ZigzagNetVis: Suggesting temporal resolutions for graph visualization using zigzag persistence – Supplemental material

Raphaël Tinarrage, Jean R. Ponciano, Claudio D. G. Linhares, Agma J. M. Traina, and Jorge Poco

A Definition of bottleneck distance

The most common distance between barcodes $\mathscr{B}$ and $\mathscr{B}'$ is the *bottleneck distance*, defined as follows. First, we say that a subset $C \subset \mathscr{B} \times \mathscr{B}'$ is a *partial correspondence* if for any bar $b \in \mathscr{B}$ (resp. $b' \in \mathscr{B}'$) there exists at most one bar $b' \in \mathscr{B}'$ (resp $b \in \mathscr{B}$) such that $(b,b') \in C$. Next, the cost of a matched pair of bars $(b,b') \in C$ is defined as the maximum difference between their endpoints. A bar $b \in \mathscr{B}$ (resp. $b' \in \mathscr{B}'$) for which there is no $b' \in \mathscr{B}'$ (resp $b \in \mathscr{B}$) such that $(b,b') \in C$ is said unmatched, and its cost is defined as its half-length. The cost of the partial correspondence C is defined as the maximal costs of its matched and unmatched bars. Last, the bottleneck distance, denoted $d_{\mathrm{B}}(\mathscr{B},\mathscr{B}')$, is defined as the minimal cost of a partial correspondence:

$$d_{\mathrm{B}}(\mathscr{B},\mathscr{B}') = \inf_C \max\left\{ \sup_{(b,b')\in C} \|b-b'\|_\infty,\ \sup_{(b,\cdot)\notin C} \|b\|_\infty,\ \sup_{(\cdot,b')\notin C} \|b'\|_\infty \right\}.$$

It is interesting to note that, by definition, this distance is either equal to $\|b-b'\|$ for a pair of bars $(b,b') \in \mathscr{B} \times \mathscr{B}'$, or to $\|b\|$ for a single bar $b \in \mathscr{B} \cup \mathscr{B}'$. That is to say, the bottleneck distance is caused either by a pair of bars, or by a bar alone. Identifying this cause allows us to derive an explainability pipeline, in Sec. 9 of the main document.

B Methodology details

When studying temporal graphs, two common uniform timeslicing methods are employed. The first one, partition timeslicing, consists of choosing a multiple αr_0 of the initial resolution and subdividing the time by stacking intervals of length αr_0. It is employed, for instance, in LargeNetVis [7]. The other one, sliding-window timeslicing, is obtained by allocating to each edge an activation window of semilength αr_0, as used in [9]. In practice, we observed that our resolution suggestion method, described in the main document (Sec. 4.2), gives better results when considering sliding-window timeslicing. The partition timeslicing suffers from *instability*, which we will exemplify here.

B.1 Instability of partition timeslicing

Consider a temporal graph that is made of only two nodes. Let the time interval $[0,T]$ be subdivided as $[0,t_1] \cup [t_1,t_2] \cup [t_2,T]$, and suppose that the edge is active only on $[0,t_1]$ and $[t_2,T]$. Let $r = \alpha r_0$ be a multiple of the initial resolution such that $r < t_2 - t_1 < 2r$. By applying partition timeslicing, a whole interval $[kr,(k+1)r]$ may be included in $[t_1,t_2]$. Since no edge is active in this interval, we obtain a graph G_k that is empty. Hence the barcode consists of two bars, the barcode being empty on the interval $[t_1,t_2]$. However, if by chance no $[kr,(k+1)r]$ is included in $[t_1,t_2]$, then the barcode will consist of only one bar. The situation is depicted in Fig. 1.

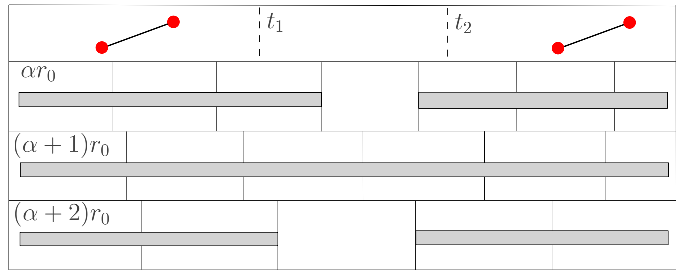

Fig. 1. From top to bottom: a temporal graph, and the barcodes obtained for three consecutive resolutions, using the partition timeslicing. The two bars of the barcode get glued and cut again, showing the instability of this timeslicing method.

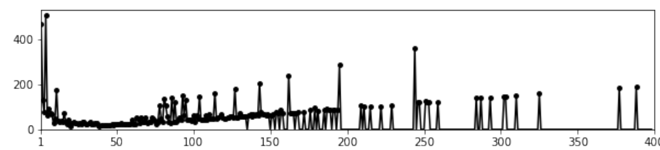

Fig. 2. Normalized suggestion curve for partition timeslicing in the Primary School network. We observed that the peaks witness an alternation of the barcode between having a large bar and two smaller bars. The x and y axes represent the resolution values and the consecutive bottleneck distance, respectively.

In practice, we have observed that increasing the resolution slightly allows us to go from one situation to the other and vice versa. A concrete example for the Primary School network is given in Fig. 2. We observed that most of the peaks correspond to the same bar of the barcode, which gets cut and merged again. In this case, our resolution selection method will detect all these critical resolution values. To analyze the graph, we would rather have detected this change in behavior only once.

The sliding-window timeslicing, in contrast, does not exhibit this problem. Indeed, the activation windows' length on each edge grows monotonically. Therefore, for our main paper's analyses, experiments, and illustrations, we choose to use the sliding window method. The only exception is when we compare our ZigzagNetVis approach with LargeNetVis [7] (see Sec. C.4 in this document). This is because, as mentioned earlier, LargeNetVis employs partition timeslicing.

B.2 Visual comparison of timeslicing methods

Given the first day of the Primary School network (recall Sec. 6 from the main document), Fig. 3 illustrates our colored barcode created using partition (Fig. 3(a)) and sliding window timeslicing (Fig. 3(b)). While the partition presents a smoother representation, the sliding window faithfully represents the activity variation over time. In the case of the partition, the number of connected components is computed only at each partition, equivalent to layouts representing the activity in grouped timestamps or timeslices, such as the LargeNetVis's Global View [7].

R. Tinarrage and J. Poco are with the School of Applied Mathematics, Fundação Getulio Vargas, Rio de Janeiro, Brazil. E-mails: {raphael.tinarrage, jorge.poco}@fgv.br.

J. Ponciano and A. Traina are with the Institute of Mathematics and Computer Sciences, University of São Paulo, São Carlos, Brazil. E-mails: {jeanponciano}@usp.br, {agma}@icmc.usp.br.

- *C. Linhares is with the Department of Computer Science and Media Technology Linnaeus University, Växjö, Sweden. E-mail: {claudio.linhares}@lnu.se.*

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxxx

In contrast, the number and evolution of the connected components are highly dynamic when using sliding windows due to the quick identification of temporal events such as grows, merges, splits, and disappearances.

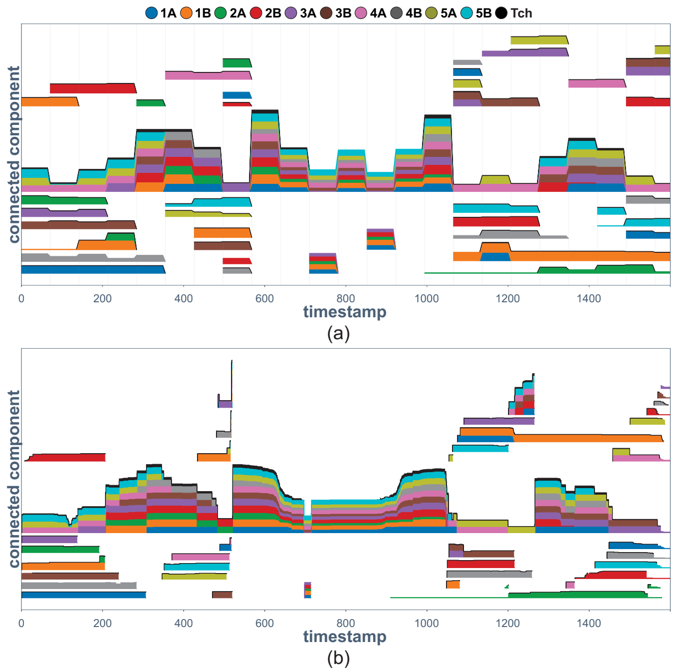

Fig. 3. Comparing colored barcodes using the first day of the Primary School network for (a) partition and (b) sliding window timeslicing. Both cases use the suggested resolution $r = 154$ and filter out components with less than 10 node members and 10 timestamps of duration.

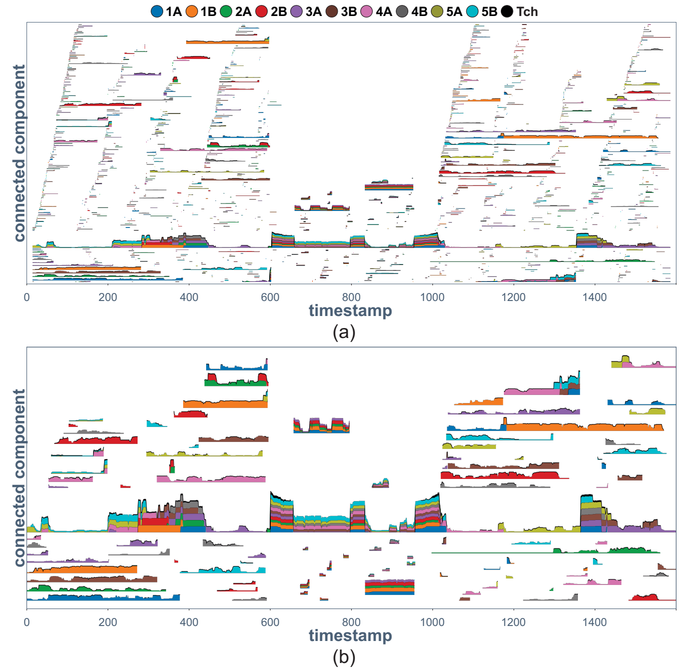

Fig. 4. Comparing the colored barcodes for the Primary School network using sliding window: (a) without filtering and (b) after filtering out components with less than 10 nodes and 10 timestamps of duration. Adopted resolution: $r = 18$.

C Additional Experiments

C.1 Visual comparison of filtering options

Fig. 4 compares the colored barcode without filtering connected components (Fig. 4(a)) with the one after filtering out components with less than 10 nodes and 10 timestamps of duration (Fig. 4(b)). The dataset used is the first day of the Primary School with resolution $r = 6$. Although the colored barcode without filtering faithfully represents all active connected components in every timestamp, this layout leads to a high level of visual clutter, hindering the analyses. Focusing on large connected components leads to important regions of interest and improves the scalability of the approach. The minimum number of nodes and duration a component must have are currently user-defined thresholds; we believe further investigation may lead to a method that automatically suggests suitable values.

C.2 Suggested resolutions

C.2.1 Resolution curves

As mentioned in Sec. 9 of the main document, Fig 10 depicts the normalized suggestion curves for the networks Primary School, High School, Hospital, InVS, Museum, Enron, Conference, and Sexual.

C.2.2 Comparison of resolutions

As also studied in Sec. 9 of the main document, Fig 11 shows the bars that differ the most when considering two selected resolutions, according to the bottleneck distance. The figure considers the networks Hospital, InVS, Museum, Enron, Conference, and Sexual. To see the bars that differ the most in the Primary School, please refer to Fig. 16 of the main document.

C.2.3 Comparison with other features

We now compare our novel method with existing techniques. In the literature, features of temporal graphs are of two sorts: either they are features of (non-temporal) graphs, adapted to the temporal case by taking their mean of their list over all the snapshots, or they directly depend on the temporal structure [10, 13].

Geometric and topological features of snapshots. Let G be a temporal graph to which we apply a sliding-window timeslicing of resolution r. Given a snapshot G_t, we consider:

- its number of nodes and edges, denoted $N(t)$ and $E(t)$,
- its density $D(t) = 2E(t)/(V(t)(V(t)-1))$,
- its number of connected components $CC(t)$,
- the mean degree $MD(t)$ of its nodes,
- its transitivity $T(t)$, defined the ratio between the number of triangles and triads (i.e., pairs of edges sharing a vertex).

Taking the mean value of such a feature over all times t yields a feature of the temporal graph G. We denote them respectively $N(r)$, $E(r)$, $D(r)$, $CC(r)$, $MD(r)$ and $T(r)$, making explicit the dependence on the resolution.

Note that, when increasing the resolution r, the features $N(r)$, $E(r)$ and $D(r)$ increase. That is, they are non-decreasing functions. In general, we except that abrupt changes in these values reflect the fact that the temporal graph exhibits a new behavior. To visualize such changes, one can plot these curves or, more efficiently, their derivative. These curves are represented in Fig. 5 for the Primary and High School networks, and in Fig. 12 for the other graphs. Since we are only interested in the peaks or qualitative behaviors of these curves, and not their absolute values, we normalize them so that their maxima equals one. In order to ease the reading, the x-axis is divided in two windows, and the curves are normalized both times.

A manual inspection of these curves allows us to compare them with our suggestion curve. For instance, in the Primary School Network, one sees that the first peak, at $r = 8$, corresponds to the global maximum of the derivative of the number of connected components. Besides, the peak at resolution 154 seems to appear simultaneously as the transitivity curve shows a constant derivative.

Similar observations can be made on the High School network. The peaks of our suggestion curve found at resolutions 46, 204, and 216

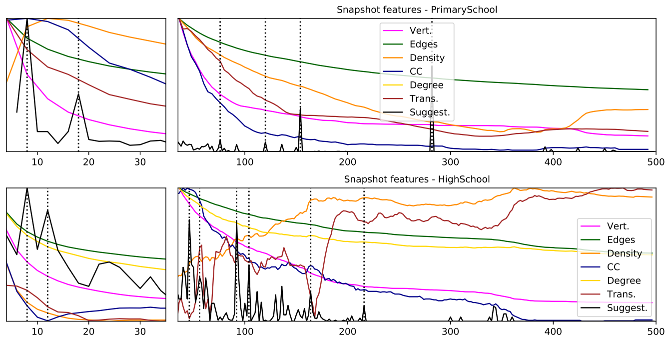

Fig. 5. Suggestion curve and derivative of the mean snapshot features for the networks Primary School (top) and High School (bottom). The x axis represents the resolution values. The curves are normalized, and a few interesting values are highlighted with a dashed line.

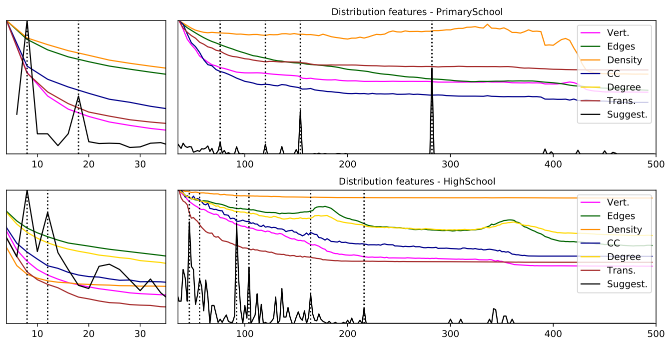

Fig. 6. Consecutive distances between the distribution of the snapshot features for the networks Primary School (top) and High School (bottom). The x axis represents the resolution values. The curves are normalized, and a few interesting values are highlighted with a dashed line.

correspond, respectively, to global maxima of the derivative of the number of connected components, transitivity, and density. Besides, the peaks at resolutions 12 and 56 correspond, respectively, to a global minimum of the derivative of the number of connected components and a significant local maximum of the transitivity.

These observations suggest that our curve captures information coming from various features of graphs. However, some peaks remain unexplained, and hence we will study other features in the paragraphs below. It must also be noted that certain features' peaks do not correspond to a peak of the suggestion curve. This may be caused by the fact that the suggestion curve, based on the homology group H_0, is blind to certain purely geometric properties of graphs, and works only in terms of connected components.

Distribution of features of snapshots. Instead of taking the average value of a feature over all the snapshots, we can compare their distribution over time. To do so, we consider the entire curves

$$f_{N,r}\colon t \mapsto N(t), \quad f_{E,r}\colon t \mapsto E(t), \quad \text{etc.}$$

Given two consecutive resolutions r and $r+2$, we compare these curves via their ℓ^2-norm

$$\|f_{N,r} - f_{N,r+2}\|_2, \quad \|f_{E,r} - f_{E,r+2}\|_2, \quad \text{etc.}$$

These are functions of r, on which we expect to observe abrupt changes in the graph's behavior. These curves are represented in Fig. 6 for the Primary and High School networks and in Fig. 13 for the other graphs.

As before, one draws correspondences between the suggestion curve and these features. For example, on the Primary School network, the first suggested resolution, $r = 8$, happens precisely during an abrupt change in the derivative of all the curves. Besides, the two peaks at resolutions 120 and 154 delimit the only interval where the curve of transitivity increases and then decreases. This last correspondence has already been observed in the last paragraph while considering the mean transitivity of the temporal graph.

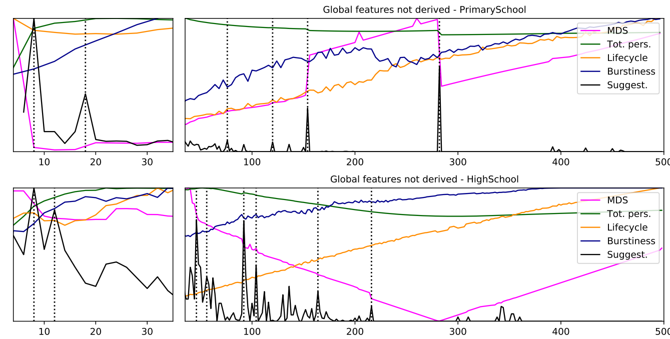

Fig. 7. Global features as functions of the resolution parameter, for the networks Primary School (top) and High School (bottom). The x axis represents the resolution values. The curves are normalized, and a few interesting values are highlighted with a dashed line.

In a similar fashion, on the High School network, one observes that the selected resolution 46 corresponds to an abrupt change in the curve built from transitivity. In a few words, comparing the distribution of the snapshot features offers compatible but also complementary information to the average values alone.

Global features. Lastly, we consider features of dynamic graphs that do not come from features of snapshots. Given a temporal graph G, timesliced at a resolution r, we consider:

- its burstiness $B(r)$ and average lifecycle $LC(r)$, defined in [12].
- its stability $S(r)$ and fidelity $F(r)$, defined in [1].

Moreover, we will also consider the *total persistence* $TP(r)$ of its corresponding zigzag persistence module, defined as the quadratic mean of the length of its bars. Finally, we will employ the curve $MDS(r)$ defined in [4]. It consists of the multidimensional scaling (MDS) in dimension 1, whose input is the set of bottleneck distances between the persistence barcodes of the temporal graph for all the resolutions considered. These curves are represented in Fig. 7 for the Primary and High School networks and in Fig. 14 for the other graphs.

On all the graphs, one observed a high correlation between our suggestion curve and the curves of MDS and total persistence. This is expected since all these features are related to the persistence barcodes of the zigzag modules. We also observe that the peak at $r = 18$ of our suggestion curve for the Primary School corresponds to a global minimum of the lifecycle. The same occurs with $r = 12$ for the High School network.

Regarding stability and fidelity, we observe that the curves $r \mapsto S(r)$ and $r \mapsto F(r)$ are convex and no abrupt change can be observed. Instead, we consider a relevant feature that can be defined from them: the point of intersection between the normalized stability and the inverse of the normalized fidelity reveals the resolution value that balances the best between these antagonists' values. The resolutions suggested by this strategy are $r = 24$ and $r = 12$ for the Primary and High School, respectively. Note that ZigzagNetVis also suggests $r = 12$ for the High School.

In conclusion of this section, the peaks of the suggestion curve can, most of the time, be mapped to peaks or bumps of other features in the literature. We stress that our analysis does not reveal the exact nature of this connection; we simply showed how the suggestion curve can be understood as related to other features.

C.2.4 Comparison with Wasserstein distance

As discussed in the Sec. 4.2 of our article, in the context of PH, a concurrent distance to the bottleneck is the Wasserstein distance. We represent in Fig. 9 the suggestion curves obtained from the Wasserstein distance, on the Primary School dataset, for orders 1, 2 and 10. The five suggested resolutions are respectively $\{18, 40, 64, 154, 282\}$, $\{8, 18, 40, 154, 282\}$ and $\{8, 18, 146, 154, 282\}$. This is to be compared with the resolutions $\{8, 18, 76, 154, 282\}$ obtained with the bottleneck

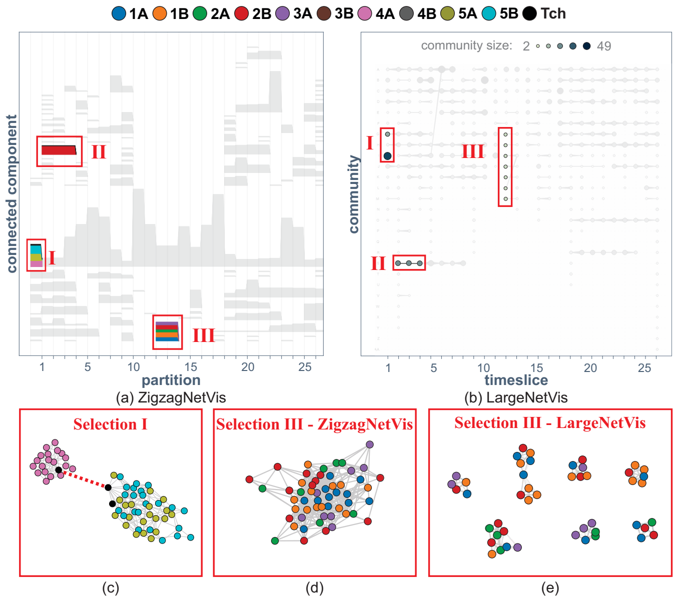

Fig. 8. Comparison between (a) ZigzagNetVis with bottom-based ordering and (b) LargeNetVis highlighting three distinct patterns (I-III) in the Primary School network.

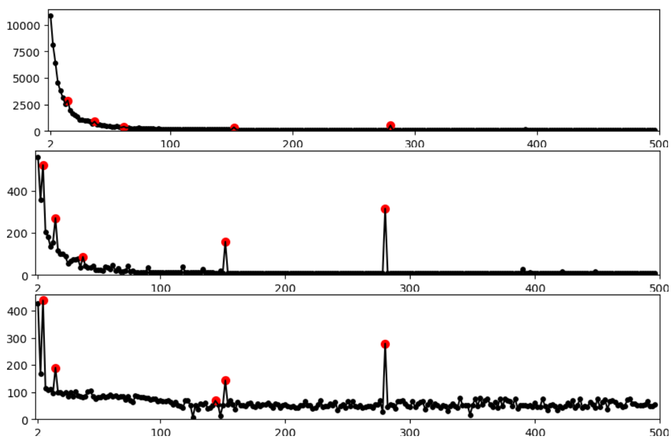

Fig. 9. Normalized suggestion curves with sliding-window timeslicing and Wasserstein distance, for the network Primary School. The distance is computed respectively with order 1, 2 and 10. The x and y axes represent the resolution values and the consecutive Wasserstein distances.

distance (Fig. 10(Primary School)). Although yielding similar resolutions, the peaks of the suggestion curves obtained with the Wasserstein distance appear smaller or less isolated than that of bottleneck distance.

C.3 Running times

Tab. 1 depicts our average running time of 10 executions for every procedure, i.e., the average running time needed to **open the dataset (1)**, **compute the suggestion curve (2)**, and **compute the colored barcode for a given resolution (3)**. The experiments were performed on a personal computer with Intel(R) Core(TM) i5-8350U x 8 CPU @ 3.60GHz, 16 GB RAM, and Ubuntu 22.10. The table considers the eight networks and maximal time values from the main document, Sec. 9.

C.4 Comparison with LargeNetVis

To further validate the colored barcode, we performed a direct comparison with LargeNetVis [7], an established approach to visualize large temporal networks. To be coherent with the partition timeslicing used

Table 1. Running times in seconds for eight distinct networks.

Network	Step 1	Step 2	Step 3
Primary School [2]	31	241	7
High School [8]	53	79	4
Hospital [14]	9	5	1
InVS [3]	3	5	1
Museum [5]	3	12	1
Enron [6]	10	78	5
Conference [5]	6	26	1
Sexual [11]	10	571	600

by LargeNetVis, we decided to use partition timeslicing in ZigzagNetVis as well (see Sec. B.2 for a visual comparison between partition and sliding timeslicing). We also forced the number of timeslices in LargeNetVis to be equal to the number of partitions in ZigzagNetVis for a fair comparison.

Using the Primary School network, Fig. 8 shows a comparison between our colored barcode and LargeNetVis' Global View, also showing node-link diagrams that support the comparison. The first highlighted pattern refers to a single connected component on ZigzagNetVis containing students and teachers from three classes (4A, 5A, 5B) (Fig. 8(a,I)) and the two equivalent communities from LargeNetVis (Fig. 8(b,I)). When analyzing the corresponding node-link diagrams (Fig. 8(c)), we see that these two communities form a single connected component thanks to a single edge (dotted in red) linking two teachers (task T1). Also, we see that students from class 4A interact with each other but not with the other two classes (5A and 5B); on the other hand, students from 5A interact with students from 5B and vice-versa (T1). This finding is supported by the fact that students in the same class interact more often with themselves than with students in other classes and that the same goes for same-grade students. Note also that LargeNetVis only allows us to identify the school classes that are found in a community through the node-link diagram. However, this information is immediate with ZigzagNetVis' colored barcode.

Regarding the second pattern (Fig. 8(a-b,II)), both layouts were able to faithfully represent the continuous level of interactions involving students in class 2B and their teacher (task T2). Once again, note that the colored barcode already shows the school class involved in the interactions. Regarding the third pattern, the colored barcode highlights a component that contains students from five different classes interacting with each other during lunch break (Fig. 8(a,III)). When analyzing the node-link diagram (Fig. 8(d)), we see several interactions between these students during that period, i.e., the component is strongly connected (task T1). In LargeNetVis, due to the nature of the community detection algorithm, this strongly connected component was divided into seven small communities, which impaired the finding of this strongly connected group (see Fig. 8(b,III) and Fig. 8(e)).

Each layout has advantages and disadvantages depending on the user task. We aimed to demonstrate that our approach compares to well-validated visualizations, producing equally relevant results.

D User study

D.1 Complete questionnaire

The questionnaire used in the user study was originally written in Brazilian Portuguese, in which all participants were fluent. The questions were translated into English in this document.

Background and experience

- What is your age group? Choose one option: (i) Between 18 and 24 years old; (ii) Between 25 and 34 years old; (iii) Between 35 and 44 years old; (iv) Between 45 and 64 years old; (v) More than 65 years old;

- Are you aware of any visual difficulties you may have?

- What area is your education in (e.g., computer science, statistics)?

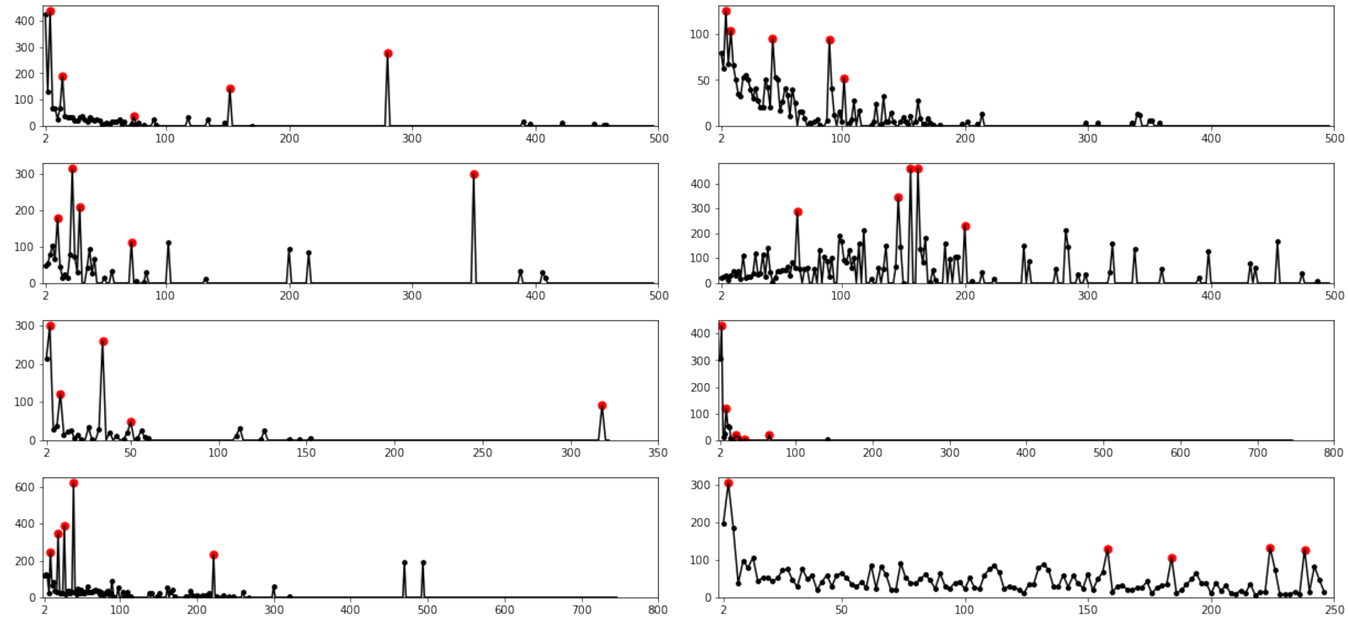

Fig. 10. Normalized suggestion curves, using sliding-window timeslicing, for the networks Primary School, High School, Hospital, InVS, Museum, Enron, Conference, and Sexual, from left-to-right then top-to-bottom. The x and y axes represent the resolution values and the consecutive bottleneck distance, respectively.

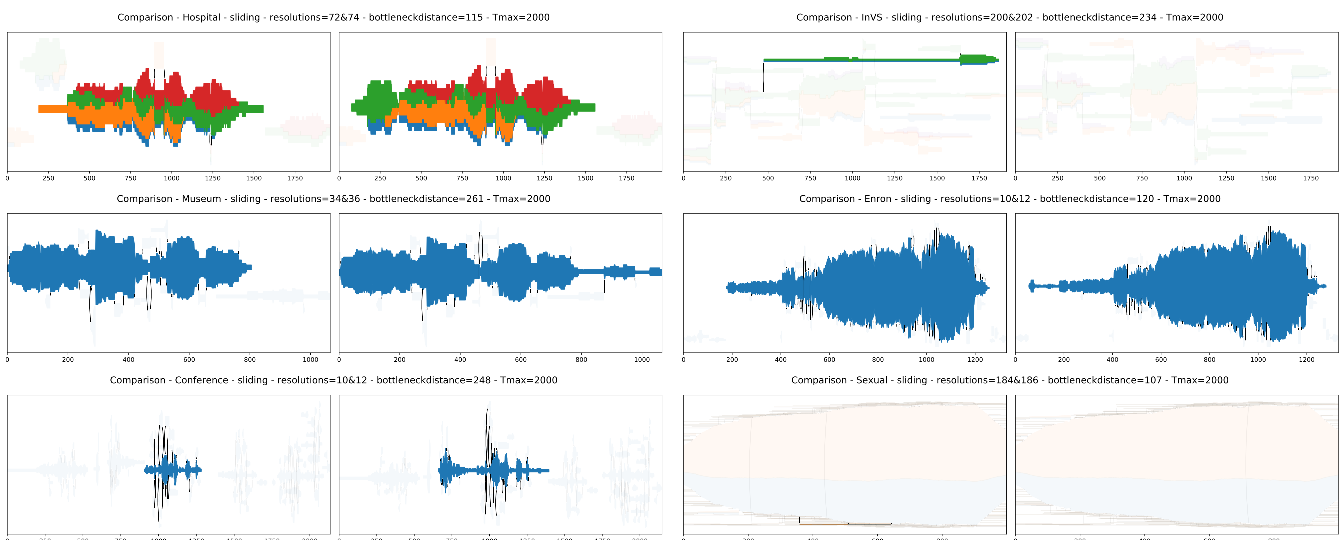

Fig. 11. Visualization of the bottleneck distance for the networks Hospital, InVS, Museum, Enron, Conference, and Sexual, from left-to-right then top-to-bottom. The x and y axes represent the resolution values and the consecutive bottleneck distance, respectively. Highlighted connected components represent the bars that differ the most when considering the two selected resolutions, according to the bottleneck distance.

- What is your most relevant academic title/function? Choose one option: (i) I'm an undergraduate student; (ii) I'm pursuing my master's degree; (iii) I'm a Ph.D. student/candidate; (iv) I'm a postdoctoral researcher; (v) I'm a professor.
- What is your degree of prior knowledge in the Information Visualization field? Choose one option: None, Basic, Intermediate, and Advanced knowledge.
- What is your degree of prior knowledge in the Network Science field? Choose one option: None, Basic, Intermediate, and Advanced knowledge.
- What is your degree of prior knowledge in the Topological Data Analysis field? Choose one option: None, Basic, Intermediate, and Advanced knowledge.
- What is your degree of prior knowledge in the Informatics in Education field? Choose one option: None, Basic, Intermediate, and Advanced knowledge.
- Briefly explain your experience with the above fields (Visualization, Network Science, Topological Data Analysis, and Informatics in Education).

Hands-on experience (ST1-ST12)

Given the Primary School network with default filters, perform the following tasks:

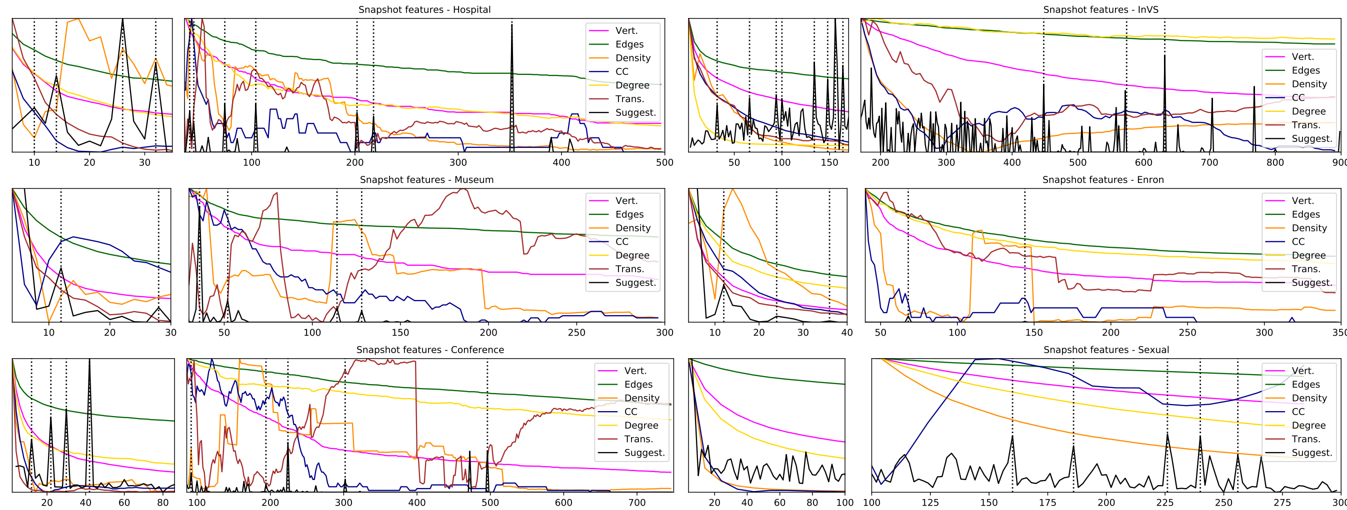

Fig. 12. Suggestion curve and derivative of the mean snapshot features for the networks Hospital, InVS, Museum, Enron, Conference, and Sexual, from left-to-right then top-to-bottom. The x-axis represents the resolution values. The curves are normalized, and a few interesting values are highlighted with a dashed line.

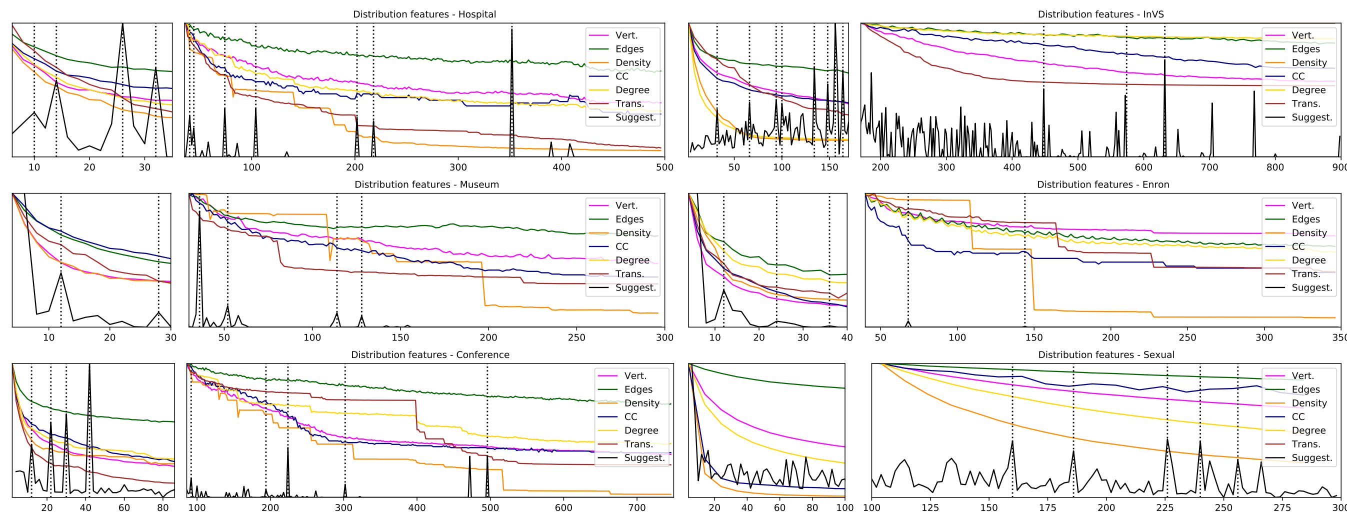

Fig. 13. Consecutive distances between the distribution of the snapshot features, for the networks Hospital, InVS, Museum, Enron, Conference, and Sexual, from left-to-right then top-to-bottom. The x-axis represent the resolution values. The curves are normalized, and a few interesting values are highlighted with a dashed line.

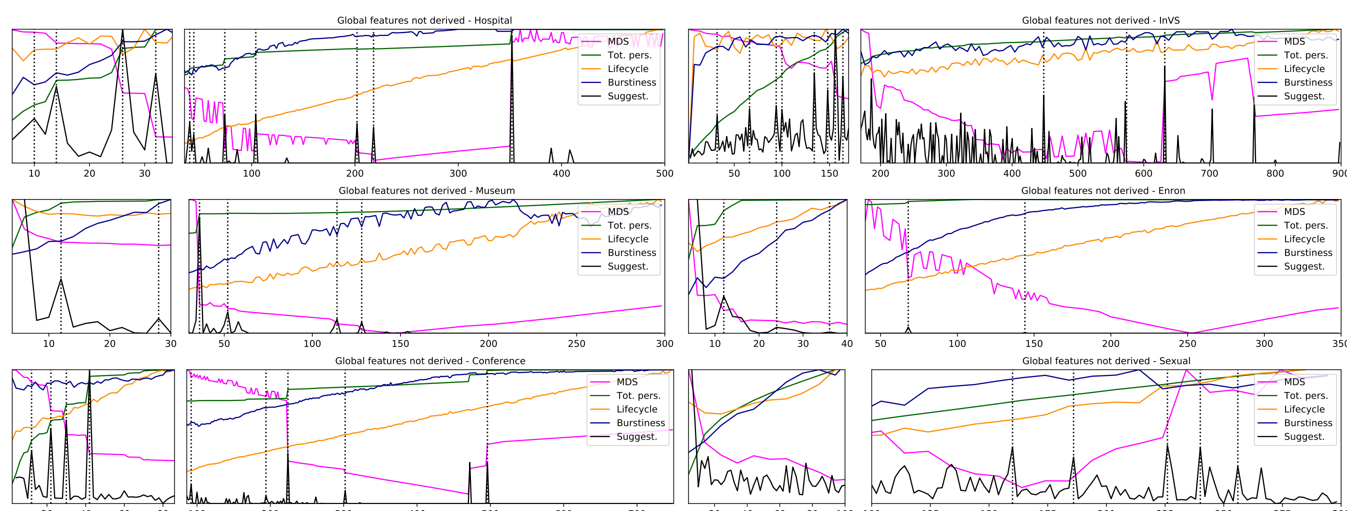

Fig. 14. Global features for the networks Hospital, InVS, Museum, Enron, Conference, and Sexual, from left-to-right then top-to-bottom. The x-axis represents the resolution values. The curves are normalized, and a few interesting values are highlighted with a dashed line.

ST1 Select the students from class 3B on the colored barcode. Action: Click on the desired barcode using the mouse's left button.

ST2 Select the students from class 1B using the color legend. Action: Click on the desired color legend label using the mouse's left button.

ST3 Remove all selection. Action: Click on a blank screen space using the mouse's right button.

ST4 Zoom in and zoom out on the colored barcode. Action: Scroll the mouse wheel up (in) and down (out).

ST5 Go back to the default zoom level. Action: Click on a blank screen space using the mouse's right button.

ST6 Select the connected component that initiates with students from classes 4A, 5A, 5B, and Teachers, and describe patterns and behaviors perceived after visually analyzing that connected component. Action: Find the connected component that contains the students 4A, 5A, 5B, and Teacher and select it using the mouse's left button.

ST7 Select the timestamps 200, 300, and 400 by moving the timestamp markers on the colored barcode. Action: Double-click on the colored barcode to show the three timestamp markers and then drag and drop each one to the desired position.

ST8 Select any connected component in any node-link diagram and describe what occurs to it in the other two timestamps. Action: First, select the flag "Select by connected components" and then click on a node in the node-link diagram using the mouse left button to select a specific component.

ST9 Remove all selections again and perform zoom-in and zoom-out in any node-link diagram. Action: Click in the node-link diagram using the mouse's right button to reset the selection and use the mouse wheel to perform zoom-in (up) and zoom-out (down).

ST10 Change the timestamp of the second node-link diagram (i.e., the one in the middle) from 300 to 350 by typing the new value. Action: Type 350 on the textbox positioned in the middle node-link diagram and click OK or press Enter to change the selected timestamp.

ST11 Expand two node-link diagrams and position them side-by-side. Action: Click on the expand button (⤡) positioned on the top-left portion of the node-link diagram, and drag and drop the opened window to position it.

ST12 Define temporal resolution and give an example. Action: No action is required in the system.

Six questions for the Primary School network (SQ1-SQ6)

Given the Primary School network with resolution $r = 76$ (SQ1-SQ3), answer:

SQ1 Consider the bars and node-link diagrams that refer to the students and teachers from the first and large connected component (i.e., the one containing students from classes 4A, 5A, 5B, and their teachers). Around timestamp 50, how is the relationship between students and teachers of different classes?

Expected answer: There is a strong relationship between the students from class 5A and 5B and their respective teachers, forming a cluster. Also, there is a strong relationship between the students of class 4A and their teacher, forming a second cluster. These two clusters connect to each other through a single edge involving two teachers, leading to one connected component.

SQ2 Select three timestamps to compare: 640, 710, and 800. What can you tell about the relationship between students and teachers when comparing those three timestamps?

Expected answer: There was a strong relationship between all students and teachers in the first timestamp (640). However, in the next timestamp (710), the students were divided into two groups, consisting of younger students (1A, 1B, 2A, 2B, and partially 3A) and older students (partially 3A, 3B, 4A, 4B, 5A, 5B), but with no interaction from the teachers. At the last timestamp (800), the students merged again and were highly connected, but again without the presence of their teachers.

SQ3 Tell the classes that, at the beginning of their activities, were increasing in size over time. Look only at the beginning of the network (from the beginning until around timestamp 100).

Expected answer: Almost every class that is divided into a single connected component has this behavior of starting with just a few students and increasing this size over time.

Given the Primary School network with resolution $r = 154$ (SQ4-SQ6), answer:

SQ4 Tell the classes that, at the beginning of their activities, were increasing in size over time. Look only at the beginning of the network (from the beginning until around timestamp 100). Notice that this question was already answered for the same network but using a different resolution.

Expected answer: In the case of this resolution, almost no class has this pattern visible except for class 2B.

SQ5 Considering the last question, according to your perception, did both resolutions lead to the same answer? In a negative case, why do you think the answers differed for the two distinct resolutions?

Expected answer: No, both resolutions had very different answers. In the case of the smaller resolution ($r = 76$), many components followed this behavior (increasing size over time). On the other hand, this pattern almost disappeared when using $r = 154$.

SQ6 Based on the currently selected network and resolution, freely explore the system and try to find new patterns or anomalies. If there are any, mention findings you consider relevant and tell us which part of the visualization helped you find them.

Expected answer: No expected answer since the participants were free to explore the system and find new patterns or anomalies.

Three questions for the High School network (SQ7-SQ9)

Given the High School network with $resolution = 46$, answer:

SQ7 Can you identify two distinct connected components containing students from the same class at the same timestamp? If so, cite a class and at which timestamp this behavior occurs.

Expected answer: Multiple cases highlight this pattern, such as 2BIO3 at any timestamp between 614 and 725, or 2BIO2 at any timestamp between 586 and 725.

SQ8 Analyze the relationship of students from the class MP2 after timestamp 810.

Expected answer: After timestamp 810, until the end of the network, there exist only three students from class MP2 with connections.

SQ9 Based on the currently selected network and resolution, freely explore the system and try to find new patterns or anomalies. If there are any, mention findings you consider relevant and tell us which part of the visualization helped you find them.

Expected answer: No expected answer since the participants were free to explore the system and find new patterns or anomalies.

Questions to compare ZigzagNetVis with other techniques

- Have you tried to do analyses similar to those described in this study? Options: Yes or no.
- (Optional if the previous answer is yes) What systems/techniques do you commonly use? Do you prefer ZigzagNetVis or the other systems/techniques you know? Why?

Likert-scale questions (LQ1–LQ10)

We used two 5-point Likert-scale questionnaires to assess the participants' preferences about ZigzagNetVis and the provided visual components (LQ1 – LQ6) and specific tasks (LQ7 – LQ10). For each question below, the participants should choose between (i) Strongly disagree; (ii) Disagree; (iii) I don't know; (iv) Agree; (v) Strongly agree.

LQ1 The colored barcode is useful and helps when analyzing the networks.

LQ2 The node-link diagrams are useful and help when analyzing the networks.

LQ3 The interaction and coordination between the views are useful and help when analyzing the networks.

LQ4 ZigzagNetVis is intuitive and easy to use.

LQ5 ZigzagNetVis is useful.

LQ6 ZigzagNetVis is fast (i.e., the provided interactions work in a satisfactory time).

LQ7 It is easy to understand the temporal evolution of connected components and particular classes, in terms of what happens with groups that interact with one another over time, when using ZigzagNetVis.

LQ8 It is easy to compare the network structure at different timestamps when using ZigzagNetVis.

LQ9 It is easy to analyze the network structure at a node level when using ZigzagNetVis.

LQ10 It is easy to analyze the network under different resolutions when using ZigzagNetVis.

Questions to collect the participants' feedback

- What are the most useful visual aids offered by the ZigzagNetVis system? Why?
- What are the most useful visual aids offered by the ZigzagNetVis system? Why?
- What other visual aids could be helpful if incorporated into the ZigzagNetVis system?
- Do you have any final comments?

D.2 Interactive features - complete analysis

Since we recorded the participants' screens, we validated the functionalities mainly used for some questions. Note that since the questions for the same network are sequential (e.g., SQ1-SQ3), in some cases, the participant did not need to interact to find new patterns; the interaction from the previous question might have helped them to answer the current one. We consider ten possible interactions categorized into general, colored barcode, and node-link diagram (Tab. 2): users can select connected components (III) or nodes with the same label (II) — in this case by clicking on the color legend (I) or the colored barcode (IV); users navigate throughout time by moving the timestamp markers (V) or by typing the new value on a node-link diagram (VII); they can pan and zoom in/out on the colored barcode (VI) and diagrams (X); and they can expand the diagrams (VIII) and use the informative tooltip (IX). Tab. 2 presents the percentage of participants who used these interactive features for questions SQ1-SQ3, SQ6, and SQ9. We chose these questions (the first three of the primary school and the two exploratory ones) as we believe they are sufficient to understand the users' behaviors on defined and exploratory tasks.

Participants interacted with the system differently. For instance, in the already mentioned SQ1, all participants made selections using connected components (justified by the question description) and zoomed in the node-link diagram to understand the relationship between classes. Since they were asked to select "around timestamp 50", some participants chose to move the timestamp markers, others typed the timestamp value on the node-link diagram, and others did both (Tab. 2(SQ1, V and VII)). For the exploratory questions, moving the timestamp markers was the best option for most participants (Tab. 2(SQ6 and SQ9, V)). Overall, they interacted more often with the colored barcode than with the diagrams. On average, the feature mainly used in the node-link diagrams was zoom (41.48%), which is justified by the small size of nodes and edges initially applied. Not least, the similar rate of usage involving selection by label (44.44%) and by component (41.48%) indicates that both were appreciated.

Table 2. Percentage of participants who used each of the ten possible interactive features to answer five questions. Features used by more than 50% of participants are highlighted in bold.

	General			Colored barcode			Node-link diagrams			
	I	II	III	IV	V	VI	VII	VIII	IX	X
SQ1	22.22	37.03	**100**	25.92	**59.25**	14.81	**59.25**	25.92	29.62	**81.48**
SQ2	25.92	29.62	40.74	22.22	22.22	25.92	**85.18**	18.51	25.92	40.74
SQ3	37.03	**70.37**	18.51	48.14	44.44	**51.85**	25.92	7.4	3.7	18.51
SQ6	7.4	**55.55**	22.22	44.44	**66.66**	33.33	3.7	14.81	7.4	37.03
SQ9	18.51	29.62	25.92	18.51	**51.85**	25.92	3.7	11.11	3.7	29.62
Avg	22.22	44.44	41.48	31.85	48.15	30.37	35.55	15.55	14.07	41.48

REFERENCES

[1] A. Chiappori and R. Cazabet. Quantitative evaluation of snapshot graphs for the analysis of temporal networks. In R. M. Benito, C. Cherifi, H. Cherifi, E. Moro, L. M. Rocha, and M. Sales-Pardo, eds., *Complex Networks & Their Applications X*, pp. 566–577. Springer International Publishing, Cham, 2022.

[2] V. Gemmetto, A. Barrat, and C. Cattuto. Mitigation of infectious disease at school: targeted class closure vs school closure. *BMC infectious diseases*, 14(1):695, Dec. 2014. doi: 10.1186/PREACCEPT-6851518521414365

[3] M. Génois, C. L. Vestergaard, J. Fournet, A. Panisson, I. Bonmarin, and A. Barrat. Data on face-to-face contacts in an office building suggest a low-cost vaccination strategy based on community linkers. *Network Science*, 3:326–347, 2015. doi: 10.1017/nws.2015.10

[4] M. Hajij, B. Wang, C. Scheidegger, and P. Rosen. Visual detection of structural changes in time-varying graphs using persistent homology. In *2018 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 125–134, 2018. doi: 10.1109/PacificVis.2018.00024

[5] L. Isella, J. Stehlé, A. Barrat, C. Cattuto, J.-F. Pinton, and W. Van den Broeck. What's in a crowd? analysis of face-to-face behavioral networks. *J Theor Biol*, 271(1):166–180, 2011. doi: 10.1016/j.jtbi.2010.11.033

[6] P. S. Keila and D. B. Skillicorn. Structure in the enron email dataset. *Comput. Math. Organ. Theory*, 2005. doi: 10.1007/s10588-005-5379-y

[7] C. D. G. Linhares, J. R. Ponciano, D. S. Pedro, L. E. C. Rocha, A. J. M. Traina, and J. Poco. LargeNetVis: Visual exploration of large temporal networks based on community taxonomies. *IEEE Trans. Visual Comput. Graphics*, pp. 1–11, 2022. doi: 10.1109/TVCG.2022.3209477

[8] R. Mastrandrea, J. Fournet, and A. Barrat. Contact patterns in a high school: A comparison between data collected using wearable sensors, contact diaries and friendship surveys. *PLOS ONE*, 10(9):1–26, 09 2015. doi: 10.1371/journal.pone.0136497

[9] A. Myers, D. Muñoz, F. A. Khasawneh, and E. Munch. Temporal network analysis using zigzag persistence. *EPJ Data Science*, 12(1):1–19, 2023.

[10] G. K. Orman, N. Türe, S. Balcisoy, and H. A. Boz. Finding proper time intervals for dynamic network extraction. *J. Stat. Mech: Theory Exp.*, 2021(3):033414, 2021.

[11] L. E. C. Rocha, F. Liljeros, and P. Holme. Simulated epidemics in an empirical spatiotemporal network of 50,185 sexual contacts. *PLoS Comput. Biol.*, 7(3):e1001109, 03 2011.

[12] L. E. C. Rocha, N. Masuda, and P. Holme. Sampling of temporal networks: Methods and biases. *Phys. Rev. E*, 96:052302, Nov 2017. doi: 10.1103/PhysRevE.96.052302

[13] A. E. Sizemore and D. S. Bassett. Dynamic graph metrics: Tutorial, toolbox, and tale. *NeuroImage*, 180:417–427, 2018.

[14] P. Vanhems, A. Barrat, C. Cattuto, J.-F. Pinton, N. Khanafer, C. Régis, B.-a. Kim, B. Comte, and N. Voirin. Estimating potential infection transmission routes in hospital wards using wearable proximity sensors. *PLoS One*, 8:e73970, 2013. doi: 10.1371/journal.pone.0073970