization types.

best observation ratio results with an average of 0.91 observations per minute. It is followed by the Sequence Chart (0.72) and the Scatterplot (0.69), suggesting that these three types support analysts well in progressing and interpreting the data. As observations may lead to findings [SSS*14], we conclude that these visualization types likely support analysts better in finding an answer to the task at hand. Conversely, the lower observation ratio for Petri nets and the IvM suggests lower efficiency for deriving immediate insights.

5.3.2. Accuracy: How accurate are information derived while using different visualizations?

While interpretation errors have been suggested as a measure for visualization accuracy [Zhu07], it is important to recognize that they can be influenced by individual factors such as the analyst's expertise, cognitive bias, or misinterpretation of prior results. Consequently, interpretation errors should be treated as an indicator rather than a definitive measure of visualization accuracy, particularly when considering evolving analysis processes.

Despite these considerations, we observed 22 errors made by 13 analysts that are attributed to specific visualization types, based on the event sequences for the Interpretation Error Perspective (cf. Sect. 4.2). 19 errors occurred in conjunction with the DFG, two with the BPMN diagram, and one with the Sequence Chart. While this results in 0.06 and 0.02 errors per minute for the DFG and the Sequence Chart respectively, for the BPMN diagram that is only used for 1.45 minutes, 1.38 errors are made per minute, resulting in a low accuracy measure. The DFG, as available in all the tools used during our study and responsible for most errors, revealed two main issues. First, errors arose from misinterpretations or miscalculations of annotations, including event frequencies and durations. Second, analysts experienced confusion due to the filtered version of the DFG, which visualizes process behavior that is not representative of the data [VDA19]. For the two errors made for BPMN diagrams, we found that they occurred because analysts interpreted BPMNs as providing a “ground truth”, i.e., they assumed that they represented the ideal process rather than a discovered model. This issue was particularly evident in the tool Celonis, where users can either manually create or generate BPMN diagrams from event logs. Interestingly, even when analysts mined the model from the data, they perceived it as prescriptive.

Our findings suggest that visualization tools insufficiently clarify the origin and meaning of certain visualization types (e.g., through annotations, visual cues, or interactive explanations). Misinterpretations of discovered models (in the DFG and BPMN) underscore

the need for better design interventions to help users distinguish between an actual process and its visualized abstraction.

6. Discussion, Limitations, and Conclusion

We introduced a novel approach for analyzing the use of visualizations in analytic processes starting from video recordings of analysis sessions. Compared to previous work, our method provides a structured view of *what*, *when*, and *how* visualizations are used by considering multiple perspectives and levels of granularity. It extends existing visualization evaluation frameworks, such as the insight-based evaluation method [SND05], long-term in-depth case studies [SP06], and analysis tools such as the VA² approach [BJK*15], by revealing patterns of tool use through statistical techniques, process- and pattern mining, and effectiveness assessments, enriched by qualitative insights. Compared to interaction log and eye-tracking analyses, which focus on controlled experiments and detailed user behavior, our approach captures broader analytical workflows in real-world scenarios without the need for controlled settings.

Our method is particularly useful for capturing realistic analysis patterns, making it a valuable resource for researchers and tool developers. Derived insights can inform tool refinement, inspire new visualization designs, and highlight training needs to improve analytical efficiency. Furthermore, by revealing visualization adoption across domains, it contributes to assessing usability, effectiveness, and opportunities for improving user support. Our study of visualization use in process mining, for instance, revealed that the DFG is the most frequently used visualization type but causes a substantial amount of interpretation errors. Other representations, such as the Causal Net and Sequence Chart, instead are more effective for deriving observations. Such findings underscore the benefits of analyzing visualization use over time and across multiple perspectives to identify both strengths and areas for improvement.

The ability of our method to capture and analyze rich, contextually grounded data, makes it well-suited for exploratory studies where research questions are not fully predefined, or where unexpected user behaviors emerge. By structuring user behavior from a process perspective, it provides a means to investigate workflows, cognitive strategies, and analytical intents, also in less controlled settings [Nor06]. In the future, it may also serve as a foundation for formalizing user strategies into structured goal (or intent) hierarchies, or event-condition-patterns, aligned with prior work on problem-solving models [JK96, MM21]. Thus, it could provide deeper insights and inform analysts about how to effectively navigate analytical tasks with VA tools.

The richness of our data, capturing multi-perspective and multi-granular event sequences, provides valuable insights into analysts' behavior. However, effectively handling such detailed and layered data remains a challenge [ZSDF*21]. Since many existing analysis techniques do not support the exploration of multiple sequences simultaneously, during our use case, we had to *manually* merge sequences by identifying overlapping events. This step of event sequence generation (cf. Sect.3.4) requires iterative adjustments to event granularity, re-identification of overlaps, and regeneration of sequences based on emerging patterns. Advanced analysis methods supporting multi-perspective data could streamline this

process by allowing adjustments during analysis. Interactive tools like Eloquence [VN18], MAQUI [LLMB18], or those reviewed in [GGJ*22], offer promising directions for future improvements.

Limitations. In our observational study, we replicated realistic analysis conditions by allowing analysts to use their preferred tools and workflows. While the presence of an interviewer during the online meeting may have influenced participant behavior, we interfered as little as possible to ensure participants could work autonomously, as they would in their work practice.

Another limitation is that we demonstrate our method only in the process mining domain. However, we believe that its applicability is not limited to this domain. Many of the behaviors we observed in the screen recordings, such as switching between tools, consulting external references (e.g., spreadsheets, documentation, web browsers), and structuring insights over time, are common in analytical workflows across domains [WLH19]. Moreover, the key analytical questions our method addresses (*what* visualizations are used, *when*, and *how*) provide a transferable framework for studying multiple-view interfaces more broadly. Still, we acknowledge that future work is needed to demonstrate the applicability of our method to different domains.

Our approach relies on video recordings and think-aloud protocols, making it widely accessible. However, the event identification and sequence generation steps (cf. Sect. 4.2) demand significant manual effort, often taking weeks. While insightful, these steps are time-consuming, requiring multiple iterations to refine coding schemes and annotations. This limits scalability, especially as data volumes grow. To improve scalability, two strategies could help. First, integrating machine learning and artificial intelligence techniques could automate event segmentation and identification. Large language models show promise for analyzing think-aloud data [Yoo24]. Similarly, activity recognition models [PT21] trained on screen recordings and eye-tracking data or additional (visual) tool support [BJK*15] could aid in identifying events automatically. Second, preprocessing techniques (e.g., mouse movement capture or eye-tracking) could be set up to enable multi-granular event identification for certain perspectives directly during data collection [AALW23]. For example, existing approaches can help to automatically detect the data objects users engage with [JA17]. Additionally, interaction logs, where available, can be used to track user interactions. However, this approach depends on access to relevant tools and requires well-defined research questions, which limits its applicability for exploratory real-world scenarios.

Conclusion. In summary, in this paper, we introduced a novel process-oriented approach for analyzing video recordings of analysis sessions across multiple perspectives and levels of granularity to reveal insights into visualization use and effectiveness. We demonstrated the approach by studying the work of process mining analysts and uncovered previously unexplored patterns in the use of process model visualizations, other process representations, and chart visualizations. Moving forward, investigating automation for event identification from screen recordings capturing analytical and cognitive steps at multiple levels of granularity and the development of more sophisticated tools for multi-perspective event sequence analysis could further strengthen our approach. While the scalability and transferability of our approach require further ex-

ploration, its ability to capture rich, contextual patterns of visualization use makes it a valuable method for studying complex, dynamic analysis workflows.

Acknowledgment

We thank Dagstuhl Seminar 23271, where the collaboration laying a key foundation for this work began. We also thank all participants of our study for their valuable time and insights.

This work was funded by the Swiss National Science Foundation as part of the ProMiSE project under Grant No.: 200021_197032, and by the Swedish Research Council under Grant No.: 2020-05000.

References

- [AALW23] ABBAD-ANDALOUSSI A., LÜBKE D., WEBER B.: Conducting eye-tracking studies on large and interactive process models using eyemind. *SoftwareX* 24 (2023), 101564. [3](#), [10](#)
- [All83] ALLEN J. F.: Maintaining knowledge about temporal intervals. *Communications of the ACM* 26, 11 (1983), 832–843. [5](#)
- [AR24] ABB L., REHSE J.-R.: Process-related user interaction logs: State of the art, reference model, and object-centric implementation. *Information Systems* 124 (2024), 102386. [doi:10.1016/j.is.2024.102386](#). [4](#)
- [AS05] AMAR R. A., STASKO J. T.: Knowledge precepts for design and evaluation of information visualizations. *IEEE Transactions on Visualization and Computer Graphics* 11, 4 (2005), 432–442. [6](#)
- [BJK*15] BLASCHECK T., JOHN M., KURZHALS K., KOCH S., ERTL T.: Va 2: a visual analytics approach for evaluating visual analytics applications. *IEEE transactions on visualization and computer graphics* 22, 1 (2015), 61–70. [2](#), [10](#)
- [BM13] BREHMER M., MUNZNER T.: A multi-level typology of abstract visualization tasks. *IEEE transactions on visualization and computer graphics* 19, 12 (2013), 2376–2385. [4](#), [6](#)
- [BWG*22] BADAKHSHAN P., WURM B., GRISOLD T., GEYER-KLINGEBERG J., MENDLING J., VOM BROCKE J.: Creating business value with process mining. *The Journal of Strategic Information Systems* 31 (12 2022), 101745. [doi:10.1016/J.JSIS.2022.101745](#). [3](#)
- [Car08] CARPENDALE S.: Evaluating information visualizations. In *Information visualization: Human-centered issues and perspectives*. Springer, 2008, pp. 19–45. [1](#), [2](#)
- [CASUCR22] CAPITÁN-AGUDO C., SALAS-URBANO M., CABANILLAS C., RESINAS M.: Analyzing how process mining reports answer time performance questions. In *International Conference on Business Process Management* (2022), Springer, pp. 234–250. [3](#)
- [Cel24] CELONIS: Process mining platform. <https://www.celonis.com/platform/process-mining/>, 2024. Accessed: 2024-10-24. [3](#), [5](#)
- [DJS*09] DOU W., JEONG D. H., STUKES F., RIBARSKY W., LIPFORD H. R., CHANG R.: Recovering reasoning processes from user interactions. *IEEE computer graphics and applications* 29, 3 (2009), 52–61. [2](#)
- [dLM15] DE LEONI M., MANNHARDT F.: Road traffic fine management process. *Eindhoven University of Technology, Dataset* 284 (2015). [3](#), [5](#)
- [GGJ*22] GUO Y., GUO S., JIN Z., KAUL S., GOTZ D., CAO N.: Survey on visual analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics* 28, 12 (2022), 5091–5112. [doi:10.1109/TVCG.2021.3100413](#). [5](#), [10](#)
- [GGZL16] GUO H., GOMEZ S. R., ZIEMKIEWICZ C., LAIDLAW D. H.: A case study using visualization interaction logs and insight metrics to understand how analysts arrive at insights. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 51–60. [doi:10.1109/TVCG.2015.2467613](#). [2](#)
- [GK94] GONG R., KIERAS D.: A validation of the goms model methodology in the development of a specialized, commercial software application. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (1994), pp. 351–357. [2](#)
- [GK12] GOODMAN E., KUNIAVSKY M.: *Observing the user experience: A practitioner's guide to user research*. Elsevier, 2012. [1](#)
- [GR12] GÜNTHER C. W., ROZINAT A.: Disco: Discover your processes. In *Demonstration Track of the 10th International Conference on Business Process Management, BPM Demos 2012* (2012), CEUR-WS. org, pp. 40–44. [5](#)
- [GTS10] GRAMMEL L., TORY M., STOREY M.-A.: How information visualization novices construct visualizations. *IEEE transactions on visualization and computer graphics* 16, 6 (2010), 943–952. [2](#)
- [GZ08] GOTZ D., ZHOU M. X.: Characterizing users' visual analytic activity for insight provenance. In *2008 IEEE symposium on visual analytics science and technology* (2008), IEEE, pp. 123–130. [6](#)
- [HMSA08] HEER J., MACKINLAY J., STOLTE C., AGRAWALA M.: Graphical histories for visualization: Supporting analysis, communication, and evaluation. *IEEE transactions on visualization and computer graphics* 14, 6 (2008), 1189–1196. [2](#)
- [HS12] HEER J., SHNEIDERMAN B.: Interactive dynamics for visual analysis: A taxonomy of tools that support the fluent and flexible use of visualizations. *Queue* 10, 2 (2012), 30–55. [6](#)
- [HZZ08] HUANG Y., ZHANG L., ZHANG P.: A framework for mining sequential patterns from spatio-temporal event data sets. *IEEE Transactions on Knowledge and Data Engineering* 20, 4 (2008), 433–448. [doi:10.1109/TKDE.2007.190712](#). [4](#)
- [IIC*13] ISENBERG T., ISENBERG P., CHEN J., SEDLMAIR M., MÖLLER T.: A systematic review on the practice of evaluating visualization. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2818–2827. [2](#)
- [IZCC08] ISENBERG P., ZUK T., COLLINS C., CARPENDALE S.: Grounded evaluation of information visualizations. In *Proceedings of the 2008 Workshop on BEyond time and errors: novel evaluation methods for Information Visualization* (2008), pp. 1–8. [1](#)
- [JA17] JIANU R., ALAM S. S.: A data model and task space for data of interest (doi) eye-tracking analyses. *IEEE Transactions on visualization and computer graphics* 24, 3 (2017), 1232–1245. [10](#)
- [JK96] JOHN B. E., KIERAS D. E.: Using goms for user interface design and evaluation: which technique? *ACM Transactions on Computer-Human Interaction (TOCHI)* 3, 4 (1996), 287–319. [2](#), [10](#)
- [KMW19] KLINKMÜLLER C., MÜLLER R., WEBER I.: Mining process mining practices: an exploratory characterization of information needs in process analytics. In *Business Process Management: 17th International Conference, BPM 2019, Vienna, Austria, September 1–6, 2019, Proceedings* 17 (2019), Springer, pp. 322–337. [2](#), [3](#)
- [LBI*11] LAM H., BERTINI E., ISENBERG P., PLAISANT C., CARPENDALE S.: Empirical studies in information visualization: Seven scenarios. *IEEE transactions on visualization and computer graphics* 18, 9 (2011), 1520–1536. [1](#), [2](#)
- [LLMB18] LAW P.-M., LIU Z., MALIK S., BASOLE R. C.: Maqui: Interweaving queries and pattern mining for recursive event sequence exploration. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 396–406. [10](#)
- [LW97] LEWIS C., WHARTON C.: Cognitive walkthroughs. In *Handbook of human-computer interaction*. Elsevier, 1997, pp. 717–732. [2](#)
- [MM21] MALINOVA M., MENDLING J.: Cognitive diagram understanding and task performance in systems analysis and design. *MIS Quarterly* 45, 4 (2021), 2101–2157. [doi:10.25300/MISQ/2021/15262](#). [10](#)
- [Mon08] MONDADA L.: Using video for a sequential and multimodal analysis of social interaction: Videotaping institutional telephone calls. *Forum Qualitative Sozialforschung / Forum: Qualitative Social Research* 9, 3 (2008). [2](#)

- [Mun09] MUNZNER T.: A nested model for visualization design and validation. *IEEE transactions on visualization and computer graphics* 15, 6 (2009), 921–928. [1](#)
- [Nie92] NIELSEN J.: Finding usability problems through heuristic evaluation. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (1992), pp. 373–380. [2](#)
- [Nor06] NORTH C.: Toward measuring visualization insight. *IEEE Computer Graphics and Applications* 26, 3 (2006), 6–9. [1](#), [10](#)
- [NSD11] NORTH C., SARAIYA P., DUCA K.: A comparison of benchmark task and insight evaluation methods for information visualization. *Information Visualization* 10, 3 (2011), 162–181. [2](#)
- [PT21] PAREEK P., THAKKAR A.: A survey on video-based human action recognition: recent updates, datasets, challenges, and applications. *Artificial Intelligence Review* 54, 3 (2021), 2259–2322. [10](#)
- [Sal15] SALDAÑA J.: *The coding manual for qualitative researchers*. Sage publications Ltd, 2015. [4](#), [5](#)
- [SKFZ21] SOURE E. J., KUANG E., FAN M., ZHAO J.: Coux: collaborative visual analysis of think-aloud usability test videos for digital interfaces. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 643–653. [1](#)
- [SND05] SARAIYA P., NORTH C., DUCA K.: An insight-based methodology for evaluating bioinformatics visualizations. *IEEE transactions on visualization and computer graphics* 11, 4 (2005), 443–456. [1](#), [2](#), [10](#)
- [SP06] SHNEIDERMAN B., PLAISANT C.: Strategies for evaluating information visualization tools: multi-dimensional in-depth long-term case studies. In *Proceedings of the 2006 AVI workshop on BEyond time and errors: novel evaluation methods for information visualization* (2006), pp. 1–7. [1](#), [10](#)
- [SSH*23] SOROKINA E., SOFFER P., HADAR I., LERON U., ZERBATO F., WEBER B.: Pem4ppm: A cognitive perspective on the process of process mining. In *International Conference on Business Process Management* (2023), Springer, pp. 465–481. [3](#)
- [SSS*14] SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G., KEIM D. A.: Knowledge generation model for visual analytics. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 1604–1613. [4](#), [9](#)
- [VBVDvdA10] VERBEEK H., BUIJS J., VAN DONGEN B., VAN DER AALST W. M.: Prom 6: The process mining toolkit. *Proc. of BPM Demonstration Track 615* (2010), 34–39. [3](#), [5](#)
- [VDA19] VAN DER AALST W. M.: A practitioner's guide to process mining: Limitations of the directly-follows graph, 2019. [3](#), [9](#)
- [vdA22] VAN DER AALST W. M.: Process mining: a 360 degree overview. In *Process Mining Handbook*. Springer, 2022, pp. 3–34. [2](#)
- [VER24] VERBI SOFTWARE: MAXQDA: Software for qualitative data analysis, 2024. Accessed: 2024-11-06. URL: <https://www.maxqda.com/de/>. [5](#)
- [vLFB*14] VON LANDESBERGER T., FIEBIG S., BREMM S., KUIJPER A., FELLNER D. W.: *Interaction Taxonomy for Tracking of User Actions in Visual Analytics Applications*. Springer New York, 2014, pp. 653–670. [doi:10.1007/978-1-4614-7485-2_26](https://doi.org/10.1007/978-1-4614-7485-2_26). [4](#)
- [VN18] VROTSOU K., NORDMAN A.: Exploratory visual sequence mining based on pattern-growth. *IEEE transactions on visualization and computer graphics* 25, 8 (2018), 2597–2610. [10](#)
- [vW13] VAN WIJK J. J.: Evaluation: A challenge for visual analytics. *Computer* 46, 7 (2013), 56–60. [2](#)
- [WAAF*24] WEBER B., ABBAD-ANDALOUSSI A., FRANCESCHETTI M., SEIGER R., VÖLZER H., ZERBATO F.: Leveraging digital trace data to investigate and support human-centered work processes. In *Evaluation of Novel Approaches to Software Engineering* (Cham, 2024), Springer Nature Switzerland, pp. 1–23. [doi:10.1007/978-3-031-64182-4_1](https://doi.org/10.1007/978-3-031-64182-4_1). [3](#)
- [WKG10] WARD M. O., GRINSTEIN G., KEIM D.: *Interactive data visualization: foundations, techniques, and applications*. AK Peters/CRC Press, 2010. [1](#)
- [WLH19] WONGSUPHASAWAT K., LIU Y., HEER J.: Goals, process, and challenges of exploratory data analysis: An interview study. *arXiv preprint arXiv:1911.00568* (2019). [10](#)
- [XOW*20] XU K., OTTLEY A., WALCHSHOFER C., STREIT M., CHANG R., WENSKOVITCH J.: Survey on the analysis of user interactions and visualization provenance. In *Computer Graphics Forum* (2020), vol. 39, Wiley Online Library, pp. 757–783. [1](#), [2](#)
- [YaKSJ07] YI J. S., AH KANG Y., STASKO J., JACKO J. A.: Toward a deeper understanding of the role of interaction in information visualization. *IEEE transactions on visualization and computer graphics* 13, 6 (2007), 1224–1231. [4](#), [6](#)
- [YM24] YESHCENKO A., MENDLING J.: A survey of approaches for event sequence analysis and visualization. *Information Systems* 120 (2024). [doi:10.1016/j.is.2023.102283](https://doi.org/10.1016/j.is.2023.102283). [2](#)
- [Yoo24] YOO Y.: Automated think-aloud protocol for identifying students with reading comprehension impairment using sentence embedding. *Applied Sciences* 14, 2 (2024), 858. [10](#)
- [Zhu07] ZHU Y.: Measuring effective data visualization. In *International symposium on visual computing* (2007), Springer, pp. 652–661. [9](#)
- [ZSDF*21] ZERBATO F., SEIGER R., DI FEDERICO G., BURATTIN A., WEBER B.: Granularity in process mining: Can we fix it? In *International Workshop on BPM Problems to Solve Before We Die* (2021), CEUR-WS, pp. 40–44. [5](#), [10](#)
- [ZSW21] ZERBATO F., SOFFER P., WEBER B.: Initial insights into exploratory process mining practices. In *Business Process Management Forum* (Cham, 2021), Springer International Publishing, pp. 145–161. [doi:10.1007/978-3-030-85440-9_9](https://doi.org/10.1007/978-3-030-85440-9_9). [2](#), [3](#)
- [ZSW22] ZERBATO F., SOFFER P., WEBER B.: Process mining practices: Evidence from interviews. In *International Conference on Business Process Management* (2022), Springer, pp. 268–285. [3](#)
- [ZZVW23] ZERBATO F., ZIMMERMANN L., VÖLZER H., WEBER B.: Promise: Process mining support for end-users. In *CAiSE Research Projects Exhibition* (2023), pp. 60–67. [5](#)
- [ZZW24] ZIMMERMANN L., ZERBATO F., WEBER B.: What makes life for process mining analysts difficult? a reflection of challenges. *Software and Systems Modeling* 23, 6 (2024), 1345–1373. [doi:10.1007/s10270-023-01134-0](https://doi.org/10.1007/s10270-023-01134-0). [3](#)