

TimeLighting: Guided Exploration of 2D Temporal Network Projections

Velitchko Filipov¹, Davide Ceneda¹, Daniel Archambault², Alessio Arleo^{1,3}

Abstract—In temporal (*event-based*) networks, time is a continuous axis, with real-valued time coordinates for each node and edge. Computing a layout for such graphs means embedding the node trajectories and edge surfaces over time in a $2D+t$ space, known as the space-time cube. Currently, these space-time cube layouts are visualized through animation or by slicing the cube at regular intervals. However, both techniques present problems such as below-average performance on tasks as well as loss of precision and difficulties in selecting timeslice intervals. In this paper, we present *TimeLighting*, a novel visual analytics approach to visualize and explore temporal graphs embedded in the space-time cube. Our interactive approach highlights node trajectories and their movement over time, visualizes node “aging”, and provides guidance to support users during exploration by indicating interesting time intervals (“when”) and network elements (“where”) are located for a detail-oriented investigation. This combined focus helps to gain deeper insights into the temporal network’s underlying behavior. We assess the utility and efficacy of our approach through two case studies and qualitative expert evaluation. The results demonstrate how *TimeLighting* supports identifying temporal patterns, extracting insights from nodes with high activity, and guiding the exploration and analysis process.

Index Terms—Human-centered computing—Visualization—Graph drawings, Empirical studies in visualization

1 INTRODUCTION

TEMPORAL (or *event-based*) networks [2] are dynamic graphs where the temporal dynamics, such as node and edge additions and removals, have real-time valued coordinates. These have been characterized and studied extensively [3], as they are used in many applications to model phenomena of commercial and academic interest, such as interactions in social media [2], communication networks [4], and contact tracing [5], to name a few. In contrast to traditional dynamic graph drawing [3], [6], where time is discretized (or *timesliced*), creating a visualization for temporal networks poses different challenges. Juxtaposition (or small multiples) presents the discrete timeslices in a side-by-side manner and requires identifying suitable timeslices, which inevitably leads to quantization errors. Such quantization errors obscure the fine temporal details that might be crucial in some domains (e.g., the exact order of personal contacts in contact tracing networks). Animation would not suffer from such artifacts, and it has been used in previous work on event-based graph drawing as a visualization metaphor to display the computed layouts [7], [8]. Animation is a more natural way to encode time; however, it is not perceptually effective for many tasks involving dynamic networks [9], [10]. In particular, animation falls short when comparing different timeslices (or states of the network) and

the changes between distant timeslices (i.e., the first and last time slices) as it requires the viewer to memorize these. Moreover, the vast majority of research on animation has been done with timesliced graphs (see, e.g., [6], [9], [10], [11]), and its application to temporal networks is still in its infancy.

Temporal networks can also be drawn in a 3D space, namely, the space-time cube ($2D+t$) [12], [13]. In this case, a drawing algorithm computes the node trajectories over time and space. Existing research [7], [8], [14] provides evidence that this drawing approach yields better quality drawings of temporal graphs, compared to their timesliced counterparts (e.g., Visone [15]), and for discrete-time graphs when many changes occur between timeslices. Despite this, research on visually depicting these trajectories and obtaining insights from their behavior is still an under-investigated topic.

On these premises, we present *TimeLighting*, a guidance-enhanced Visual Analytics (VA) solution for exploring node trajectories in the space-time cube keeping the *full* temporal resolution of the network. Our approach supports understanding temporal patterns and behaviors and, in general, extracting insights from datasets with complex temporal dynamics. *TimeLighting* has several key differences compared to existing methods. Unlike traditional dynamic graph drawing techniques that discretize time and often result in quantization errors, our approach maintains the full temporal resolution, preserving fine temporal details crucial in domains like contact tracing. Additionally, we integrate multiple elements to guide and ease the exploration process, highlighting key intervals for further inspection. We further present a movement score, based on the length of each node’s trajectory over time that helps to rank and identify the more and less stable parts of the graph, facilitating a better understanding of the nodes’ dynamics.

The design of *TimeLighting* is inspired by the “time-

• ¹ Centre for Visual Analytics Science and Technology (CVASt), TU Wien, Vienna (AT). E-mail: {name.surname}@tuwien.ac.at.

• ² School of Computing, Newcastle University, Newcastle (UK). E-mail: daniel.archambault@newcastle.ac.uk.

• ³ Visualization Cluster, Eindhoven University of Technology, Eindhoven (NL). E-mail: a.arleo@tue.nl.

• A preliminary version of this research has been presented at the International Symposium on Graph Drawing and Network Visualization (GD) 2023 [1], selected by the Program Chairs, and invited for publication in TVCG. Part of the text has been reused in this publication.




Fig. 1: The TimeLighting metaphor. Rays of light (dashed lines), travel from $t = 0$ through the space-time cube and reach the observer at t_{max} . They interact with nodes' trajectories and carry this information to the projection plane.

coloring" operation [13], where time is mapped to color to visualize the evolution of nodes and edges through the space-time cube; and is loosely motivated by *transfer functions* used in direct volume rendering [16] to emphasize features of interest in the data. Conceptually, in our approach, we "shine" light through the space-time cube down along its time axis (hence the system's name – TimeLighting) in a manner that resembles the behavior of transfer functions for volume rendering. As the light interacts with the node trajectories, they are visualized and colored differently according to the node's age (i.e., applying the time-coloring operation), generating a 2D visualization of the 3D embedding (see also Fig. 1). The resulting visualization is an explorable 2D map of the nodes' densities, with visible individual node movement and "aging" over time. We complement this visualization with several interactive controls to explore the data and introduce a movement score to rank the nodes based on the length of their trajectory. We designed TimeLighting introducing multiple elements of visual guidance [17] to enhance (and possibly ease) the network exploration process. Finally, we describe two case studies demonstrating how our guidance-enhanced approach supports users in achieving the system design tasks. In summary, our key contributions are:

- A novel visual analytics approach, TimeLighting, to visualize and explore temporal graphs embedded in the space-time cube (see Section 4).
- Two case studies presenting the utility of TimeLighting to extract insights from temporal graphs and how the guidance-enhanced features support this (see Section 5).
- A two-fold expert evaluation of our proposed approach and the supplementing guidance features provided (see Section 6).

2 RELATED WORK

Visualization of Dynamic Networks. The typical definition of a dynamic network implies the notion of *timeslicing*: the temporal dynamics are represented as a sequence of snapshots (the *timeslices*), each one representing the status of the graph over an interval of time, placed at equally

spaced moments [6]. Several algorithms have been proposed for drawing dynamic networks by using timeslices [18]. Branke [19] presents an overview of methods in dynamic graph drawing in the early days of the field. Beck et al. [6] broadly characterize techniques into timeline (or "time-to-space") and animation (or "time-to-time") [20]. Kerracher et al. [21] introduce a two-dimensional design space of dynamic network visualization along two categories: graph structural and graph temporal encodings. Filipov et al. [22], [23], [24] explore this design space by comparing adjacency matrices and node-link diagrams, with different temporal encodings (i.e., juxtaposition, superimposition, animation, and timeline). The qualitative study suggests the combination of animation and node-link representation as the better-performing option for the tasks tested.

Farrugia and Quigley [9] measure the performance of animation and timeline dynamic network visualization on node-link diagrams. They find that timeline representations were faster than animation and preferred by users. Archambault et al. [10] found that timeline approaches were faster with no difference in error rate for most tasks, but some evidence for animation in helping with appearance tasks. Archambault and Purchase [11] looked at the mental map's impact in low-stability settings (i.e., significant changes between subsequent timeslices) for node-link dynamic graph drawings by comparing animation and timeline. However, small multiples were consistently faster than animation with no significant difference in error rate for most tasks in this series of experiments.

For matrix representations, Bach et al. explore visualizing dynamic networks by stacking chronologically ordered adjacency matrices, effectively using the third dimension as a time axis [25]. Windhager et al. [26] explore the use of space-time cubes to provide an overview of multifaceted data over time.

Many approaches for online dynamic graph drawing have been proposed where a single timeslice is modified online with no knowledge of future changes [27], [28], [29], [30]. However, we consider offline dynamic graph drawing whereby the algorithm has full access to the information across the entire time period of interest during drawing.

All the above techniques use the timeslice as a basis for visualizing the dynamic network. Visualizing temporal networks differs significantly from how typically we draw and visualize dynamic graphs in the graph drawing and visualization community as no such sequence of timeslices necessarily exists. The problem with timeslicing is that many networks of scientific interest do not have natural timeslices, rather edges and nodes come in with their own specific time coordinates. Therefore, choosing a regular sampling and duration often results in the loss of temporal information. In response, visualization techniques have been proposed for interesting timeslice selection. Wang et al. [31] present a technique for non-uniform time slicing (that is, selecting slices of different duration) based on histogram equalization to produce timeslices with the same number of events. Lee et al. [32], [33] experimented with methods for interactive time slicing on large touch displays.

Algorithms now exist that draw temporal networks directly in the space-time cube [7], [8]. These approaches do not impose timeslices on the temporal network but directly

embed the network in the space-time cube. However, the only way to visualize such $2D + t$ drawings on a 2D plane was to select timeslices or present an animation of the data over time. The work we present here bridges this gap for temporal network visualization.

Guidance for Graph Exploration Due to the challenge of analyzing complex events such as those modeled by temporal networks, researchers also investigated approaches to provide support and ease the analysis for users. The resulting approaches fall under the definition of “guidance” [17]. Guidance is characterized as *active* support in response to a *knowledge gap* which hinders the completion of an interactive visual data analysis session. Over the years, several approaches have been devised, providing different types of guidance [34], [35]. May et al. [36] describe a method to enhance the exploration of large graphs using glyphs. While the user explores a given area of interest (the focus), the system automatically highlights the path to other possibly off-screen interesting nodes (the context). Gladisch et al. [37] provide support during the navigation through large hierarchical graphs by suggesting what to explore next. Despite the work in this area, applying guidance to temporal networks is uncharted territory. Given a temporal network drawn in a space-time cube, our goal is to provide guidance to support the identification of interesting time intervals and nodes requiring further attention and analysis.

3 DESIGN CONSIDERATIONS

In this section, we discuss the most relevant aspects that influenced the design of `TimeLighting`, namely, the data characteristics, the users and tasks [38], and the time-coloring paradigm used to provide guidance.

3.1 Data, Users, Tasks.

Data. `TimeLighting` aims to support the visualization, analysis, and exploration of interval temporal network data [2]. In a temporal network $D = (V, E)$, V is the set of nodes, E is the set of edges. Each node and edge comes with time-dependant (expressed as $t \in \mathbb{R}$ in the following) *attributes*. These take the form of functions in the $V \times t$ and $E \times t$ domains for nodes and edges, respectively. For simplicity, and in accordance with existing literature [7], [8], we consider all attribute functions as piece-wise linear functions. The *appearance* A_x attribute, for example, models the intervals in time in which nodes and edges exist:

$$\begin{aligned} A_v &: V \times t \rightarrow [\text{true}, \text{false}] \\ A_e &: E \times t \rightarrow [\text{true}, \text{false}] \end{aligned}$$

A_v and A_e map to the node and edge insertion and deletion *events*, respectively. For this reason, these graphs are also called *event-based* networks, and the terms, temporal and event-based, will be used interchangeably in the remainder of this paper. The *position* $P_v : V \times t \rightarrow \mathbb{R}^2$ attribute describes the nodes’ position over time. The following is an example of how each node’s ($v \in V$) position is computed by an event-based layout algorithm describing the movement over time and space:

$$P_v(t) = \begin{cases} (5, 6) \rightarrow (12, 11) & \text{for } t \in [0, 1] \\ (8, 3) \rightarrow (1, 7) & \text{for } t \in [5, 7] \\ \dots & \\ (0, 0) & \text{otherwise} \end{cases}$$

Therefore, each node draws a *trajectory* in time and space. While a timesliced drawing would require computing the node coordinates at each point in time, in a temporal layout these trajectories must be computed and optimized in the three dimensions of the space-time cube. It is a complex and computationally intensive operation: in this paper, we use the `MultiDynNoS` [7] event-based layout algorithm to generate the drawings of the temporal graphs.

Users. Temporal graph drawing is a difficult concept to comprehend and put into practice. With `TimeLighting` we attempt to lower the bar of the necessary expertise to deal with these drawings by introducing guidance (see Section 3.3). However, in our system, we still target users with expertise in graph drawing. We, therefore, primarily aim to support the target audience of dynamic network visualization experts and support them in the analysis of temporal graphs. We further consider their feedback and discuss directions for future work on how to adapt the system to support novice users in extracting insights from temporal graphs.

Tasks. We derive `TimeLighting` tasks from the task taxonomy for network evolution analysis by Ahn et al. [39] and the taxonomy of operations on the space-time cube by Bach et al. [13]. The specific tasks are discussed in the following, along with a short description and justification.

T1 - Overview. The system should provide an overview of the temporal information at a glance. Providing an overview is typically the first necessary step in any VA process. This task considers the entire graph as an entity of the analysis [39] and is related to the orthogonal “time” flattening operation [13].

T2 - Tracking Events. Understanding the temporal dynamics and the events’ frequency helps the user isolate interesting occurrences in time. Events also cause the nodes’ trajectories to bend, i.e., make the node change direction. Understanding the *shape* [39] of changes in nodes’ movement over time (e.g., speed, repetition, etc.) would provide further insights during the exploration of the data.

T3 - Investigate relationships. Each edge occurrence perturbs the trajectories. Identifying which relationships have the most impact or how often they occur might help explain the formation of clusters, or, in general, the phenomenon at hand. This task aligns with the “point extraction” task described in Bach et al.’s taxonomy [13] since we are selecting an individual node and observing how it behaves over time (i.e., changing neighborhood over time).

3.2 A 2D Projection of a Temporal Graph

To highlight interesting nodes and trajectories and more generally to visualize the network’s dynamics, we took inspiration from the time-coloring operation described by Bach et al. [13] in their survey. Time-coloring is a content transformation operation applied to the space-time cube. Given a timesliced graph, the procedure consists in coloring

each timeslice based on a uniform linear color scale to facilitate the identification of the “age” of the data points.

To produce our 2D projection of a temporal graph, we took inspiration from the time coloring operation. However, to extend this technique to temporal graphs, we need to find a way to visually represent two features of each node: their *movement* (i.e., its behavior during its appearance intervals) and their *aging*. This means analyzing how their movements are distributed in time from the point of view of the observer, in a continuous time-axis scenario.

To solve this problem, we propose the following solution. If we projected light through the cube along the time axis, the interaction between the light and the node trajectories will be visible from the observer, watching from the other side of the cube (at t_{max} , see Fig. 1). The edges of the trajectories closer to the light will appear brighter, with the side closer to the observer will appear darker, reflecting their aging process respectively (older—transparent; younger—opaque). Concerning the node’s movement, we perform *time flattening* [13]. The results are 2D trajectories that “narrate” both the age and movement of the nodes over time. Inevitably, some temporal information is obfuscated in a 2D projection: however, we integrate it in the visualization through auxiliary views, interactions, and *guidance* [17] to support the users during exploration.


3.3 Guidance

In addition to visualizing the temporal network, we designed guidance to support its exploration and analysis. In general, the degree of support provided to the user may vary significantly, and at least three guidance degrees can be identified [17]:

- *Orienting* helps users keep an overview of the problem and alternative analytical paths they can choose to proceed;
- *Directing* guidance provides users with a set of options and orders them according to their importance for solving current tasks;
- *Prescribing* guidance (as the name suggests) prescribes a series of actions to take to conclude the task.

Given the similarities between our problem and the transfer functions used in volume rendering (although we apply them to a 2D visualization), the general idea at the base of our guidance-enhanced approach is to support and ease the identification of time intervals and nodes with specific desirable characteristics. The aim is to investigate their relationships and how they interact (i.e., analyzing the moment of the trajectories), making them stand out from the rest for the user’s convenience. Considering our set of design tasks, we highlight the nodes that have the longest trajectories. Long trajectories represent nodes that are associated with many events and with high movement in the currently selected temporal interval (to support T2). This type of guidance can be classified as directing, as trajectories are ranked based on their length, and, is necessary to ease or solve the system design tasks. In addition, we highlight temporal intervals in which nodes defined by the user interact with each other, thus providing guidance to T3.

4 TIMELIGHTING

In this section, we describe our system in detail and how we implemented it considering the design requirements described in the previous section. *TimeLighting* is a guidance-enhanced VA system comprised of two linked views and a focus+context approach. An overview of the prototype is shown in Fig. 2. For more details and resources we refer to the online supplementary material [40]. The source code is available on  GitHub.

4.1 Main View

In the main view of *TimeLighting* (see Fig. 2-B), we show a 2D projection of the complete temporal graph (i.e., an overview—T1). We discuss the details of our approach in the following. We employ several different encodings to highlight features of nodes and edges.

Node Positions are represented as trajectories encoded as a trail of circles (see Fig. 3). First, we place each interval’s start and end position within every node’s P_v attribute (see Section 3). These come directly from the computed drawing and have an orange stroke to make them distinguishable from the sampled nodes. To ease the comprehension of the movement flow (T2), between the start and end positions of each interval, we place several sampled nodes, where the coordinates are interpolated. The user can fine-tune the number of interpolated positions by choosing an appropriate “sampling frequency”. The resulting sampled nodes are positioned along the node trajectory but are encoded as smaller circles with no stroke to differentiate them from non-sampled nodes. We calculate and visualize the node *aging* as follows: for each node visualized on screen (nodes both in P_v and interpolated ones are considered), its age is computed as the difference between its time coordinate and the time of the node’s first appearance. We use a linear opacity scale to visually represent the node’s aging process. This encoding makes it easier to understand the information about the progression and movement of each node over time, providing an overview of the evolution of the network. We use relative aging in this context as it is focused on the individual node’s trajectory. Hovering over the node makes it possible to see its position in the timeline (in the context of the full temporal extent of the network) as a yellow bar, corresponding to its time coordinate (see Fig. 2-C). Typically, nodes are visualized in gray. However, users can change the visual appearance of the nodes to reflect the cumulative amount of movement (see Section 3.3—*Guidance*). Activating this type of guidance changes the coloring from gray-scale to a continuous color scale making nodes with higher movement visually distinct from the more stationary ones (T2/T3).

Edges are represented as solid straight lines that connect pairs of nodes belonging to two distinct trajectories. Edges have a pivotal role because these interactions eventually cause node movements and the creation of temporal clusters (T3). Edges might also appear or disappear within P_v movement intervals. This justifies the choice of introducing sampled nodes: edges can appear between all the points belonging to a trajectory, including interpolated ones. This allows the user to keep an overview of the finer temporal details, as we can display edges closest to their exact time

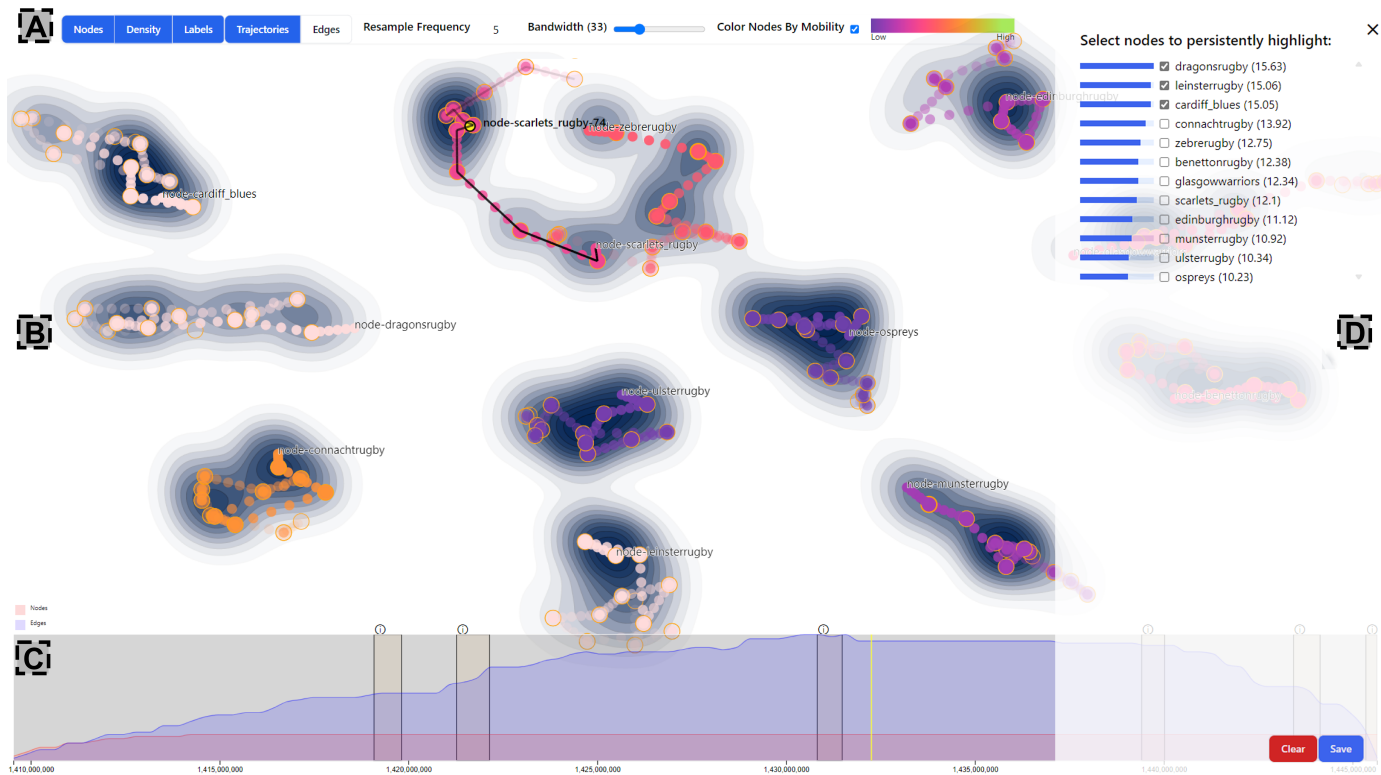


Fig. 2: TimeLighting overview. The view is comprised of the (A) toolbar and sidebar, (B) main view, (C) event timeline, and (D) the side bar depicting the movement scores. The yellow bar in the timeline shows the absolute age of the hovered node (visible in the top-center area).

coordinates. In turn, this could generate visual clutter as each edge is shown once for every node pair between each trajectory, and, depending on the sampling frequency, a trajectory can be comprised of several nodes. We mitigate this by showing edges on-demand and related to one trajectory at a time similar to an ego-centric approach (the selected one). The user, by hovering on any node belonging to a trajectory, will make all edges incident to those nodes appear. Edge aging is encoded similarly to node aging. Fig. 6-B illustrates an example of how edges connect the sampled nodes within the current temporal selection for the currently hovered node (Munster Rugby).

Movement is visualized using a polyline connecting each position in the nodes' P_v attribute. It represents how (and how much) the node's position changes over time due to the bends in the trajectories computed by the layout process. To clarify, the **movement** score of a node is calculated as the total distance traveled by the node over the time period for which it is present, this is the sum of the Euclidean distances between consecutive positions of the node's P_v attribute. This score directly correlates with the trajectory length; hence, longer trajectories indicate higher movement scores.

This movement (trajectory) encoding supplements the nodes' opacity as opacity alone might not be sufficient in showing how trajectories evolve (T2/T3). We calculate the age of each trajectory segment and apply a linear opacity scale (as with the nodes and edges). The trajectory's age is calculated as the mean of the ages of each pair of nodes in that given polyline segment. This information is also shown on-demand by hovering over a node, and, either the edges

or the movement of a node can be seen (top bar selection— see Fig. 2-A).

Density is represented as a contour map (see the dark-blue areas in the center of Fig. 2), providing a quick visual indicator that emphasizes locations where a larger number of nodes have existed (T2). This kind of encoding also provides a first glance at the trajectories' "shape", which eases keeping an overview of the events (T1). To calculate the density map, we translate the original set of nodes for each point in time into a series of objects with x, y coordinates and relative age. The x, y coordinates determine the contours of the density map. The age acts as a weighing function such that older nodes would contribute less to the density map compared to more recent nodes. The "bandwidth" sets the standard deviation of the Gaussian kernel, with lower values showing a sharper picture and higher values more distributed, but also more blurred, representation.

4.2 Other Views and Guidance

Top Bar. At the top of the prototype (see Fig. 2-A), a set of controls allows users to configure parameters and the appearance of the visualization. All individual encodings can be toggled on or off (nodes, labels, and density map).

Furthermore, depending on the specific analysis or exploration task, the nodes' trajectories or edges between them can be toggled as visible on mouseover. Additionally, we have a configurable re-sampling frequency that will recompute the nodes and their coordinates and update the visualization accordingly when changed. A modifiable kernel size allows the user to fine-tune the the contour

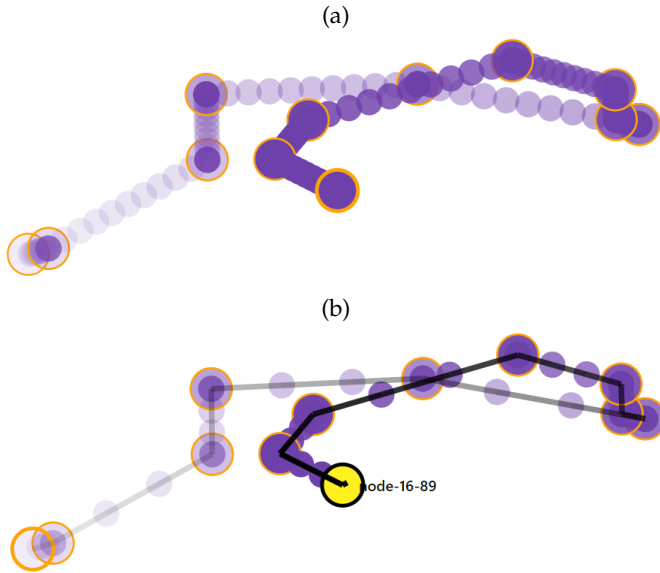


Fig. 3: Trajectories: (a) a higher sampling frequency is selected and aging is visible thanks to the change in opacity; (b) sampling is lowered (3 points per segment) and the trajectory is superimposed on mousing over the yellow node. Orange stroke and size difference indicate coordinates coming from the data compared to interpolated nodes.

density map. Finally, there is a checkbox that controls the state of the node coloring. By default, the nodes are colored gray and their opacity represents their age. By toggling the checkbox the coloring is updated to reflect their movement over time using a continuous color scale, with opacity still representing their age.

Side Panel shows a list of nodes ordered according to their movement scores (see Fig. 2-D). This ordering acts as guidance (see Section 3) by highlighting nodes that had many interactions with other nodes and that could, therefore, be considered interesting. Each node in this list is also accompanied by a small bar chart visualizing the differences in the movement scores. From this panel, the user can select and “lock” trajectories in the main view. A locked trajectory is always shown regardless of the current temporal interval selection in the timeline described in the following. The purpose of locked trajectories is to support users with a selection mechanism to focus on and examine specific node behaviors in detail over time, even as other aspects of the network and visualization change. This is particularly useful for tracking the movement and interactions of the locked (selected) trajectories without losing their context in the overall network. For example, if a researcher is analyzing a social network to understand the behavior of highly mobile individuals, they might lock the trajectories of these key nodes. This allows them to observe how these individuals interact with others over time, how relationships are formed and dissolved, and how their movement influences the rest of the network. By locking these trajectories, the researcher can maintain a continuous overview of these nodes’ activities, providing insights into patterns and anomalies that might be missed if the trajectories were not consistently visible. Locked nodes are colored in bright red




Fig. 4: Example of a locked trajectory. The circles in red are fully within the temporal selection of the users, whereas the less saturated ones are outside of the selected interval.

if they are in the current temporal selection, while nodes out of the current temporal selection are colored in a less saturated hue (see Fig. 4). The encoding and ordering of node trajectories in the side panel serve as visual guidance to support the tracking of specific events (i.e., guidance to ease T2). Additionally, when loading a new graph, the three nodes with the highest movement are locked by default (this can later be refined or changed).

Timeline, shown in (Fig. 2-C), allows users to select and explore specific temporal intervals as well as keep an overview (T1) of the number of nodes and edges that are visible over time. This overview is obtained by considering the number of nodes and edges for which the appearance attribute (A_v and A_e) is *true*. We visually encode this information as two overlapping area charts showing the changes occurring to the network as it evolves (in red for the nodes and blue for the edges).

The timeline serves two purposes: first, the user can brush to select a specific interval within the available data. As a result, only the subgraph existing during the newly selected time interval will be shown in the main view (temporal filtering). This also affects movement coloring, relative age, and density calculation, as well as, limits the edges shown on-screen to those existing in the current selection (these do not apply to locked trajectories). Second, **TimeLighting** uses the timeline to provide guidance and suggest specific time intervals for further inspection. Specifically, the system highlights intervals in time when all the currently locked trajectories interact with each other. These intervals are represented as orange rectangles drawn on top of the timeline (see Fig. 2-C). Clicking one of these intervals will snap onto that temporal selection, helping the user to keep track of and investigate relationships (i.e., guidance to ease T3).

4.3 Interactions

Highlighting: The individual nodes in the main view can be interacted with by mouseover, which highlights the currently selected node (yellow) as well as showing either the trajectory or edges to adjacent nodes depending on the current configuration. Furthermore, when mouseovering a specific node its current location in time will be highlighted in the timeline using a yellow bar (see Fig. 2-C).

Zoom & Pan: The main view also supports common interactions such as using the mouse scroll wheel to zoom and rescale the main view and panning or dragging the main view to reposition it. This interaction is intended to be used when only subsets of the network in the main view should be investigated in more detail.

Sidebar Guidance: In the sidebar, we sort all nodes by their movement score in a descending manner and lock the top

three. This default selection is intended to be a type of guidance to point to the more interesting nodes in a data set (i.e., nodes with high movement). The supplementing bar chart representations for each node in the list are also designed to support this idea. In cases where a larger selection of nodes (more than three) is of more interest, the bar charts can be used to suggest which nodes should be selected (i.e., higher movement scores).

Timeline Guidance: The timeline can be used to identify intervals where a large number of changes to the number of nodes or edges are occurring. These are considered very dynamic intervals in the evolution of the network. The area charts showing the number of edge and node events over the entire period of time serve as a visual indicator to identify interesting intervals (i.e., high event frequency). Furthermore, the orange rectangles (see Fig. 2-C) that are superimposed on the timeline guide specific intervals where all of the locked trajectories exist at the same time, supporting a more detailed analysis of the current selection.

5 CASE STUDIES

In this section, we discuss two case studies on two real temporal networks. We show how insights can be extracted from the data using *TimeLighting* and how the design tasks are achieved and supported with guidance. We build our case studies on the *Rugby* [41] and *Reality Mining* [4] datasets and are described in the following.

5.1 Case Study 1: *Rugby* Dataset

The *Rugby* dataset is a collection of 3151 tweets posted during the Pro12 rugby competition of the 2014-2015 season [41], specifically from September 2014 to October 2015. The network has a node for each team participating in the competition (12 teams in total), and an edge exists between two teams when a tweet from one mentions another. While nodes will stay visible from the moment they appear until the end, edges appear at the exact moment the tweet was posted. To improve the visibility of the edges during the layout process (as tweets do not have a “duration”), edges are given a 24-hour duration. For example, if an edge \bar{e} has a timestamp t , then $A_{\bar{e}} = [t - 12h, t + 12h]$. Multiple edges between the same teams are merged if their appearances overlap by a duration of less than one day. This simplification has already been applied in previous work using this dataset [8], and discretization of this dataset with similar resolution would require 417 timeslices. This dataset is particularly interesting as we have ground truth data to validate our findings.

The First and Second Half of the Season: Examining Trajectories. We begin our use case with an *Overview* task (T1), examining the trajectories in the first (see Fig. 5-A) and the second half of the 2014-2015 season (see Fig. 5-B). We can immediately observe from the timeline the trend of the events. There is a steady increase in the number of tweets from the beginning of the season that peaks around the beginning of the second round of the league. This peak remains until the season’s final and significantly impacts the nodes’ movement. We can also see that nodes move less in the second half of the season compared to the first half.

Specifically, in Fig. 5-B the majority of the nodes are within the purple to orange range of the scale—lower movement—whereas in Fig. 5-A they are in the yellow to green range—encoding higher movement. In the first half, instead, tweets are sparser, meaning that the influence of an edge on the movement of nodes is kept (as there is no inertia) until another one changes its trajectory. Continuing the analysis of the network, tweets happen at a much higher rate in the second half of the season, and, since this network is a clique (all teams eventually play against one another), they tend to be “locked” in place by the attractive forces exerted by the other nodes. It must be considered that the layout algorithm attempts to optimize (and reduce) node movement, placing the nodes in an area of the plane where they will likely remain. This behavior can be seen in the density map too, where hot spots are larger and more numerous (i.e., nodes tend to linger more in the same areas) in the second half compared to the first half. Nonetheless, the amount of attractive force will depend on the public interest (i.e., the number of tweets) about individual matches.

Tracking the Two Least Winning Teams. In this second case study, we *track* (T2) the trajectories and *investigate* (T3) the relationships between the two least winning teams of the season (according to the historical information available), namely “Zebre” (Z) and “Benetton” (B). We begin by selecting them in the sidebar so that the whole trajectory is locked permanently on the screen. Guidance shows us the different moments in time when the two teams interact, and we focus on the period when the two teams play against each other around the midpoint of the season. The teams played two matches (during the first and second leg of the tournament) in adjacent rounds (12 and 13). The status of the interface is reported in Fig. 6-C. It is possible to see how the Z trajectory bends significantly towards B at this point, and a similar effect is visible the other way around. This attraction strength can be interpreted as the “hype” of the matches building up, as Z and B are the only two Italian teams in the competition. Finally, we compare the relationships between the last two teams in the ranking and the first two, “Glasgow Warriors” (G), the winners, and “Munster Rugby” (M), referring to the time around the tournament finals (see Fig. 6-D). If we focus on M, it is easy to identify the time the final was played, with mostly all teams connected to it. The B trajectory is strongly influenced by M, as it was one of the final matches before the final; Z, instead, is not largely influenced and drifts away.

5.2 Case Study 2: *Reality Mining* Dataset

In this second case study, we focus on the *Reality Mining* dataset. The data comes from a study [4], where the mobile phones of the study participants were equipped with logging software. The study took place between September 2004 and June 2005, with the participants being students and staff from a major university. The complete dataset includes a wide variety of events, such as short messages, voice calls, local Bluetooth connections, and association with cell towers. For our case study, we only consider voice calls and text messages. Therefore, our nodes are the identifiers of the individual contacts (as they are referred to in the original

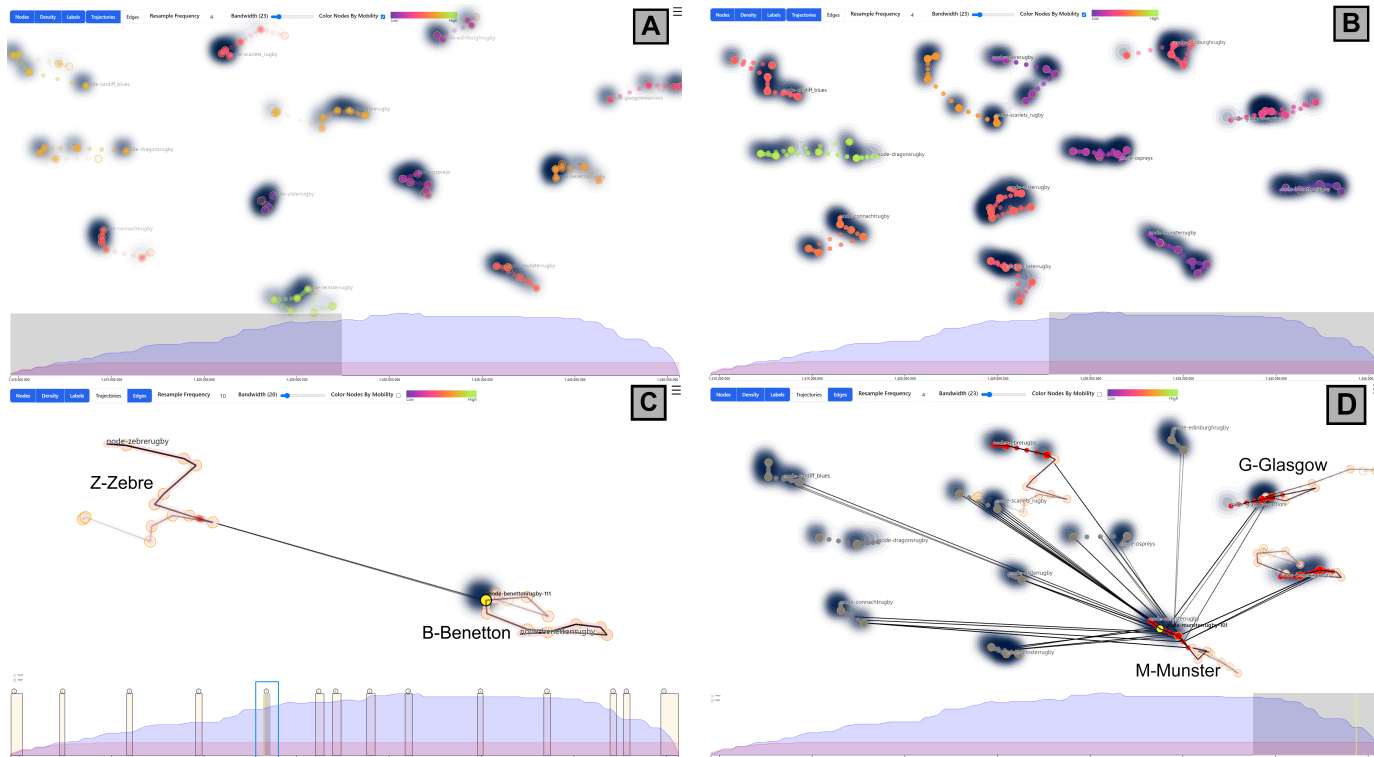


Fig. 5: Illustrations from Case Study 1. We can observe how the movement of nodes changes in the two halves of the season. (A) The first half of the season; (B) The second half of the season. Timeline brushing is used to filter out events. Trajectory sampling is set at 4 points. (C) The two teams’ trajectories are visible as guidance intervals in the timeline. The selected one (blue rectangle) relates to the match between the two teams in the first round of the competition. (D) Highlights the final match and the connections between the 2nd best team and other teams of interest in the case study.

dataset), and each edge between them is either a call or a short message. We use the real duration of the calls and set a fixed duration for the short text messages. Moreover, node appearance and presence are taken from the original data. For readability, we limit the number of nodes to 170, which results in a total of 333 events (edges). We modified the `MultiDynNoS` algorithm parameters to increase the movement of nodes. This network is more complex than *Rugby*, and therefore we take advantage of the guidance and filtering tools of `TimeLighting` to build our story.

Investigating communities. The status of the system can be seen in Fig. 6-A. We set the resampling frequency to 5 (i.e., the number of sampled nodes between the start and end positions, see Section 3) to reduce the clutter and the kernel size for the density to 15 to improve cluster separation. Nodes are colored according to their movement. From the timeline, it is possible to spot that the majority of traffic happens around the first third of the timeline, it peaks and then subsides. We argue that it is the beginning of group projects or courses, as sub-communities form and a few entities participate in more than group one across the entire timeline, as we will see in the following (T1). It is easy to spot the different sub-communities: while some are more tightly connected, i.e., the large group in the center, the group on the left and top right belong perhaps to smaller courses or friendship circles (see Fig. 6-A). If we filter out the moments around the peak, we can see the cores previously mentioned communities forming (see Fig. 6-B). Moving

the selected time interval forward, we see “beyond” the overview, meaning that as the semester progresses students come and go (i.e., appear and disappear from the dataset) and that groups reshuffle. Nodes move about significantly: the rate of movement can be inferred by the relative distance between the resampled nodes. A larger distance means more movement, as it must be covered in the same time interval indicated in the timeline.

Elements changing groups. We now track the two nodes with the greater movement (node 5390—highest and node 3947—third highest). The node with the second largest movement followed a similar trajectory to 5390, and we could not make meaningful observations. We, therefore, focus on 3947 and select the portion of the timeline from the beginning to the peak. Our first finding is that, at the very beginning, 5390 and 3947 interacted with the same nodes from one group, but not between each other directly. Node 3947 is the first to appear and shortly after departs from the group. Node 5390 appears later, and remains in the group until the peak before leaving as in Fig. 6-C (T2, T3). Following the proposed guidance, it is possible to find other instances where this behavior repeats. If we focus on the area in the timeline indicating the peaks where multiple guidance intervals are located (see Fig. 6-D), we see that the two nodes interact with node 93. This occurrence repeats close to the “tail” of the dataset: around this time, node 3947 is interacting with the nodes remaining in the cluster, and interacts again with node 93 and, coincidentally, with

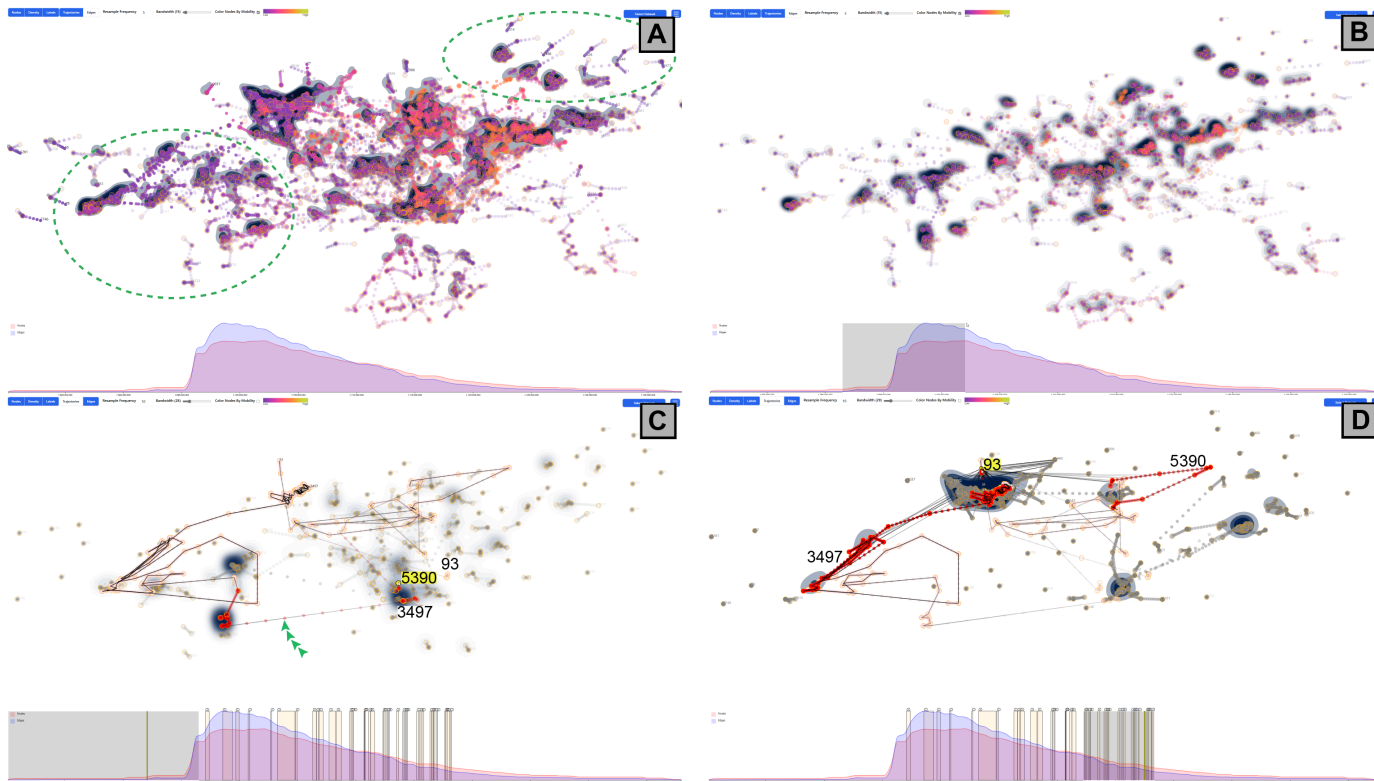


Fig. 6: Illustrations from Case Study 2. The complete projection of the dataset is shown in (A), with the two sub-communities referenced in Section 5.2 have been highlighted with green ellipses. In (B) the cores of these communities are shown as they begin to form. In (C) the nodes trajectories have been labeled for readability. The green arrow indicates the position of node 3497 as 5390 appears. Finally, in (D) we highlight the positions of nodes 5390 and 3497.

5390. Afterward, these last two will permanently disappear, followed shortly after by 3490. Further investigation shows that 93 was the last node to appear between the three and followed a direct trajectory similar to 5390 with a similar duration. We argue that these elements (5390 and 3497) were possibly tutors, as they visited several groups during the semester, with 93 being either a common friend or colleague. This example further shows how the system features and guidance are useful in untangling more complex temporal networks.

6 EVALUATION

We evaluate our system through two heuristics: the ICE-T methodology by Wall et al. [42] to assess the *value* of our proposed visualization, and the guidance evaluation heuristic by Ceneda et al. [43]. This methodology aims to provide a well-rounded appraisal of the design decisions regarding both the visual encodings and the guidance features, thus also suggesting how well they support TimeLighting design tasks (as in Section 3). In the following, we outline the setting and procedures we follow, and present our results. For further reading about the evaluation and results, we refer to our online supplementary material [40].

6.1 Methodology

Our evaluation is based on two concepts, which we aim to assess: the *value of visualization* and *guidance quality*.

Value of Visualization. The concept of *value* of visualization was introduced by Stasko et al. [44] as the ability to convey a “*true understanding*” of the data, beyond what usability task-based tests are typically able to uncover. Temporal networks are interesting to the visualization community as they carry more nuanced information over timeslicing. Therefore, we aim to visualize this extra knowledge appropriately and expressively. For this reason, we employ the ICE-T heuristic evaluation methodology proposed by Wall et al. [42] - as is an evaluation framework that stems from Stasko’s principles. The objective is to determine how appropriate and effective the visual representation and interactions are in responding to data-driven questions, generating insights, and inspiring confidence in the analysis results. The framework is divided in four visualization *components*: *Insight*, *Confidence*, *Essence*, and *Time*. Each component is divided into a set of guidelines, which are again divided into a set of heuristics, formulated as actionable and rateable statements. **Guidance Quality.** Ceneda et al. [43] in their work investigate what makes guidance *effective*. They identify 8 guidance quality criteria: each criterion is then divided into a set of heuristics, formulated as rateable insights. The paper proposes a different set of heuristics if they have to be rated by visualization users or guidance experts: depending on the evaluation perspective and purpose. In our evaluation, we use the user heuristics.

Both the evaluation methodologies described above aim to evaluate, in a qualitative manner, different aspects of a visualization approach. Moreover, they follow a similar pro-

to col, each of the study participants fill out a questionnaire evaluating each heuristic on a 1-7 Likert scale. Hence they can be effectively combined to provide an assessment of the TimeLighting overall value. We refer the interested reader to the respective papers for further reading [42], [43].

6.2 Setting and Procedure

We recruited 6 experts who voluntarily agreed to participate with no compensation. All participants had a strong background in dynamic network visualization as the intended audience of our proposed approach requires such expertise. Each participant was individually interviewed.

The interview sessions were structured as open-ended exploratory analysis scenarios lasting 1 hour. Each began with a 15-minute onboarding, detailing the visual encodings, interaction techniques, and guidance features of TimeLighting. In the remainder, participants were asked to freely interact with the system on three different graphs of increasing topological and temporal density. We prepared a set of optional tasks (see supplemental material [40]), mainly intended to bootstrap the exploration (as they were not required by the evaluation methodologies we use). The interviews were conducted online using a remote video conferencing application and the participants were asked to share their screen contents as they interacted with the prototype. We kept a record of each interview, detailing how the participants interacted with our approach. We also recorded the insights they gained, comments, and feedback they had about the visual encodings, interaction techniques, or guidance features. Participant voices and on-screen interactions were recorded with their consent.

At the end of each interview, we asked the participants to fill out a questionnaire structured as a combination of the ICE-T survey [42] and guidance evaluation methodology [43]. The questionnaire was provided to the participants as an online form to complete at their convenience.

We prepared three datasets for the evaluation, of increasing complexity. All three datasets have been previously used in visualization evaluations - with InfoVis being particularly familiar to our intended audience and study participants.

Rugby: is a temporal network derived from tweets exchanged between January 2014 and 2015 about the “Guinness Pro12” rugby competition [8], [41]. The teams represent the nodes and the relationships are derived from tweets - it is the same dataset as in case study 1 (see Section 5.1).

Reality Mining: It is the same dataset used in Case Study 2 (see Section 5.2), albeit in a simplified version. Specifically, compared to the case study, we only include the first 200 events.

InfoVis: is a co-authorship network for papers published in the InfoVis conference from 2008 to 2018 [45]. Authors of a paper are connected in a clique at the time of publication.

6.3 Results

Value of Visualization. As described by Wall et al. [42], we averaged the questionnaire results per component. The component scores are in the range of 4.61 to 5.48 and the individual guidelines scores range from 3 to 5.67 (see Section 7). A score of 5 is considered the minimum for a successful evaluation with the ICE-T method [42].

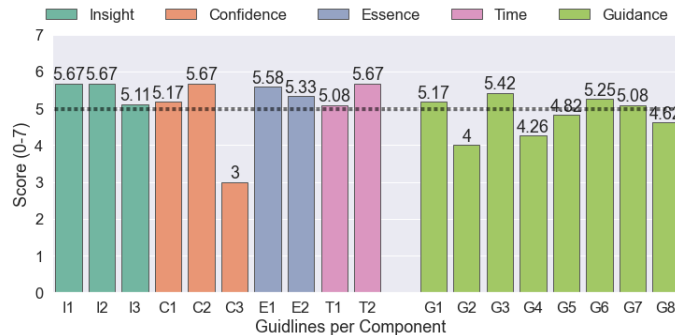


Fig. 7: The evaluation results of the individual guidelines from the ICE-T [42] and guidance evaluation methodologies [43]. The scores are calculated as the average of the heuristics related to each guideline using a Likert scale (0-7). The dashed line indicates a score of 5, which corresponds to the minimum score for an appropriate and effective visualization. The complete definitions of the shortcodes for the guidelines are available in our online supplementary material [40].

Insight component (Score 5.48): The participants rated TimeLighting highly in terms of its ability to facilitate insight generation and provide new perspectives on temporal networks, scoring 5.67 (I1–The visualization facilitates answering questions about the data) and 5.67 (I2–The visualization provides a new or better understanding of the data), respectively. Additionally, the participants rated our approach with a score of 5.10 (I3–The visualization provides opportunities for serendipitous discoveries) for providing opportunities for serendipitous discoveries, indicating its efficacy in revealing unexpected patterns within the data. Overall, TimeLighting succeeds in effectively facilitating exploration and gaining an in-depth understanding of temporal network dynamics. Our proposed approach and the supporting interactions provide opportunities for insightful discoveries and gaining a better understanding of the underlying network patterns and trends.

Confidence component (Score 4.61): Overall, confidence guidelines were rated lower than the others. Participants rated 5.17 the visualization’s capacity to help avoid making incorrect inferences (C1–The visualization helps avoid making incorrect inferences), but they believe it facilitates broader learning about the domain with a score of 5.67 (C2–The visualization facilitates learning more broadly about the domain of the data). Finally, they expressed a lower score of 3 (C3–The visualization helps understand data quality) regarding its ability to understand data quality. This is somewhat contrary to the insights and observations made during the evaluation sessions, where the participants utilized the movement scores and mouseover interactions to identify unexpected behavior of the network. These scores show that this is a potential area for improvement to enhance how TimeLighting could better support an immediate understanding of data quality and uncertainty in temporal networked data.

Essence component (Score 5.46): TimeLighting received a score of 5.58 for its capacity to provide a big-picture of the data (E1–The visualization provides a big-picture perspective of the data) and was ranked 5.33 on the ability to provide

an understanding beyond individual data cases (**E2**–*The visualization provides an understanding of the data beyond individual data cases*), indicating its effectiveness in conveying overarching trends and patterns present in the data.

Time component (Score 5.38): Participants rated the visualization favorably in terms of its capabilities in offering rapid comprehension of the data with 5.08 (**T1**–*The visualization affords rapid parallel comprehension for efficient browsing*) as well as enabling to quickly lookup and seek for specific information (data points) with 5.67 (**T2**–*The visualization provides mechanisms for quickly seeking specific information*). The time component was rated highly by all the participants and this further highlights `TimeLighting` efficiency in supporting users' exploration of temporal data.

Guidance component (Score 4.83): We averaged the score per guidance quality, with all of them scoring between 4.00 and 5.42—suggesting a positive final assessment [43]. The scores of the individual quality metrics highlight the strengths of our proposed features. The proposed guidance was perceived as flexible as it adapted to the selection (i.e., selecting specific periods/windows of time or specific nodes to investigate in more detail) of the user (5.17, **G1**–*Flexible: the guidance degree is adequate and adapts to the situation*). The participants ranked the visibility of the guidance quite highly and commented that it was easily distinguishable from other visual encodings (5.42, **G3**–*Visible: the guidance and its status are easy to identify and well integrated with the rest of the visual environment*). Furthermore, the participants positively ranked the proposed guidance to use appropriate and expressive encodings (5.25, **G6**–*Expressive: the visual encoding of guidance is appropriate*), as well as the timeliness of the provided guidance (5.08, **G7**–*Timely: the guidance is provided on time when needed*). Although expected (as we did not design specifically the guidance to address this), some aspects received lower ratings, providing directions for future work. Most notably, the capability of the guidance to interactively and dynamically adapt to the user's preferences (4, **G2**–*Adaptive: the guidance system considers the user's preferences habits and the current task requirements for generating suggestions*). The ability to steer the guidance and evaluate alternative paths of analysis was also ranked low. This implies that presenting alternative paths of analysis or providing provenance of interactions performed and being able to revert the guidance to previous states are aspects that we can improve upon (4.22, **G4**–*Controllable: the guidance can be steered, allow the user to evaluate alternatives, switch on/off, revert the guidance*). Furthermore, participants did not feel like the guidance was self-explanatory and understandable, and this is also reflected in the score (4.82, **G5**–*Explainable: the guidance and the way it was generated can be easily understood*). Investigating how to improve the explainability of the provided guidance in our proposed approach, and more generally, to other application domains and use cases presents a pressing direction for future research. Finally, based on the outcome of our evaluation we identified that we can improve the capability of the proposed guidance features in providing relevant suggestions depending on the context and recovery from error states (4.62, **G8**–*Relevant: the guidance helps prevent or recover from errors guides towards a goal helps discover the unexpected provides multiple options and perspectives*).

Summarizing the feedback and results from the questionnaire, we arrive at a final score of **5.11** which we consider to be a successful evaluation of `TimeLighting`.

7 DISCUSSION

The feedback and comments from the evaluation of our proposed approach gathered by expert interviews offer valuable insights into its usability, effectiveness, and potential improvements.

Participants found the concept of flattening the space-time cube and the `TimeLighting` metaphor intuitive, aiding them in comprehending temporal graph data (**T1**). They effectively performed open-ended exploration tasks (**T2** & **T3**), generating and answering data-driven questions and reasoning about the network's behavior over time. The timeline and area charts were frequently used to select specific intervals of interest (“when”) and the density maps highlighted areas in the network that seemed intriguing for further inspection. The movement scores and node coloring were appreciated for identify specific entities of the network (“where”) to investigate.

The proposed guidance features (i.e., movement scores, node coloring, and the highlighted timeline intervals) were instrumental in identifying elements of the network and time intervals for in-depth analysis. These features supported **T2** & **T3** by assisting in understanding temporal dynamics and event frequency. Participants appreciated the ability to control and toggle the elements on or off and explore the data from different perspectives. Finally, the mouseover interactions for exploring the immediate ego network of a specific node and the edge's age encoding (**T3**) were essential in understanding connectivity and development over time. These interactions facilitated the identification of impactful edges and their frequency, supporting the identification of cluster formation and reasoning about behavior.

Limitations: The lack of clarity about node movement and layout decisions impacted the interpretability and trust in the guidance features and analysis. Enhancing the explainability and interpretability of the underlying layout algorithm is crucial for improving confidence in the results and interpretation as well as the overall usability of our approach. Furthermore, some scalability concerns emerged with larger datasets, suggesting the need for performance optimization as well as investigating different layout algorithms for a better user experience. Specifically, when depicting networks that have a large number of node trajectories (>100) over extended time periods, the visualization's capability of conveying more subtle changes diminished, however, this was also remedied to some extent by narrowing down the interval selection and the resampling frequency. Overlapping node trajectories further complicated the analysis, indicating a need for improved layout algorithms to ensure trajectory separability (e.g., Fig. 5, all nodes' trajectories are clearly separated this may be a specific property for some of the datasets).

8 CONCLUSION AND FUTURE WORK

In this paper, we presented and described `TimeLighting`, a guidance-enhanced VA approach to support the analysis

and exploration of temporal networks embedded in a space-time cube. We augmented the visualization approach with guidance to better support users in the visual analysis of the data. We demonstrated the effectiveness of our approach with two use cases and a heuristic evaluation conducted with experts. The case studies depict a scenario where the aim is to explore and extract insights from temporal networks. The results of the expert evaluation highlight the strengths of TimeLighting and outline areas for improvement and refinement. While certain aspects of our interactions and encodings can be improved, the results also demonstrate the benefits of visualizing temporal network data using TimeLighting and providing guidance features to highlight “when” and “where” there are interesting data points for further inspection. Addressing these findings in future iterations of our approach will enhance the scalability for larger datasets, interpretability of the movement score, and the confidence that participants have in their analysis results. The potential of visualization in guiding and supporting the exploration and analysis of temporal graphs is an opportunity for the entire graph drawing and visualization communities.

ACKNOWLEDGMENTS

The authors have applied a Creative Commons Attribution (CC-BY) license to any Author Accepted Manuscript version arising from this submission. This work was supported by WWTF [10.47379/ICT19047], FFG DoRIAH [#880883], FWF ArtVis [10.55776/P35767] and SANE [10.55776/I6635]. The authors acknowledge TU Wien Bibliothek for financial support through its Open Access Funding Programme.

REFERENCES

- [1] V. Filipov, D. Ceneda, D. Archambault, and A. Arleo, “Time-Lighting: Guidance-Enhanced Exploration of 2D Projections of Temporal Graphs,” in *Graph Drawing and Network Visualization (GD '23)*. Springer, 2024, pp. 231–245.
- [2] P. Holme and J. Saramäki, “Temporal Networks,” *Physics Reports*, vol. 519, no. 3, pp. 97–125, 2012.
- [3] V. Filipov, A. Arleo, and S. Miksch, “Are We There Yet? A Roadmap of Network Visualization from Surveys to Task Taxonomies,” in *Computer Graphics Forum*, vol. 42, 2023.
- [4] N. Eagle and A. S. Pentland, “Reality Mining: Sensing Complex Social Systems,” *Personal and ubiquitous computing*, vol. 10, no. 4, pp. 255–268, 2006.
- [5] M. Sondag, C. Turkay, K. Xu, L. Matthews, S. Mohr, and D. Archambault, “Visual Analytics of Contact Tracing Policy Simulations During an Emergency Response,” *Computer Graphics Forum*, vol. 41, no. 3, pp. 29–41, 2022.
- [6] F. Beck, M. Burch, S. Diehl, and D. Weiskopf, “A Taxonomy and Survey of Dynamic Graph Visualization,” *Computer Graphics Forum*, vol. 36, no. 1, p. 133–159, 2017.
- [7] A. Arleo, S. Miksch, and D. Archambault, “Event-based Dynamic Graph Drawing without the Agonizing Pain,” *Computer Graphics Forum*, vol. 41, no. 6, pp. 226–244, 2022.
- [8] P. Simonetto, D. Archambault, and S. Kobourov, “Event-based Dynamic Graph Visualisation,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 7, pp. 2373–2386, 2018.
- [9] M. Farrugia and A. Quigley, “Effective Temporal Graph Layout: A Comparative Study of Animation versus Static Display Methods,” *Journal of Information Visualization*, vol. 10, no. 1, pp. 47–64, 2011.
- [10] D. Archambault, H. Purchase, and B. Pinaud, “Animation, Small Multiples, and the Effect of Mental Map Preservation in Dynamic Graphs,” *IEEE Transactions on Visualization and Computer Graphics*, 2011.
- [11] D. Archambault and H. C. Purchase, “Can Animation Support the Visualisation of Dynamic Graphs?” *Information Sciences*, p. 495–509, 2016.
- [12] T. Hägerstrand, “What About People in Regional Science, Regional Science Association Papers,” Vol. xxiv, 1970.
- [13] B. Bach, P. Dragicevic, D. Archambault, C. Hurter, and S. Carpendale, “A Descriptive Framework for Temporal Data Visualizations Based on Generalized Space-Time Cubes,” *Computer Graphics Forum*, vol. 36, no. 6, pp. 36–61, 2017.
- [14] P. Simonetto, D. Archambault, and S. Kobourov, “Drawing Dynamic Graphs Without Timeslices,” in *Graph Drawing and Network Visualization*, 2017, pp. 394–409.
- [15] M. Baur, M. Benkert, U. Brandes, S. Cornelsen, M. Gaertler, B. Köpf, J. Lerner, and D. Wagner, “VisOne Software for Visual Social Network Analysis,” in *Graph Drawing and Network Visualization*. Springer, 2001, pp. 463–464.
- [16] P. Ljung, J. Krüger, E. Groller, M. Hadwiger, C. D. Hansen, and A. Ynnerman, “State of the Art in Transfer Functions for Direct Volume Rendering,” *Computer Graphics Forum*, vol. 35, no. 3, pp. 669–691, 2016.
- [17] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, and C. Tominski, “Characterizing Guidance in Visual Analytics,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 1, pp. 111–120, 2016.
- [18] U. Brandes and M. Mader, “A Quantitative Comparison of Stress-minimization Approaches for Offline Dynamic Graph Drawing,” in *Graph Drawing and Network Visualization*, 2011, pp. 99–110.
- [19] J. Branke, *Dynamic Graph Drawing*. Springer, 2001, pp. 228–246.
- [20] W. Aigner, S. Miksch, H. Schumann, and C. Tominski, *Visualization of Time-Oriented Data*, 2nd ed. Springer, 2023.
- [21] N. Kerracher, J. Kennedy, and K. Chalmers, “The Design Space of Temporal Graph Visualisation,” in *EuroVis - Short Papers*. The Eurographics Association, 2014.
- [22] V. Filipov, A. Arleo, and S. Miksch, “Exploratory User Study on Graph Temporal Encodings,” in *2021 IEEE 14th Pacific Visualization Symposium (PacificVis)*, 2021, pp. 131–135.
- [23] V. Filipov, A. Arleo, M. Bögl, and S. Miksch, “On Time and Space: An Experimental Study on Graph Structural and Temporal Encodings,” in *Graph Drawing and Network Visualization*, ser. Lecture Notes in Computer Science, vol. 13764. Springer, 2022, pp. 271–288.
- [24] V. Filipov, A. Arleo, M. Bögl, and S. Miksch, “On Network Structural and Temporal Encodings: A Space and Time Odyssey,” *IEEE Transactions on Visualization and Computer Graphics*, 2023.
- [25] B. Bach, E. Pietriga, and J.-D. Fekete, “Visualizing Dynamic Networks with Matrix Cubes,” in *Proc. of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '14. Association for Computing Machinery, 2014, p. 877–886.
- [26] F. Windhager, S. Salisu, R. A. Leite, V. Filipov, S. Miksch, G. Schreder, and E. Mayr, “Many Views Are Not Enough: Designing for Synoptic Insights in Cultural Collections,” *IEEE Computer Graphics and Applications*, vol. 40, no. 3, pp. 58–71, 2020.
- [27] Y. Frishman and A. Tal, “Online Dynamic Graph Drawing,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 4, pp. 727–740, 2008.
- [28] R. F. Cohen, G. Di Battista, R. Tamassia, I. G. Tollis, and P. Bertolazzi, “A Framework for Dynamic Graph Drawing,” in *Proceedings of the Eighth Annual Symposium on Computational Geometry*, ser. SCG '92, 1992, p. 261–270.
- [29] J. Ellson, E. R. Gansner, E. Koutsofios, S. C. North, and G. Woodhull, *Graphviz and Dynagraph — Static and Dynamic Graph Drawing Tools*. Springer, 2004, pp. 127–148.
- [30] K. Gray, M. Li, R. Ahmed, and S. Kobourov, “Visualizing Evolving Trees,” in *Graph Drawing and Network Visualization: 30th International Symposium, GD 2022, Tokyo, Japan, September 13–16, 2022, Revised Selected Papers*, 2023, p. 319–335.
- [31] Y. Wang, D. Archambault, H. Haleem, T. Moeller, Y. Wu, and H. Qu, “Nonuniform Timeslicing of Dynamic Graphs Based on Visual Complexity,” in *2019 IEEE Visualization Conference (VIS)*, 2019, pp. 1–5.
- [32] A. Lee, D. Archambault, and M. Nacenta, “Dynamic Network Plaid: A Tool for the Analysis of Dynamic Networks,” in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–14.
- [33] A. Lee, D. Archambault, and M. A. Nacenta, “The Effectiveness of Interactive Visualization Techniques for Time Navigation of

- Dynamic Graphs on Large Displays," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 528–538, 2021.
- [34] D. Ceneda, T. Gschwandtner, and S. Miksch, "A Review of Guidance Approaches in Visual Data Analysis: A Multifocal Perspective," *Computer Graphics Forum*, vol. 38, no. 3, pp. 861–879, 2019.
- [35] D. Ceneda, A. Arleo, T. Gschwandtner, and S. Miksch, "Show Me Your Face: Towards an Automated Method to Provide Timely Guidance in Visual Analytics," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 12, pp. 4570–4581, 2022.
- [36] T. May, M. Steiger, J. Davey, and J. Kohlhammer, "Using Signposts for Navigation in Large Graphs," *Computer Graphics Forum*, vol. 31, no. 3pt2, pp. 985–994, 2012.
- [37] S. Gladisch, H. Schumann, and C. Tominski, "Navigation Recommendations for Exploring Hierarchical Graphs," in *Advances in Visual Computing: 9th International Symposium, ISVC 2013, Rethymnon, Crete, Greece, July 29-31, 2013. Proceedings, Part II 9, 2013*, pp. 36–47.
- [38] S. Miksch and W. Aigner, "A Matter of Time: Applying a Data–Users–Tasks Design Triangle to Visual Analytics of Time-Oriented Data," *Computers & Graphics*, vol. 38, pp. 286–290, 2014.
- [39] J.-w. Ahn, C. Plaisant, and B. Shneiderman, "A Task Taxonomy for Network Evolution Analysis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 20, no. 3, p. 365–376, 2014.
- [40] "TimeLighting Miro Board," <https://miro.com/app/board/uXjVNUoEkKM=/>, [Accessed 06/02/2024].
- [41] "PRO-12 Rugby Competition 2014-2015 Season Standings," <http://rd.pro12rugby.com/matchcentre/table.php?includeref=11189&season=2014-2015>, online, Accessed: 2023-06-03.
- [42] E. Wall, M. Agnihotri, L. Matzen, K. Divis, M. Haass, A. Endert, and J. Stasko, "A Heuristic Approach to Value-Driven Evaluation of Visualizations," *IEEE Transactions on Visualization and Computer Graphics*, vol. 25, no. 1, pp. 491–500, 2019.
- [43] D. Ceneda, C. Collins, M. El-Assady, S. Miksch, C. Tominski, and A. Arleo, "A Heuristic Approach for Dual Expert/End-User Evaluation of Guidance in Visual Analytics," *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–11, 2023.
- [44] J. Stasko, "Value-driven Evaluation of Visualizations," in *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, ser. BELIV '14. Association for Computing Machinery, 2014, p. 46–53.
- [45] P. Isenberg, F. Heimerl, S. Koch, T. Isenberg, P. Xu, C. D. Stolper, M. Sedlmair, J. Chen, T. Möller, and J. Stasko, "vispubdata.org: A Metadata Collection about IEEE Visualization (VIS) Publications," *IEEE Transactions on Visualization and Computer Graphics*, vol. 23, no. 9, pp. 2199–2206, 2016.



Velitchko Filipov is a post-doc researcher in the research group (CVASt) at the Institute of Visual Computing and Human-Centered Technology, TU Wien. He received his Ph.D. degree from TU Wien, Austria, in 2024 defending his thesis called "Networks in Time and Space: Visual Analytics of Dynamic Network Representations". His research interests include visual analytics of dynamic graphs and networks with a focus on novel representations and interactions.



Davide Ceneda received the PhD degree from TU Wien, Austria, in 2020 defending a thesis called "Guidance-Enriched Visual Analytics". He is a post-doc with the CVASt group at the Institute of Visual Computing and Human-Centered Technology, TU Wien. His research interests include visualization, human perception, visual analytics and guidance.



Daniel Archambault received the PhD degree from the University of British Columbia, in 2008 and is now a professor of Visualization/Data Science at Newcastle University where he co-leads the scalable computing group. A main theme of his research is network visualization and evaluating the perceptual effectiveness of such approaches.



Alessio Arleo is an Assistant Professor at the Eindhoven University of Technology, The Netherlands. He obtained a PhD in 2018 at the University of Perugia, Italy, investigating distributed computing platforms applied to graph drawing. He has been a Post-Doc in the Visual Analytics research unit (CVASt) at the Institute of Visual Computing and Human-Centered Technology, at TU Wien. His research interests include temporal network visualization, information diffusion, and distributed graph algorithms.