Linköping Studies in Science and Technology
Dissertation No. 2543

TEMPORAL REASONING AND TRACES

VISUAL ANALYSIS OF COMPLEX EVENT SEQUENCES

LINKÖPING UNIVERSITY

PEILIN YU

AVAILABLE EVERYWHERE
25.09.2026

Linköping Studies in Science and Technology
Dissertations, No. 2543

TRACES

Temporal Reasoning and Visual Analysis of Complex Event Sequences

PEILIN YU

Linköping University
Department of Science and Technology
Division of Media and Information Technology
Linköping University, SE–601 74 Norrköping, Sweden

Norrköping, 2026

Typeset using XƎTEX

Printed by LiU-Tryck, Linköping 2026

Edition 1:1

ISBN 978–91–8118–687–1 (print)
ISBN 978–91–8118–688–8 (PDF)
https://doi.org/10.3384/9789181186888
ISSN 0345-7524

ABSTRACT

Event sequence data, such as system logs, interaction logs, eye-tracking recordings, and electronic health records, are collected widely across application domains to record how activities and processes unfold over time. Such data are rarely ready for direct analysis. Raw sequences are often characterized by large event alphabets, varying lengths, high fragmentation, and substantial noise, and a persistent gap separates low-level logged events from the higher-level constructs that produced them, such as exploration strategies and cognitive processes. Fully automated methods risk discarding infrequent yet analytically significant information, whereas purely manual inspection becomes impractical as data grow in volume and heterogeneity. These challenges motivate visual analytics (VA) approaches that couple computational methods with interactive visualization while keeping the human analyst in the loop. This thesis investigates how VA can support analysts in transforming, exploring, comparing, extracting, and interpreting complex event sequence data derived from interaction logs.

The thesis comprises four papers that form a methodological progression along the VA process. First, a VA approach is presented for improving the quality of noisy and fragmented event sequences in the domain of en-route air traffic control. The approach supports transparent, context-aware transformations and provides visual cues about potential information loss, enabling analysts to assess the impact of each transformation step and improve data quality progressively while preserving analytically significant information. Once sequences are prepared, the analytical focus shifts to what they reveal about user behavior. Accordingly, the second paper proposes a VA approach for analyzing interaction logs from an interactive touchscreen exhibit in an informal learning environment. Coordinated visual representations support navigation across multiple levels of granularity, enabling analysts to inspect individual behaviors, compare users, and identify recurring exploration strategies. As such data extend from controlled studies to unstructured, in-the-wild collections, individual inspection alone becomes insufficient. The third paper therefore introduces a VA workflow that integrates a tailored similarity measure, hierarchical clustering, and process mining to guide the stepwise extraction, characterization, and validation of interaction patterns in large, unstructured interaction logs. Finally, the fourth paper addresses the gap between observable behavior and cognition through a VA approach grounded in a systems thinking framework. Interaction sequences are transformed into multidimensional time series, and motif discovery is coupled with similarity analysis and interactive clustering to characterize pupils' problem-solving behavior in a digital learning environment for the carbon cycle, supporting hypothesis generation about underlying reasoning processes. The proposed approaches are implemented in interactive systems and examined through usage scenarios, case studies, and expert workshops conducted in close collaboration with domain experts.

Together, these contributions advance VA for event sequence analysis across diverse contexts, including air traffic control, interactive public exhibits, and digital learning environments. They demonstrate how transparent, interactive, and domain-informed workflows can render complex event sequence data analytically accessible and help analysts connect fine-grained interaction traces to higher-level interpretations of user behavior.

SVENSK SAMMANFATTNING

Händelsesekvensdata, såsom systemloggar, interaktionsloggar, ögonrörelseinspelningar och elektroniska patientjournaler, samlas in i stor omfattning inom olika tillämpningsområden för att dokumentera hur aktiviteter och processer utvecklas över tid. Sådana data är sällan direkt redo att analyseras. Råa sekvenser kännetecknas ofta av stora händelsealfabet, varierande längd, hög fragmentering och betydande brus. Dessutom råder en bestående klyfta mellan loggade lågnivåhändelser och de övergripande konstruktioner som gett upphov till dem, såsom utforskningsstrategier och kognitiva processer. Helt automatiserade metoder riskerar att förkasta sällsynt men analytiskt viktig information, medan rent manuell granskning blir opraktisk när data växer i volym och heterogenitet. Dessa utmaningar motiverar visual analytics (VA) metoder, det vill säga visuella analysmetoder som kopplar samman beräkningsmetoder med interaktiv visualisering samtidigt som den mänskliga analytikern förblir delaktig i analysprocessen. Denna avhandling undersöker hur VA kan stödja analytiker i att transformera, utforska, jämföra, extrahera och tolka komplexa händelsesekvensdata härledda ur interaktionsloggar.

Avhandlingen omfattar fyra projekt som bildar en metodologisk progression längs VA-processen. I den första artikeln presenteras en VA-metod för att förbättra kvaliteten hos brusiga och fragmenterade händelsesekvenser inom flygtrafikledning. Metoden stödjer transparenta, kontextmedvetna transformationer och ger visuella ledtrådar om potentiell informationsförlust, vilket gör det möjligt för analytiker att bedöma effekten av varje transformationssteg och successivt förbättra datakvaliteten samtidigt som analytiskt viktig information bevaras. När sekvenserna väl är förberedda förskjuts det analytiska fokuset mot vad de avslöjar om användarbeteende. Följaktligen föreslås i den andra artikeln en VA-metod för att analysera interaktionsloggar från en interaktiv pekskärmsutställning i en informell lärandemiljö. Koordinerade visuella representationer stödjer navigering mellan flera detaljnivåer, vilket gör det möjligt för analytiker att granska enskilda beteenden, jämföra användare och identifiera återkommande utforskningsstrategier. När sådana data utvidgas från kontrollerade studier till ostrukturerade samlingar insamlade i verkliga miljöer räcker enskild granskning inte längre till. I den tredje artikeln introduceras därför ett VA-arbetsflöde som integrerar ett skräddarsytt likhetsmått, hierarkisk klustring och processutvinning för att vägleda en stegvis extrahering, karakterisering och validering av interaktionsmönster i stora, ostrukturerade interaktionsloggar. Slutligen behandlar den fjärde artikeln klyftan mellan observerbart beteende och kognition genom en VA-metod förankrad i ett ramverk för systemtänkande. Interaktionssekvenser transformeras till flerdimensionella tidsserier, och motivupptäckt kombineras med likhetsanalys och interaktiv klustring för att karakterisera elevers problemlösningsbeteende i en digital lärandemiljö om kolets kretslopp, vilket stödjer hypotesgenerering om underliggande resonemangsprocesser. De föreslagna metoderna är implementerade i interaktiva system och prövas genom användningsscenarier, fallstudier och expertworkshoppar genomförda i nära samarbete med domänexperter.

Tillsammans för dessa bidrag VA för händelsesekvensanalys framåt i skilda sammanhang, däribland flygtrafikledning, interaktiva offentliga utställningar och digitala lärandemiljöer. De visar hur transparenta, interaktiva och domäninformerade arbetsflöden kan göra komplexa händelsesekvensdata analytiskt tillgängliga och hjälpa analytiker att koppla finkorniga interaktionsspår till tolkningar av användarbeteende på högre nivå.

PROLOGUE

First and foremost, I thank my beloved supervisors, **Katerina Vrotsou** and **Aida Nordman**, without whom none of this would have been possible.

Thank you for the journey this far: for the guidance, the support, the joy, and the suffering, all of which made this an unforgettable experience and a lifelong lesson.

You have now been at least five percent of my life, assuming I live to a hundred.

This thesis closes one chapter of that journey.

For the unwritten: may we live long and prosper.

LIST OF PUBLICATIONS

Paper I Peilin Yu, Aida Nordman, Lothar Meyer, Supathida Boonsong, and Katerina Vrotsou. "Interactive Transformations and Visual Assessment of Noisy Event Sequences: An Application in En-Route Air Traffic Control." In: *2023 IEEE 16th Pacific Visualization Symposium (PacificVis)*. 2023, pp. 92–101. DOI: `10.1109/PacificVis56936.2023.00017`

Paper II Peilin Yu, Aida Nordman, Marta Koc-Januchta, Konrad Schönborn, Lonni Besançon, and Katerina Vrotsou. "Revealing Interaction Dynamics: Multi-Level Visual Exploration of User Strategies with an Interactive Digital Environment." In: *IEEE Trans. Vis. Comput. Graph.* 31.1 (2025), pp. 831–841. DOI: `10.1109/TVCG.2024.3456187`

Paper III Peilin Yu, Aida Nordman, Takanori Fujiwara, Marta Koc-Januchta, Konrad Schönborn, Lonni Besançon, and Katerina Vrotsou. "Visual Extraction of Interaction Patterns Guided by Hierarchical Clustering and Process Mining." In: *IEEE Trans. Vis. Comput. Graph.* 32.1 (2026), pp. 1142–1152. DOI: `10.1109/TVCG.2025.3634837`

Paper IV Peilin Yu, Aida Nordman, Mina Mani, Konrad Schönborn, Marta Koc-Januchta, Gunnar Höst, Måns Gezelius, Lonni Besançon, and Katerina Vrotsou. "From Interaction to Systems Thinking: A Visual Analytics Approach for Exploring Pupils' Behavior in a Digital Learning Environment." Submitted to IEEE Trans. Vis. Comput. Graph. 2026. DOI: `10.35542/osf.io/9zp7k_v1`

CONTENTS

Abstract **iii**

Svensk sammanfattning **v**

Prologue **vii**

List of Publications **ix**

Contents **xi**

1 Trailhead **1**
1.1 Motivation 2
1.2 Visual Analytics as an Approach 3
1.3 Research Scope and Domain Context 3
1.4 Research Questions 4
1.5 Thesis Contributions 5
1.6 Thesis Structure 6

2 Landscape **7**
2.1 Event Sequence Data 7
2.2 Event Sequence Visualization and Visual Analytics 9
2.3 Algorithmic Methods for Event Sequence Analysis 13
2.4 Bridging Event Sequences and Higher-Level Constructs 16
2.5 Application Contexts 16

3 Waypoints **19**
3.1 An Integrated Pipeline 19
3.2 Positioning the Contributions 19

4 Fragmented Trace **21**
4.1 Background and Domain 21
4.2 Methodology 22
4.3 Contributions and Conclusions 24

5 Exploratory Trace **27**
5.1 Background and Domain 27
5.2 Methodology 28
5.3 Contributions and Conclusions 29

6 Unstructured Trace **33**
6.1 Background and Domain 33

6.2 Methodology . . . 34
6.3 Contributions and Conclusions . . . 34

7 Systemic Trace **37**
7.1 Background and Domain . . . 37
7.2 Methodology . . . 38
7.3 Contributions and Conclusions . . . 40

8 Synthesis **43**

9 Horizons **45**
9.1 Contributions . . . 45
9.2 Conclusions . . . 45
9.3 Limitations and Future Work . . . 47
9.4 Onward . . . 48

Epilogue **49**

Bibliography **51**

Paper I **63**

Paper II **75**

Paper III **89**

Paper IV **103**

CHAPTER 1

TRAILHEAD

In the summer of 1854, a severe cholera outbreak killed more than 600 people in the Soho district of London within a matter of weeks. The prevailing medical theory of the time held that cholera spread through "bad air". John Snow, a physician who was skeptical of this theory, traced cholera deaths by plotting bars on a street map of the neighborhood (Fig. 1.1), showing a striking spatial concentration of cases near a single water source. This finding eventually led to the removal of the pump handle, and the outbreak subsided. Snow's map not only described the epidemic but also made a causal argument visible in a way that words and tables of numbers could not, resulting in a life-saving intervention.

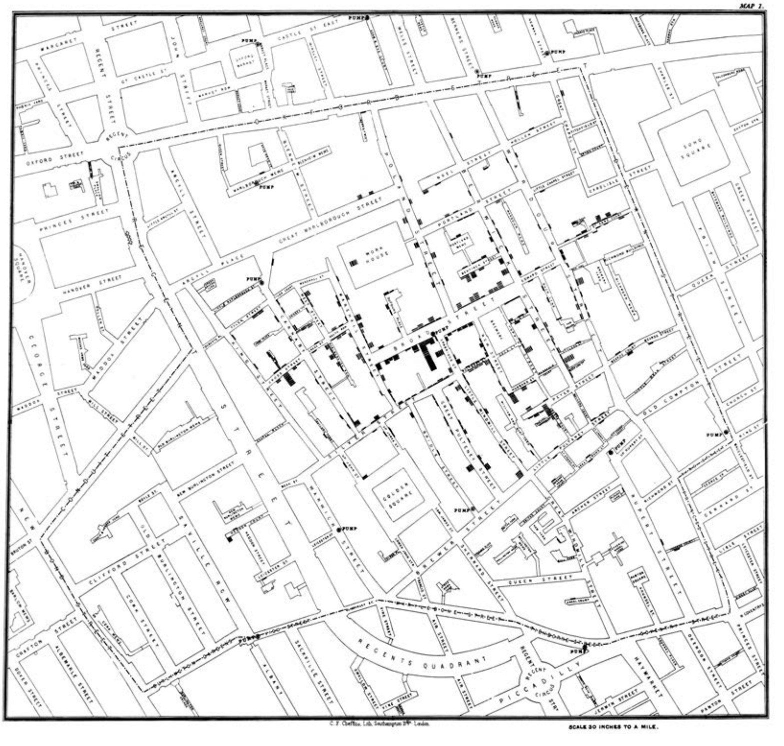

Figure 1.1: John Snow's 1854 cholera map of the Broad Street outbreak in London [95]. Black stacked bars indicate cholera deaths, revealing a clear concentration of cases around the contaminated Broad Street pump. This image is in the public domain.

1.1 Motivation

Although more than a century and a half separates Snow's map from today's digital environments, its central premise endures: carefully designed visual representations can reveal structure hidden within complex data and turn observation into actionable understanding. Digital systems have become an inescapable part of how people work, learn, and explore information. Behind these everyday interactions are traces of human activity: a tap on a screen, a click on an interface, a dragged object, a gaze shift, a selected option. Over time, these traces accumulate into interaction logs and event sequences that offer a detailed yet often complex account of how people engage with digital devices and environments. Such sequence data now appear across health informatics, internet applications, and Industry 4.0, where their scale, dimensionality, and heterogeneity make them both valuable and difficult to interpret [45].

Unlike aggregated measures, such as total time, click counts, or final task outcomes, interaction logs preserve the temporal structure of behavior. This temporal structure is crucial, as behavior in digital environments is rarely a set of isolated actions; users revisit options, compare alternatives, recover from mistakes, test assumptions, and adapt their actions as a task unfolds. Interaction logs can therefore reveal not only what users did, but how they moved through a system. For system designers, they can support interface redesign, surface usability problems, and expose discrepancies between intended and actual use [13, 109]. For educators and researchers, logs from digital learning environments can provide insight into otherwise hidden cognitive processes, such as how pupils approach problems, where they struggle, and which strategies they develop over time [22, 25]. For experts in domains such as air traffic control, medical monitoring, and cybersecurity, logs can support the analysis of operator performance, competency assessment, and the design of future automation support [39, 62, 79].

Across these diverse application domains, a common analytical challenge emerges: moving beyond knowing *what* a user did toward understanding *how* and *why* they did it. To address this, a wide range of techniques has been developed for the visual representation, exploration, and comparison of event sequence data [45, 99]. Despite this progress, transforming raw interaction logs into meaningful accounts of human behavior remains an open and evolving challenge, rooted in the data's intrinsic properties, domain differences, the limitations of existing analytical approaches, and the rapid emergence of new, increasingly complex digital environments. Four intertwined difficulties, which persist across application domains and jointly motivate this thesis, are examined as follows: the nature of the data, the insufficiency of fully automated approaches, the impracticality of purely manual analysis, and the gap between event sequences and their interpretation.

The nature of the data. Event sequences tend to be long and contain many distinct event types, spanning dozens to hundreds. Spurious or uninformative events frequently fragment them, and the sequences exhibit high variability in both event frequency and duration. These properties complicate the identification of recurring behavioral patterns, both within individual sequences and across user groups.

The insufficiency of fully automated approaches. Computational methods such as pattern mining, clustering, and sequence modeling handle scale efficiently [29, 45, 111], but they are fundamentally domain-agnostic. A frequent subsequence need not be behaviorally meaningful, and a rare event need not be irrelevant. Without domain knowledge, automated preprocessing risks discarding precisely what should be preserved, and no downstream analysis can recover what has been lost [54]. Moreover, purely automated pipelines offer little transparency: the analyst often receives a processed result with no visibility into what was changed, what was removed, and why [45, 91, 108].

The impracticality of purely manual analysis. Domain experts bring irreplaceable knowledge and interpretive judgment to the analysis, but inspecting raw sequences of many events across multiple sessions without computational support is impractical, time-consuming, and difficult to reproduce [54]. Manual analysis is also vulnerable to inconsistency, as findings produced without systematic and traceable tools become difficult to compare, verify, reproduce, or build upon [81, 94, 109].

The gap between event sequences and interpretation. Interaction logs typically capture low-level behavior, such as which object was clicked, which region was inspected, or which connection was made. Yet the questions that researchers, educators, and system designers seek to answer are generally high-level, such as what strategy a user employs, or what an identified pattern reveals about expertise, understanding, or learning. The mapping from observable actions to such higher-level constructs, and ultimately to reasoning, is neither trivial nor automatic. It instead requires analytical framing, domain expertise, and human judgment that no algorithm can fully substitute for [21, 74, 91].

1.2 Visual Analytics as an Approach

Taken together, the four challenges outlined above motivate the use of *visual analytics* [55, 97]. Visual analytics has been defined as the science of analytical reasoning facilitated by interactive visual interfaces [97], integrating automated analysis with interactive visualization to support reasoning over complex data. Its defining characteristic is the tight coupling of computational methods and visual representations through interaction, which keeps the human analyst in the loop throughout the analytical process. This coupling directly addresses the four challenges above. Domain knowledge can be incorporated into automated transformation before analytically significant information is lost; the operations applied to the data can be made transparent and traceable; the scale of the data can be handled computationally while interpretation remains under expert control; and the path from low-level events to higher-level interpretation can be supported in a way that no fully automated pipeline can provide.

Building on this paradigm, this thesis investigates how visual analytics can support the analysis of complex event sequence data throughout the analytical pipeline, from transformation to interpretation. This investigation is grounded in, and examined through, concrete application domains, each contributing distinct interaction log data with different challenges and analytical demands.

1.3 Research Scope and Domain Context

The scope of this thesis is the visual analysis of complex event sequence data derived from interaction logs. The main focus is on interactive, human-in-the-loop analysis that incorporates automated methods to support analysts in transforming, exploring, comparing, extracting, and interpreting complex interaction sequences. This analysis is developed and examined across two application domains that span a spectrum of settings: a safety-critical professional activity on one end, and, on the other, two learning contexts ranging from informal public engagement to structured formal education. The domains are not studied in isolation; they serve as complementary settings through which a common set of analytical challenges, spanning data transformation to higher-level interpretation, is investigated.

The first setting concerns ***air traffic control*** (ATC). En-route control involves maintaining safe and efficient traffic flow within a controlled airspace sector, with the primary objective of ensuring minimum separation between aircraft. One of the central tasks performed by *air traffic controllers* (ATCOs) in this context is *conflict detection and resolution* (CD&R), which requires continuous visual scanning and interaction with complex control interfaces to detect impending conflicts between aircraft, formulate a resolution strategy, and issue the appropriate clearances [65]. Understanding how ATCOs carry out this task, and which behavioral patterns characterize their scanning and resolution strategies, is of direct interest to ATC researchers and trainers seeking to assess competency and inform the design of future automation support.

The second setting concerns **informal learning** [5] in a public science center setting, where an interactive exhibit, "*Sweden in Numbers*" (SiN), was investigated. The exhibit is implemented on a large touchscreen table that enables visitors to explore complex demographic and societal data about Sweden through an interactive map. This setting involves diverse groups of people with varying backgrounds, motivations, and prior knowledge engaging freely with a digital environment. Understanding how such visitors explore the exhibit, which strategies they employ, and what their interaction patterns reveal about engagement and cognitive processing is of direct interest to educational researchers, exhibit designers, and science center stakeholders.

The third setting concerns **formal learning** in a digital learning environment, *Tracing Carbon* (TC) [68]. In this environment, pupils demonstrate their understanding of the global carbon cycle by drawing directed connections between carbon reservoirs. This activity is task-directed and embedded in a structured educational objective. The analytical challenge is not merely to describe what pupils did, but to interpret the meaning of their exploration patterns in relation to higher-level cognitive constructs and reasoning, making this domain of direct interest to educators and learning researchers seeking insight into how understanding develops.

1.4 Research Questions

Motivated by the challenges outlined above and situated in the three application settings described, this thesis develops visual analytics approaches to support the transformation, exploration, comparison, extraction, and interpretation of complex event sequence data. Four research questions are formulated as follows:

RQ1. **How can visual analytics approaches support the transformation of noisy event sequence data while preserving analytically significant information?**

Raw event sequences are rarely ready for direct analysis. Noise, fragmentation, and irrelevant events can obscure meaningful patterns, yet fully automated preprocessing may remove infrequent but important events without analysts' awareness. Therefore, this question focuses on interactive, transparent, and domain-informed transformation, enabling analysts to assess the impact of each transformation step and to improve the quality of downstream analysis.

RQ2. **How can visual analytics methods support the multi-level exploration and comparison of complex event sequences across different levels of granularity?**

Once event sequence data have been prepared, the analytical challenge shifts to understanding what the sequences reveal. Meaningful behavior manifests at different scales, yet individual events are too fine-grained to reveal overarching strategies; aggregated summaries hide the actions that constitute them; and recurring strategies only become apparent when sequences are viewed in relation to one another. Therefore, this question focuses on multi-level representation and comparison of event sequences, supporting analysts in exploring and inspecting user behavior across levels of granularity, from individual events to entire tasks and sessions, and in comparing sequences across users, tasks, or sessions.

RQ3. **How can algorithmic workflows support guided, human-in-the-loop characterization and identification of complex interaction patterns?**

As event sequence data grow from controlled studies to larger, more heterogeneous collections, individual inspection becomes inefficient and insufficient. Algorithmic methods can extract candidate patterns at scale, yet their results are only useful if analysts can understand, steer, and act upon them. Therefore, this question focuses on interactive workflows that couple algorithmic pattern extraction with coordinated visualization, guiding analysts toward potentially meaningful structures while keeping algorithmic results accessible, interpretable, and verifiable.

RQ4. **How can visual analytics systems support the interpretation and mapping of complex interaction patterns in relation to higher-level constructs, such as strategies, cognitive processes, and reasoning?**

Describing what users did is only part of the analytical goal; the questions that ultimately matter to domain experts concern what these interactions may imply about how users explore, learn, or reason. Such higher-level constructs are of central interest, yet they are not directly observable in interaction logs nor directly computable by any algorithm, and a persistent gap separates the recorded low-level events from the cognitive processes they may reflect. Therefore, this question focuses on theory-informed interpretation of interaction patterns, enabling analysts to connect low-level events to higher-level constructs, such as exploration strategies, learning processes, or reasoning frameworks grounded in application domains.

1.5 Thesis Contributions

In response to these questions, this thesis contributes visual analytics approaches for analyzing interaction event sequence data across three settings. The four included papers form a methodological progression: improving noisy raw sequences, revealing interaction strategies, extracting behavioral patterns, and finally interpreting these patterns in relation to higher-level constructs, such as learning frameworks and reasoning. Four main contributions are as follows:

C1. **Interactive transformation of noisy event sequences and quality assessment of the transformation results.** The thesis contributes a visual analytics approach that enables analysts to progressively transform noisy, complex event sequence data to improve their quality while preserving analytically significant information (RQ1). This contribution is realized in Paper I.

C2. **Multi-level exploration and comparison of interaction strategies.** The thesis contributes coordinated visualization methods that support the exploration and comparison of interaction sequences across multiple levels of granularity, enabling analysts to investigate individual behaviors, compare them, and identify recurring interaction strategies (RQ2 & RQ4). This contribution is mainly realized in Papers II and IV, and is further supported by Paper III.

C3. **User-driven extraction and characterization of interaction patterns.** The thesis contributes interactive analytical workflows that integrate algorithmic methods with interactive visualization to guide analysts in stepwise extraction, characterization, and validation of behavioral patterns from unstructured, complex interaction sequences (RQ3 & RQ2). This contribution is realized in Papers III and IV.

C4. **Theory-grounded interpretation of interaction patterns in relation to higher-level cognitive constructs.** The thesis contributes visual analytics designs that couple domain-informed behavioral abstractions with established theoretical frameworks, enabling analysts to connect low-level interaction events to higher-level cognitive constructs, such as reasoning processes, learning strategies, and systems thinking, and to formulate evidence-based hypotheses about them (RQ4). This contribution is realized in Papers II and IV, and is supported by Paper III.

1.6 Thesis Structure

The remainder of this thesis is organized as follows:

- Chapter 2 surveys the background and related work on event sequence data, visual analytics, algorithmic pattern extraction, and the higher-level constructs through which interaction behavior is interpreted.
- Chapter 3 provides a brief synopsis of the four included papers and their relationships.
- Chapters 4 to 7 present the four papers in turn, progressing from the transformation of noisy event sequences, through multi-level exploration and guided pattern extraction, to theory-grounded interpretation.
- Chapter 8 synthesizes the papers and relates them to the research questions and contributions.
- Chapter 9 concludes with limitations and directions for future work.

CHAPTER 2

LANDSCAPE

What intersecting fields shape the landscape in which this thesis is situated? This chapter traces the connections among event sequence data and their properties, the visual analytics approaches developed to explore and compare them, the algorithmic methods that scale analysis beyond manual inspection, and the constructs, such as strategies, cognitive processes, and reasoning, through which low-level interactions acquire higher-level meaning.

2.1 Event Sequence Data

This section characterizes event sequence data as the object of study of this thesis, establishing what such data are, the levels of granularity at which they can be analyzed, and the challenges that arise when analyzing them.

2.1.1 Definition and Prevalence

Event sequence data record ordered occurrences of discrete events, commonly associated with an entity, case, user, object, or process instance. Formally, an event denotes a single discrete observation captured in time, and a sequence is an ordered list of such events tied to a common entity [45]. Each event typically includes a timestamp and often carries additional attributes, such as duration, user identifier, outcome, or other contextual metadata. Such data are collected widely across domains, including healthcare, finance, manufacturing, security, education, and interactive systems, and event sequence analysis has been positioned as an interdisciplinary field spanning information visualization, visual analytics, and human-centered interaction research [45]. The value of such data lies in their richness: they capture not only what happened, but also when, in what order, for how long, and with what associated attributes. Analyzing such rich data, however, is not straightforward; a first consideration, therefore, concerns the level of granularity at which the analysis is conducted.

2.1.2 Levels of Granularity

Event sequence data can be analyzed at several levels of granularity (Fig. 2.1), and the level chosen shapes both the questions that can be asked and the patterns that become visible. At the finest level, analysis focuses on the individual **event**, the fundamental unit at which low-level detail is inspected. Contiguous events form **subsequences**, which preserve local temporal order and are often the level at which recurring motifs and behavioral patterns are sought. A complete **sequence** constitutes the full ordered record associated with a single entity, such as one user, case, or session, and is the level at which an entity's overall progression is examined. At the coarsest level, a **cohort** is analyzed to summarize common behaviors across a population or to compare groups of entities [45]. These levels are complementary rather than mutually exclusive: meaningful analysis frequently requires moving between them, relating fine-grained events to the higher-level structures they compose, so that individual actions can be understood in the context of broader behavior. This motivates multi-granular analytical approaches, in which low-level events, intermediate interaction structures, and aggregate behavior are examined both separately

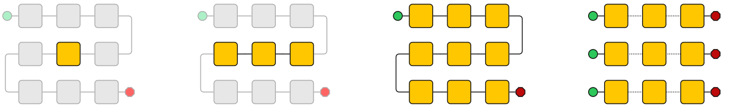

(a) Event. (b) Subsequence. (c) Sequence. (d) Cohort.

Figure 2.1: Levels of granularity at which interaction event sequence data can be analyzed, from a single event (a), through a subsequence (b) and a complete sequence (c), to a cohort of sequences (d). Behavioral structure becomes visible at different levels, which motivates the multi-level analysis pursued in this thesis. Illustrations adapted from Guo et al. [45].

and in relation to one another. However, analysis across these levels of granularity requires overcoming several inherent challenges associated with event sequence data.

2.1.3 Analytical Challenges

The characteristics that make event sequence data expressive are also those that make them difficult to analyze in raw form. Sequences tend to be long, span a large number of event types, and are highly fragmented and noisy [45]. These properties lead to large variations in how patterns manifest, making the identification of recurring behavior both within and across sequences difficult. When interaction logs are the subject of analysis, these difficulties are further pronounced, as such logs are typically unstructured and collected across many users, without a predefined schema of expected behavior.

As a consequence, raw sequences are rarely ready for direct analysis, and simplification via data transformation is often required to improve their comparability and applicability for downstream analysis [29]. Improving data quality is consequently recognized as an important step in the analysis process, which deserves greater attention [31, 36]. Fully automated simplification, however, carries distinct risks: infrequent, short-duration events may be discarded even though they are meaningful for the analytical task and may not be recovered once removed. This tension, between rendering event sequences tractable and preserving analytically significant detail, defines the first challenge of working with such data and motivates interactive, transparent, domain-informed transformation.

Transformation, however, resolves only part of the difficulty, as several challenges persist even in well-processed sequences. Analytically meaningful structures are rarely delimited in the data, as interaction logs capture low-level events, such as clicks, taps, or interface operations, whereas the tasks, sessions, and interaction strategies that analytical questions concern lack explicit boundaries in the log and must be constructed during analysis. As collections grow larger and more heterogeneous, manual inspection alone becomes insufficient, and algorithmic extraction becomes necessary. Finally, the constructs of interest, such as strategies, cognitive processes, and reasoning, are latent and must be inferred, a mapping that is neither automatic nor unique.

2.2 Event Sequence Visualization and Visual Analytics

In response to the challenges established in the previous section, this section positions visualization and visual analytics as analytical means, first introducing the two fields and the process by which they combine automated computation with human reasoning.

2.2.1 Information Visualization and Visual Analytics

Information visualization concerns the design of interactive visual representations of abstract data to amplify cognition and support exploratory analysis [19]. Its central premise is that the human visual system, supported by carefully designed representations and interactions, can externalize reasoning and reveal structures that are difficult to detect through automated computation alone. A foundational principle of the field is the Visual Information-Seeking Mantra, which states that any visualization system should provide an overview first, then zooming and filtering, and finally details on demand [90], and continues to guide the design of current systems.

As data have grown in scale and analytical complexity, visual representation alone is often insufficient, motivating the integration of automated computation with human reasoning. Visual analytics emerged from this need, defined as the science of analytical reasoning facilitated by interactive visual interfaces [97]. Within this paradigm, the analyst is not a passive recipient of a visualization but an active participant who interprets computational output, contributes domain knowledge, and iteratively steers the analysis [55]. Rather than treating automated analysis and visual representation as separate stages, visual analytics couples them through interaction, so that computational methods and human judgment reinforce one another throughout this process [55].

A recurring design consideration within this paradigm is the balance between automation and human control [27]. Fully automated pipelines can obscure the assumptions and transformations applied to the data, whereas purely manual analysis is impractical at scale. Human-in-the-loop designs address this gap by enabling analysts to investigate, validate, and refine algorithmically generated results with domain knowledge. This coupling allows automation to handle scale while human judgment retains control over interpretation, and it is this balance that situates the work of this thesis and foregrounds the visual analysis of complex event sequence data.

2.2.2 Visual Analytics Process

The process underlying visual analytics is commonly described as a closed loop in which data, automated models, and visualizations are coupled through user interaction [55], as illustrated in Fig. 2.2. Within this loop, raw data are first transformed and then mapped to both computational models and visual representations; the analyst inspects the visualizations, steers the underlying models, and iteratively refines both, while the resulting insight accumulates toward knowledge and feeds back to inform subsequent transformation and analysis [55]. Interaction is central to this process, serving as the means by which domain knowledge enters the analytical pipeline, allowing automated methods to scale analysis while keeping human judgment in control of what is preserved, examined, and interpreted.

This process model is adopted to frame the work of this thesis. Each of its four stages is annotated with the high-level *stage goals* the analysis aims to achieve at that stage. This annotation draws inspiration from existing typologies of event sequence analysis, in particular a multi-level task framework [120] and the design space and analytical tasks synthesized in a recent survey [45]. The stage goals are kept at a high level of abstraction, describing the kind of analytical work each stage serves rather than the specific methods utilized. In brief, the *data* stage concerns the transformation and quality assessment of

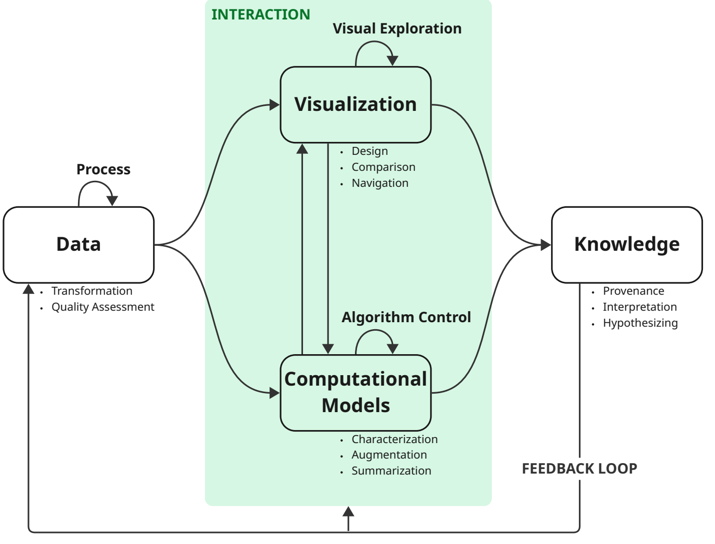

Figure 2.2: The visual analytics process adopted in this thesis, with each stage annotated by its high-level *stage goals*. Transformed data are passed to both interactive *visualization* and automated *computational models*, which are most tightly coupled through user interaction (the highlighted region); insight accumulates toward *knowledge*, and *feedback loops* return to refine earlier stages of the analysis. Illustrations adapted from an established visual analytics process model [55]; the stage-goal annotation draws inspiration from a multi-level task framework [120] and a survey of event sequence analysis [45], while the inclusion of interpretation and hypothesizing at the knowledge stage is particular to this thesis.

event sequences; the *computational models* stage concerns characterization, augmentation, and summarization; the *visualization* stage concerns design, comparison, and navigation; and the *knowledge* stage concerns provenance together with the interpretation and hypothesizing that connect observed patterns to higher-level constructs. Interaction occurs throughout the entire loop rather than residing in a single stage. The highlighted region in Fig. 2.2 marks the tight coupling between the *visualization* and *computational models* stages, where interaction is most concentrated, as analysts inspect visual representations, steer automated methods, and iteratively refine both. The feedback loops indicate that interaction also reaches the remaining stages, for instance when data transformations are revised and reapplied in light of what the visualizations reveal, or when accumulated knowledge motivates a return to an earlier stage of the analysis. Interpretation and hypothesizing at the *knowledge* stage extend beyond the scope of existing task typologies. They are included because linking low-level interactions to higher-level constructs is central to this thesis's aims.

The visual analytics process model provides the conceptual spine of the review that follows, and the analytical tasks discussed are situated within its loop.

2.2.3 Event Sequence Visualization

Analytical tasks on event sequence data include identifying recurrent patterns, detecting anomalies, and comparing behavioral trajectories across individuals or groups. These tasks have motivated substantial research on visualization, with recent surveys providing systematic overviews of the field [30, 45, 87, 111].

At a conceptual level, two complementary ways of visualizing event sequences are distinguished, namely **instance-based** and **model-based** representations [111]. *Instance-based* representations depict actual observed sequences as recorded, laying out their constituent events individually in temporal or sequential order so that each occurrence remains visible and inspectable. *Model-based* representations, in contrast, aggregate multiple sequences into abstract specifications conveying a summary of behavior rather than individual instances. The two are complementary: instance-based representations support the inspection of specific traces, whereas model-based representations support the discovery of process logic and aggregated behavioral summaries.

Within these conceptual levels, existing approaches are largely categorized by different **visual encodings**. *Timeline-based* representations remain the most prevalent paradigm, mapping events along a temporal or sequential axis using visual channels such as position, color, labels, and bar length [28, 58, 78, 104, 107]. They are among the most intuitive representations, as they present events chronologically in a directly readable manner. *Matrix-based* designs support fine-grained comparison of co-occurring event pairs across stages [11, 119], and *flow-based* and *Sankey-based* encodings emphasize transition structure over strict temporal positioning [41, 106]. While additional encoding strategies exist, selecting a design approach entails fundamental trade-offs among maintaining chronological order, highlighting transitional structure, and communicating aggregated results. Several examples are illustrated in Fig. 2.3.

The intuitiveness of instance-based, and particularly timeline-based, representations comes at a cost, since long sequences or those spanning a large number of event types quickly become visually cluttered, increasing cognitive load and making patterns difficult to discern. This has motivated a range of **abstraction** and **simplification techniques**. *Aggregation-based* approaches collapse similar sequences into compact summaries [23, 44], while *pattern-mining-based* systems explicitly extract frequent subsequences to guide visual inspection [18, 103]. Complementary efforts simplify visualizations through user-guided substitution and filtering [71, 78], or by applying hierarchical edge bundling [50]. Such transformations, however, risk discarding infrequent yet semantically important events.

Encoding and abstraction thus play complementary roles, with encoding determining which aspects of temporal and transitional structure are made visible, and abstraction governing how much underlying detail is preserved. An alternative to abstraction is to retain detail while managing complexity through interaction, following the principle of an overview first, then zooming and filtering, and details on demand [90]. The appropriate combination of encoding, abstraction, and interaction depends on the structure of the data, the analytical task, and the intended users. This balance is especially consequential when infrequent events carry analytical significance.

2.2.4 Event Sequence Visual Comparison

Visual comparison of event sequences provides the interpretive bridge between algorithmic similarity and human reasoning, rendering the relationships uncovered by similarity measures, examined in the following section, inspectable and actionable. The type of comparison is defined by the granularity of the items being compared, ranging from individual events and subsequences to whole sequences and cohorts [99], while the visual arrangement typically follows one of three principal strategies: **juxtaposition**, **superposition**, and **explicit encoding** [37, 38]. Several examples are illustrated in Fig. 2.4.

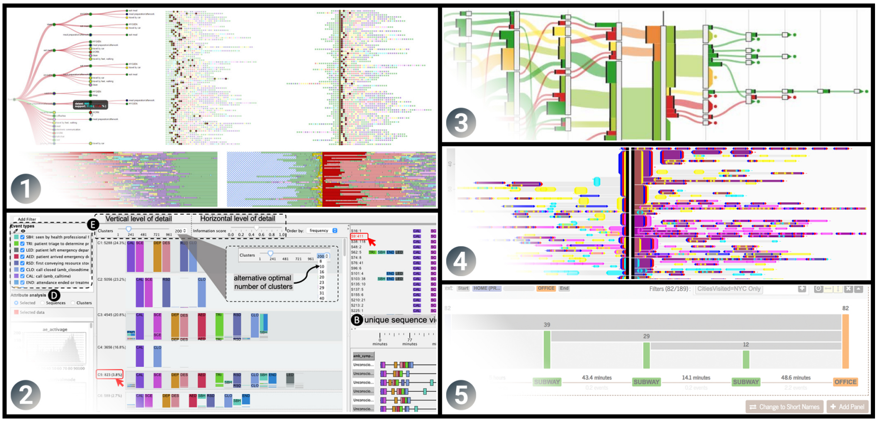

Figure 2.3: Examples of event sequence visualization techniques. (1) ELOQUENCE [103] displays patterns as they are mined and enables steering of the algorithm's execution using a tree representation. (2) SEQUEN-C [67] provides a multi-level overview presenting a variable number of sequence clusters derived from a hierarchical aggregation tree. (3) OUTFLOW [106] visualizes alternative progression paths using a color-coded flow representation mapped to outcomes. (4) EVENTFLOW [78] creates an aggregated view of a dataset by grouping records with identical event sequences. (5) MAQUI [60] applies a hierarchy-based visualization to represent multiple frequent patterns, combined with a timeline to convey temporal information. All images are reprinted with permission. ©IEEE.

Juxtaposition arranges sequences side by side along a shared axis, preserving the context of events, individual sequences, or cohorts. However, this technique relies on the analyst's memory to establish correspondences between separated objects, a demand that grows as the number or length of sequences increases. Its effectiveness therefore depends on careful layout, proximity, ordering, and visual alignment [98]. *Superposition* overlays sequences, pathways, or aggregate patterns within a shared frame, enabling point-wise comparison and reducing eye movement. It is effective for small collections or when aggregate summaries suffice [43, 44], but quickly produces occlusion and overplotting as the number of overlapping elements grows. Its effectiveness therefore depends on prior alignment, filtering, or aggregation, and empirical work indicates that even small differences in alignment can affect perception [117]. *Explicit encoding* represents differences or similarities directly through visual channels such as position, color, size, or shape. A significant limitation is the loss of absolute context, as focusing on relationships or differences may obscure or omit the raw values.

The choice of comparison design depends on the analytical goal, data scale, event alphabet, and the need to preserve individual detail. Each strategy carries certain trade-offs, so the selection and combination should be made with care. Through such comparison, recurring interaction strategies become discernible across users, linking the visual representation of individual sequences examined earlier to the higher-level interpretation of strategies pursued later in this thesis.

Figure 2.4: Examples of strategies for visually comparing event sequences. (1) IDMVis [118] employs a juxtaposition design for diabetes treatment decision support. (2) U4 [80] employs a superposition design to display the sequential patterns of anomalous action sequences. (3) MizBee [73] adopts an additive design, in which an explicit encoding of the relationships between an object and others is superposed on a view of the object itself. All images are reprinted with permission. ©IEEE.

2.3 Algorithmic Methods for Event Sequence Analysis

The techniques discussed so far place the analyst in direct visual contact with the sequences, an approach that is effective when collections are modest in size and well structured. As event sequence data grow, however, from controlled studies toward larger, more heterogeneous, and less structured collections, visual inspection and comparison at this scale become inefficient and, ultimately, insufficient. Meeting this challenge requires moving from the visualization stage to the computational models stage of the visual analytics process, where automated methods operate over the collection as a whole. Three families of such methods recur across event sequence analysis. They are examined in turn: measures that quantify how similar two sequences are, clustering and summarization techniques that group similar sequences to uncover common behavior, and mining approaches that extract patterns and models. Throughout, a central concern is how these methods can be integrated into the analytical loop while remaining interpretable and steerable.

2.3.1 Event Sequence Similarity Measures

At the foundation of much algorithmic event sequence analysis lies the notion of similarity, since quantifying how alike two sequences are is the primitive on which clustering, ranking, retrieval, outlier detection, and strategy identification all depend [99]. A wide range of similarity measures has been developed [69, 96], and these are grouped along several axes, for example, by data type and underlying methodology, including **lock-step** and **elastic** measures [2, 26, 88].

Lock-step measures, of which the Euclidean distance is the canonical example [34], perform a direct point-wise comparison between values at corresponding positions in two sequences. They are computationally efficient, but a limitation is that they are sensitive to temporal drift, speed variation, misalignment, and noise; thus, two sequences describing the same behavior at slightly different rates may be judged dissimilar [85]. More flexible combinatorial approaches relax this restriction, such as the longest common subsequence, which identifies the greatest number of matching elements occurring in the same relative order [7]. In interaction log data, where similar behaviors may unfold at different rates or over different durations, the limitations of such position-dependent measures are particularly pronounced. *Elastic measures* overcome these limitations through non-linear alignment, allowing sequences to be locally stretched or compressed so that similar patterns are exposed despite temporal variation. Dynamic time warping [8] is a prominent technique that computes an optimal alignment path between two sequences and defines their distance as the cost of that alignment [2, 56]. Through such alignment, it is robust to variations in speed, length, or timing [2]. String metrics, such as the Levenshtein edit distance, instead define dissimilarity as the minimum cumulative cost of transforming one sequence into another through insertions, deletions, and substitutions [61]. Edit distances are well-suited to discrete categorical sequences, but they discard temporal information such as event durations and may mix meaningful events with noise unless augmented with event weights or domain-specific costs.

Furthermore, each event in an interaction log typically carries multiple attributes, such as its type, duration, and contextual information, forming multidimensional sequences to which measures defined over a single attribute must be extended or adapted. Additionally, many distance measures cannot be applied directly to categorical event data, since assigning numeric values to categories imposes an ordering and spacing that the data do not inherently possess, and the analytical weight of an event is often domain-dependent, with some event types carrying greater significance than others. Custom distance functions that address these challenges are therefore often required and are pursued in this thesis.

These measures collectively make clear that similarity is not a universal property [99]. Its definition depends on the data type, sequence length, temporal structure, level of abstraction, and analytical goals. Selecting an appropriate measure requires explicit decisions about what constitutes a meaningful difference, whether in event type, order, timing, duration, pace, or noise. Once a measure has been chosen and applied, the pairwise similarities it produces yield a distance matrix over all sequences, which serves as the input to techniques such as clustering, comparison, and summarization.

2.3.2 Event Sequence Clustering and Summarization

While similarity is defined between pairs of sequences, summarization addresses the collection of sequences as a whole, grouping them to reveal common behaviors, anomalies, and progression patterns [2, 45]. Clustering is the technique most commonly applied for this purpose, partitioning a collection so that similar sequences fall together and distinct behaviors become separable. Partitional methods such as k-means are widely used [35], but they require the number of clusters to be specified in advance, which is limiting when little is known about the data.

Hierarchical clustering is better suited to such settings, since it does not presuppose a fixed number of clusters and instead constructs a nested hierarchy of groupings [83]. Two complementary strategies exist for constructing this hierarchy: *agglomerative* approaches, which begin from individual sequences and progressively merge the most similar groups in a bottom-up manner, and *divisive* approaches, which begin from a single group and recursively split it top-down [83]. Hierarchical clustering produces a dendrogram, a tree structure that depicts the relationships between sequences and clusters across levels of granularity. The dendrogram conveys clustering quality and cluster sizes at each level, which supports the identification of outliers and undesirable groupings. In all cases, the collection is ultimately reduced to a compact set of representative structures, but at the cost of individual detail. This reveals a limitation common to both clustering and summarization techniques: the loss of infrequent yet analytically important events without the analyst's awareness.

A more fundamental difficulty is that no single clustering algorithm is universally optimal [32]. The outcome of a clustering depends jointly on the chosen similarity measure, the algorithm's parameter settings, and the granularity at which the hierarchy is cut [56]. The appropriate configuration is therefore not intrinsic to the data, but contingent on the analytical task, and the resulting groupings cannot be assumed to be valid without inspection. This parameter sensitivity, together with the risk of aggregation discarding significant events, means that clustering and summarization are most effective as interactive analytical support whose results an analyst can navigate, validate, and refine, rather than as automated procedures yielding final outputs.

2.3.3 Event Sequence Pattern Extraction and Process Mining

Once sequences have been grouped, the analytical task shifts to describing what each group contains, for which pattern and process mining provide complementary means. Sequential pattern mining extracts frequently occurring subsequences from a cohort, surfacing recurring motifs that characterize common behavior, and it is often coupled with interactive visualization so that the mining process can be steered and its results inspected [45]. Exploratory, user-guided mining of long and complex sequences has been supported through progressive pattern-growth strategies that allow support thresholds to be tuned and candidate matches to be inspected in context [102, 103]. Soft patterns relax the matching process to tolerate noise and intervening gaps [40]. Motif extraction offers a further means of identifying representative behavior, recovering recurring subsequences whose repetition marks them as characteristic of a sequence or a cohort [64, 82]. A persistent difficulty across these techniques is the trade-off between the completeness of the discovered patterns and the cognitive load they impose, since mining a collection exhaustively tends to yield many overlapping or redundant patterns that overwhelm rather than inform the analyst [29, 45]. More fundamentally, pattern mining reveals what recurs, but it does not by itself convey the overall structure of behavior within a cohort, nor does it distinguish meaningful patterns from those arising as artifacts of the mining constraints.

Process mining offers a complementary, model-based perspective, deriving structured models directly from event logs. It consists of three stages: **discovery**, **conformance**, and **enhancement** [1, 86], two of which are particularly relevant to this thesis. *Process discovery* constructs a process model from an event log without prior information, using, for example, a directly-follows graph whose nodes denote activities and whose directed arcs denote precedence relations, yielding a summary of the dominant behavior within a cohort. *Conformance checking* then compares individual sequences against such a model, assigning each a measure of how well it aligns and thereby supporting the detection of outliers and the validation of a cohort's coherence. Process mining and visual analytics developed largely in isolation, with emphasis placed on algorithmic performance and comparatively little support for human-centered reasoning, though their intersection has since attracted growing attention, as interactive visualization has proven essential for examining alternative models, incorporating domain expertise, and building trust in algorithmic outcomes [42, 74]. Applied to clustered interaction data, process discovery offers a model-based characterization of each cluster, while model conformance provides a principled means of assessing the representativeness of that characterization.

In combination, these algorithmic methods make it feasible to analyze collections of event sequences at a scale beyond manual inspection. Yet this feasibility carries a cost that is central to this thesis, since delegating analysis to automated methods introduces a gap in interpretability and trust. The output of a clustering algorithm or a mined model is not self-evidently correct, as a frequent subsequence need not be meaningful, a cluster boundary need not be justified, and a discovered model may fit some sequences poorly. When such results are presented without visibility into how they were produced, the analyst is left unable to judge their validity or to bring domain knowledge to bear. This concern motivates integrating algorithmic methods into interactive, transparent visual analytics workflows, in which extraction and characterization are rendered accessible, interpretable, and verifiable.

2.4 Bridging Event Sequences and Higher-Level Constructs

The *visualization* and *computational models* stages support transforming, exploring, comparing, and characterizing event sequences, but they do not interpret what users did. A broader aim of analyzing interaction data is to understand what those actions imply, since a sequence of recorded interactions is rarely of interest in itself but as evidence of the strategy, intent, or cognitive process that produced it. This shift, from observable low-level events to the higher-level constructs they signify, is central to this thesis.

A higher-level construct is a conceptual category, typically drawn from learning theories, such as Bloom's taxonomy [4, 14] and systems thinking [57, 77], onto which observed behavior is mapped so that it becomes interpretable. Such constructs are not directly recorded in an interaction log; they are inferred from it. The most immediate of these is the notion of an exploration strategy, a purposeful way of moving through a system that recurs across users or tasks. Interaction logs preserve the ordering, timing, and recurrence of actions from which such strategies can be reconstructed. Yet, the mapping from a sequence of clicks to a named strategy is neither automatic nor unique [15]. The same subsequence may be read as different strategies depending on the theoretical lens through which it is interpreted and the context in which it occurs, since surrounding events and strategies shape what a given pattern signifies.

This thesis approaches this gap by treating the mapping from interaction patterns to higher-level constructs as an integrated analytical activity. Discovered exploration strategies are related to the cognitive dimensions of learning frameworks, so that patterns of interaction can be read in terms of the processes they may reflect, and this interpretation is grounded in the judgment of domain experts rather than asserted by the system. How this mapping is constructed depends on the setting in which the analysis takes place.

2.5 Application Contexts

The analysis this thesis addresses requires both visual and algorithmic methods, since visualization reaches its scalability limits as collections grow, while algorithmic methods produce opaque results. A central effort of this thesis is therefore to integrate the two in a human-in-the-loop process [55, 74]. Such integration takes concrete form only within specific settings, each imposing its own data characteristics, analytical questions, and constraints on what an analysis must preserve. Two such domains ground this thesis: a safety-critical professional domain and interactive learning environments.

2.5.1 Event Sequences in Air Traffic Control

In the ATC context, and en-route control in particular, the primary task is to maintain safe separation between aircraft through CD&R. Understanding how controllers build and maintain situational awareness during this task has direct implications for competence assessment, training, and the design of future automation, which motivates the analysis of their working practices [65]. Because a controller's competence is closely tied to visual scanning behavior during CD&R, eye-tracking is commonly combined with interaction logging to capture where attention is directed, in what order, and for how long [72, 105].

The visual analysis of eye-tracking data has been studied extensively, but predominantly under the assumption that the field of view can be divided into a set of static areas of interest, from which sequences are then derived according to how those areas are visited [12, 17, 84]. This assumption has enabled the study of visual scan patterns in settings such as tower control. En-route control, however, is different, since the field of view comprises moving aircraft and dynamic interface elements and is therefore continuously changing. Methods that presuppose static areas of interest are consequently inapplicable, and alternative approaches are required that extract discrete event sequences anchored to specific objects and events of interest rather than to fixed screen regions. A further characteristic of the data from this domain is that they are long, noisy, and highly fragmented, yet contain infrequent events, such as the clearances a controller issues, that are central to the task and must not be discarded during any simplification. This combination of dynamic areas of interest and rare but critical events embedded in noisy sequences defines the analytical demands that the transformation work of this thesis addresses.

2.5.2 Event Sequences in Learning Environments

Interactive touchscreen exhibits have become prevalent in informal learning settings such as museums and science centers, where they are designed to prolong engagement and support understanding and knowledge transfer, often by aligning with educational frameworks that classify learning processes [5]. The design of such environments is frequently framed through the notion of *exploranation*, which combines exploratory and explanatory visualization to scaffold communication and learning [112]. To study how visitors engage with these exhibits, their interactions are logged, producing event sequences that reflect the unique exploration approaches individuals adopt.

Analyzing such logs to understand learning is an established pursuit, and visual analytics and learning analytics techniques have been applied across online, classroom, and informal environments to reveal learning pathways, support reflection and monitoring, relate exploration behavior to outcomes, and inform educational interventions [22, 70, 76]. The unstructured nature of data generated in informal learning environments increases the complexity of this analytical task. Visitors are self-directed; sessions can be brief and highly variable; learning outcomes are typically unlabeled; and multiple visitors may interact concurrently or consecutively. While interaction logs have informed studies of visitor engagement, exhibit use, and redesign [48, 66], existing approaches largely remain at the level of observable patterns, such as frequencies, paths, clusters, or transitions. There remains a significant gap in coupling these observable behaviors with the underlying cognitive constructs.

In learning contexts, interaction is of analytical interest precisely because it is thought to externalize otherwise hidden cognitive processes [15]. A widely used framework is Bloom's taxonomy, which classifies cognitive processes into levels progressing from remembering and understanding, through applying and analyzing, to evaluating and creating, and which was formulated to render learning observable and measurable [4, 14]. Such frameworks supply the vocabulary in which the significance of an interaction pattern is expressed. Existing work, however, provides limited support for making this connection. Much prior work treats discovered patterns as the analytical result rather than as evidence to be interpreted against a theoretical frame, and few approaches couple recorded interactions with the cognitive processes those interactions are intended to reflect. Bridging this gap, therefore, requires more than extracting patterns; it requires methods that render the correspondence between observed interaction and higher-level constructs explicit, inspectable, and open to the domain expert's judgment.

This also intersects with the field of learning analytics [89, 93], which focuses on measuring and analyzing learner data to understand and optimize learning. Rather than treating learning analytics as a separate application domain, this thesis positions it as the target research community to which the interaction-log analysis presented here contributes, since interpreting exploration behavior in relation to cognitive and motivational frameworks is precisely the kind of question that community pursues. Taken together, these challenges underscore the need for visual analytics approaches that render interaction sequences interpretable as evidence of higher-level cognitive constructs.

CHAPTER 3

WAYPOINTS

How does an analysis travel from a log of recorded actions to a claim about how a user reasons? No single stage covers that distance. This chapter lays out the route the four works chart together and marks where each of them stands.

3.1 An Integrated Pipeline

The work of this thesis is organized around the visual analytics process introduced in Sec. 2.2.2, in which the data, computational models, visualization, and knowledge stages are coupled through user interaction (Fig. 2.2). A central premise is that the value of the process lies in the coupling between stages under the analyst's control.

The data stage determines what can be analyzed, since the transformations applied to raw logs determine what any subsequent method can reveal. The computational models extract and characterize structure at a scale that manual inspection cannot reach. The visualizations, in turn, make these results inspectable and support the exploration and comparison of behavior across users, tasks, and levels of granularity. Interaction binds these stages. Observing how the models shape the resulting visualizations allows the analyst to understand what the computational methods are doing, while acting on those visualizations adjusts the parameters or transforms the data on which the models operate. Each loop aims to bring the analysis closer to a result the analyst can interpret.

3.2 Positioning the Contributions

The contributions of this thesis enter the process at different stages, with each of the four works addressing a different aspect of the coupling described above. Taken together, they map onto the pipeline as a methodological progression, summarized in Tab. 3.1.

Table 3.1: Overview of the four works, showing the pipeline stages each primarily addresses. Filled markers denote a primary contribution at a stage; open markers denote a supporting role.

Paper	Data	Visualization	Comp. Models	Knowledge
Paper I	•	◦	◦	◦
Paper II		•		•
Paper III		•	•	◦
Paper IV		•	•	•

Paper I. Paper I operates at the **data** stage, realizing the stage goals of *transformation* and *quality assessment*. It further supports the **computational models** and **visualization** stages, since reduced sequence complexity improves both mining outcomes and the clarity of subsequent representations, and also the **knowledge** stage by preserving analytic provenance. The contribution of this work is a visual analytics approach that supports an interactive, stepwise, and context-aware application of data transformations to improve the quality of complex interaction event sequences collected in safety-critical domains. The approach is examined in the domain of air traffic control, using event sequences derived by merging eye-tracking and simulator data from a human-in-the-loop simulation experiment with 14 air traffic controllers. As a result, the complexity of the sequences can be progressively simplified while subtle, infrequent, yet crucial events are retained, which leads to better downstream mining outcomes.

Paper II. Paper II operates at the **visualization** stage, realizing the stage goals of *design* and *comparison*, and contributes to the **knowledge** stage as the surfaced strategies are *interpreted*. Shifting attention from the preparation of sequences to what the prepared sequences reveal, this work contributes a visual analytics approach that enables multi-granular exploration of users' interaction strategies within an interactive digital environment and couples these low-level interactions with higher-level cognitive learning processes. To this end, both the investigation of individual behavior and multi-user comparison are supported, and a dynamic time warping measure is employed to identify similarities across sequences, thereby surfacing the strategies users adopt. The approach is examined through an experimental study in which 14 science center visitors complete tasks of increasing complexity, designed to elicit different levels of cognitive processing. As a result, a set of task-solving strategies is identified, which reflect different levels of the cognitive processes described in Bloom's taxonomy [4, 14].

Paper III. Paper III operates at the **computational models** and **visualization** stages, realizing the stage goals of *characterization*, *augmentation*, and *summarization* while rendering these results *navigable* and *comparable*, and supports the **knowledge** stage as the characterized clusters are *interpreted* against prior findings. Extending the multi-level perspective to unstructured collections, this work contributes a visual analytics approach that integrates hierarchical clustering and process mining to support analysts in exploring large, unstructured interaction sequence data collected in an uncontrolled digital environment. A tailored similarity measure based on dynamic time warping enables comparison of interaction sequences, and the resulting clusters are characterized through process mining, so that algorithmic groupings can be inspected, validated, and refined with domain knowledge. As a result, five distinct exploration strategies are identified from a combination of in-the-wild and previously collected experimental logs, and validated by relating them to previously established behavioral patterns.

Paper IV. Paper IV operates at the **computational models**, **visualization**, and **knowledge** stages, realizing the computational models goal of *augmentation*, the visualization goals of *design* and *comparison*, and the knowledge goals of *interpretation* and *hypothesizing*. Turning from pattern extraction to the reasoning it may reflect, this work contributes a visual analytics approach that couples motif discovery over multidimensional behavioral series with multi-level visual representations, supporting the exploration of pupils' problem-solving behaviors within interaction logs. A systems thinking framework is adopted as an operational interpretive lens, structuring the analysis across hierarchical levels so that hypotheses linking observable interaction events to higher-level cognitive constructs, such as Analysis, Synthesis, and Implementation, can be formulated. As a result, previously uncharacterized behavior shared across pupils is uncovered, and evidence-based hypotheses about candidate indicators of systems thinking reasoning are formulated, with the underlying behavioral findings confirmed by domain experts against an earlier analysis of the same dataset.

CHAPTER 4

FRAGMENTED TRACE

Event sequence data are fragmented footprints in time,
each step preserved,
the journey yet to take shape.

4.1 Background and Domain

Event sequence data collected from real-world systems are commonly noisy and fragmented, consist of events with irregular durations, and have a large number of event types [45]. Such characteristics make raw event sequences difficult to analyze directly, obscuring meaningful behavioral structures and complicating downstream tasks such as pattern discovery, clustering, sequence comparison, and model construction [31, 36, 54]. At the same time, infrequent events may carry substantial analytical value, particularly when they represent critical user actions or domain-specific behaviors.

Existing approaches address this challenge through manual preprocessing, which incorporates domain knowledge but is often labor-intensive and difficult to reproduce [91], or through automated transformation, which scales well but operates in bulk with limited transparency. As a result, analysts may unintentionally remove important information, have little opportunity to assess the consequences of transformation decisions, and lack support for determining whether data quality has actually improved. This challenge motivates visual analytics approaches that place analysts within the transformation process, supporting it through interactive, transparent, and domain-informed workflows that enable iterative refinement of event sequences while assessing the impact of each operation. Such support requires mechanisms to explore transformation candidates, understand their local and global effects, identify potential information loss, and retain a record of the steps applied so that the process remains transparent and reproducible.

These requirements become concrete in domains that are safety-critical and cognitively demanding, and where interaction data are correspondingly fine-grained. One such domain is en-route air traffic control (ATC), where air traffic controllers (ATCOs) maintain safe separation between aircraft within a controlled airspace sector. This task centers on conflict detection and resolution (CD&R), a cognitively demanding activity in which the ATCO continuously assesses the traffic situation, formulates resolution decisions, and executes appropriate clearances. Understanding how ATCOs perform this task has direct implications for competence assessment, training design, and the development of automation solutions.

To study ATCO work patterns empirically, human-in-the-loop simulation experiments are conducted in which ATCOs perform CD&R scenarios. Eye-tracking captures their visual scanning behavior, operations are logged through the simulator software, and both datasets are merged into a Human-Machine Interaction (HMI) log. Eye-gaze coordinates are discretized into events by computing intersections with graphical objects on the radar screen, including aircraft track symbols, clickable information labels, and conflict detection tool overlays (Fig. 4.1). Intersecting gaze points yield *look* events with durations corre-

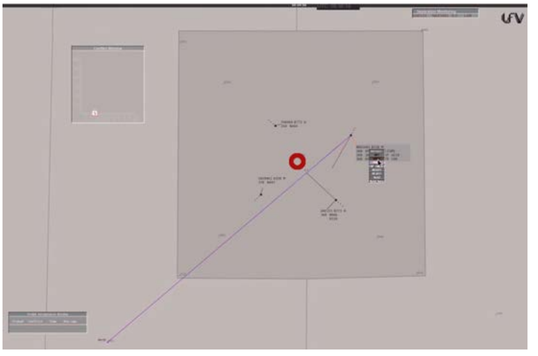

Figure 4.1: The radar screen presents an en-route ATC simulation of a CD&R scenario involving four aircraft. Three auxiliary windows surround the display: a conflict window in the top-left corner, a probe acceptance window in the bottom-left corner, and a separation monitoring window in the top-right corner. A red circle marks the ATCO's eye-gaze point.

sponding to the intersection interval, while gaze points for which no intersection can be determined yield *unknown* events. *Interaction* events are logged at the moment of mouse action, encoding information queries (e.g., label clicks), tool activations (e.g., separation information), and clearance commands (e.g., direction or flight level changes). These events are treated as instantaneous and, despite their rarity, are analytically critical as markers of CD&R decisions.

In the ATC setting, the general characteristics of real-world event sequence data manifest as four interrelated challenges that preclude direct analysis:

- **Noise**: A substantial number of events are *unknown*. Many *look* events shorter than 100 ms add further noise since reliable perception cannot be assumed at such durations [75].
- **Fragmentation**: Short *unknown* events interrupt coherent *look* subsequences, producing patterns such as `look_A` → `unknown` → `look_A`.
- **Alphabet size**: Sequences contain over 100 distinct event types.
- **Variability**: Event durations and frequencies vary substantially, from event types occurring once per sequence to others recurring hundreds of times.

This work therefore investigates how the quality of long, noisy, complex interaction event sequences can be progressively improved through the interactive, context-aware application of transformation operations, while preserving infrequent yet crucial information. The approach targets users such as ATC instructors, researchers, and domain experts seeking to understand controllers' work activities from interaction and eye-tracking logs.

4.2 Methodology

Pattern-based data transformation. A pattern-based transformation framework is introduced to progressively improve the quality of noisy event sequences. Analysts define event patterns of inter-

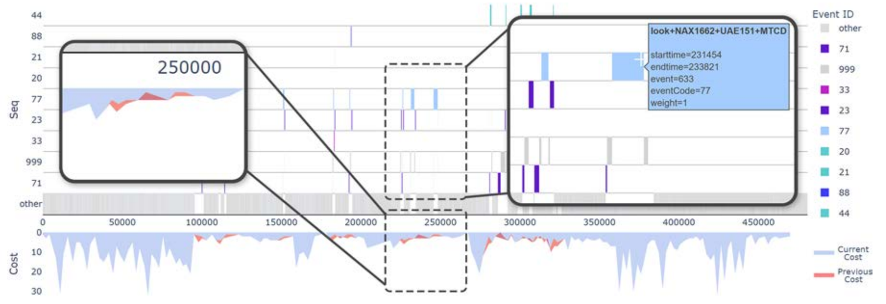

Figure 4.2: Main components of the visual analytics system. The Overview+Detail Panel (top) presents the full event sequence along a timeline, highlighting every occurrence of the currently selected pattern. The Sequence Weight Assessment View (bottom) displays event weights across the sequence, serving as a cue for judging how sequence quality has changed since the preceding transformation step. The image is reprinted with permission. ©IEEE.

est, retrieve all matches, and replace or aggregate the matched subsequences into semantically more meaningful ones. Three complementary means of defining patterns are supported:

- *Manual patterns* are specified explicitly by the analyst based on domain knowledge.
- *Template patterns* use wildcards to apply structural simplifications uniformly across the alphabet, for example removing `unknown` events in all instances of `X` → `unknown` → `X`, where `X` is a known event type.
- *Algorithmically mined patterns* from external pattern mining systems [102, 103] can be imported as candidate transformation targets, surfacing recurring subsequences the analyst may not have anticipated.

Since strict sequential matching misses many meaningful subsequences in fragmented data, the matching process adopts the principle of soft pattern matching [40], permitting several non-pattern events to interrupt a match. This flexibility broadens what can be transformed but makes the consequences of each operation harder to anticipate, motivating the quality cues described next.

Information loss visual cues. To enable analysts to steer transformations while preserving rare but analytically critical events, quality assessment is integrated directly into the transformation workflow at two scales.

A composite event-weighting scheme encodes the analytical importance of each event as a scalar quantity, combining a domain-informed event-type weight with the event duration. In the ATC context, event categories are ranked by significance from *look* events to *clearances*, with *unknown* events weighted at zero, so that the resulting weights reflect analytical significance rather than statistical frequency alone. For *local assessment*, a sliding time window produces a continuous weight profile across the sequence, rendered as overlapping layers reflecting the original, preceding, and current states (Fig. 4.3b). A steep downward displacement of the current layer signals that a recent operation removed high-weight events in that region. For *global assessment*, a measure extending the Levenshtein

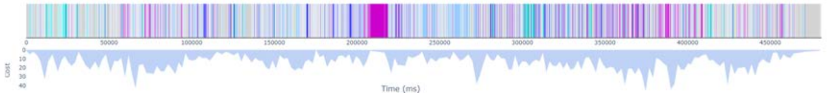

(a) Original sequence: timeline and event weight graph of the sequence before any transformations.

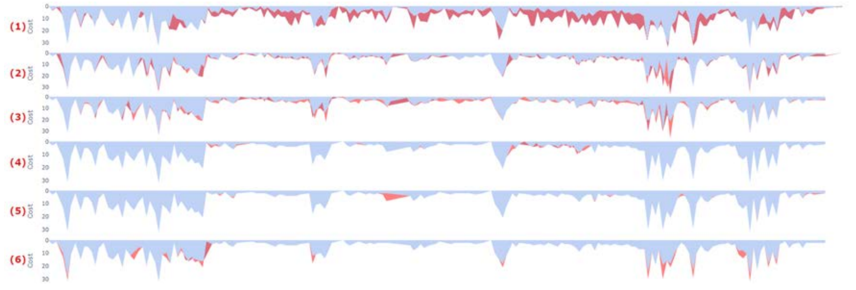

(b) Event weight graphs after each successive operation.

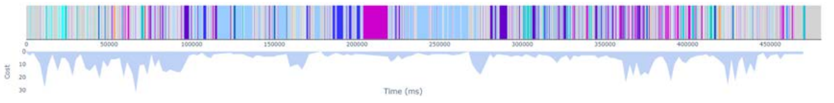

(c) Simplified sequence: timeline and event weight graph after transformations have been applied.

Figure 4.3: Stepwise transformation of a single ATCO sequence in the usage example. The event weight graphs reveal how each operation reshapes the distribution of analytically important events, with peaks corresponding to high-weight activity retained throughout. The image is reprinted with permission. ©IEEE.

distance [46, 61] is computed between the original and current sequences at each step, weighting edit operations by event importance so that removing a critical event incurs a higher penalty than removing a *look* or *unknown* event. Together, these cues provide a multi-scale account of transformation quality, supporting informed decisions at each step of an iterative process.

Analytic provenance and reproducibility. Since such a process unfolds over many steps, an action history logs each transformation to ensure analytic provenance [81, 109], supporting transparency, reversibility, and reproducibility. Analysts can undo operations and revert to earlier sequence states, enabling exploration of alternative transformation strategies. The complete history can be exported as a structured log alongside the transformed data and reused as a transformation blueprint for datasets with similar characteristics.

4.3 Contributions and Conclusions

This work introduces a visual analytics approach to support context-aware transformation and quality assessment of noisy event sequences. The primary contributions are summarized as follows:

- **Pattern-based transformation**: An interactive transformation workflow is presented, in which analysts define, match, and replace or aggregate event subsequences, supporting manual, template-based, and algorithmically mined patterns combined with soft pattern matching to identify and remove noise from the data.

- **Context-aware visual cues**: Visual cues, grounded in a domain-informed event-weighting scheme (Fig. 4.2), are introduced to communicate the localized and global information loss incurred by each transformation, enabling analysts to assess quality at each step.

- **Analytic provenance**: All transformation operations are recorded in an exportable action history, supporting a transparent, reversible, and reproducible transformation process that can be reused as a blueprint for datasets with similar characteristics.

- **Usage example**: The applicability of the approach is demonstrated using HMI logs from 14 ATCOs, yielding sequences with substantially reduced fragmentation while preserving infrequent yet crucial information (Fig. 4.3). Mining sequential patterns before and after transformation showed that the transformed sequences yield coherent and interpretable patterns of controller activity.

Taken together, these contributions address **RQ1**, which concerns transforming noisy event sequence data while preserving analytically significant information. The pattern-based transformation framework operationalizes interactive and domain-informed refinement of raw sequences; the context-aware visual cues make the consequences of each operation assessable at both local and global scales; and the recorded provenance renders the entire process transparent and reproducible. This design allows analysts to progressively improve data quality while retaining rare but analytically critical events, thereby realizing the thesis contribution on *interactive transformation of noisy event sequences and quality assessment of the transformation results* (C1).

While transformation improves the analytical value of noisy event sequences, the question of what those sequences reveal remains open. Once sequences of adequate quality are available, the analytical challenge shifts from preparing the data to exploring and comparing user behavior, from individual events to entire tasks and sessions, and across users.

CHAPTER 5

EXPLORATORY TRACE

Behind every trace lies a strategy,
behind every strategy, a mind
reasoning at a level of its own choosing.

5.1 Background and Domain

Informal learning environments, such as museums and science centers, increasingly leverage interactive visualizations to communicate complex concepts and foster public STEM literacy [9, 16]. At the core of these public exhibits is the concept of *exploranation*, a paradigm blending exploratory and explanatory visualization to scaffold user learning [51, 112]. To structure and evaluate user engagement, designers often draw on educational frameworks such as Bloom's taxonomy, which defines a cognitive hierarchy ranging from lower-order retrieval and understanding to higher-order evaluation and creation [4, 14]. However, the design of interaction techniques for science communication in public contexts remains understudied [10], as is how visitors navigate these environments and construct knowledge. Interaction logs capture detailed traces of visitor behavior, yet mapping such low-level log data to higher-level cognitive learning processes is non-trivial.

To investigate this challenge, a real-world exploranative exhibit was selected: "*Sweden in Numbers*" (SiN), an exhibit at the Norrköping Visualization Center [101]. It visualizes Swedish demographic and societal data via a choropleth map on a large, multi-user interactive touchscreen table (Fig. 5.1). Selecting a geographic region triggers a pop-up information panel (henceforth "infopanel") displaying localized statistical summaries. Within these infopanels, visitors can adjust a temporal *Year* slider or apply categorical filters, such as *Gender*, across four thematic dimensions: *Age*, *Education*, *Income*, and *Population*. The resulting interaction sequences are heterogeneous and multi-granular, varying substantially across individuals and reflecting a wide range of problem-solving strategies and exploration behaviors.

Existing event sequence analysis approaches support visualization, comparison, and summarization through techniques such as timelines, aggregation, clustering, and similarity analysis [45]. However, they typically operate on low-level interaction events, such as screen taps or the opening and closing of interface components, and offer limited support for relating observable behaviors to learning processes.

This work therefore investigates how analysts can be supported in exploring, aligning, and comparing multi-granular interaction sequences and interpreting them in relation to cognitive learning processes. An experiment consisting of five tasks of increasing complexity, aligned with Bloom's taxonomy, was designed and conducted, and the resulting interaction data were analyzed to identify exploration patterns corresponding to distinct cognitive levels. The proposed approach is implemented in a visual analytics system, VISID (**VIS**ualization of **I**nteraction **D**ynamics), and targets users such as educational researchers, exhibit designers, and science center stakeholders. In the remainder of this chapter, individuals interact-

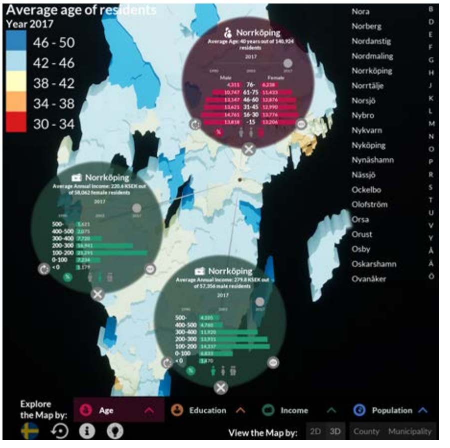

Figure 5.1: Three infopanels opened concurrently for a single municipality in the "Sweden in Numbers" exhibit, showing *Age* statistics (red) alongside *Income* statistics (green) filtered by "Male" and "Female". Such concurrent panels illustrate the interface-level behavior captured in the interaction logs. The image is reprinted with permission. ©IEEE.

ing with the exhibit are referred to as "visitors" or "participants", whereas users of VISID are referred to as "analysts".

5.2 Methodology

Theoretical framework and collaborative design. Bloom's taxonomy [4, 14] is incorporated as the theoretical framework for characterizing cognitive learning processes in an observable and measurable manner. This framework has been widely adopted in educational research and assessment, and its levels correspond to qualitatively different kinds of engagement with information. Five levels are used to design an experiment consisting of five tasks of increasing complexity, each stimulating a specific cognitive dimension, as follows:

1. **Remember**: The least demanding stage, this level focuses on basic factual information retrieval.
2. **Understand**: At this level, engagement involves making simple comparisons, for example, comparing income between genders.
3. **Apply**: This stage requires participants to perform more complex comparisons across multiple data points to reach a general conclusion.
4. **Analyze**: This level involves differentiating between non-explicit elements by organizing data in a suitable order to perform mental calculations and integrate those findings.
5. **Evaluate and Create**: This is the most complex and open-ended dimension, involving testing assumptions and generating hypotheses to identify potential factors associated with broader changes in the data.

This design makes the targeted cognitive processes explicit and provides a basis for interpreting observed interaction strategies. The approach was developed through a year-long collaborative design process involving a multidisciplinary team of visualization researchers, visual learning and communication experts, and an engineer central to the exhibit's development. Across 12 workshops, following a user-centered design process [24], domain insights were distilled into a set of high-level design requirements mapped to analytical tasks. A pilot study provided an initial testbed, revealed logging limitations that bounded the analytical scope, and guided the iterative refinement of the resulting system.

Multi-granular data abstraction. To couple interaction data with cognitive dimensions, the interaction logs are abstracted into three levels of granularity. The *interface level* tracks the lifecycle of high-level components, such as infopanels, indicating whether participants explored them simultaneously or sequentially. The *interaction level* records discrete, active engagement with these components. The *attribute-change level* captures fine-grained contextual updates, such as gender filtering, recording the exact data states examined. Together, these levels reveal both how the interface is used and which data are being explored.

Coordinated multi-level visualization and comparison. The approach is realized in a visual analytics system, VISID, which implements two coordinated views (Fig. 5.2). The INDIVIDUAL VIEW supports detailed examination of a single participant's chronological exploration through overview+detail and details-on-demand interactions. This includes an *Interaction View* that couples interaction event sequences with superposed attribute-change sequences, and a *Holistic View* that visualizes the concurrent lifetimes of open infopanels at the interface level. The COMPARISON VIEW supports cross-participant comparison through a nested juxtaposed layout [38]. Pairwise sequence similarity is computed using dynamic time warping [8]. The resulting scores sort all sequences relative to a selected baseline participant and are further visualized in a *Similarity Plot* as a heatmap and a clustergram, highlighting emerging behavioral groups and outliers. Analysts thus move iteratively between cohort-level similarity structure and individual-level inspection to identify shared exploration strategies and investigate diverging participants.

5.3 Contributions and Conclusions

This work introduces a visual analytics approach for multi-level and multi-granular visual exploration of interaction sequences, coupling low-level interaction patterns with higher-level cognitive constructs derived from Bloom's taxonomy [4, 14]. The primary contributions are summarized as follows:

- **Problem and data description**: The problem and the data underlying the analysis of users' exploration strategies in informal education settings are characterized and distilled into a set of design requirements and analytical tasks that support a data-driven approach.
- **Interactive system**: The proposed approach is implemented in a visual analytics system, VISID, that supports multi-granular visual exploration at individual level and comparison of event sequence data, built through a user-centered design process.
- **Application scenario**: A real-world dataset is collected from an experimental study involving 14 science center visitors, from which task-solving exploration strategies and their connections to cognitive learning processes defined by Bloom's taxonomy [4, 14] are identified.

Figure 5.2: Main components of VISID. The INDIVIDUAL VIEW (top) features (1) *Interaction View* that visualizes a participant's attribute changes for *Gender* and interaction event sequences. Each event is uniformly colored, with the background indicating statistical categories. (2) *Holistic View* displays interface events showing the lifetime duration and concurrently opened infopanels for Task 3. The COMPARISON VIEW (bottom) offers (3) *Sequence Comparison Visualization* sorting interaction sequences by similarity to the baseline participant P7. Participants who are most similar to P7 also adopt a similar strategy. Each event is colored by *Municipality*, with attribute-change sequences for *Year* and *Gender* superposed. (4) The similarity score bars and delta values depict the similarity and dissimilarity with respect to the baseline participant. (5) *Cluster View* shows potential clusters of similar participants. The image is reprinted with permission. ©IEEE.

Taken together, these contributions first address **RQ2**, which concerns the multi-level exploration and comparison of complex event sequences across different levels of granularity. Interaction sequences are abstracted into three levels of granularity, from interface-level events through interaction events to fine-grained attribute changes, and two coordinated views operationalize these levels: the INDIVIDUAL VIEW preserves all three within a single participant's chronological exploration, while the COMPARISON VIEW lifts them to the cohort scale through similarity-based ranking. This design allows analysts to move fluidly between levels, thereby realizing the thesis contribution on *multi-level exploration and comparison of interaction strategies* (C2).

This work also addresses **RQ4**, which concerns the interpretation and mapping of complex interaction patterns in relation to higher-level constructs. Bloom's taxonomy is adopted both to design the experiment's tasks of increasing complexity, each stimulating a specific cognitive dimension, and to serve as an interpretive framework for the observed behavior. By analyzing the resulting interaction sequences from a multi-level and multi-granular perspective in VISID, analysts can trace how distinct interaction patterns, such as "cascading" and "nested-loop", correspond to higher-level cognitive processes. This coupling demonstrates that combining multi-granular representations with theoretical frameworks can reveal strategies hidden in complex event sequence data and support their interpretation in terms of cognitive processes, thereby realizing the thesis contribution on *theory-grounded interpretation of interaction patterns in relation to higher-level cognitive constructs* (C4).

Despite the value of this controlled and theory-driven analysis, the 14 participants knowingly worked through tasks whose goals, difficulty, and expected cognitive demands were predefined. This raises the question of how readily these findings transfer to settings where such structure cannot be assumed. Visitors arrive without assigned tasks and with unknown intents; the number of sequences far exceeds what individual inspection can cover, and whether the identified strategies remain recognizable or entirely new ones emerge is itself an open question. The analytical challenge therefore shifts from analyzing a controlled set of sequences to extracting meaningful patterns from large and unstructured data.

CHAPTER 6

UNSTRUCTURED TRACE

No single eye can follow every path,
so the machine gathers what is alike,
and the mind decides what comes to light.

6.1 Background and Domain

In the third paper of the thesis, the same exploranative environment, "*Sweden in Numbers*" (SiN), from the previous work is investigated. However, rather than controlled experimental data from predefined tasks, this work examines unstructured interaction data collected in the wild over five months, yielding 117 interaction sessions. This transition introduces several analytical challenges. Visitors' goals become unpredictable, their sequences exhibit greater variability in duration and structure, and behavioral patterns are no longer organized around predefined tasks. Moreover, the increased data volume renders detailed inspection of individual sequences impractical. Compounding this, interaction patterns rarely manifest as exact repetitions: visitors may perform similar actions in different orders, spend varying amounts of time on specific activities, and pursue different goals while exhibiting comparable underlying behaviors. Identifying meaningful patterns thus requires methods that accommodate such variability while preserving behaviorally significant similarities.

Algorithmic methods offer mechanisms for processing large collections of sequences, yet each has limitations for interactive analysis. Clustering techniques group similar sequences, but the resulting clusters depend heavily on the selected similarity measures and parameters, particularly regarding the number of clusters and the quality of each cluster [3, 32, 33]. Process mining techniques summarize sequence collections through process representations that reveal common behavioral structures, but the resulting models often become visually complex and difficult to interpret. Furthermore, both families of methods frequently emphasize generating final outputs, leaving analysts limited opportunity to adjust parameters, evaluate intermediate results, and observe how such changes affect the extracted patterns. This motivates visual analytics approaches that embed algorithmic pattern extraction in an interactive, stepwise workflow in which analysts are guided to progressively explore, evaluate, refine, validate, and interpret computationally generated results.

This work therefore investigates how algorithmic workflows can be integrated to support guided, human-in-the-loop characterization and identification of complex interaction patterns. It extends the multi-level exploration perspective to the cohort level and examines potential connections between extracted behavioral patterns and higher-level cognitive constructs. The proposed approach is implemented in a visual analytics system, *re*VISID (***re***visiting the **VIS**ualization of **I**nteraction **D**ynamics), and targets users such as educational researchers, exhibit designers, and science center stakeholders. As in the previous chapter, individuals interacting with the exhibit are referred to as "visitors" or "participants", whereas users of *re*VISID are referred to as "analysts".

6.2 Methodology

Data abstraction and sequence comparison. To emphasize navigation and interaction strategies over the specific content explored, the interaction logs are abstracted into four high-level event variables, identified through prior collaborations with domain experts: *Open* infopanel, *Revisit* infopanel, *Year Change*, and *Gender Change*. Based on these variables, pairwise sequence similarity is computed. Since similar behavioral strategies often manifest as sequences with varying lengths, execution speeds, and temporal structures, a tailored elastic alignment measure based on multidimensional dynamic time warping [92] is employed, adapted to the discrete and categorical nature of interaction events. The computation yields a pairwise distance matrix that serves as the foundation for subsequent clustering.

Hierarchical clustering and representation. Agglomerative hierarchical clustering [47] is applied to the distance matrix because it does not require a predefined number of clusters, which is suitable when the underlying behavioral profiles are unknown. The resulting hierarchy is treated as an exploratory space through which analysts progressively investigate behavioral groups at different levels of abstraction. Navigating this hierarchy requires a suitable representation. Since a full dendrogram and heatmap become difficult to interpret at scale, a compact partial dendrogram is proposed that can be explored at different levels of the hierarchy, with triangular glyphs summarizing the size and linkage structure of each cluster, providing visual cues for deciding where to subdivide further and investigate clusters. A complementary matrix view displays computed measures of intra-cluster conciseness and inter-cluster dissimilarity, enabling analysts to evaluate and refine clusters across hierarchy levels (Fig. 6.1.1).

Process mining integration. Even once clusters have been evaluated and refined, their grouping alone does not explain what the sequences within them have in common. Process mining is therefore used to characterize the behavioral structures within each cluster. A process model is generated per cluster and rendered as an arc diagram with fixed node ordering, enabling compact visual comparison of behavioral structures across clusters (Fig. 6.1.2). Conformance checking connects these cluster-level abstractions back to individual sequences. A fitness score quantifies how well each sequence conforms to its cluster's process model, and the resulting ranking allows analysts to identify representative sequences and outliers (Fig. 6.1.4). Finally, an interactive cluster adjustment supports expert-driven refinement. By comparing process models, analysts can identify clusters that exhibit the same exploration strategy yet are separated by the clustering algorithm, validate the grouping through conformance scores, and merge them based on domain knowledge (Fig. 6.1.3). This establishes an iterative refinement loop in which clusters are repeatedly inspected, restructured, and reassessed until a set of meaningful, interpretable, and domain-relevant interaction strategies is obtained.

6.3 Contributions and Conclusions

This work introduces a visual analytics approach to extract and characterize interaction patterns from large, unstructured interaction logs by integrating hierarchical clustering and process mining into a guided, human-in-the-loop workflow. The primary contributions are summarized as follows:

- **Stepwise, guided clustering exploration**: A workflow tailored for the stepwise, interactive exploration of hierarchical clustering structures is presented, allowing analysts to navigate complex hierarchies of interaction data.
- **Integration of process mining**: Process mining techniques are integrated to characterize derived clusters, extracting process models visualized as arc diagrams that summarize the behavior within

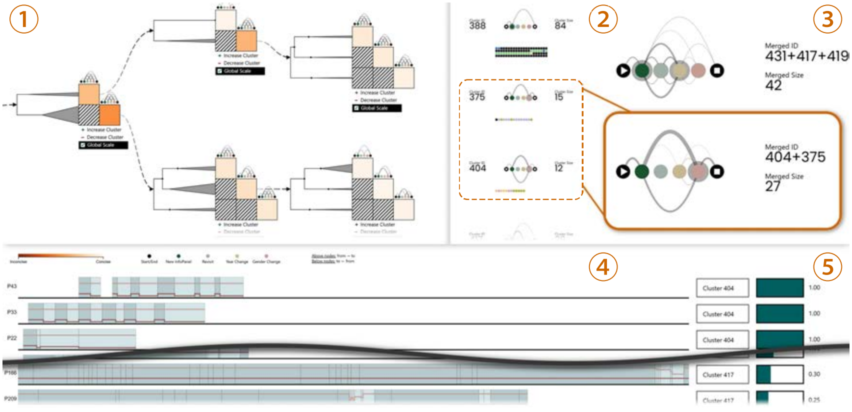

Figure 6.1: Main components of *re*VISID. The CLUSTER AND PROCESS VIEW features (1) *Cluster Pane*, which summarizes the top levels of the hierarchical clustering structure into a partial dendrogram with triangular glyphs alongside a lower-triangular similarity matrix. (2) *Process Description Pane* displays cluster descriptions derived through process mining and facilitates cluster description comparison. (3) Through the *Process Merging Pane*, analysts may manually merge process models based on expert-guided assessments. (4) SEQUENCE VIEW presents individual sequences corresponding to selected clusters, supporting detailed behavioral analysis. (5) Analysts can evaluate how well each individual sequence conforms to a selected process model based on conformance scores. The image is reprinted with permission. ©IEEE.

each cluster, and employing conformance checking to evaluate how well individual sequences align with those models.

- **Tailored similarity measure**: A customized dynamic time warping similarity measure is developed for discrete interaction event sequences, accounting for the asynchronous nature of complex user interactions in unstructured, in-the-wild logs.
- **Interactive system**: The proposed approach is implemented in a visual analytics system, *re*VISID, supporting the identification and extraction of exploration strategies through an iterative process of cluster refinement, model comparison, and inspection of individual sequences (Fig. 6.1).
- **Case studies**: The applicability of the approach is demonstrated through two case studies: first, analyzing the in-the-wild logs to inform exhibit design, and then combining them with interaction data from the controlled experiment of the previous work to identify five distinct exploration strategies validated against previously established behavioral patterns.

Taken together, these contributions first address **RQ3**, which investigates how algorithmic workflows support guided, human-in-the-loop characterization and identification of complex interaction patterns. The proposed approach reveals intermediate analytical results throughout the pipeline and supports their inspection, validation, and iterative refinement through interactive visualization. By combining hierarchical clustering with process-based cluster characterization, conformance analysis, and expert-driven cluster merging, the workflow enables analysts to validate computational results and incorporate domain knowledge throughout the analytical process. This directly realizes the thesis contribution on *user-driven extraction and characterization of interaction patterns* (C3).

This work also addresses **RQ2**, which concerns the multi-level exploration and comparison of complex event sequences across different levels of granularity. The clustering hierarchy supports analysis at the cohort level, where analysts inspect, subdivide, and merge behavioral groups; the process models provide an intermediate abstraction that characterizes the behavior shared within each cluster; and the individual sequences within each cluster, ordered by conformance scores, allow analysts to distinguish representative behaviors from deviating or outlier sequences. This design allows analysts to move fluidly between summarized behavioral patterns and their underlying interaction sequences, thereby contributing to *multi-level exploration and comparison of interaction strategies* (C2).

Finally, this work addresses **RQ4** by supporting the interpretation and mapping of complex interaction patterns in relation to higher-level constructs. The extracted strategies were interpreted through the lenses of Bloom's taxonomy [4, 14] and a three-factor situational interest model [63], relating minimal engagement to lower-order cognitive processing and sustained sessions with repeated attribute changes to deeper engagement. This illustrates how extracted interaction patterns can be mapped to meaningful domain concepts, thereby contributing to *theory-grounded interpretation of interaction patterns in relation to higher-level cognitive constructs* (C4).

The extracted strategies describe how visitors moved through the exhibit, yet they leave open how those visitors actually reason about what they explore, including interconnections, feedback, and emergence. The analytical challenge therefore shifts from extracting and characterizing patterns to interpreting what those patterns imply about how users reason, which calls for grounding the analysis in an explicit theoretical framework of cognition.

CHAPTER 7

SYSTEMIC TRACE

The trace was never the destination,
only the surface on which
reasoning becomes visible.

7.1 Background and Domain

In STEM education, understanding the carbon cycle is essential for scientifically informed reasoning about climate issues, and comprehending its systemic nature depends on systems thinking (ST), a fundamental element of learning and teaching science [49] that comprises skills for interpreting complex, dynamic systems [100]. Development of ST is described as a three-level hierarchy of increasingly complex cognitive stages [6]:

- *Analysis*: Pupils identify a system's key elements and processes.
- *Synthesis*: Pupils connect these elements and recognize their dynamic, cyclical interactions.
- *Implementation*: Pupils interpret hidden system features, generalize about system functioning, and apply temporal reasoning to past events and future outcomes.

Complementing this hierarchy, a related conceptualization, "the iceberg" [57, 77], describes visible *Events* as arising from deeper *Patterns* and *Systemic Structures* built on underlying *Mental Models*, stressing that ST requires perceiving interrelationships, non-linearity, feedback, and emergence.

Reasoning of this kind is difficult to observe directly: although modern educational software captures granular digital traces of pupil activity, these traces are often analyzed as isolated, discrete actions rather than as evidence of conceptual reasoning. Since ST reasoning is non-linear, emergent, and feedback-driven, it cannot be directly parsed from such data; mapping fine-grained interaction sequences to theoretically grounded ST constructs thus constitutes a central challenge for the analysis of educational interaction data.

Tracing Carbon (TC) [68] is a web-based learning environment for pupils in grades 7 to 9 in the Swedish educational context and provides a suitable testbed for addressing this challenge (Fig. 7.1). TC features interactive tasks centered on the carbon cycle, including a relational arrow-drawing task in which pupils explicitly map carbon transfers between reservoirs, externalizing how they structurally organize relationships within a complex system. Since TC was designed around the ST hierarchy, the environment aligns with the analytical lens applied in this work.

This work therefore investigates how low-level interaction sequences and higher-level cognitive hierarchies can be bridged, supporting analysts in exploring task-solving strategies and generating hypotheses about the relationship between observable behavior and ST reasoning. The proposed approach is implemented in a visual analytics system, BEHAVE (**B**ehavioral **E**xploration and **H**ierarchical **A**nalytics

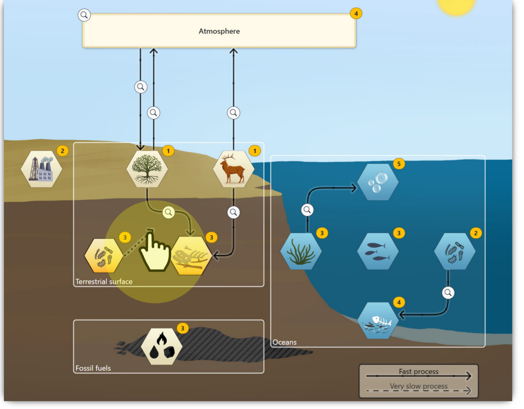

Figure 7.1: A screenshot of the TC digital learning environment. The interface displays 12 carbon reservoirs distributed across the atmosphere, terrestrial surface, oceans, and fossil fuels. Pupils work by drawing arrows between reservoirs to denote carbon transfer processes such as photosynthesis, respiration, and decomposition, where directional links convey both the flow and its relative speed (fast versus very slow). The filled yellow circle highlights an active interaction, indicating the pupil's cursor while a link is being drawn between components. Logging these interactions captures how pupils construct relationships within the carbon cycle and how they approach the problem. The image is reprinted with permission. ©IEEE.

in **V**isual **E**nvironments), and targets users such as educational researchers and learning scientists. In the remainder of this chapter, individuals interacting with the learning environment are referred to as "pupils", whereas users of BEHAVE are referred to as "analysts".

7.2 Methodology

Behavioral descriptor abstraction. In this work, ST serves as both a conceptual framework and an interpretive lens for analyzing pupils' interaction logs from TC. To support interpretation in relation to ST, discrete interaction sequences are first mapped to a set of ST-related events, or behavioral building blocks, derived in collaboration with domain experts:

- **Tracing**: Connecting reservoirs sequentially, with each arrow starting where the previous one ended.
- **Mini-cycles**: Creating a cycle between two reservoirs.
- **Same-start**: Drawing repeated arrows from a single reservoir to several others.
- **Same-end**: Drawing repeated arrows from multiple reservoirs to a single one.
- **Within-category**: Connecting two reservoirs that belong to the same category.

A temporal measure of **sparsity** complements these events by capturing how active a pupil is over a given period, and a Boolean **correctness** measure indicates whether a connection is correct.

To capture how these behaviors evolve over time, the sequences are transformed into multidimensional time series by computing normalized signals for each descriptor over a sliding window. The resulting representation reveals how a pupil's problem-solving strategies shift throughout a task.

Motif discovery. Building on these series, a motif discovery pipeline, adapted from consensus and multidimensional matrix profile techniques [53, 110], identifies recurring behavioral patterns shared across pupils. Each identified behavioral subsequence is compared across all sequences, and candidate patterns are ranked by the number of pupils who exhibit them, prioritizing patterns shared across the group over those recurring within a single pupil's sequence. The matches are further consolidated, deduplicated, and filtered to yield a compact, non-redundant set of consensus motifs. These motifs constitute a library of candidate behavioral patterns that potentially reflect distinct ST-related strategies or conceptual difficulties for analysts to inspect and validate.

ST-aligned multi-level visualization. The discovered motifs, behavioral abstractions, and individual sequences are explored through three coordinated panels in BEHAVE, each aligned with one of the three ST-informed levels of events, recurring patterns, and underlying structures (Fig. 7.2).

The Cluster Panel supports identifying cohorts with similar behavioral profiles by k-prototypes clustering [52] over a flexible combination of survey variables, behavioral descriptors, and sequence-level statistics, enabling cohort construction from behavioral, performance-related, or survey-based perspectives. The Motif and Comparison Panel supports exploration of recurring temporal patterns through its central *Motif Library* and a compact sequence comparison visualization, enabling analysts to examine where, when, and how frequently motifs occur within sequences and across cohorts. Finally, the Individual Panel provides fine-grained inspection of a selected sequence through a Marey-chart-based representation with two complementary modes: *sequential mode* spaces interactions uniformly to emphasize their order, while *durational mode* scales them to real time to reveal pacing and idle intervals. The panel further displays the derived multidimensional time series and the ST-related event sequence, relating abstract behavioral tendencies back to the original interaction trajectory.

Through these panels, cluster-driven and motif-driven workflows converge on detailed individual inspection, allowing analysts to move between abstract summaries and behavioral evidence and to generate hypotheses about the relationship between observable behaviors and ST reasoning.

Case study and expert validation. The approach was evaluated through a case study and a structured expert validation. The evaluation adopted an insight-oriented, qualitative perspective [20, 59], focused on whether the approach supports experts in exploring interaction patterns and generating hypotheses about pupils' ST reasoning.

The case study yielded several insights. Four cohorts with distinct strategy profiles and differing correctness outcomes were identified, indicating that systematic exploration strategies are not equally productive. Previously uncharacterized trial-and-error behavior was uncovered and shown to be shared across multiple pupils, alongside frequent shifts between strategies during problem solving. These observations further drove hypothesis generation regarding ST reasoning levels, for example, linking low use of *within-category* connections to difficulties in identifying system boundaries (Synthesis level), and cross-system replication of a completed food chain to generalization skills (Implementation level).

The expert validation was grounded in triangulation with a prior empirical analysis of the same dataset, conducted through workshop sessions that combined open-ended exploration with think-aloud

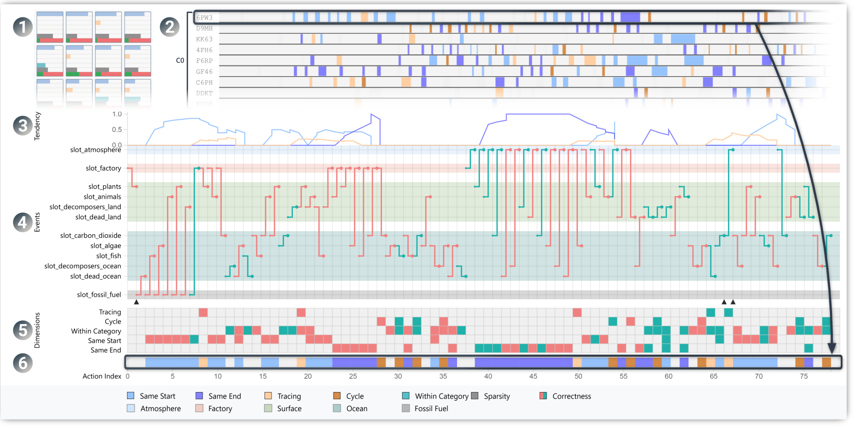

Figure 7.2: Main components of BEHAVE. The Motif and Comparison Panel features (1) the *Motif Library*, containing all discovered motifs, and (2) a compact visualization of mutually exclusive ST-related events across all interaction sequences, shown in *durational mode* to support cross-sequence comparison. The Individual Panel presents details for one selected pupil. It features (3) multidimensional time series that summarize behavioral tendencies; (4) a Marey-chart-based representation of interaction events; (5) a representation of the behavioral building blocks; and (6) a compact visualization of the ST-related event sequence in *sequential mode*. The topmost sequence in the Motif and Comparison Panel is the same one shown in the Individual Panel, as marked by the annotated arrow on the right. The image is reprinted with permission. ©IEEE.

commentary. The experts confirmed that BEHAVE reproduced the four previously identified task-solving strategies and further judged that it extended prior work by revealing how these behaviors are organized in both order and time, showing how strategies combine, persist, and shift within individual sequences. The validated contribution lies in making complex interaction patterns systematically visible and interpretable, thereby enabling the formulation and refinement of evidence-based hypotheses about pupils' ST reasoning.

7.3 Contributions and Conclusions

This work introduces a theory-driven visual analytics approach to bridge the gap between low-level interaction logs and higher-level cognitive processes. The primary contributions are summarized as follows:

- **ST-grounded visual analytics approach**: A visual analytics approach is presented that organizes the analysis of pupil interaction logs along three ST-informed levels: events, recurring patterns, and underlying structures, employing ST both as a design lens and an interpretive framework.
- **Tailored motif discovery**: Discrete interaction sequences are converted into multidimensional time series, from which consensus motifs shared across pupils are detected and serve as candidates for recurring behavioral patterns.

- **Interactive system**: The proposed approach is implemented in a visual analytics system, BEHAVE (Fig. 7.2), containing two complementary analytical workflows, one driven by clusters and one by motifs, both converging on detailed inspection of individual sequences and facilitating hypothesis-driven exploration of how interactions relate to ST reasoning.

- **Case study**: The usefulness and applicability of the proposed approach are demonstrated through pupil interaction data collected in the TC environment, uncovering previously unseen behavioral patterns and generating new hypotheses concerning the relationship between problem-solving strategies and ST reasoning. Domain experts confirmed the resulting behavioral findings by triangulating them against an earlier analysis of the same dataset.

Taken together, these contributions first address **RQ2**, which concerns the multi-level exploration and comparison of complex event sequences across different levels of granularity. Three coordinated panels operationalize distinct analytical scales: the Cluster Panel supports cohort-level characterization and comparison; the Motif and Comparison Panel, through its *Motif Library*, enables cross-sequence inspection of recurring patterns and comparison across pupils; and the Individual Panel preserves the temporal structure of event sequences through a Marey-chart-based representation in complementary sequential and durational modes. This design allows analysts to move fluidly from cohort-level tendencies through shared motifs to fine-grained individual behavior, thereby realizing the thesis contribution on *multi-level exploration and comparison of interaction strategies* (C2).

This work also addresses **RQ3**, which investigates how algorithmic workflows support guided, human-in-the-loop characterization and identification of complex interaction patterns. The motif discovery pipeline surfaces recurring behavioral patterns algorithmically, while interactive clustering, motif exploration, and coordinated views keep computational outcomes open to interpretation and validation through domain expertise. This directly realizes the thesis contribution on *user-driven extraction and characterization of interaction patterns* (C3).

Finally, this work addresses **RQ4** by supporting the interpretation and mapping of complex interaction patterns in relation to higher-level constructs. The ST hierarchy is adopted as both a design lens and an interpretive framework, establishing an explicit bridge between low-level interaction events and higher-level reasoning constructs. Discrete interaction sequences are abstracted into domain-informed behavioral building blocks that situate observable actions within the ST reasoning hierarchy. Consensus motif discovery, combined with multi-level visual representations, then supports the formulation of evidence-based hypotheses about candidate behavioral indicators of ST reasoning levels, such as Analysis, Synthesis, and Implementation. This directly realizes the thesis contribution on *theory-grounded interpretation of interaction patterns in relation to higher-level cognitive constructs* (C4).

Together with the preceding publications, this work completes the progression from preparing event sequence data toward interpreting the reasoning behind them. The remainder of the thesis therefore turns from the four papers considered in isolation to what they jointly establish.

CHAPTER 8

SYNTHESIS

This thesis has investigated how visual analytics can support the transformation, exploration, comparison, extraction, and interpretation of complex interaction event sequence data. This chapter revisits the four research questions formulated in Chapter 1 and considers how the included papers jointly answer them. The correspondence between papers, research questions, and contributions is summarized in Fig. 8.1.

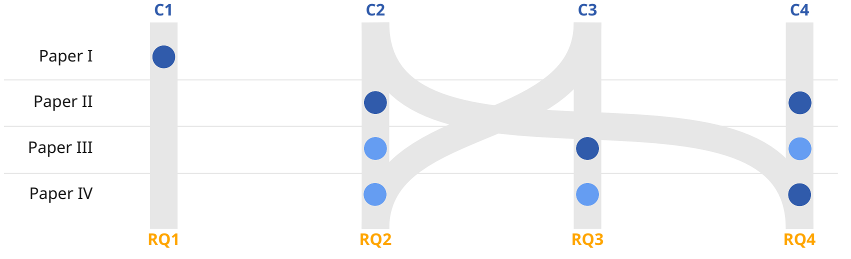

Figure 8.1: Correspondence between the four included papers, the four thesis contributions, and the four thesis research questions. Each paper realizes one or more contributions, and each contribution answers one or more research questions. A dark blue circle denotes a primary contribution, and a light blue circle denotes a supporting role. Paper I contributes to C1, which answers RQ1. Papers II to IV contribute to C2, C3, and C4, with several papers jointly realizing the same contribution and several contributions addressing the same research question.

The first research question concerns how visual analytics approaches can support the transformation of noisy event sequence data while preserving analytically significant information. This question is addressed by Paper I, which realizes contribution C1. Paper I treats transformation as an interactive, context-aware process rather than a purely automated preprocessing step. Pattern-based transformations, specified manually, through templates, or from algorithmically mined sequential patterns, progressively reduce the complexity of raw sequences; visual cues convey the potential information loss of a transformation; and an exportable action history and undo functionality keep the process transparent and reproducible, supporting analytic provenance.

The second research question concerns how visual analytics methods can support the multi-level exploration and comparison of complex event sequences across different levels of granularity. This question is primarily addressed by Papers II and IV, and further supported by Paper III; together, these papers realize contribution C2. Paper II supports multi-granular investigation across the interface, interaction, and attribute-change levels, together with individual inspection and comparison across participants. Paper III scales this capability to in-the-wild data by supporting guided exploration of a navigable cluster hierarchy while retaining access to individual sequences. Paper IV develops cohort-level, pattern-level, and individual-level representations, with a Marey-chart-based visualization revealing the pacing and structure of interactions in sequential and durational modes.

The third research question concerns how algorithmic workflows can support guided, human-in-the-loop characterization and identification of complex interaction patterns. This question is primarily addressed by Paper III and complemented by Paper IV; together, these papers realize contribution C3. Paper III couples a tailored dynamic time warping similarity measure, hierarchical clustering, process mining, and conformance checking with visual summaries that allow the analyst to inspect, steer, and refine each intermediate result. Paper IV transforms discrete interaction sequences into multidimensional time series and discovers cross-pupil consensus motifs, which serve as candidate recurring patterns for subsequent inspection and comparison rather than as final results.

The fourth research question concerns how visual analytics systems can support the interpretation and mapping of complex interaction patterns in relation to higher-level constructs. This question is primarily addressed by Papers II and IV, and further supported by Paper III; together, these papers realize contribution C4. Paper II adopts Bloom's taxonomy to structure tasks of increasing complexity and to interpret the resulting exploration strategies, while Paper IV employs systems thinking as both a design lens and an interpretive framework, allowing analysts to hypothesize about observable behavior in relation to higher-level constructs. Paper III supports the contribution by showing that such interpretation remains attainable under uncontrolled conditions.

Together, these answers trace a cumulative path from preparing raw event sequences to interpreting the behavior they record, laying the groundwork for a closing account of the thesis contributions, limitations, and future directions.

CHAPTER 9

HORIZONS

What remains once the traces have been transformed, explored, compared, extracted, and interpreted? This chapter gathers the contributions of the four works, draws the conclusions they jointly support, and marks the boundaries at which those conclusions end and future research begins.

9.1 Contributions

This thesis makes four primary contributions to the field of visual analytics for event sequence analysis.

C1. **Interactive transformation of noisy event sequences and quality assessment of the transformation results.** The thesis contributes a visual analytics approach that enables analysts to progressively transform noisy, complex event sequence data to improve their quality while preserving analytically significant information.

C2. **Multi-level exploration and comparison of interaction strategies.** The thesis contributes coordinated visualization methods that support the exploration and comparison of interaction sequences across multiple levels of granularity, enabling analysts to investigate individual behaviors, compare them, and identify recurring interaction strategies.

C3. **User-driven extraction and characterization of interaction patterns.** The thesis contributes interactive analytical workflows that integrate clustering, similarity analysis, process mining, and motif discovery with interactive visualization to guide analysts in stepwise extraction, characterization, and validation of behavioral patterns from unstructured, complex interaction sequences.

C4. **Theory-grounded interpretation of interaction patterns in relation to higher-level cognitive constructs.** The thesis contributes visual analytics designs that couple domain-informed behavioral abstractions with established theoretical frameworks, enabling analysts to connect low-level interaction events to higher-level cognitive constructs, such as reasoning processes, learning strategies, and systems thinking, and to formulate evidence-based hypotheses about them.

Taken together, these four contributions are closely interdependent and distributed across the stages of the visual analytics process model. Each addresses a stage on which the next depends: transformation bounds what any subsequent method can reveal (C1); multi-level representation determines which behaviors become visible (C2); algorithmic support extends this reach to collections beyond manual inspection (C3); and theoretical grounding determines what the resulting patterns can be taken to mean (C4). Across all four, the unifying principle is the tight coupling of computational methods, interactive visualization, and domain expertise.

9.2 Conclusions

This thesis has been concerned with the visual analysis of interaction event sequence data, the detailed logs that interactive systems record as users interact, explore, and reason with them. Such logs are routinely collected across a wide range of digital environments. Yet, their analysis has often relied on aggregate statistics, isolated metrics, or fully automated methods that conceal the sequential, temporal, and individual character of the underlying behavior. Across three application settings, namely en-route air traffic control, interactive science center exhibits, and digital learning environments, this thesis has developed visual analytics approaches that together span the entire analytical process (Fig. 2.2). The conclusions drawn below trace this same progression, beginning with the nature of the data itself.

Interaction logs should be regarded as rich but incomplete behavioral evidence. They capture detailed traces of what users did in an interactive system, such as the ordering, timing, recurrence, and duration of actions. In their raw, exported form, however, such logs can rarely be analyzed to explain why users acted a certain way, what strategies they employed, or what they understood while interacting with a system. Meaningful analysis therefore requires methods that preserve the temporal and contextual structure of interaction data while enabling analysts to relate observable actions to domain-relevant interpretations. This requirement shapes every stage of the analytical process, beginning with the preprocessing of the data itself.

Data preprocessing must be understood as a central analytical activity rather than a separate step. The findings of this thesis show that event sequence transformation can strongly affect which patterns become visible in subsequent analysis. Removing noise, merging events, or simplifying repeated interactions may improve the readability and comparability of sequences, yet such operations also risk discarding subtle or infrequent events that carry importance in a particular domain. Transformation should therefore be transparent, reversible, and open to expert assessment. Interactive transformation, visual cues about data quality, and analytic provenance are essential, as they enable analysts to understand not only the resulting data but also *how* and *why* they were transformed. However, improved data quality is only a precondition; the analytical challenge that follows concerns the behavior the sequences encode.

Understanding that behavior, this thesis concludes, requires exploration and analysis at multiple levels of granularity. Behavioral patterns manifest at different scales, as individual actions, local subsequences, complete trajectories, recurring motifs, or differences between user groups, and no single representation is sufficient for all these perspectives. This thesis therefore demonstrates the value of coordinated views that support movement between abstraction levels, between individual and cohort comparison, and between computational summaries and the underlying event sequences. Such coordination helps analysts connect high-level patterns with the concrete interactions from which they emerge. As collections of logs grow larger, more heterogeneous, and more unstructured, however, even well-designed visualizations reach their limits, and algorithmic support becomes necessary.

Computational methods can reveal structure in such collections that would otherwise be difficult to identify through purely visual approaches. Sequence similarity measures, clustering, process mining, motif discovery, and conformance checking can reduce the complexity of large collections of logs and direct attention toward prominent, reproducible, and potentially meaningful patterns. This thesis also shows, however, that algorithmic methods do not produce definitive behavioral explanations on their own. The computational results are often sensitive to parameter choices, and equally reasonable settings can yield different patterns; analysts must therefore be able to explore alternative configurations and tune parameters interactively rather than accept the computed outcome as final. Moreover, when only the final output is presented, the algorithm acts as a black box: the analyst can see what was extracted but not *how* or *why*, leaving no basis for judging whether a cluster, process model, or motif reflects genuine behavior or an artifact of the method. The approaches in this thesis therefore make the extraction process itself transparent and adjustable, enabling analysts to understand, steer, and act upon it. Computational outputs thus become entry points for visual exploration and expert-driven sense-making rather than final analytical results, a sense-making that extends beyond the data themselves toward higher-level constructs that interaction logs do not record directly.

This ability to connect interaction patterns to higher-level constructs becomes particularly important in learning contexts. In these settings, educators and researchers seek to understand whether users reason systematically, engage with intent, or develop understanding. Yet, such cognitive processes leave no direct traces in interaction logs. The analysis of interaction logs from interactive exhibits and digital learning environments demonstrates that visual analytics can help experts investigate possible relationships between interaction patterns and cognitive or educational frameworks. Theories, such as Bloom's taxonomy and systems thinking, serve as interpretive frameworks that guide exploration and hypothesis formation and, in some cases, also as design lenses shaping behavioral abstractions and the structure of visual analytics systems. At the same time, the thesis emphasizes that visual analytics systems cannot directly infer users' intentions, knowledge, or cognitive abilities; instead, they support the formulation of evidence-based hypotheses that can subsequently be examined with additional educational, qualitative, or experimental evidence.

In sum, these conclusions point to the broader contribution of visual analytics in this area: bridging computational pattern extraction and contextual human interpretation. Across these domains, the proposed approaches make interaction data accessible to analysts who need to understand not only what users did, but also how they did it and what their actions may imply. By integrating data transformation, pattern discovery, interactive visualization, and domain expertise, this thesis demonstrates a pathway from raw event sequences to interpretable behavioral insights.

9.3 Limitations and Future Work

The contributions of this thesis were developed and validated under specific empirical, methodological, and epistemic conditions, and the answers to the four research questions should be understood within these boundaries. At the same time, each boundary charts a concrete direction for continued research.

Domain dependency and broadening the empirical base. All approaches were developed and validated on datasets from specific application contexts, whose scale and diversity bound the empirical claims. The thesis demonstrates the applicability and analytical value of the proposed approaches, while statistical generalization of the discovered behaviors and patterns lies outside its scope. The reported results are dependent on domain-specific choices of transformations, similarity measures, behavioral descriptors, and visual encodings. Future work could apply the approaches to larger and more diverse cohorts, tasks, and domains, distinguishing findings specific to a particular environment from general properties of interaction behavior and documenting how workflow components must be adapted to new settings. Such transfer, however, raises a more fundamental question of what the logs themselves can capture.

Missing context and complementary data sources. Interaction logs capture only actions and lack context such as attention, prior knowledge, and intent; distinct cognitive processes, however, can leave indistinguishable traces. Complementary modalities were deliberately excluded, since the contributions target logs generated during ordinary use and additional instrumentation would have altered the study settings and raised further ethical considerations across the participant groups involved. The thesis thus analyzes interactions as a proxy for behavior, and behavior, in turn, is considered a proxy for cognition. Future work could complement logs with modalities, such as screen recordings, think-aloud protocols, and task outcomes, presented through coordinated representations that help analysts judge whether similar traces reflect similar intentions. Richer data, in turn, place greater demands on the analyst, foregrounding the costs of the human-in-the-loop commitment.

Analyst effort, scalability, and the balance with automation. Keeping the analyst in the loop demands sustained effort and domain expertise, makes subjective judgments part of the analytical result, and sharpens the scalability challenge. Long sequences, large event alphabets, and more users or clusters induce visual clutter and cognitive load, which the employed abstraction, filtering, and overview+detail techniques mitigate, though at the cost of a complete global overview. Future systems could reduce this effort while retaining the analyst's interpretive role, for example through onboarding and explanation-on-demand tailored to users' expertise and visualization literacy. Whether such mechanisms improve analytical outcomes is an empirical question, which leads to evaluation.

Evaluation and comparative evidence. The evaluation is grounded in case studies, usage scenarios, and expert feedback, which are appropriate for investigating analytical utility yet leave comparative effectiveness relative to alternative methods unexamined. Comparative studies could quantify how alternative encodings, similarity measures, and clustering methods affect analysts' performance, while benchmark datasets or independently coded annotations could assess the computational components against a known reference. Even rigorous quantitative evaluation, however, only establishes that patterns are extracted reliably; what they mean remains an interpretive question, marking the final boundary of this thesis.

Interpretive boundaries and the validation of hypotheses. This thesis does not claim that visual analytics systems can infer cognition, learning, or intent from interaction data; identified relationships between behavioral patterns and constructs, such as systems thinking, remain interpretive hypotheses. Because the behavioral descriptors were derived in relation to the applied theoretical framework, evidence from the logs alone cannot test a mapping it already presupposes. A central future direction is therefore to test the generated hypotheses against independently measured criteria, such as pre- and post-assessments, think-aloud protocols, and longitudinal studies, examining whether particular motifs or strategy profiles are associated with learning outcomes or transfer across tasks.

9.4 Onward

Together, these directions describe a cycle rather than an endpoint: interaction data are analyzed, hypotheses about higher-level constructs are formed and tested, and the resulting insights reshape the environments themselves, whose subsequent use lays down new traces to follow.

The cholera map (Fig. 1.1) that opens this thesis did not merely record where people had died; it made a hidden cause visible, and in doing so it changed what could be done about it. Interaction logs pose the same challenge at a different scale. Read with computational support and human judgment together, they reveal not only what users did, but something of how they explored, learned, and reasoned along the way.

EPILOGUE

I extend my appreciation to those whom I knew, know, and have yet to know, in the past, present, and future:

- All my colleagues at Linköping University and the Information Visualization unit. A place where people from many countries, disciplines, and backgrounds pull in the same direction, where doors remain open, where help is offered before it is requested, and where the work itself remains at the center. May it continue to grow and flourish as exactly that.
- My collaborators, who made the work productive and the process easy, generously contributed their time, effort, and expertise, and whose encouragement, humor, and laughter made the whole collaboration a pleasure.
- My friends, who throughout these years have been together at conferences, courses, seminars, labs, retreats, meetings, offices, and corridors, as well as activities, meals, fikas, afterworks, and board games that ran later than anyone had planned.
- My spouse, a mutual choice.
- My family, not by choice.
- My cat, who chose me.
- And finally, the Swedish Research Council, which supported this work through grant 2020-05000.

[1] Wil van der Aalst. "Data Science in Action." In: *Process Mining: Data Science in Action.* Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 3–23. DOI: 10.1007/978-3-662-49851-4_1.

[2] Saeed Aghabozorgi, Ali Seyed Shirkhorshidi, and Teh Ying Wah. "Time-series clustering – A decade review." In: *Inf. Syst.* 53 (2015), pp. 16–38. ISSN: 0306-4379. DOI: 10.1016/j.is.2015.04.007.

[3] Mohammed Ali, Ali Alqahtani, Mark W. Jones, and Xianghua Xie. "Clustering and Classification for Time Series Data in Visual Analytics: A Survey." In: *IEEE Access* 7 (2019), pp. 181314–181338. DOI: 10.1109/ACCESS.2019.2958551.

[4] Lorin W Anderson and David R Krathwohl. *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives: complete edition.* Addison-Wesley, 2001.

[5] Lucija Andre, Tracy Durksen, and Monique L Volman. "Museums as avenues of learning for children: A decade of research." In: *Learn. Environ. Res.* 20.1 (2017), pp. 47–76. DOI: 10.1007/s10984-016-9222-9.

[6] Orit Ben-Zvi Assaraf and Nir Orion. "Development of system thinking skills in the context of earth system education." In: *J. Res. Sci. Teach.* 42.5 (2005), pp. 518–560. DOI: 10.1002/tea.20061.

[7] L. Bergroth, H. Hakonen, and T. Raita. "A survey of longest common subsequence algorithms." In: *Proceedings Seventh International Symposium on String Processing and Information Retrieval. SPIRE 2000.* 2000, pp. 39–48. DOI: 10.1109/SPIRE.2000.878178.

[8] Donald J. Berndt and James Clifford. "Using dynamic time warping to find patterns in time series." In: *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining.* AAAIWS'94. Seattle, WA: AAAI Press, 1994, pp. 359–370.

[9] Lonni Besançon, Konrad Schönborn, Erik Sundén, He Yin, Samuel Rising, Peter Westerdahl, Patric Ljung, Josef Wideström, Charles Hansen, and Anders Ynnerman. "Exploring and explaining climate change: Exploranation as a visualization pedagogy for societal action." In: *VIS4GOOD, a workshop on Visualization for Social Good, held as part of IEEE VIS 2022.* Institute of Electrical and Electronics Engineers (IEEE). 2022.

[10] Lonni Besançon, Anders Ynnerman, Daniel F. Keefe, Lingyun Yu, and Tobias Isenberg. "The State of the Art of Spatial Interfaces for 3D Visualization." In: *Comput. Graph. Forum* 40.1 (2021), pp. 293–326. DOI: 10.1111/cgf.14189.

[11] Carla Binucci, Annika Bonerath, Walter Didimo, Henry Förster, Seok-Hee Hong, Maria Eleni Pavlidi, and Alessandra Tappini. "Clusterix: A Hybrid Visualization Model for Hierarchically Clustered Networks." In: *Comput. Graph. Forum* (2026). DOI: 10.1111/cgf.70469.

[12] Tanja Blascheck, Kuno Kurzhals, Michael Raschke, Michael Burch, Daniel Weiskopf, and Thomas Ertl. “State-of-the-art of visualization for eye tracking data.” In: *Proc. EuroVis (STARs)*. Eurographics Assoc., 2014. DOI: 10.2312/eurovisstar.20141173.

[13] Tanja Blascheck, Lindsay MacDonald Vermeulen, Jo Vermeulen, Charles Perin, Wesley Willett, Thomas Ertl, and Sheelagh Carpendale. “Exploration Strategies for Discovery of Interactivity in Visualizations.” In: *IEEE Trans. Vis. Comput. Graph.* 25.2 (2019), pp. 1407–1420. DOI: 10.1109/TVCG.2018.2802520.

[14] Benjamin S Bloom, Max D Engelhart, Edward J Furst, Walker H Hill, David R Krathwohl, et al. *Taxonomy of educational objectives: The classification of educational goals. Handbook 1: Cognitive domain*. Longman, 1956.

[15] Alejandro Bogarín, Rebeca Cerezo, and Cristóbal Romero. “A survey on educational process mining.” In: *WIREs Data Min. Knowl. Discov.* 8.1 (2018), e1230. DOI: 10.1002/widm.1230.

[16] Katy Börner, Adam Maltese, Russell Nelson Balliet, and Joe Heimlich. “Investigating aspects of data visualization literacy using 20 information visualizations and 273 science museum visitors.” In: *Inf. Vis.* 15.3 (2016), pp. 198–213. DOI: 10.1177/1473871615594652.

[17] Michael Burch. *Eye tracking and visual analytics*. River Publ., 2022.

[18] Bram C.M. Cappers and Jarke J. van Wijk. “Exploring Multivariate Event Sequences Using Rules, Aggregations, and Selections.” In: *IEEE Trans. Vis. Comput. Graph.* 24.1 (2018), pp. 532–541. DOI: 10.1109/TVCG.2017.2745278.

[19] Stuart K Card, Jock Mackinlay, and Ben Shneiderman. *Readings in information visualization: using vision to think*. Morgan Kaufmann, 1999.

[20] Sheelagh Carpendale. “Evaluating Information Visualizations.” In: *Information Visualization: Human-Centered Issues and Perspectives*. Berlin, Heidelberg: Springer, 2008, pp. 19–45. ISBN: 978-3-540-70956-5. DOI: 10.1007/978-3-540-70956-5_2.

[21] Davide Ceneda, Theresia Gschwandtner, Thorsten May, Silvia Miksch, Hans-Jörg Schulz, Marc Streit, and Christian Tominski. “Characterizing Guidance in Visual Analytics.” In: *IEEE Trans. Vis. Comput. Graph.* 23.1 (2017), pp. 111–120. DOI: 10.1109/TVCG.2016.2598468.

[22] Qing Chen, Xuanwu Yue, Xavier Plantaz, Yuanzhe Chen, Conglei Shi, Ting-Chuen Pong, and Huamin Qu. “ViSeq: Visual Analytics of Learning Sequence in Massive Open Online Courses.” In: *IEEE Trans. Vis. Comput. Graph.* 26.3 (2020), pp. 1622–1636. DOI: 10.1109/TVCG.2018.2872961.

[23] Yuanzhe Chen, Panpan Xu, and Liu Ren. “Sequence Synopsis: Optimize Visual Summary of Temporal Event Data.” In: *IEEE Trans. Vis. Comput. Graph.* 24.1 (2018), pp. 45–55. DOI: 10.1109/TVCG.2017.2745083.

[24] Olivier Christmann, Sylvain Fleury, Jérôme Migaud, Vincent Raimbault, Benjamin Poussard, Thibaut Guitter, Geoffrey Gorisse, and Simon Richir. “Visualizing the invisible: User-centered design of a system for the visualization of flows and concentrations of particles in the air.” In: *Inf. Vis.* 21 (2022), pp. 311–320. DOI: 10.1177/14738716221085965.

[25] Carleton Coffrin, Linda Corrin, Paula de Barba, and Gregor Kennedy. “Visualizing patterns of student engagement and performance in MOOCs.” In: *Proceedings of the Fourth International Conference on Learning Analytics And Knowledge*. LAK ’14. Indianapolis, Indiana, USA: ACM, 2014, pp. 83–92. ISBN: 9781450326643. DOI: 10.1145/2567574.2567586.

[26] Hui Ding, Goce Trajcevski, Peter Scheuermann, Xiaoyue Wang, and Eamonn J Keogh. “Querying and mining of time series data: Experimental comparison of representations and distance measures.” In: *Proc. VLDB Endow.* 1.2 (2008), pp. 1542–1552.

[27] Veronika Domova and Katerina Vrotsou. “A Model for Types and Levels of Automation in Visual Analytics: A Survey, a Taxonomy, and Examples.” In: *IEEE Trans. Vis. Comput. Graph.* 29.8 (2023), pp. 3550–3568. DOI: 10.1109/TVCG.2022.3163765.

[28] Fan Du, Catherine Plaisant, Neil Spring, and Ben Shneiderman. “EventAction: Visual analytics for temporal event sequence recommendation.” In: *VAST*. 2016, pp. 61–70. DOI: 10.1109/VAST.2016.7883512.

[29] Fan Du, Ben Shneiderman, Catherine Plaisant, Sana Malik, and Adam Perer. “Coping with Volume and Variety in Temporal Event Sequences: Strategies for Sharpening Analytic Focus.” In: *IEEE Trans. Vis. Comput. Graph.* 23.6 (2017), pp. 1636–1649. DOI: 10.1109/TVCG.2016.2539960.

[30] Omar ElTayeby and Wenwen Dou. “A Survey on Interaction Log Analysis for Evaluating Exploratory Visualizations.” In: *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization*. BELIV ’16. Baltimore, MD, USA: ACM, 2016, pp. 62–69. ISBN: 9781450348188. DOI: 10.1145/2993901.2993912.

[31] Florian Endel and Harald Piringer. “Data Wrangling: Making data useful again.” In: *IFAC-PapersOnLine* 48.1 (2015), pp. 111–112. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2015.05.197.

[32] Vladimir Estivill-Castro. “Why so many clustering algorithms: a position paper.” In: *SIGKDD Explor. Newsl.* 4.1 (June 2002), pp. 65–75. ISSN: 1931-0145. DOI: 10.1145/568574.568575.

[33] Brian S. Everitt, Sabine Landau, Morven Leese, and Daniel Stahl. *Cluster Analysis, 5th Edition*. Wiley-Blackwell, Jan. 2011. ISBN: 9780470749913. DOI: 10.1002/9780470977811.

[34] Christos Faloutsos, Mudumbai Ranganathan, and Yannis Manolopoulos. “Fast subsequence matching in time-series databases.” In: *ACM SIGMOD Rec.* 23.2 (1994), pp. 419–429.

[35] Edward W Forgy. “Cluster analysis of multivariate data: efficiency versus interpretability of classifications.” In: *Biometrics* 21 (1965), pp. 768–769.

[36] Tim Furche, Georg Gottlob, Leonid Libkin, Giorgio Orsi, and Norman W. Paton. “Data Wrangling for Big Data: Challenges and Opportunities.” In: *Advances in Database Technology — EDBT 2016*. Advances in Database Technology. Univ. Konstanz, 2016, pp. 473–478. ISBN: 978-3-89318-070-7. DOI: 10.5441/002/edbt.2016.44.

[37] Michael Gleicher. “Considerations for Visualizing Comparison.” In: *IEEE Trans. Vis. Comput. Graph.* 24.1 (2018), pp. 413–423. DOI: 10.1109/TVCG.2017.2744199.

[38] Michael Gleicher, Danielle Albers, Rick Walker, Ilir Jusufi, Charles D. Hansen, and Jonathan C. Roberts. "Visual comparison for information visualization." In: *Inf. Vis.* 10.4 (2011), pp. 289–309. DOI: 10.1177/1473871611416549.

[39] Sarah Goodwin, Arnaud Prouzeau, Ryan Whitelock-Jones, Christophe Hurter, Lee Lawrence, Umair Afzal, and Tim Dwyer. "VETA: Visual eye-tracking analytics for the exploration of gaze patterns and behaviours." In: *Vis. Inform.* 6.2 (2022), pp. 1–13. ISSN: 2468-502X. DOI: 10.1016/j.visinf.2022.02.004.

[40] David Gotz. "Soft patterns: Moving beyond explicit sequential patterns during visual analysis of longitudinal event datasets." In: *Proceedings of the IEEE VIS 2016 Workshop on Temporal & Sequential Event Analysis*. 2016.

[41] David Gotz and Harry Stavropoulos. "DecisionFlow: Visual Analytics for High-Dimensional Temporal Event Sequence Data." In: *IEEE Trans. Vis. Comput. Graph.* 20.12 (2014), pp. 1783–1792. DOI: 10.1109/TVCG.2014.2346682.

[42] Theresia Gschwandtner. "Visual Analytics Meets Process Mining: Challenges and Opportunities." In: *Data-Driven Process Discovery and Analysis*. Ed. by Paolo Ceravolo and Stefanie Rinderle-Ma. Cham: Springer International Publishing, 2017, pp. 142–154. ISBN: 978-3-319-53435-0.

[43] Shunan Guo, Zhuochen Jin, David Gotz, Fan Du, Hongyuan Zha, and Nan Cao. "Visual Progression Analysis of Event Sequence Data." In: *IEEE Trans. Vis. Comput. Graph.* 25.1 (2019), pp. 417–426. DOI: 10.1109/TVCG.2018.2864885.

[44] Shunan Guo, Ke Xu, Rongwen Zhao, David Gotz, Hongyuan Zha, and Nan Cao. "EventThread: Visual Summarization and Stage Analysis of Event Sequence Data." In: *IEEE Trans. Vis. Comput. Graph.* 24.1 (2018), pp. 56–65. DOI: 10.1109/TVCG.2017.2745320.

[45] Yi Guo, Shunan Guo, Zhuochen Jin, Smiti Kaul, David Gotz, and Nan Cao. "Survey on Visual Analysis of Event Sequence Data." In: *IEEE Trans. Vis. Comput. Graph.* 28.12 (2022), pp. 5091–5112. DOI: 10.1109/TVCG.2021.3100413.

[46] Patrick A. V. Hall and Geoff R. Dowling. "Approximate String Matching." In: *ACM Comput. Surv.* 12.4 (Dec. 1980), pp. 381–402. ISSN: 0360-0300. DOI: 10.1145/356827.356830.

[47] John A Hartigan. *Clustering algorithms*. Wiley, 1975.

[48] Uta Hinrichs, Holly Schmidt, and Sheelagh Carpendale. "EMDialog: Bringing Information Visualization into the Museum." In: *IEEE Trans. Vis. Comput. Graph.* 14.6 (2008), pp. 1181–1188. DOI: 10.1109/TVCG.2008.127.

[49] Cindy E Hmelo-Silver, Surabhi Marathe, and Lei Liu. "Fish swim, rocks sit, and lungs breathe: Expert-novice understanding of complex systems." In: *J. Learn. Sci.* 16.3 (2007), pp. 307–331. DOI: 10.1080/10508400701413401.

[50] Danny Holten. "Hierarchical Edge Bundles: Visualization of Adjacency Relations in Hierarchical Data." In: *IEEE Trans. Vis. Comput. Graph.* 12.5 (2006), pp. 741–748. DOI: 10.1109/TVCG.2006.147.

[51] Gunnar Höst, Karljohan Palmerius, and Konrad Schönborn. "Nano for the Public: An Exploranation Perspective." In: *IEEE Comput. Graph. Appl.* 40.2 (2020), pp. 32–42. DOI: 10.1109/MCG.2020.2973120.

[52] Zhexue Huang. “Extensions to the k-means algorithm for clustering large data sets with categorical values.” In: *Data Min. Knowl. Discov.* 2.3 (1998), pp. 283–304.

[53] Kaveh Kamgar, Shaghayegh Gharghabi, and Eamonn Keogh. “Matrix Profile XV: Exploiting Time Series Consensus Motifs to Find Structure in Time Series Sets.” In: *Proc. ICDM.* 2019, pp. 1156–1161. DOI: 10.1109/ICDM.2019.00140.

[54] Sean Kandel, Jeffrey Heer, Catherine Plaisant, Jessie Kennedy, Frank van Ham, Nathalie Henry Riche, Chris Weaver, Bongshin Lee, Dominique Brodbeck, and Paolo Buono. “Research directions in data wrangling: Visualizations and transformations for usable and credible data.” In: *Inf. Vis.* 10.4 (2011), pp. 271–288. DOI: 10.1177/1473871611415994.

[55] Daniel Keim, Gennady Andrienko, Jean-Daniel Fekete, Carsten Görg, Jörn Kohlhammer, and Guy Melançon. “Visual Analytics: Definition, Process, and Challenges.” In: *Information Visualization: Human-Centered Issues and Perspectives.* Berlin, Heidelberg: Springer, 2008, pp. 154–175. ISBN: 978-3-540-70956-5. DOI: 10.1007/978-3-540-70956-5_7.

[56] Eamonn Keogh and Jessica Lin. “Clustering of time-series subsequences is meaningless: implications for previous and future research.” In: *Knowl. Inf. Syst.* 8.2 (Aug. 2005), pp. 154–177. ISSN: 0219-3116. DOI: 10.1007/s10115-004-0172-7.

[57] Daniel H Kim. *Introduction to systems thinking.* Vol. 16. Pegasus Commun., 1999.

[58] Milos Krstajic, Enrico Bertini, and Daniel Keim. “CloudLines: Compact Display of Event Episodes in Multiple Time-Series.” In: *IEEE Trans. Vis. Comput. Graph.* 17.12 (2011), pp. 2432–2439. DOI: 10.1109/TVCG.2011.179.

[59] Heidi Lam, Enrico Bertini, Petra Isenberg, Catherine Plaisant, and Sheelagh Carpendale. “Empirical Studies in Information Visualization: Seven Scenarios.” In: *IEEE Trans. Vis. Comput. Graph.* 18.9 (2012), pp. 1520–1536. DOI: 10.1109/TVCG.2011.279.

[60] Po-Ming Law, Zhicheng Liu, Sana Malik, and Rahul C. Basole. “MAQUI: Interweaving Queries and Pattern Mining for Recursive Event Sequence Exploration.” In: *IEEE Trans. Vis. Comput. Graph.* 25.1 (2019), pp. 396–406. DOI: 10.1109/TVCG.2018.2864886.

[61] Vladimir I Levenshtein et al. “Binary codes capable of correcting deletions, insertions, and reversals.” In: *Soviet physics doklady.* Vol. 10. 8. Soviet Union. 1966, pp. 707–710.

[62] C. Liebers, S. Agarwal, M. Krug, K. Pitsch, and F. Beck. “VisCoMET: Visually Analyzing Team Collaboration in Medical Emergency Trainings.” In: *Comput. Graph. Forum* 42.3 (2023), pp. 149–160. DOI: 10.1111/cgf.14819.

[63] Lisa Linnenbrink-Garcia, Amanda M. Durik, AnneMarie M. Conley, Kenneth E. Barron, John M. Tauer, Stuart A. Karabenick, and Judith M. Harackiewicz. “Measuring Situational Interest in Academic Domains.” In: *Educ. Psychol. Meas.* 70.4 (2010), pp. 647–671. DOI: 10.1177/0013164409355699.

[64] Zheng Liu, Jeffrey Xu Yu, Xuemin Lin, Hongjun Lu, and Wei Wang. “Locating Motifs in Time-Series Data.” In: *Advances in Knowledge Discovery and Data Mining.* Berlin, Heidelberg: Springer, 2005, pp. 343–353.

[65] Jonas Lundberg. “Situation awareness systems, states and processes: a holistic framework.” In: *Theor. Issues Ergon. Sci.* 16.5 (2015), pp. 447–473. DOI: 10.1080/1463922X.2015.1008601.

[66] Joyce Ma, Kwan-Liu Ma, and Jennifer Frazier. "Decoding a Complex Visualization in a Science Museum – An Empirical Study." In: *IEEE Trans. Vis. Comput. Graph.* 26.1 (2020), pp. 472–481. DOI: 10.1109/TVCG.2019.2934401.

[67] Jessica Magallanes, Tony Stone, Paul D Morris, Suzanne Mason, Steven Wood, and Maria-Cruz Villa-Uriol. "Sequen-C: A Multilevel Overview of Temporal Event Sequences." In: *IEEE Trans. Vis. Comput. Graph.* 28.1 (2022), pp. 901–911. DOI: 10.1109/TVCG.2021.3114868.

[68] Mina Mani, Måns Gezelius, Jonas Löwgren, Gunnar Höst, Marta M. Koć-Januchta, Lena A. E. Tibell, and Konrad J. Schönborn. "Designing with multiple perspectives: an interactive learning environment for developing systems thinking in a carbon cycle context." In: *Interact. Learn. Environ.* 34.1 (2026), pp. 218–243. DOI: 10.1080/10494820.2025.2494161.

[69] Sabrina Mantaci, Antonio Restivo, and Marinella Sciortino. "Distance measures for biological sequences: Some recent approaches." In: *Int. J. Approx. Reason.* 47.1 (2008), pp. 109–124. ISSN: 0888-613X. DOI: 10.1016/j.ijar.2007.03.011.

[70] Wannisa Matcha, Nora'ayu Ahmad Uzir, Dragan Gašević, and Abelardo Pardo. "A Systematic Review of Empirical Studies on Learning Analytics Dashboards: A Self-Regulated Learning Perspective." In: *IEEE Trans. Learn. Technol.* 13.2 (2020), pp. 226–245. DOI: 10.1109/TLT.2019.2916802.

[71] Matthew Louis Mauriello, Ben Shneiderman, Fan Du, Sana Malik, and Catherine Plaisant. "Simplifying Overviews of Temporal Event Sequences." In: *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems.* CHI EA '16. San Jose, California, USA: ACM, 2016, pp. 2217–2224. ISBN: 9781450340823. DOI: 10.1145/2851581.2892440.

[72] Lothar Meyer, KJ Klang, S Boonsong, C Westin, A Nordman, J Lundberg, B Josefsson, and K Vrotsou. "Mapping the decision-making process of conflict detection and resolution in en-route control: An eye-tracking based approach." In: *Proc. SESAR Innovation Days.* 2022.

[73] Miriah Meyer, Tamara Munzner, and Hanspeter Pfister. "MizBee: A Multiscale Synteny Browser." In: *IEEE Trans. Vis. Comput. Graph.* 15.6 (2009), pp. 897–904. DOI: 10.1109/TVCG.2009.167.

[74] Silvia Miksch, Claudio Di Ciccio, Pnina Soffer, and Barbara Weber. "Visual Analytics Meets Process Mining: Challenges and Opportunities." In: *IEEE Comput. Graph. Appl.* 44.6 (2024), pp. 132–141. DOI: 10.1109/MCG.2024.3456916.

[75] Robert B. Miller. "Response time in man-computer conversational transactions." In: *Proceedings of the December 9-11, 1968, Fall Joint Computer Conference, Part I.* AFIPS '68 (Fall, part I). San Francisco, California: ACM, 1968, pp. 267–277. ISBN: 9781450378994. DOI: 10.1145/1476589.1476628.

[76] Zeynab Mohseni, Italo Masiello, Rafael M. Martins, and Susanna Nordmark. "Visual Learning Analytics for Educational Interventions in Primary and Secondary Schools: A Scoping Review." In: *J. Learn. Anal.* 11.2 (June 2024), pp. 91–111. DOI: 10.18608/jla.2024.8309.

[77] Jamie P Monat and Thomas F Gannon. "What is systems thinking? A review of selected literature plus recommendations." In: *Am. J. Syst. Sci.* 4.1 (2015), pp. 11–26.

[78] Megan Monroe, Rongjian Lan, Hanseung Lee, Catherine Plaisant, and Ben Shneiderman. "Temporal Event Sequence Simplification." In: *IEEE Trans. Vis. Comput. Graph.* 19.12 (2013), pp. 2227–2236. DOI: 10.1109/TVCG.2013.200.

[79] Phong H. Nguyen, Rafael Henkin, Siming Chen, Natalia Andrienko, Gennady Andrienko, Olivier Thonnard, and Cagatay Turkay. "VASABI: Hierarchical User Profiles for Interactive Visual User Behaviour Analytics." In: *IEEE Trans. Vis. Comput. Graph.* 26.1 (2020), pp. 77–86. DOI: 10.1109/TVCG.2019.2934609.

[80] Phong H. Nguyen, Cagatay Turkay, Gennady Andrienko, Natalia Andrienko, Olivier Thonnard, and Jihane Zouaoui. "Understanding User Behaviour through Action Sequences: From the Usual to the Unusual." In: *IEEE Trans. Vis. Comput. Graph.* 25.9 (2019), pp. 2838–2852. DOI: 10.1109/TVCG.2018.2859969.

[81] Chris North, Remco Chang, Alex Endert, Wenwen Dou, Richard May, Bill Pike, and Glenn Fink. "Analytic provenance: process+interaction+insight." In: *CHI '11 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '11. Vancouver, BC, Canada: ACM, 2011, pp. 33–36. ISBN: 9781450302685. DOI: 10.1145/1979742.1979570.

[82] P. Patel, E. Keogh, J. Lin, and S. Lonardi. "Mining motifs in massive time series databases." In: *2002 IEEE International Conference on Data Mining, 2002. Proceedings.* 2002, pp. 370–377. DOI: 10.1109/ICDM.2002.1183925.

[83] Xingcheng Ran, Yue Xi, Yonggang Lu, Xiangwen Wang, and Zhenyu Lu. "Comprehensive survey on hierarchical clustering algorithms and the recent developments." In: *Artif. Intell. Rev.* 56.8 (Aug. 2023), pp. 8219–8264. ISSN: 1573-7462. DOI: 10.1007/s10462-022-10366-3.

[84] Michael Raschke, Tanja Blascheck, and Michael Burch. "Visual Analysis of Eye Tracking Data." In: *Handbook of Human Centric Visualization*. New York, NY: Springer, 2014, pp. 391–409. DOI: 10.1007/978-1-4614-7485-2_15.

[85] Chotirat Ratanamahatana, Eamonn Keogh, Anthony J. Bagnall, and Stefano Lonardi. "A Novel Bit Level Time Series Representation with Implication of Similarity Search and Clustering." In: *Advances in Knowledge Discovery and Data Mining*. Berlin, Heidelberg: Springer, 2005, pp. 771–777.

[86] Stefanie Rinderle-Ma, Florian Stertz, Juergen Mangler, and Florian Pauker. "Process Mining—Discovery, Conformance, and Enhancement of Manufacturing Processes." In: *Digital Transformation: Core Technologies and Emerging Topics from a Computer Science Perspective*. Berlin, Heidelberg: Springer, 2023, pp. 363–383. DOI: 10.1007/978-3-662-65004-2_15.

[87] Christoffer Rubensson, Harleen Kaur, Timotheus Kampik, and Jan Mendling. "Timeline-based process discovery." In: *Inf. Syst.* 133 (2025), p. 102568. ISSN: 0306-4379. DOI: 10.1016/j.is.2025.102568.

[88] Mohammed Saeed and Roger Mark. "A novel method for the efficient retrieval of similar multiparameter physiologic time series using wavelet-based symbolic representations." en. In: *AMIA Annu. Symp. Proc.* 2006 (2006), pp. 679–683.

[89] Neil Selwyn. "Re-imagining 'Learning Analytics' … a case for starting again?" In: *Internet High. Educ.* 46 (2020), p. 100745. ISSN: 1096-7516. DOI: 10.1016/j.iheduc.2020.100745.

[90] B. Shneiderman. "The eyes have it: a task by data type taxonomy for information visualizations." In: *Proc. VL*. 1996, pp. 336–343. DOI: 10.1109/VL.1996.545307.

[91] Ben Shneiderman and Catherine Plaisant. "Interactive Visual Event Analytics: Opportunities and Challenges." In: *Computer* 52.1 (2019), pp. 27–35. DOI: 10.1109/MC.2018.2890217.

[92] Mohammad Shokoohi-Yekta, Bing Hu, Hongxia Jin, Jun Wang, and Eamonn Keogh. "Generalizing DTW to the multi-dimensional case requires an adaptive approach." In: *Data Min. Knowl. Discov.* 31 (2017), pp. 1–31. DOI: 10.1007/s10618-016-0455-0.

[93] George Siemens. "Learning Analytics: The Emergence of a Discipline." In: *Am. Behav. Sci.* 57.10 (2013), pp. 1380–1400. DOI: 10.1177/0002764213498851.

[94] Claudio T. Silva, Juliana Freire, and Steven P. Callahan. "Provenance for Visualizations: Reproducibility and Beyond." In: *Comput. Sci. Eng.* 9.5 (2007), pp. 82–89. DOI: 10.1109/MCSE.2007.106.

[95] John Snow. "On the Mode of Communication of Cholera." en. In: *Edinb. Med. J.* 1.7 (Jan. 1856), pp. 668–670.

[96] Matthias Studer and Gilbert Ritschard. "What Matters in Differences Between Life Trajectories: A Comparative Review of Sequence Dissimilarity Measures." In: *J. R. Stat. Soc. Ser. A Stat. Soc.* 179.2 (Feb. 2016), pp. 481–511. ISSN: 0964-1998. DOI: 10.1111/rssa.12125.

[97] James J Thomas. *Illuminating the path: The research and development agenda for visual analytics*. IEEE Comput. Soc., 2005.

[98] Edward R Tufte. *Beautiful evidence*. Vol. 1. Graphics Press, 2006.

[99] Sanne van der Linden, Evie de Fouw, Stef van den Elzen, and Anna Vilanova. "A survey of visualization techniques for comparing event sequences." In: *Comput. Graph.* 115 (2023), pp. 522–542. ISSN: 0097-8493. DOI: 10.1016/j.cag.2023.05.016.

[100] Roald P. Verhoeff, Marie-Christine P. J. Knippels, Melde G. R. Gilissen, and Kerst T. Boersma. "The Theoretical Nature of Systems Thinking. Perspectives on Systems Thinking in Biology Education." In: *Front. Educ.* Volume 3 - 2018 (2018). ISSN: 2504-284X. DOI: 10.3389/feduc.2018.00040.

[101] Visualiseringscenter C. *Visualiseringscenter C*. https://visualiseringscenter.se. 2026.

[102] Katerina Vrotsou and Aida Nordman. "A Window-based Approach for Mining Long Duration Event-sequences." In: *EuroVA@ Eurographics/EuroVis*. 2020, pp. 55–59.

[103] Katerina Vrotsou and Aida Nordman. "Exploratory Visual Sequence Mining Based on Pattern-Growth." In: *IEEE Trans. Vis. Comput. Graph.* 25.8 (2019), pp. 2597–2610. DOI: 10.1109/TVCG.2018.2848247.

[104] Katerina Vrotsou, Anders Ynnerman, and Matthew Cooper. "Are we what we do? Exploring group behaviour through user-defined event-sequence similarity." In: *Inf. Vis.* 13.3 (2014), pp. 232–247. DOI: 10.1177/1473871613477852.

[105] Carl Westin, Katerina Vrotsou, Aida Nordman, Jonas Lundberg, and Lothar Meyer. "Visual scan patterns in tower control: Foundations for an instructor support tool." In: *SESAR Innovation Days*. 2019.

[106] Krist Wongsuphasawat and David Gotz. "Outflow: Visualizing patient flow by symptoms and outcome." In: *IEEE VisWeek Workshop on Visual Analytics in Healthcare, Providence, Rhode Island, USA*. American Medical Informatics Association. 2011, pp. 25–28.

[107] Krist Wongsuphasawat, John Alexis Guerra Gómez, Catherine Plaisant, Taowei David Wang, Meirav Taieb-Maimon, and Ben Shneiderman. "LifeFlow: visualizing an overview of event sequences." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: ACM, 2011, pp. 1747–1756. ISBN: 9781450302289. DOI: 10.1145/1978942.1979196.

[108] Yingcai Wu, Nan Cao, David Gotz, Yap-Peng Tan, and Daniel A. Keim. "A Survey on Visual Analytics of Social Media Data." In: *IEEE Trans. Multimedia* 18.11 (2016), pp. 2135–2148. DOI: 10.1109/TMM.2016.2614220.

[109] Kai Xu, Alvitta Ottley, Conny Walchshofer, Marc Streit, Remco Chang, and John Wenskovitch. "Survey on the Analysis of User Interactions and Visualization Provenance." In: *Comput. Graph. Forum* 39.3 (2020), pp. 757–783. DOI: 10.1111/cgf.14035.

[110] Chin-Chia Michael Yeh, Nickolas Kavantzas, and Eamonn Keogh. "Matrix Profile VI: Meaningful Multidimensional Motif Discovery." In: *Proc. ICDM*. 2017, pp. 565–574. DOI: 10.1109/ICDM.2017.66.

[111] Anton Yeshchenko and Jan Mendling. "A survey of approaches for event sequence analysis and visualization." In: *Inf. Syst.* 120 (2024), p. 102283. ISSN: 0306-4379. DOI: 10.1016/j.is.2023.102283.

[112] Anders Ynnerman, Jonas Löwgren, and Lena Tibell. "Exploranation: A New Science Communication Paradigm." In: *IEEE Comput. Graph. Appl.* 38.3 (2018), pp. 13–20. DOI: 10.1109/MCG.2018.032421649.

[113] Peilin Yu, Aida Nordman, Takanori Fujiwara, Marta Koc-Januchta, Konrad Schönborn, Lonni Besançon, and Katerina Vrotsou. "Visual Extraction of Interaction Patterns Guided by Hierarchical Clustering and Process Mining." In: *IEEE Trans. Vis. Comput. Graph.* 32.1 (2026), pp. 1142–1152. DOI: 10.1109/TVCG.2025.3634837.

[114] Peilin Yu, Aida Nordman, Marta Koc-Januchta, Konrad Schönborn, Lonni Besançon, and Katerina Vrotsou. "Revealing Interaction Dynamics: Multi-Level Visual Exploration of User Strategies with an Interactive Digital Environment." In: *IEEE Trans. Vis. Comput. Graph.* 31.1 (2025), pp. 831–841. DOI: 10.1109/TVCG.2024.3456187.

[115] Peilin Yu, Aida Nordman, Mina Mani, Konrad Schönborn, Marta Koc-Januchta, Gunnar Höst, Måns Gezelius, Lonni Besançon, and Katerina Vrotsou. "From Interaction to Systems Thinking: A Visual Analytics Approach for Exploring Pupils' Behavior in a Digital Learning Environment." Submitted to IEEE Trans. Vis. Comput. Graph. 2026. DOI: 10.35542/osf.io/9zp7k_v1.

[116] Peilin Yu, Aida Nordman, Lothar Meyer, Supathida Boonsong, and Katerina Vrotsou. "Interactive Transformations and Visual Assessment of Noisy Event Sequences: An Application in En-Route Air Traffic Control." In: *2023 IEEE 16th Pacific Visualization Symposium (PacificVis)*. 2023, pp. 92–101. DOI: 10.1109/PacificVis56936.2023.00017.

[117] Yixuan Zhang, Sara Di Bartolomeo, Fangfang Sheng, Holly Jimison, and Cody Dunne. "Evaluating Alignment Approaches in Superimposed Time-Series and Temporal Event-Sequence Visualizations." In: *Visualization Conference*. 2019, pp. 1–5. DOI: 10.1109/VISUAL.2019.8933584.

[118] Yixuan Zhang, Kartik Chanana, and Cody Dunne. "IDMVis: Temporal Event Sequence Visualization for Type 1 Diabetes Treatment Decision Support." In: *IEEE Trans. Vis. Comput. Graph.* 25.1 (2019), pp. 512–522. DOI: 10.1109/TVCG.2018.2865076.

[119] Jian Zhao, Zhicheng Liu, Mira Dontcheva, Aaron Hertzmann, and Alan Wilson. "MatrixWave: Visual Comparison of Event Sequence Data." In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. Seoul, Republic of Korea: ACM, 2015, pp. 259–268. ISBN: 9781450331456. DOI: 10.1145/2702123.2702419.

[120] Kazi Tasnim Zinat, Saimadhav Naga Sakhamuri, Aaron Sun Chen, and Zhicheng Liu. "A Multi-Level Task Framework for Event Sequence Analysis." In: *IEEE Trans. Vis. Comput. Graph.* 31.1 (2025), pp. 842–852. DOI: 10.1109/TVCG.2024.3456510.

Papers

The papers associated with this thesis have been removed for copyright reasons. For more details about these see:

https://doi.org/10.3384/9789181186888

FACULTY OF SCIENCE AND ENGINEERING

Linköping Studies in Science and Technology. Dissertation No. 2543, 2026

Department of Science and Technology

Linköping University
SE-581 83 Linköping, Sweden

www.liu.se

TRACES

/ ˈtreɪsɪz/ (n)

1. A surviving mark, sign, or evidence of the former existence, influence, or action of some agent or event.

LINKÖPING UNIVERSITY