**Liebers, Carina; Agarwal, Shivam; Krug, Maximilian; Pitsch, Karola; Beck, Fabian**


## VisCoMET: Visually Analyzing Team Collaboration in Medical Emergency Trainings


Date of secondary publication: 29.08.2023

Version of Record (Published Version), Article

Persistent identifier: urn:nbn:de:bvb:473-irb-902556

**Primary publication**

Liebers, Carina; Agarwal, Shivam; Krug, Maximilian; Pitsch, Karola; Beck, Fabian: VisCoMET: Visually Analyzing Team Collaboration in Medical Emergency Trainings. In: Computer Graphics Forum. 42 (2023), 3, pp. 149-160. DOI: 10.1111/cgf.14819

Eurographics Conference on Visualization (EuroVis) 2023
D. Archambault, R. Bujack, and T. Schreck
(Guest Editors)

COMPUTER GRAPHICS *forum*
*Volume 42 (2023), Number 3*

# VisCoMET: Visually Analyzing Team Collaboration in Medical Emergency Trainings

C. Liebers[†1], S. Agarwal[2], M. Krug[1], K. Pitsch[1], and F. Beck[2]

[1]University of Duisburg-Essen, Germany
[2]University of Bamberg, Germany

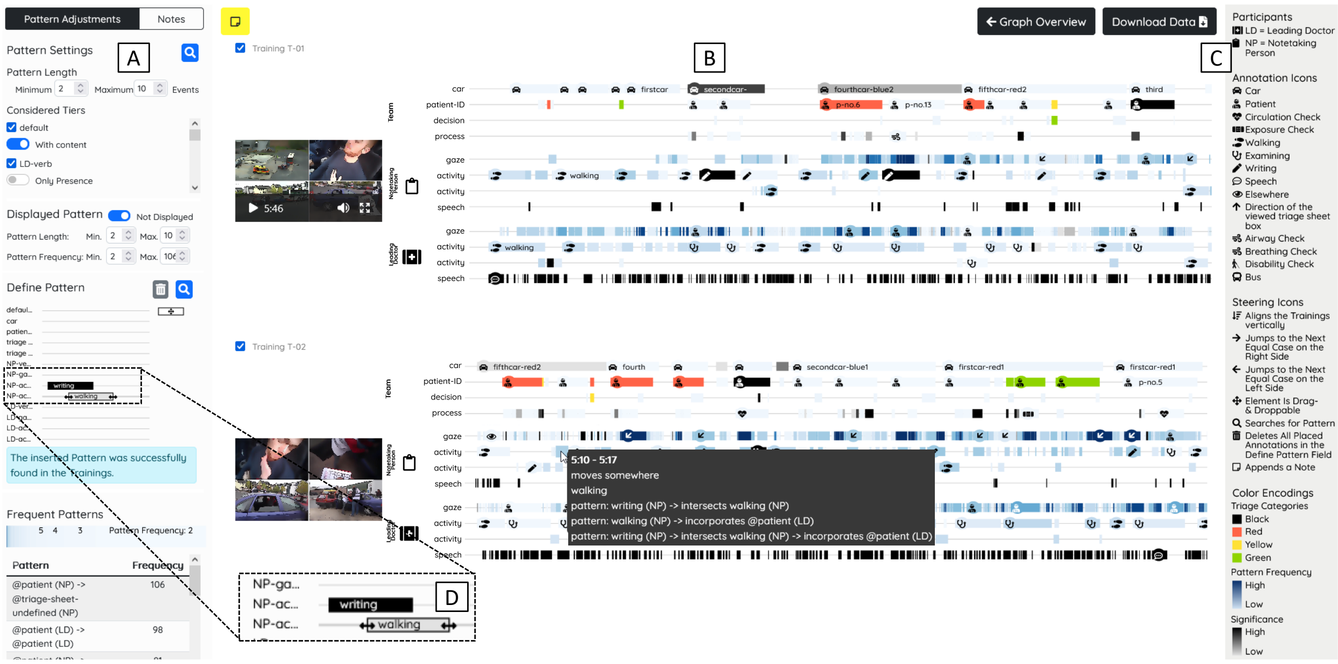


**Figure 1:** *Timeline-based details view of VisCoMET with (A) a panel offering options to detect frequent patterns in the data and specifying particular patterns for search, (B) timeline visualizations showing the parallel events of selected training sessions with their corresponding videos, (C) a legend explaining icons and color scales, and a (D) zoomed-in cutout of a defined pattern.*


## Abstract

*Handling emergencies requires efficient and effective collaboration of medical professionals. To analyze their performance, in an application study, we have developed VisCoMET, a visual analytics approach displaying interactions of healthcare personnel in a triage training of a mass casualty incident. The application scenario stems from social interaction research, where the collaboration of teams is studied from different perspectives. We integrate recorded annotations from multiple sources, such as recorded videos of the sessions, transcribed communication, and eye-tracking information. For each session, an information-rich timeline visualizes events across these different channels, specifically highlighting interactions between the team members. We provide algorithmic support to identify frequent event patterns and to search for user-defined event sequences. Comparing different teams, an overview visualization aggregates each training session in a visual glyph as a node, connected to similar sessions through edges. An application example shows the usage of the approach in the comparative analysis of triage training sessions, where multiple teams encountered the same scene, and highlights discovered insights. The approach was evaluated through feedback from visualization and social interaction experts. The results show that the approach supports reflecting on teams' performance by exploratory analysis of collaboration behavior while particularly enabling the comparison of triage training sessions.*



## CCS Concepts

• ***Human-centered computing*** → ***Visualization techniques; Information visualization;***


† e-mail: carina.liebers@uni-due.de

## 1. Introduction

At a mass casualty incident, rescuers have to quickly gain an overview and categorize the casualties regarding their medical condition. The result determines the priority of treatment and transport to the hospital. This process is known as *triage* [IM07]. Its correct and quick execution is vital, as delayed treatment might have massive consequences for the patient's life and requires an efficient collaboration of the health professionals [PBD20]. The triage process is practiced in field training, where the participating health professionals perform the triage on supernumeraries acting as patients. Multiple teams perform the same training scenario, which makes the sessions comparable and allows for an in-depth analysis of their performance and interactions.

We study this scenario from the perspective of social interaction research as an example of synchronous professional team collaboration in a semi-structured work process. A central question is how the team members collaborate and interact, which might influence the triage's outcome. To understand these effects, social interaction researchers use recorded training videos [PBD20]. However, a multitude of aspects need to be covered to describe collaborative behavior, like analyzing what the team members said, where they looked, and whom they interacted with. Social interaction researchers lack an overview to ease comparing and identifying interaction patterns of different teams. Although there exist tools supporting annotating videos [WBR*06], we are not aware of any approaches that visually analyze annotated social interactions across comparable sessions.

To support social interaction researchers in exploring the data and generating research insights, we propose VisCoMET, a visual analytics approach to investigate team collaboration and compare team performance in medical triage trainings. The approach visualizes the data on two levels in distinct views. First, as shown in Fig. 1B, a timeline visualization displays the events of a session per person and allows a detailed comparison with other selected sessions. The timeline provides means to identify frequent patterns in the event sequence data and to search for specific sequences (Fig. 1A). Second, an aggregated overview shows the individual training sessions using glyphs as nodes with rich details on team interactions, while displaying similarities between different sessions as links (Fig. 4B).

VisCoMET has been developed through an application study in close collaboration between visualization and social interaction researchers in an iterative process (three visualization researchers and two social interaction researchers, all being co-authors of this paper). In a jointly performed application example, we demonstrate that our approach not only allows analyzing the parallel events (simultaneous events in the context of a team) within one training session efficiently, but also supports the comparison of event patterns and overall team performance across different sessions. The approach was further evaluated by external visualization and social interaction experts, which helped refine the visual interface and confirmed the discoverability of relevant insights.

The supplemental materials contain details of the expert feedback and a video demonstrating the interactive use of the system. We plan to publicly showcase VisCoMET after acceptance.

## 2. Related Work

Our work relates to visualization of multi-stream empirical data, observing persons and their interactions with each other, with their work environment, or with a computer system. It is also an event sequence visualization, similar to timeline-based representations revealing interactions between parallel events, or aggregating and comparing collections of events.

**Visualization of Work Behavior and Collaboration.** VisCoMET visualizes collaborative work behavior, annotated and transcribed from videos. While video visualization generally summarizes video content for diverse applications [BCD*12], some approaches specifically visualize video annotations on timelines similar to our approach [BCNS15, BNFD16, FWHH11, JKKW17, Kip12]. Generally, it is common to use different rows for different persons, objects, actions, or other data streams, and colored markers in these row to indicate activity regarding the respective entity. Partly, these approaches are specialized for specific use cases. For instance, John et al. [JKKW17] visualize movies annotated with information on active characters, their activities, and used objects in a timeline, as well as support the comparison of different movies in juxtaposed timelines. Further works introduced approaches to annotate videos [BCNS15, BNFD16, JKKW17, Kip12], including manual, semi-automatic, or automatic methods. Moreover, there are qualitative research tools supporting video annotations like Atlas.ti [ATL23], NVivo [Alf22], and ELAN [WBR*06]. Here, we assume that the videos have already been annotated.

Observing users interacting with a computer system is also related, and we refer to surveys of visually analyzing interaction data [ED16, XOW*20] and eye movements [BKR*17] for an overview. Those approaches are most similar to ours that show multiple aspects of data on a timeline. For instance, Blascheck et al. [BJK*16] integrate eye-tracking data with interaction logs and think-aloud transcripts on session-specific timelines. Like in our approach, an event pattern search allows identifying situations of interest. Dou et al. [DJS*09] use WireVis [CGK*07] to explore interactions of users of a visual analytics system. Minelli et al. [MMLB14] visualize interactions of software developers across different views of their development environment. Instead of focusing on single sessions, reVISit [NWC*21] aggregates study participants across sessions. However, these approaches study interactions between a human and a computer, not of a team of professionals.

Some approaches consider synchronous collaboration scenarios. Targeting also a medical case, Echeverria et al. [EMMBS19] analyze teamwork in nursing training and generate a timeline representation of different sessions. They use networks to visualize communication among the nurses and patients. Agarwal et al. [AASB20] showcase a remote collaboration scenario where two users solve a puzzle game within a mixed-reality environment, visualized as on a timeline showing users, game objects, and their interactions. Similarly, MRAT [NSW*20] helps analyze multi-user mixed-reality sessions, visualizing events on a timeline and within the mixed-reality scene; the approach is applied to a simplified triage scenario. Weibel et al. [WFE*12] use ChronoViz [FWHH11] to study video, eye tracking data, communication, and other data recorded for two pilots in a plane. Studying team meetings, Yu et al. [YYA*10] present a workflow for automatically recognizing and visualizing social interactions on a timeline. These approaches, however, rarely go beyond the analysis of single sessions and do not facilitate comparing sessions through more than juxtaposed representations.

**Event Sequence Visualizations.** Regarding the design space of event visualizations by Guo et al. [GGJ*22], our visualization, as *data scale* dimension, considers a *sequence collection* (i.e., parallel sequences and multiple sessions); as *automated sequence analysis* dimension, it applies a simple form of *pattern discovery* (i.e., pattern search); as *visual representation* dimension, it is *timeline-based*; and as *interaction technique* dimension, it leverages *filter/query* and *alignment*. Regarding *analysis tasks*, our approach mostly focuses on *visual comparison*, supporting the *comparison of individual event sequences* (i.e., within one session) and the *comparison of sequence collections* (i.e., across multiple sessions). The survey [GGJ*22] also names *health informatics* as a main application area, but discusses mostly electronic health records. Regarding the design space of timelines [BLB*16], our main visualization is a *linear* representation with *chronological* scale in a *faceted* layout (i.e., showing different facets as rows)—a common combination.

Among timeline-based event visualizations, showing parallel streams of events and interactions among them is most related. For instance, LifeLines [PMR*96] visualizes biographies of persons, separating events into several rows. Similarly, ClinicalPath [LLP*22] uses several rows to displays an overview of patient's test results. Bombalytics [AWB20] shows matches in a bomb laying game; players and game objects form the rows of a timeline, connected if a player interacts with an object. Moreover, parallel execution traces in supercomputing can be shown on a multi-row timeline, with a row for every process and links between them for messages [IBJ*14]. Timeline Trees [BBD08] show parallel events on a timeline across a hierarchy of objects, while small thumbnails hint at interactions. Highlighting interactions, some dynamic graph visualizations plot links between nodes represented as rows of a timeline [vdEHBvW13]. Although these works also mark events and interactions on a timeline, they focus on events within one collection (i.e., session), not comparing multiple ones.

For a high-level comparison of event sequences collections using a graph-based overview, we find fewer related examples in the literature. Nguyen et al. [NTA*18] summarize different sessions as small bar charts indicating an anomaly score and other numeric attributes, then allowing to inspect session details in a timeline representation. Blascheck et al. [BSBE17] generate radial transition graphs that aggregate the gaze movements of a user between multiple areas of interest. Contrasting different sessions, these radial representations are used for difference graphs comparing two sessions and a matrix-based comparison showing high-level similarities among sessions. Although approaches exist to summarize or compare event sequences through common patterns (e.g., [CXR18, GXZ*17]), they are not suitable in our scenario as we deal with interlinked parallel streams of events.

## 3. Analysis Context

As background, we introduce the triage training and recorded data, summarize the status quo of the analysis process in social interaction research, and report gathered requirements for our approach.


**Figure 2:** *Anonymized frame of a training video with four cameras: (A) first-person camera with eye-tracking data of the leading doctor and (B) the notetaking person, (C, D) scene overviews.*

**Triage Training Videos.** At a mass casualty incident, first, a triage is performed to dispatch ambulance and pre-hospital care resources. The triage is an assessment process, determining the treatment and transport priority of injured persons [EK14]. It is executed by a team of an emergency physician and a paramedic. The physician assesses each patient's medical condition and leads the team (*leading doctor (LD)*) while the paramedic documents the results (*notetaking person (NP)*).

We study ten training sessions that apply the *ABCDE* triage process (Airway, Breathing, Circulation, Disability, and Exposure). It is part of the *mSTaRT* algorithm, which extends the commonly used *START* algorithm [BKS96] to prioritize injured persons [KHK*06]. The triage categories are assigned colors to express treatment's urgency: red (immediate assistance necessary), yellow (can be postponed), green (follow-up treatment), blue (unlikely survival), and black (dead patients). The training was performed at one location, with a scenario simulating thirteen patients in several vehicles being involved in an accident. The ground truth triage categories of patients are known to the organizers of the training and the teams were tasked to correctly triage all of them in a given timespan. The teams were free to choose the order of examining patients. The training sessions were recorded with four cameras: two for the first-person perspectives of team members and two for an overview of the entire scene (see Fig. 2). In addition, the first-person videos of one or both team members include eye-gaze data.

**Status Quo: Data Annotation and Analysis.** The goal of social interaction researchers is to discover problems and strategies in interpersonal interaction and communication. The status quo analysis process—as performed by the contributing social interaction researchers—is the following: First, they (often together) identify aspects of interest in the communication, interaction, and collaboration by watching video excerpts of the recorded videos. Next, the researchers iteratively transcribe, analyze, and annotate the recordings. An annotation scheme is established and annotations are systematically applied. During analysis, they switch between the annotations and the video while documenting their findings. This manual analysis becomes cumbersome for multiple videos, leading to a need for more support comparing the annotations across sessions and analyzing similarities between the teams.

The social interaction researchers annotate their videos with the annotation tool ELAN [HVUH*21, WBR*06]. Their annotations consist of a type, a time interval, the annotation text, and other descriptions. These are grouped in different channels with similar characteristics, called *tiers*. The triage scenario has thirteen tiers defined: Each *gaze*, *speech*, and two *activity* tiers per person (LD and NP), as well as *car*, *patient-ID*, *decision*, *triage true category*, and *process*. The *gaze* tiers contain annotations of where a person looked, the *activity* tiers annotations of the member's actions, and the *speech* tiers annotations that characterize the spoken phrases. Other tiers refer to the triage scenario: *car* lists annotations of vehicles, while *patient-ID* contains annotations of the patients referenced by an identifier; *decision* marks triage decisions, *triage true category* the ground truth for the patient classifications, and *process* the performed triage step.

**Requirements.** At the project's outset, the contributing visualization researchers interviewed the collaborating social interaction researchers to gather requirements, and prioritized the results (details of the process are in the supplementary material). The requirements cover visualizing individual sessions, comparing the information within and across sessions, and providing specific support for the analysis workflow.

> **R1**: Session Details
> The approach should visualize all details of a session: a) the video, b) the annotations and their properties, and c) the triage decisions in contrast to the ground truth.

> **R2**: Comparison
> The approach should enable comparison a) of parallel events (across tiers) in one session and b) similar temporal event patterns between multiple sessions, c) enable finding similarities and differences among the sessions, and d) convey an overview of teams' activities.

> **R3**: Workflow
> The approach should a) enable the look-up of annotations in the corresponding video, b) enable documenting findings, and c) support the definition and highlighting of interaction patterns between team members.

## 4. Visualization Approach

Our visual analytics approach, VisCoMET, is specifically tailored to social interaction researchers analyzing the medical training scenario described above. Its development included four extensive feedback meetings with the social interaction researchers, where they gave feedback on the current interface and further improvements. We designed VisCoMET with two levels of abstraction, one for inspecting details of selected sessions and another for getting an overview of all sessions. On the detail level, one or multiple sessions can be viewed on a timeline. The overview level displays a node-link graph diagram to compare all training sessions while visually summarizing session details. Hence, requirements regarding the visualization of session details (**R1**) are mostly met through the detail view, while the overview visualization particularly facilitates comparison (**R2**). However, the two views also interlace features for the respective other requirements. The detail view reflects most requirements regarding the workflow (**R3**). We implemented it as a web application using Vue.js [You] and D3.js [Bos].

### 4.1. Timeline-based Detail View

The detail view shown in Fig. 1 integrates one or several session timeline visualizations (Fig. 1B) with a panel for pattern search (Fig. 1A) and a legend (Fig. 1C). Each timeline visualization of a session displays the annotations chronologically as events (**R1b**, **R1c**) together with the corresponding video on the left (**R1a**). In case several sessions are selected for comparison (**R2b**), these visualizations are juxtaposed vertically.

**Layout.** On the vertical axis, per session, it shows twelve tiers as grouped rows: One group for each team member and, to communicate common actions prominently, one general group at the top concerning both members (Fig. 3A). A meaningful order for the rows in each group was determined jointly during development, for instance, listing *car*, *patient-ID*, and *decision* at the very top to provide an overview of the training process. Within a row, session time is shown on horizontal axis, annotations are displayed as rectangles; their position and width encode their start time and duration (Fig. 3B, **R1b**). A gray-to-black color scale is used to encode the annotation categories depending on their importance as specified by the social interaction researchers (the darker, the more important). For example, NP's main role is writing and LD's main activity is examining, indicated by the darkest colors in their tiers. Each annotation type is assigned an icon image. The icons are displayed inside the rectangle, as well as a label revealing its type if it is wide enough. Encoding two tiers in one row, we visualize a patient's ground truth classification (*true triage category*) in the *patient-ID*'s annotations with respective triage colors ■ ■ ■ ■ (**R1c**). Directly below, the team's triage decision is indicated by colored rectangles in the *decision* row, which supports comparing a triage decision to the ground truth (**R2a**).

**Interactions.** If multiple sessions are loaded, they are normalized to the same timescale. The timeline can be zoomed via mouse scrolling, and—using the additional space—more icons and labels become visible (Fig. 3D). To support comparing different time ranges across sessions (**R2b**) zooming is independent for each session. Hovering over an annotation reveals details in a tooltip (see Fig. 3C, **R1b**). A selection highlights all annotations of the same type throughout all sessions (**R2a, R2b**). Upon selection, the corresponding video jumps to the annotation's start to support the verification of an event (**R3a**). In case an annotation from a general row gets selected, like a patient case or a vehicle, an alignment button appears above (Fig. 3C). Clicking aligns the first occurrence of the same annotation type in all displayed sessions (Fig. 3C). A vertical line indicates the alignment. If a session has multiple annotations of the same type, arrows ← → enable jumping

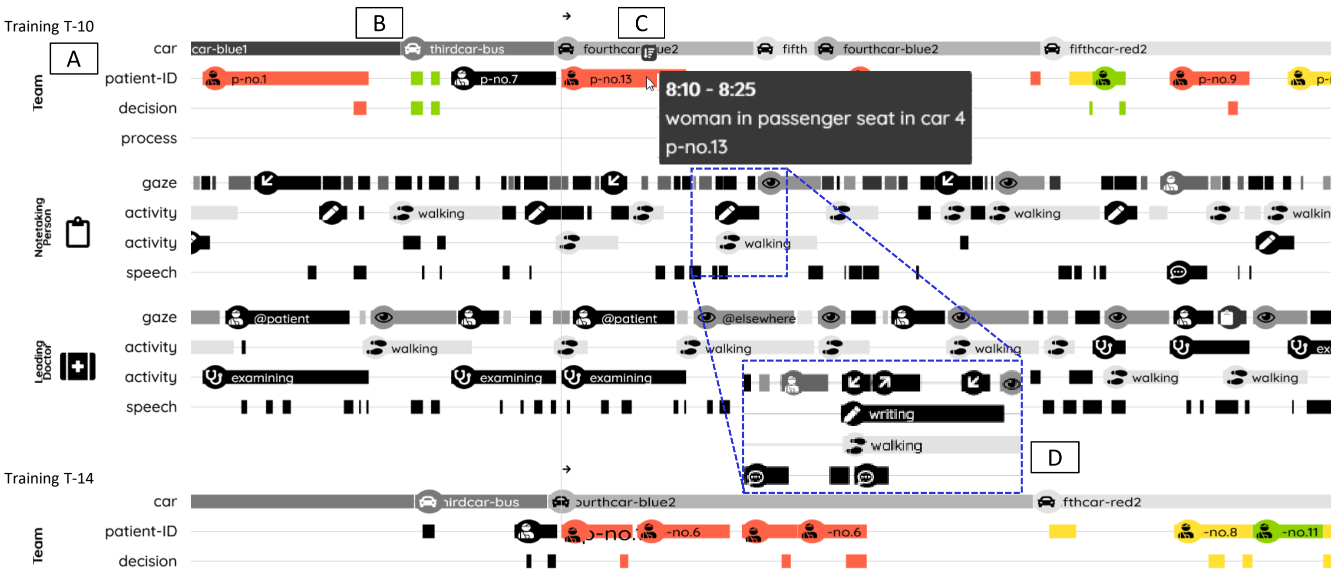


**Figure 3:** *Timeline visualization of a session (top) with (A) tier labels, (B) visually encoded annotations on a vertical timescale with (C) annotation details inside a tooltip and a button for alignment, and (D) a close-up of a section in the timeline. An excerpt of a second timeline visualization (bottom) is aligned with the event of visiting patient no. 13 above.*

to the previous or next occurrence. These features support comparing similar cases among the juxtaposed sessions (**R2a, R2b**). To reduce the vertical distance for comparison, users can temporarily hide sessions using checkboxes (Fig. 1, **R2b**). They can document their findings by selecting the note icon at the top (**R3b**); the notes are visually represented with a marker in the visualization.

**Event Patterns.** We define four types of event patterns, which cover important cases of parallelism across rows in our scenario: (i) consecutive events, (ii) partly overlapping events, (iii) events starting at the same time, and (iv) an event happening within the span of another event (not starting at the same time). Focusing on them allowed us to flatten the parallel events to a linear sequence per session by coding the cases (i, ii, iv) through symbols or generating merged events (iii). Using the flattened sequence, we apply the SPAM algorithm [AFGY02] for detecting all occurrences of a pattern (**R2a, R2b**).

The panel on the left (Fig. 1A) allows specifying a minimum and maximum pattern length and provides further options regarding the considered tiers for the pattern search and displayed patterns. At the bottom, a pattern list shows the most frequent patterns detected. For instance, the most frequent examples shown in Fig. 1 are sequences of two consecutive events. We map the frequencies to a scale from light blue (less frequent) to dark blue (more frequent). The color-coded frequency distribution is shown above the list. The frequencies of detected patterns are mapped to the annotations in the timeline visualization, encoding the most frequent sequence in the respective blue color (Fig. 1B, **R3c**). All patterns an annotation belongs to are shown in its tooltip. Hovering a pattern in the list enlarges all related annotations in the timelines (**R3c**). The panel includes a tile to define a pattern for search (Fig. 1D, **R3c**). It shows the same tiers as the timelines and encodes annotations alike. Dragging the generic annotation at the top-right and dropping it onto a row adds a default annotation type. The type can be changed by clicking on a button. Users can adapt the position and size of the rectangle to model the four cases of parallelism between each pair of annotation events. After searching, an answer is displayed.

### 4.2. Graph-based Overview

On start, VisCoMET presents a graph-based overview (Fig. 4) visualizing each session as a node and encoding similarities between the sessions as edges. Characteristics of each person and team behavior per session are depicted in a nested glyph visualization. The design serves as a visual fingerprint enabling the identification of different team behavior in sessions at a glance, and detailed comparison through interactions. The view also serves as the basis to select sessions for detailed inspection in the timeline visualization.

**Session Nodes.** Each node, as depicted in detail in Fig. 5, aggregates the team's activities of a session **(R2d)**. It consists of the outer session node and two inner team member sub-nodes. The outer one visualizes general team attributes, the inner sub-nodes display equivalent data individually for the leading doctor (*LD*, left) and notetaking person (*NP*, right). The outer node and inner sub-nodes employ rings divided into arcs, which encode a numeric value each. The values are min-max-normalized per ring across all sessions. The two outer rings of each session node summarize viewing of the patients and simultaneous activities. The outer gaze ring visualizes how often team members looked at the patient per triage case (sum of gaze events, Fig. 5A). The inner activity ring Fig. 5B, visualizes the number of co-occurrences of each pair of two activities: *walking & walking*, *walking & examining*,

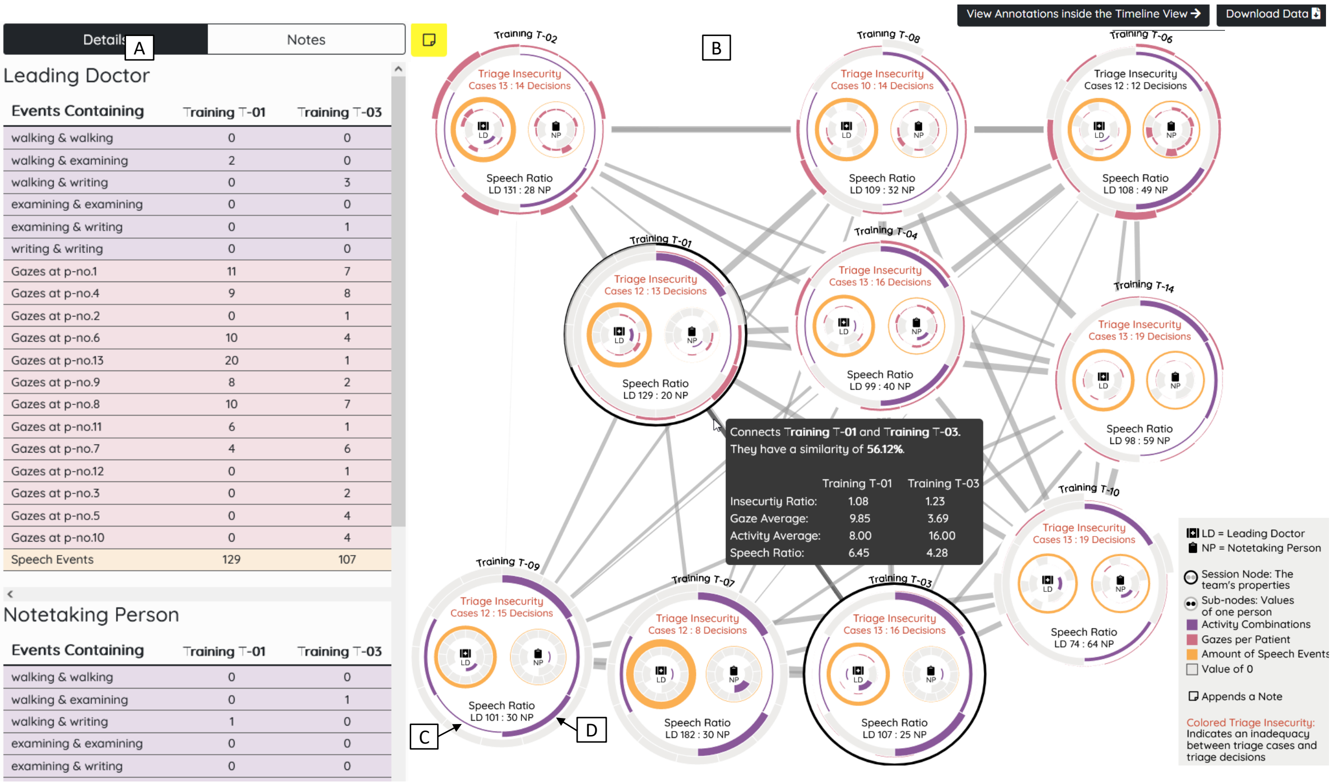


**Figure 4:** *Graph-based overview consisting of (A) a panel listing details of selected session nodes, (B) a node-link graph showing aggregated characteristics of sessions in the nodes and the similarities between sessions as links.*

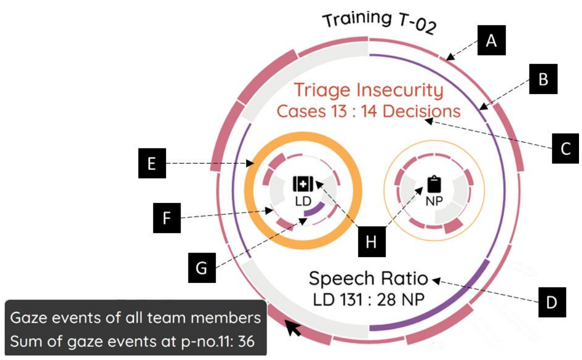


**Figure 5:** *A session node showing at the outer ring overall team characteristics: (A) gaze events per patient, (B) frequencies of parallel activity pairs, (C) the insecurity ratio, (D) the speech ratio. The inner sub-nodes per team member visualizes (E) speech events, (F) gaze events per patient, (G) number of parallel, paired activities, and (H) icon identifiers. On hover, a tooltip displays details.*

*walking & writing*, *examining & examining*, *examining & writing*, and *writing & writing*. The *triage insecurity* ratio (Fig. 5C) communicates how decisive the team was, contrasting the number of triage cases (patient cases) and the team's triage decisions. It is depicted in ■ red if the two numbers do not match—a deviation indicates that either the team did not completely fulfill the task or more triage decisions than cases indicates re-assessment. The *speech ratio* contrasts the number of speech event between the team members (Fig. 5D). Since a balance is not expected, the ratio is not color-coded. Nested in the session nodes, each team member sub-node contains the three rings visualizing ■ *speech* (Fig. 5E) in addition to ■ *gaze* (Fig. 5F) and ■ *activity* (Fig. 5G). The outer ring encodes the number of speech events for the assigned person in a single arc (ring). The other two are designed consistently with the respective rings in the outer session node. Only, co-occurrence of the same activity is impossible for one person and hence not shown here.

The initial idea of using glyphs for visualizing team properties was inspired by TargetVue [CSL*15]; the design was influenced by Guo et al. [GJC*19], CircleView [KSS04], and FluxFlow [ZCW*14] (usage of nested nodes and rings), as well as by CloudDet [XWY*19] and CircleView [KSS04] (using arcs).

**Similarity Edges and Layout.** The graph's edges encode the similarity between the training sessions (**R2c**). We consider as main characteristics of a session the ■ gaze distribution per patient, the ■ activity pairs, the insecurity ratio, and the speech ratio. To calculate the similarity between the nodes, we apply the Euclidean distance on vectors summarizing the main characteristics of the sessions. We construct each vector from the activity combinations sum, the gaze events average, and the relative values of the two ratios. The distances are normalized and inverted to represent similarity. To compute the graph layout, we use the D3.js implementation of a force-directed layout and adapt the forces to repel connected nodes according to their Euclidean distance.

**Details on Demand.** Hovering displays more details in a tooltip, for instance, for the arcs within the glyphs (Fig. 5). Selecting an edge opens details on the similarity between the two sessions. Selected nodes are emphasized through a black border, and tables in the detail panel display their values (Fig. 4A), where background colors correspond to the ones used in the glyphs. One table is dedicated to each team member and one to their collaboration. The columns display the selected training sessions to contrast characteristics between trainings (**R2b**). Sessions are added to the table by selection and removed on deselection. Users can compare the selected sessions in-depth in the timeline-based details view (Fig. 1) when clicking the respective button on the top-right (Fig. 4).

## 5. Application Example

To demonstrate how VisCoMET could support social interaction researchers, we conducted among the interdisciplinary team of authors a joint analysis of the entire dataset with annotated training sessions. Starting with the graph overview and then moving to the timelines, we present specific insights that were gained and match them with the initially elicited requirements.

**Examining Patients.** The scenario expects the leading doctor to examine the patients, which implies looking at them. In even numbered-session, both team members wore an eye-tracker, in odd-numbered sessions only the leading doctor. However, it was surprising to discover zero gaze events at any patient in sessions T-07 and T-09 (Fig. 4 bottom; **R2c**), which hints at recording issues. In the remaining sessions, only one team (T-02) looked at every patient (all pink-colored arcs are present in the outermost ring in Fig. 4 top-left). Among the other teams, some missed looking at many patients (e.g., teams in T-06, T-08, and T-14 skipped at least three patients (Fig. 4); **R2c**). Moreover, focusing on pink-colored arcs in sub-nodes, we find that, in three sessions (T-04, T-06, and T-08), the notetaking person looked at the majority of patients with a higher frequency than the leading doctor (**R2c**). The behavior might indicate that the notetaking person is unable to follow the expected vocal instructions from the leading doctor, demanding the notetaking person to frequently look at the patient for self-assessment.

**Team Roles.** As lead doctors decide the triage class, we assumed them to have more speech events than notetaking persons, who document the decisions. Almost all sessions matched this expectation, except for training T-10 (similar width of yellow circles in node T-10, as shown in Fig. 4B at the bottom right; **R2b**, **R3c**), where both spoke similarly frequent. A notetaking person talking a lot may indicate either a highly collaborative team sharing the roles or conflicting roles. To check whether unexpected behavior might have impacted the team's results, we inspected the triage insecurity ratio of the session, which also differed strongly from the optimum (Fig. 4, insecurity ratio 13:19 in node T-10; **R3c**).

**Re-classification of Patients.** Based on previous observations, we selected sessions T-02, T-08, and T-10 for comparison in the timeline view. From the *patient-ID* tier, we can observe the examination sequence. From Fig. 6, we read that, in training T-08, *patient no. 7* was examined twice (**R1c**). First, the team made two triage decisions (Fig. 6A); being maybe unsure about the classification, the team revisited the patient and clarified the decision (Fig. 6B). A similar re-examination occurred in training T-10 for *patient no. 8* (Fig. 6C and the tooltip; **R1c**). A gap in the patient tier and frequent speech event annotations after visiting the patient indicate that the team might have discussed the case in more detail (**R1b**).

**Ill-timed Documentation of Triage Decisions.** We expected the triage decisions to be documented at the end of examining a patient. However, this varied within sessions. For instance, in trainings T-02 and T-10, we saw no documentation of a decision for an examined patient (Fig. 6C, D; **R1b**, **R1c**). Documentation may have been missed or delayed by the notetaking person, requiring the team to come back for re-examination (Fig. 6C). Similarly, we observed that two triage decisions were made by the same team while examining *patient no. 13* (Fig. 3C; **R1c**). Due to an unusually increased number of speech events (**R1b**) by notetaking person around these decisions (Fig. 3, *NP-speech* in T-10), we watched the video playback at this timestep (**R1a**) and found that the notetaking person said *"too many red patients"* before documenting the decision. This statement was particularly interesting as it provides a concrete indication of why this team might come back to previous patients—they might be confused due to feeling pressured by not having the correct categories in their earlier decisions.

**Simultaneous Activities.** Regarding simultaneous activities that might influence the triage result, we searched for the pattern *writing intersected by walking* for the notetaking person (Fig. 1D; **R3c**) because we expected difficulties performing these two tasks together. It was found once in training T-02 and twice in training T-10. It occurred after inspecting the first patient who did not receive a triage decision in training T-02. In training T-10, this happened after the triage decision for patient no. 13 (Fig. 3D). Therefore, this activity combination might not influence the decision but concern the documentation of patient examination results. We concluded that the notetaking person probably uses the time between patient cases to document other medical properties. Furthermore, we observed that the notetaking person exhibited a complex combination of activities *walking & writing* four times in three training sessions, T-03, T-07, and T-10 (Fig. 4D; **R2a**). These were later found in the timeline visualization using the pattern search (Fig. 1D and Fig. 3D).

**Unusual Collaboration Patterns.** To investigate collaboration issues, we inspected parallel activities that we expected to be hindering the collaboration (**R2c**, **R2d**, **R3c**), for instance, activities not fitting the role or being difficult to perform in parallel. In training T-09 (Fig. 4C), we found an occurrence of both team members examining the patient (*examining & examining*), although it is a task exclusively assigned to the leading doctor (**R2d**). For such an instance, it would be interesting to perform a *qualitative single case analysis* and to explore how the team members jointly organize examining a patient, how the notetaking person takes this unusual interactional role, and how team members reach a mutual decision.

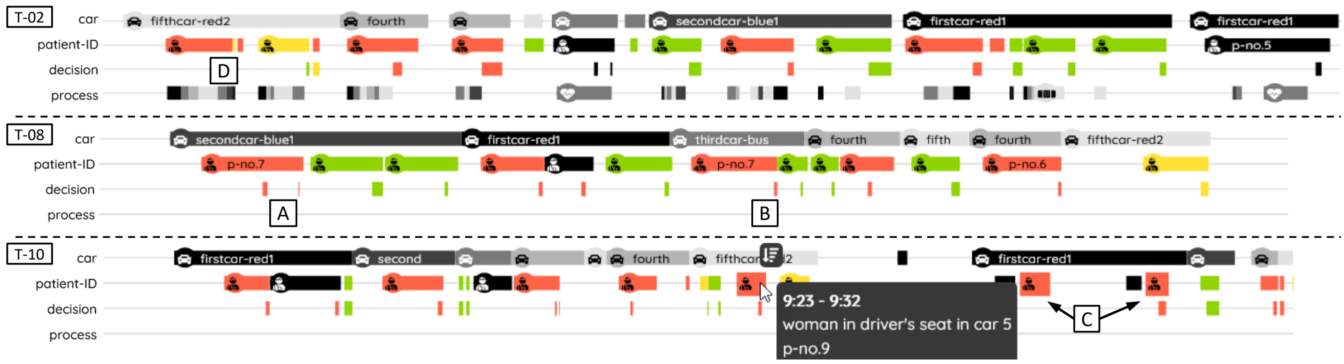


**Figure 6:** *Snippets of the timeline view comparing the first five tiers of training T-02, T-08, and T-10. (A) and (B) display incidents of re-examinations of patient no. 7, (C) shows a delayed triage decision at the same patient, and (D) an absent decision.*

## 6. Expert Feedback

We evaluated our approach with two groups of experts. In a first study with visualization experts, we investigated the quality of visual encodings. After adapting our system slightly with the feedback, we observed social interaction researchers, our target audience, using VisCoMET in a realistic analysis setting.

**Participants and Interview Design.** Five visualization experts (V1–V5) participated. Two held a PhD (V1, V2) and had experience in visualization research for more than four years. From the other three postgraduate students, V3 and V4 had worked for 3–4 years in the field, and V5 between 1–2 years. For the second study, we recruited three social interaction graduate students (S1–S3) doing social interaction research on average for 4,33 years (SD: 058); they were familiar with the used triage training dataset. Regarding their highest educational qualification, two stated to have a bachelor's, while one indicated a master's degree. All experts were recruited from our close professional network, but did not contribute to the development of the approach.

The studies were conducted in one-to-one remote interviews of about 60 minutes each. For visualization experts, the system was shared from the interviewer's machine and remotely controlled by the participants. Simulating the workflow of social interaction researchers analyzing the data either individually or in pairs, we conducted one interview with a single expert (S3) and another with two experts collaborating (S1, S2); the system was shared and controlled from the interviewer's machine, while participants stated the interactions to be performed by the interviewer.

For both expert groups, the interviewer presented a brief introduction to the triage training process and explained the annotated dataset. The visualization experts were asked to think-out-aloud to record the first-impressions on the visual encoding. The exploration was guided through questions in an online questionnaire using the graph and the timeline-based detail view. Then they were asked specific questions about visual encoding that worked well and weaknesses of the two views. Further, they should rate different features in the approach on a 7-point Likert scale. Lastly, the experts could provide other scenarios where the presented visualization could be useful, as well as their estimation of the system's usability. Since we wanted to simulate the workflow of the social interaction researchers, we did not impose any structure in their analysis process during the interview. The interviewer facilitated exploratory analysis and answered clarification questions when asked. Study artifacts of the visualization expert interviews, the introductory presentation, questionnaire, the experts' responses, and transcripts of the social interaction expert interviews are submitted as supplementary material.

For the interviews with visualization experts, we used the dataset of two annotated (T-01 and T-02) triage training sessions and anonymized the corresponding videos. Since the social interaction researchers had the required data permissions, we used fourteen annotated training sessions without needing to anonymize the videos.

**Feedback of Visualization Experts.** All experts agreed that the timeline successfully displays the annotations, while enabling their comparison and a detailed analysis of the team members' interactions. Experts said the event durations and their sequences are shown well (V5), enabling the comparison of multiple sessions (V2). All rated the usage of the icons representing activities to be intuitive, and agreed that color-coding patients based on their triage category supports comparison. In contrast, the gray color scale encoding importance of annotations received mixed feedback, with two experts (V1, V2) either *disagreed* or *slightly disagreed* that it supports differentiating between activities. Two experts (V2, V3) expressed difficulty interacting with thin rectangles representing short events. Two experts (V1, V5) suggested clearer highlighting events by enlarging the corresponding rectangles. Experts mentioned that the pattern panel was difficult to understand (V2) and too complicated for non-domain experts (V4).

In the graph overview, two experts (V2, V3) liked the encoding in session nodes' representing team's activities, while three (V1, V2, V3) appreciated the design of sub-nodes for encoding behavior of each team member. The encoding of the attributes as segments in different rings was positively emphasized (V5) as it aids comparison between team members (V1, V4). Additionally, V5 found the design of the graph useful to encode pairwise session similarities as edge thickness in the graph. Three experts (V1, V2, V5) found the tooltips necessary to estimate the similarity. Two experts (V1, V2) found the dense information encoding in a node unintuitive and difficult to remember. V1 further highlighted that the circular design might also represent a time-varying attribute (clock metaphor) and could therefore be confusing. V3 remarked that it is difficult to remember the rings' encodings due to the different number of rings representing activity of a team (two outermost rings in session node) and a team member (three rings in sub-node).

**Feedback of Social Interaction Experts.** At first, the experts had problems understanding the graph overview, overwhelmed by information-rich encodings (S1) and deviating initial expectations (S2). However, after an explanation and using the view, they appreciated its features to compare training sessions and to *"see what's exciting"* in the dataset (S1). They stated that, while it displays characteristics and might support the evaluation, it requires an understanding of the training sessions as it does not explain why the cases are similar or different (S1, S2). S3 found the overview particularly useful at the beginning of the analysis to identify research questions as it displays similar and different cases and emphasized VisCoMET as time-saving. All assessed the graph overview to be especially useful as it enables selections due to the made annotations (S1), offers more insights and knowledge through statistical values (S2), enables finding first impressions on what would be interesting, and new analysis options through the systematic visualization (S3). They finally confirmed the understandability and usefulness of the view: *"Once it is explained to you, you can understand it[...], despite the existing complexity"* (S1); *"The longer I look at it, the more I can actually start to make comparisons"* (S2).

During exploration, the experts used the rings and arcs within the session nodes to compare team activities. S1 and S2 focused on gaze events and noticed some unfilled arcs. They proposed to look into the videos to understand if these are founded in missing data or training situations (S1). Searching for different speech behavior between teams, S3 identified T10 as training with balanced speech events, T1 as similar, and T14 as the most different training for comparison, using the speech arcs. Regarding the insecurity ratio, S1 and S2 mentioned it can be influenced by many aspects, such as a patient condition change during the training (S3, S2), the given examination time per patient (S1), reassessments (S2), and unclarity in the annotation process as events might not be clearly defined (S1, S2). Thus they stated that the ratio can be off *"when a scene becomes too chaotic and the team has to react to the new events"* (S1) and might also *"reflect the [...] speed of reaction"* in a decision (S2). However, they concluded that it indicates the smoothness of a team's process (S2) and might hint at where to look (S1).

Compared to the graph overview, the timeline view did not raise many initial questions, probably since it resembles the known ELAN tool. Mostly its comparative features were emphasized by the experts. They stated to easily perceive the time span and event types of the annotations (S2), where a triage process accelerates and the process was well done (S2), and compare the order of events to recognize systematics in the annotations with the tool (S1). S1 liked jumping between re-classification events exploring when they were annotated, while S3 emphasized the usage for searching for relevant passages: *"I find super interesting [... the] comparison and picking out [...] sequences, which I [want to] look at deeper."*

The experts gained noteworthy insights with the timeline view. For instance, S3 compared the speech behavior of the team members within the pre-selected training session and noticed that the notetaking person in training T-10 gets increasingly talkative, while the leading doctor had mostly constant speech events. As no decision was made in such a talkative passage, the expert hypothesized that they might discuss organizational topics, whereby holes in their dialogue might be a third, not annotated person speaking. The video embedding was found helpful, as it prevents searching for passages through its linkage (S2). This procedure was often used to verify hypotheses and understand passages identified as conspicuous (S1, S2, S3). Searching for similar cases using an activity sequence, where the notetaking person looked at the sheet and did not talk while the leading doctor was engaged with the patient, was successful in trainings T-1 and T-14. All researchers found repetitive patient cases in their inspection of the training session. They discussed reasons that might lead to such repetitions, like reassessments and the arrangement of cars, thus directing a team twice to one vehicle when the team *"couldn't get to the other side properly, or they couldn't open the car from one side"* (S3). However, the experts found repetitive patient cases interesting, as they could hint at difficult patient cases when multiple teams show such conspicuity (S1). S3 suggested using the alignment between different sessions to investigate re-examinations at particular cars.

Overall, the social interaction experts proposed several use cases of VisCoMET in their work. As a first analysis step, it could be used to *"revise the annotations"* as it offers an overview of the annotations (S1) to find weak points of annotation characteristics and be able to redefine them (S2). Afterward, VisCoMET might support *collaborative data analysis* when reasoning about uncertainties in the annotation process (S2), or comparing similar and different sequences (S3). When doing a *conversation analysis*, S1 would use it to find peculiarities in the data. Aside from enhancing current analysis methods, S1 highlighted that, as it provides statistics and comparison functionalities, *case collections* can be created. The expert suggested using the tool to find similar cases, analyze them, and build a system from the gained insights: *If you have already annotated cases, you can actually recognize such systematics.*

## 6.1. Limitations

Our two studies included a limited number of experts. However, analyzing their responses already showed signs of repeating arguments, which is a sign of saturation (i.e., that the limited number did not affect the analysis results). The experts were recruited through our professional network, which might have influenced them to reply more positively to the study. We tried to compensate by specifically asking them to highlight the limitations of the approach.

# 7. Discussion and Future Work

Based on concrete insights from the application example and expert feedback, we discuss the approach along with ideas for future work.

**Learning Curve.** Despite some visualization experts were concerned that the graph overview might be difficult to interpret, the social interaction experts who used VisCoMET quickly familiarized with and correctly read the visualization. We conclude that, while the overview may initially seem overwhelming, the view becomes useful after short time, and the learning curve is generally moderate. Despite limitations to read the concrete values, the glyphs act as visual fingerprints that allow fast recognition of similarities and differences in trainings, as in Keim et al. [KO07].

**Scalability.** When larger data sets are used, the approach's scalability may become a limiting factor. Due to limited vertical space and the amount of tiers to be displayed, the timeline visualization is only designed to compare a few selected sessions (max. 4) at a time. However, filtering the tiers or aggregated representation of tiers would save space and allow comparing more sessions. Still, the timeline does not enable reordering tiers, so users may have to compare similar rows over a distance. However, a timeline representation integrating multiple sessions is difficult to design as an appropriate time mapping across sessions is challenging and might depend on the specifics of the use case. In the graph view, typical force-directed layouts will not easily scale beyond about 50 nodes. In our case, this is further reduced to approximately 20 nodes because the nodes contain space-consuming details that would otherwise get too small or nodes would overlap. Visualizing teams of more than two persons as a session node is possible (potentially up to five members). However, bigger teams would lead to the nodes growing larger. Both challenges can be partly addressed by interactive selection and filtering of attributes to be shown in the nodes.

**Interactive Pattern Recommendation.** When analyzing social interactions, a major challenge is defining concrete constraints for meaningful event sequences. To not exceed reasonable session duration in our studies, we have not yet tested this empirically and cannot comment on the specific performance of our pattern search. However, displaying all existing patterns, even when focusing on the most frequent ones, leads to a multitude of sequences, not all representing meaningful social interactions. Hence, the approach could benefit from interactive navigation of the discovered patterns, e.g., through a node-link visualization in CoreFlow [LKD*17]. However, automatic pattern search only finds a narrowly defined set of patterns. Hence, our approach could be extended to include user-defined ratings of patterns regarding their meaningfulness and to making pattern recommendations based on these ratings.

**Levels of Annotation.** Since there already is an established and efficient process for annotating, the scope of this work has been set on visualizing already annotated videos. Still, it might be worthwhile in the future to extend the scope and support the annotation process with multiple levels of annotation. For instance, a first rough annotating pass might generate hypothesis and identify relevant research questions, followed by one or multiple fine-grained analyses regarding different aspects of social interactions. The approach would then need to visually discern these annotation levels and, ideally, support adapting and refining the annotations within the system. This might result in a better visualization support throughout different phases of a social interaction research project.

**Insights and Analysis Tasks.** The specific insights from the application example and the expert feedback provide a clearer view of relevant analysis tasks for VisCoMET and related approaches: First, *analyzing diverse and multimodal interactions*, such as, between humans (e.g., conversation) or with the environment (e.g., triage sheet), to understand collaboration in a dynamic scenario. Second, *investigating the timing* of performed actions in collaborative decisions. The timing analysis looks into coordination within a team (e.g., late documentation of a triage result), which might reveal insights on team performance (e.g., re-visiting patients due to confusion). Third, *comparing the member roles* of different teams. The analysis helps understand the differences in the coordination strategies among different teams. For instance, notetaking persons speaking more than others in some teams, although their task is to document results. Insights into such differences could help the teams undergoing training to improve their collaaboration strategy.

**Generalizability.** Our approach is customized to a specific triage scenario and the workflow of the social interaction researchers. Although it cannot be directly applied to analyze other collaborative scenarios, various elements of our approach can be easily reused. These might serve as a basis to develop a framework for visualizing synchronous social interactions. Reusable concepts are the definition of tiers being assigned to different teams members, or tiers containing specific annotation types. Furthermore, annotations should be always linked to a video, while our pattern detection will work analogously for other annotation types and tiers. Although the general concept of a graph overview of sessions is reusable, the specification of session characteristics to be shown in the nodes and the metric for calculating session similarities are scenario-dependent.

## 8. Conclusions

Teamwork in triage trainings is challenging to study as decisions are made in a short time considering many factors, and the interactions among the health professionals might relate to different, not all explicitly expressed aspects. To understand these interactions, analyzing the annotated videos requires comparison among parallel events within one session, as well as comparison across different teams and sessions. Our approach supports facilitating this comparison, aside from making the annotated events of the sessions efficiently readable. Its timeline view is designed to relate the parallel events, align them with similar events in juxtaposed timelines of other sessions, and find patterns of interest effectively. The graph-based overview complements this on a higher level of abstractions, giving quick insights about the characteristics of each session through using glyphs in its nodes, as well as similarities among them. VisCoMET has been useful to find diverse collaborative behaviors, anomalies, deviations in the teams' performances, and to identify related patterns. We found evidence that the system can support social interaction research in several ways: To revise made annotations, as visualization and comparison system in a collaborative data analysis session, and to identify case collections of similar interaction cases. With this work, we have made a step towards facilitating social interaction research to compare multiple sessions efficiently instead of only analyzing single sessions.

## Acknowledgments

This work is partly funded by *Mercator Research Center Ruhr* (*MERCUR*; project: “Vergleichende Analyse dynamischer Netzwerkstrukturen im Zusammenspiel statistischer und visueller

Methoden"). Open Access funding enabled and organized by Projekt DEAL.

## References


[AASB20] AGARWAL S., AUDA J., SCHNEEGASS S., BECK F.: A design and application space for visualizing user sessions of virtual and mixed reality environments. In *Vision, Modeling, and Visualization* (2020), The Eurographics Association. doi:10.2312/vmv.20201194. 2

[AFGY02] AYRES J., FLANNICK J., GEHRKE J., YIU T.: Sequential PAttern Mining using A Bitmap Representation. In *Proceedings of the eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2002), ACM, pp. 429–435. doi:10.1145/775047.775109. 5

[Alf22] ALFASOFT GMBH: NVivo. Qualitative Datenanalyse leicht gemacht, Nov 2022. URL: https://www.nvivo.de/. 2

[ATL23] ATLAS.TI SCIENTIFIC SOFTWARE DEVELOPMENT GMBH: Die #1 software für qualitative Datenanalysen, Feb 2023. URL: https://atlasti.com/de. 2

[AWB20] AGARWAL S., WALLNER G., BECK F.: Bombalytics: Visualization of competition and collaboration strategies of players in a bomb laying game. *Computer Graphics Forum 39* (2020), 89–100. doi:10.1111/cgf.13965. 3

[BBD08] BURCH M., BECK F., DIEHL S.: Timeline Trees: Visualizing sequences of transactions in information hierarchies. In *Proceedings of the 9th International Working Conference on Advanced Visual Interfaces* (2008), ACM, pp. 75–82. doi:10.1145/1385569.1385584. 3

[BCD*12] BORGO R., CHEN M., DAUBNEY B., GRUNDY E., HEIDEMANN G., HÖFERLIN B., HÖFERLIN M., LEITTE H., WEISKOPF D., XIE X.: State of the Art Report on Video-Based Graphics and Video Visualization. *Computer Graphics Forum 31*, 8 (2012), 2450–2477. doi:10.1111/j.1467-8659.2012.03158.x. 2

[BCNS15] BIANCO S., CIOCCA G., NAPOLETANO P., SCHETTINI R.: An Interactive Tool for Manual, Semi-Automatic and Automatic Video Annotation. *Computer Vision and Image Understanding 131* (2015), 88–99. doi:10.1016/j.cviu.2014.06.015. 2

[BJK*16] BLASCHECK T., JOHN M., KURZHALS K., KOCH S., ERTL T.: VA$^2$: A Visual Analytics Approach for Evaluating Visual Analytics Applications. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (2016), 61–70. doi:10.1109/tvcg.2015.2467871. 2

[BKR*17] BLASCHECK T., KURZHALS K., RASCHKE M., BURCH M., WEISKOPF D., ERTL T.: Visualization of eye tracking data: A taxonomy and survey. *Computer Graphics Forum 36*, 8 (2017), 260–284. doi:10.1111/cgf.13079. 2

[BKS96] BENSON M., KOENIG K. L., SCHULTZ C. H.: Disaster triage: Start, then save—a new method of dynamic triage for victims of a catastrophic earthquake. *Prehospital and disaster medicine 11*, 2 (1996), 117–124. 3

[BLB*16] BREHMER M., LEE B., BACH B., RICHE N. H., MUNZNER T.: Timelines revisited: A design space and considerations for expressive storytelling. *IEEE Transactions on Visualization and Computer Graphics 23*, 9 (2016), 2151–2164. doi:10.1109/tvcg.2016.2614803. 3

[BNFD16] BIRESAW T. A., NAWAZ T., FERRYMAN J., DELL A. I.: ViTBAT: Video tracking and behavior annotation tool. In *Proceedings of the 13th IEEE International Conference on Advanced Video and Signal Based Surveillance* (2016), IEEE, pp. 295–301. doi:10.1109/AVSS.2016.7738055. 2

[Bos] BOSTOCK M.: D3 Data-Driven Documents. https://d3js.org. 4

[BSBE17] BLASCHECK T., SCHWEIZER M., BECK F., ERTL T.: Visual Comparison of Eye Movement Patterns. *Computer Graphics Forum 36*, 3 (2017), 87–97. doi:10.1111/cgf.13170. 3

[CGK*07] CHANG R., GHONIEM M., KOSARA R., RIBARSKY W., YANG J., SUMA E., ZIEMKIEWICZ C., KERN D., SUDJIANTO A.: WireVis: Visualization of categorical, time-varying data from financial transactions. In *Proceedings of the 2007 IEEE Symposium on Visual Analytics Science and Technology* (2007), IEEE, pp. 155–162. doi:10.1109/vast.2007.4389009. 2

[CSL*15] CAO N., SHI C., LIN S., LU J., LIN Y.-R., LIN C.-Y.: TargetVue: Visual Analysis of Anomalous User Behaviors in Online Communication Systems. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (2015), 280–289. doi:10.1109/TVCG.2015.2467196. 6

[CXR18] CHEN Y., XU P., REN L.: Sequence Synopsis: Optimize Visual Summary of Temporal Event Data. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2018), 45–55. doi:10.1109/TVCG.2017.2745083. 3

[DJS*09] DOU W., JEONG D. H., STUKES F., RIBARSKY W., LIPFORD H. R., CHANG R.: Recovering reasoning processes from user interactions. *IEEE Computer Graphics and Applications 29*, 3 (2009), 52–61. doi:10.1109/mcg.2009.49. 2

[ED16] ELTAYEBY O., DOU W.: 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* (2016), ACM, pp. 62–69. doi:10.1145/2993901.2993912. 2

[EK14] ELLEBRECHT N., KAUFMANN S.: Boosting efficiency through the use of IT?: Reconfiguring the management of mass casualty incidents in Germany. *International Journal of Information Systems for Crisis Response and Management 6*, 4 (2014), 1–18. doi:10.4018/IJISCRAM.2014100101. 3

[EMMBS19] ECHEVERRIA V., MARTINEZ-MALDONADO R., BUCKINGHAM SHUM S.: Towards Collaboration Translucence: Giving Meaning to Multimodal Group Data. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), ACM, pp. 1–16. doi:10.1145/3290605.3300269. 2

[FWHH11] FOUSE A., WEIBEL N., HUTCHINS E., HOLLAN J. D.: ChronoViz: A System for Supporting Navigation of Time-Coded Data. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems* (2011), ACM, p. 299–304. doi:10.1145/1979742.1979706. 2

[GGJ*22] GUO Y., GUO S., JIN Z., KAUL S., GOTZ D., CAO N.: Survey on Visual Analysis of Event Sequence Data. *IEEE Transactions on Visualization and Computer Graphics 28*, 12 (2022), 5091–5112. doi:10.1109/TVCG.2021.3100413. 3

[GJC*19] GUO S., JIN Z., CHEN Q., GOTZ D., ZHA H., CAO N.: Visual Anomaly Detection in Event Sequence Data. In *Proceedings of the 2019 IEEE International Conference on Big Data* (2019), IEEE, pp. 1125–1130. doi:10.1109/BigData47090.2019.9005687. 6

[GXZ*17] GUO S., XU K., ZHAO R., GOTZ D., ZHA H., CAO N.: EventThread: Visual Summarization and Stage Analysis of Event Sequence Data. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2017), 56–65. doi:10.1109/TVCG.2017.2745320. 3

[HVUH*21] HELLWIG B., VAN UYTVANCK D., HULSBOSCH M., HULSBOSCH M., TACCHETTI M., SOMASUNDARAM A., GEERTS J., SLOETJES H.: *ELAN – Linguistic Annotator*, 6.2 ed. Max Planck Institute for Psycholinguistics, 2021. https://www.mpi.nl/corpus/manuals/manual-elan.pdf, last accessed: 4 June 2023. 4

[IBJ*14] ISAACS K. E., BREMER P.-T., JUSUFI I., GAMBLIN T., BHATELE A., SCHULZ M., HAMANN B.: Combing the communication hairball: Visualizing large-scale parallel execution traces using logical time. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 2349–2358. doi:10.1109/TVCG.2014.2346456. 3

[IM07] ISERSON K. V., MOSKOP J. C.: Triage in Medicine, Part I: Concept, History, and Types. *Annals of Emergency Medicine 49*, 3 (2007), 275–281. doi:https://doi.org/10.1016/j.annemergmed.2006.05.019. 2

[JKKW17] JOHN M., KURZHALS K., KOCH S., WEISKOPF D.: A Visual Analytics Approach for Semantic Multi-Video Annotation. In *Proceedings of the 2nd Workshop on Visualization for the Digital Humanities* (2017). https://www.informatik.uni-leipzig.de/~stjaenicke/vis4dh/Vis4DH2017_paper_5.pdf. 2

[KHK*06] KANZ K. G., HORNBURGER P., KAY M., MUTSCHLER W., SCHÄUBLE W.: mSTaRT-Algorithmus für Sichtung, Behandlung und Transport bei einem Massenanfall von Verletzten. *Notfall+ Rettungsmedizin 9*, 3 (2006), 264–270. doi:10.1007/s10049-006-0821-3. 3

[Kip12] KIPP M.: Multimedia Annotation, Querying, and Analysis in Anvil. In *Multimedia Information Extraction*. Wiley, 2012, pp. 351–367. doi:10.1002/9781118219546.ch21. 2

[KO07] KEIM D. A., OELKE D.: Literature Fingerprinting: A New Method for Visual Literary. *IEEE Symposium on Visual Analytics Science and Technology* (2007), 115–122. doi:10.1109/VAST.2007.4389004. 10

[KSS04] KEIM D. A., SCHNEIDEWIND J., SIPS M.: CircleView – A New Approach for Visualizing Time-related Multidimensional Data Sets. In *Proceedings of the 2004 Working Conference on Advanced Visual Interfaces* (2004), ACM, pp. 179–182. doi:10.1145/989863.989891. 6

[LKD*17] LIU Z., KERR B., DONTCHEVA M., GROVER J., HOFFMAN M., WILSON A.: CoreFlow: Extracting and Visualizing Branching Patterns from Event Sequences. *Computer Graphics Forum 36*, 3 (2017), 527–538. doi:https://doi.org/10.1111/cgf.13208. 10

[LLP*22] LINHARES C. D. G., LIMA D. M., PONCIANO J. R., OLIVATTO M. M., GUTIERREZ M. A., POCO J., TRAINA C., TRAINA A. J. M.: ClinicalPath: a Visualization tool to Improve the Evaluation of Electronic Health Records in Clinical Decision-Making. *IEEE transactions on visualization and computer graphics* (2022), 1. doi:10.1109/TVCG.2022.3175626. 3

[MMLB14] MINELLI R., MOCCI A., LANZA M., BARACCHI L.: Visualizing Developer Interactions. In *In Proceedings of the Second IEEE Working Conference on Software Visualization* (2014), IEEE, pp. 147–156. doi:10.1109/VISSOFT.2014.31. 2

[NSW*20] NEBELING M., SPEICHER M., WANG X., RAJARAM S., HALL B. D., XIE Z., RAISTRICK A. R. E., AEBERSOLD M., HAPP E. G., WANG J., SUN Y., ZHANG L., RAMSIER L. E., KULKARNI R.: MRAT: The Mixed Reality Analytics Toolkit. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020), ACM, p. 1–12. doi:10.1145/3313831.3376330. 2

[NTA*18] NGUYEN P. H., TURKAY C., ANDRIENKO G., ANDRIENKO N., THONNARD O., ZOUAOUI J.: Understanding User Behaviour through Action Sequences: From the Usual to the Unusual. *IEEE Transactions on Visualization and Computer Graphics 25*, 9 (2018), 2838–2852. doi:10.1109/TVCG.2018.2859969. 3

[NWC*21] NOBRE C., WOOTTON D., CUTLER Z., HARRISON L., PFISTER H., LEX A.: reVISit: Looking under the Hood of Interactive Visualization Studies. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (2021), ACM. doi:10.1145/3411764.3445382. 2

[PBD20] PITSCH K., BACHMANN P., DUDDA M.: 'Triage' in Mass Casualty as Situated Interaction. Algorithm and Participation. In *Proceedings of 18th European Conference on Computer-Supported Cooperative Work – Exploratory Papers, Reports of the European Society for Socially Embedded Technologies* (2020), EUSSET. doi:10.18420/ecscw2020_p05. 2

[PMR*96] PLAISANT C., MILASH B., ROSE A., WIDOFF S., SHNEIDERMAN B.: LifeLines: visualizing personal histories. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (1996), ACM, pp. 221–227. doi:10.1145/238386.238493. 3

[vdEHBvW13] VAN DEN ELZEN S., HOLTEN D., BLAAS J., VAN WIJK J. J.: Dynamic network visualization with extended massive sequence views. *IEEE Transactions on Visualization and Computer Graphics 20*, 8 (2013), 1087–1099. doi:10.1109/tvcg.2013.263. 3

[WBR*06] WITTENBURG P., BRUGMAN H., RUSSEL A., KLASSMANN A., SLOETJES H.: ELAN: A Professional Framework for Multimodality Research. In *Proceedings of the 5th International Conference on Language Resources and Evaluation* (2006), pp. 1556–1559. 2, 4

[WFE*12] WEIBEL N., FOUSE A., EMMENEGGER C., KIMMICH S., HUTCHINS E.: Let's look at the cockpit: Exploring mobile eye-tracking for observational research on the flight deck. In *Proceedings of the Symposium on Eye Tracking Research and Applications* (2012), ACM, pp. 107–114. doi:10.1145/2168556.2168573. 2

[XOW*20] XU K., OTTLEY A., WALCHSHOFER C., STREIT M., CHANG R., WENSKOVITCH J.: Survey on the analysis of user interactions and visualization provenance. *Computer Graphics Forum 39*, 3 (2020), 757–783. doi:10.1111/cgf.14035. 2

[XWY*19] XU K., WANG Y., YANG L., WANG Y., QIAO B., QIN S., XU Y., ZHANG H., QU H.: CloudDet: Interactive Visual Analysis of Anomalous Performances in Cloud Computing Systems. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2019), 1107–1117. doi:10.1109/TVCG.2019.2934613. 6

[You] YOU E.: Vue.js. The Progressive JavaScript Framework . https://vuejs.org. 4

[YYA*10] YU Z., YU Z., AOYAMA H., OZEKI M., NAKAMURA Y.: Capture, recognition, and visualization of human semantic interactions in meetings. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications* (2010), IEEE, pp. 107–115. doi:10.1109/PERCOM.2010.5466987. 2

[ZCW*14] ZHAO J., CAO N., WEN Z., SONG Y., LIN Y.-R., COLLINS C.: #FluxFlow: Visual Analysis of Anomalous Information Spreading on Social Media. *IEEE Transactions on Visualization and Computer Graphics 20*, 12 (2014), 1773–1782. doi:10.1109/TVCG.2014.2346922. 6