

Evaluation of Two Interaction Techniques for Visualization of Dynamic Graphs

Paolo Federico^(✉) and Silvia Miksch

Institute of Software Technology and Interactive Systems, TU Wien,
Favoritenstrasse 9-11, 1040 Vienna, Austria
{federico,miksch}@ifs.tuwien.ac.at

Abstract. Several techniques for visualization of dynamic graphs are based on different spatial arrangements of a temporal sequence of node-link diagrams. Many studies in the literature have investigated the importance of maintaining the user's mental map across this temporal sequence, but usually each layout is considered as a static graph drawing and the effect of user interaction is disregarded. We conducted a task-based controlled experiment to assess the effectiveness of two basic interaction techniques: the adjustment of the layout stability and the highlighting of adjacent nodes and edges. We found that generally both interaction techniques increase accuracy, sometimes at the cost of longer completion times, and that the highlighting outclasses the stability adjustment for many tasks except the most complex ones.

Keywords: Network visualization · Dynamic graphs · Interaction · Evaluation · User-study · Time-oriented data

1 Introduction

Dynamic graphs can be used to model different complex real-world phenomena and, therefore, are collected and analysed in various disciplines. Visualization is an indispensable mean to make sense of this type of time-oriented data and gain valuable insights about the phenomena they represent. In recent years, the research about visualization of dynamic graphs has seen a rapid growth, with many novel approaches, techniques, and systems, as surveyed by recent reviews. Likewise, many evaluation studies have investigated those visualizations, to understand which are the key design factors, how they are perceived by users, and how they can support users in analysing data and solving their tasks. The evaluation of dynamic graph visualization has focused mainly on two aspects: comparing animation versus static views, and assessing the importance of the mental map preservation. These studies have often found conflicting results, or a high variability of results depending on different tasks. The use of interaction, in order to control the amount of mental map preservation, or to switch from animation to static views, has been proposed as a means to increase the applicability of a given visualization to diverse tasks. Nevertheless,

few evaluation studies have focused on the role of interaction in dynamic graph visualization: usually static views are considered as noninteractive, while for animated views the most contemplated interaction is the playback control. To fill this gap, we focused on the mental map preservation and its interactive control, which we empirically evaluated in comparison with another common interaction such as highlighting. Thus, the contribution of this paper is a controlled task-based experiment to quantitatively evaluate two interaction techniques for dynamic graph visualization, namely the interactive control of the mental map and the interactive highlighting of adjacent nodes and links. In the following, we review the related work; list the hypotheses, the design, and the settings of our study; present the results and discuss their implications.

2 Related Work

Herman et al. [23] provide an early survey on graphs in information visualization, focusing on layout algorithms for both the general case and special subclasses (e.g., planar graphs, trees) as well as on techniques for interactive navigation, in particular focus+context and clustering. The state of the art report by von Landesberger et al. [41] offers an updated and extensive review of the field; it has a particular focus on issues and solutions for large scale graphs, but contains sections on dynamic graphs as well as interactions. Kerracher et al. [25] explore, and outline a structure of, the design space of dynamic graph visualization. Archambault et al. [3] also review the field, discussing in particular multivariate and temporal aspects of networks. A comprehensive survey on visualizing dynamic graphs is found in Beck et al. [10].

The *Layout Stability*. Many existing algorithms for drawing dynamic graphs ensure the layout stability in order to preserve the user's mental map of the graph [30]. This stability can be seen as an additional aesthetic criterion for dynamic graphs, prescribing that the placement of nodes should change as little as possible [16]. The utility of this dynamic aesthetics has been highly disputed in literature and several evaluations have been conducted, from both the algorithmic and the perceptual perspective. As for the algorithmic evaluation, Brandes and Mader [14] compare three approaches: *aggregation* (fixed nodes positions are obtained from the layout of an aggregate of all graphs in the sequence, achieving maximum stability), *anchoring* (nodes are attracted by reference positions), and *linking* (nodes are attracted by instances of themselves in adjacent time slices of the sequence). Their results suggest that the generally preferable approach is linking, that is also the most computationally demanding; a faster alternative is anchoring to an aggregate layout initialized with the previous one in the sequence. Many user studies have been performed to empirically assess the importance of mental map preservation for readability, memorability, or interpretability of dynamic graphs. Archambault and Purchase [6] review many of them. In an early study about readability of direct acyclic graphs (DAGs), Purchase et al. [34] found that the layout stability is beneficial for several tasks. Conversely, a similar study about readability of DAGs by Zaman et al. [46] found

no significant effect of the layout stability. Purchase and Samra [35] tested several interpretation tasks for directed graphs and found that extremes (no stability or maximum stability) are better than a medium stability. Conversely, Saffrey and Purchase [37], by investigating readability and interpretability of directed graphs, found that the layout stability does not provide any advantage and can be even harmful for certain tasks. While all the evaluations mentioned so far were conducted on timeline based visualization, Ghani et al. [20] studied the effects of layout stability in readability of animated node-link diagrams, finding that a fixed layout (maximum stability) outperforms a force-directed layout with no stability. Analogously, by studying the memorability of animated node-link diagrams, Archambault and Purchase [5] found that maximum layout stability was the best condition.

User Interaction in Dynamic Graph Visualization. User interaction is, by definition, a crucial aspect of Information Visualization [15, page 6]. Various motivated calls have been issued to establish a “Science of Interaction” to complement Information Visualization [32]. Yi et al. [45] propose a taxonomy of interaction in Information Visualization based on the notion of user intent; Lam [27] introduces a theoretical framework to understand and possibly reduce the costs of interaction. Nevertheless, the importance of user interaction in dynamic graph visualization is generally underestimated in literature [10]; timeline approaches are generally considered as sequences of static (i.e., non interactive) drawings, while the most discussed interaction for animation approaches deals with animation control (e.g., play/pause, or time seeker). Wybrow et al. [44] review interaction techniques for multivariate graphs and propose a classification based on the Information Visualization Reference Model [15], distinguishing among view level, visual-representation level, and data level. Notable examples include a technique for selecting and manipulating subgraphs [29] or a network-aware navigation integrating “Link sliding” (guided panning when dragging along a visible link) and “Bring and Go” (bringing adjacent nodes nearby when pointing to a node), with the latter having the best performance [31]. Another example of evaluating interaction in dynamic graph visualization is provided by Rey and Diehl [36], who investigate the effects of two interaction techniques for animated visualization: interactive control of the animation speed and a tooltip showing details on demand. They found that the speed control does not provide a significant benefit, and the tooltip is outperformed by a visualization having labels always visible.

Given the high variability in the importance of the mental map (depending on tasks, user preferences, and possibly other factors), some scholars have proposed an interactive control of the layout stability, to let the user fine-tune it [8, 19]. According to the taxonomy of interaction by Yi et al. [45], interactive control of stability can be understood as a *Reconfigure* interaction, corresponding to the user’s intent: “Show me a different spatial arrangement”. It falls into the class of user-controlled adjustments of the layout, which are very common visual-level interactions for graphs [44]. Bach et al. [8] evaluated this stability slider in the context of a specific technique (GraphDiaries) featuring staged animations of

node-link diagrams, but they did not consider it as an independent factor in the study design. Smuc et al. [38] also evaluated a graph visualization featuring a stability slider, but without a specific focus on it.

Layout stability has been also described as a form of spatial highlighting, where position is used to identify different instances of the same node over time [7]. *Highlighting*, in the stricter sense, is a *brushing* interaction technique, originally developed for scatter plots [11], and then extended and applied to other visualization techniques. Brushing is a “change in the encoding of one or more items essentially immediately following, and in response to, an interaction with another item” [39, p. 235]. In particular, in the case of highlighting, the change may affect hue, brightness, or color. Brushing operates within a view or across multiple views; in the latter case, the interaction technique is better known as *linking and brushing* [24]. Highlighting makes some information stand out from other information; it effectively exploits pre-attentive processing [42], which is the human capability to process visual information prior to, or in the early stage of, focusing conscious attention. Linking and brushing techniques support two user’s intents: *Select*, i.e., “Mark something as interesting” and *Connect*, i.e., “Show me related items”, according to the interaction taxonomy by Yi et al. [45]. In the context of graph visualization, highlighting of adjacent nodes upon selection of a certain node (for example, by mouse hovering) is a common interaction technique, also known as *connectivity highlighting* [21]. An experiment by Ware and Bobrow has shown that interactive highlighting can efficiently support visual queries on graphs [43]. In the case of timeline visualization of dynamic graphs, the highlighting technique can be extended in order to fulfil the need of visually linking and synchronizing multiple instances of the same graph entities in different time slices [10], by considering adjacency not only across the graph structure, but also along the time dimension.

Evaluation of Other Aspects in Graph Visualization. Besides the importance of preserving the mental map in node-link diagrams, another issue which attracts the interest of scholars is the comparison between animation approaches and timeline approaches, the latter being usually based on small multiples in a juxtaposition arrangement. The controlled experiment by Farrugia and Quigley [17] found that static drawings outperform animation in terms of task completion time. Archambault et al. [4], in an analogous user study, found that small multiples are generally faster, but more error-prone for certain tasks; moreover, mental map preservation has little influence on both response time and error rate. Boyandin et al. [13] also conducted a comparative evaluation of animation versus small multiples in the context of flow maps. They found that with animation there were more findings of changes in adjacent time slices, where with small multiples there were more findings about longer time periods. Moreover, they suggest that switching from one view to the other might lead to an increase in the numbers of findings; we see this consideration in analogy with the mental map case, where the introduction of a stability slider might allow the user to adapt the layout to a particular task and possibly increase the overall visualization effectiveness.

Task Taxonomies. A profound understanding of analytical tasks is a necessary prerequisite to design novel visualization techniques as well as evaluate existing ones. Ahn et al. [1] propose a task taxonomy for dynamic graphs along three different axes: graph entities, temporal features, and properties (structural and domain-specific). According to Bach et al. [8], each task can be understood as a question containing references to two dimensions and requiring an answer in the third one. In this way, they distinguish between topological tasks, temporal tasks, and behavioural tasks. Archambault and Purchase [3] structure their taxonomy along two dimension, mostly aiming at assessing the importance of the mental map. They distinguish between local and global tasks, and between distinguishable and undistinguishable tasks (depending on whether graph entities need to be distinguished from each other or can be aggregated). A task taxonomy for multivariate networks can be found in [33].

Open Challenges. Summarizing our review of related work on visualization of dynamic graphs, we can observe that many techniques have been designed and evaluated, but the research community lacks final and well-established results about highly disputed issues, such as the importance of the layout stability, or the comparison between animation and timeline approaches. In both cases, it has been suggested that enabling users to interactively switch between different views might broaden the number of tasks that they can efficiently solve. Hence, there is a commonly recognised need of understanding the role of interaction in the context of dynamic graph visualization, but only few studies have specifically focused on the evaluation of interaction techniques.

3 Our Evaluation

Addressing the aforementioned need, we performed a user study to explore interaction in the context of dynamic graphs visualization. In particular, we considered a timeline visualization with the juxtaposition approach, where several node-link diagrams are arranged along a horizontal timeline (Fig. 1). We evaluated two interaction techniques. The first one is the interactive control of the layout stability, which is executed by the means of a slider control (thus, for the sake of brevity, we will refer to it as the *Slider*). The second interaction is the highlighting of adjacent nodes, adapted for dynamic graphs (in the following, *Highlighting*). In this section we detail the study design, the stimuli, the tasks, and the settings of our empirical experiment.

Study Design. We designed our user study as a quantitative controlled task-based evaluation, with two observed dependent variables: time and error. We considered two factors, i.e. independent variables: the presence of the *Slider* interaction (2 levels: off/on), and the presence of the *Highlighting* interaction (2 levels: off/on). In other words, we considered four different interfaces: no-interaction, only *Slider*, only *Highlighting*, and both interactions. We chose this design in order to compare the two interaction with each other and with the non-interactive baseline, and also to assess how the two interactions work together.

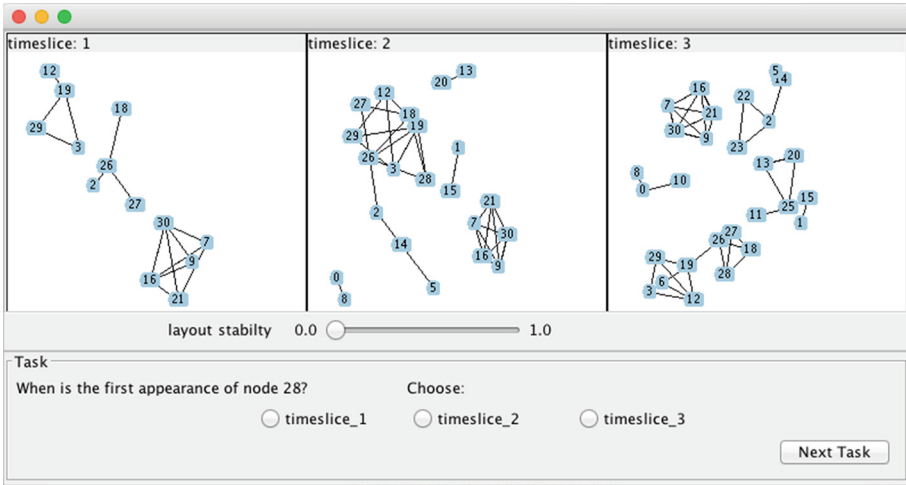


Fig. 1. The remote evaluation software displays stimuli, provides instructions, and measures time and error.

We tested 6 *task* types. The full factorial design led to a total amount of $N = Task \times Slider \times Highlighting = 6 \times 2 \times 2 = 24$ conditions. To mitigate the effects of personal skills and preferences, we chose a within-subject design; each subject tested 24 conditions, by solving a different task for each of the six task type and for each of the four interfaces. In order to lower the cognitive effort of switching between different interfaces, we grouped conditions by interface. To mitigate the effects of learning and fatigue, we used a Latin square arrangement of the interfaces and we randomized the order of the tasks within each interface. Moreover, we randomized the initial slider position.

Stimuli. We selected as baseline a spring-embedder layout as implemented in the Prefuse visualization toolkit [22] (Fig. 1). According to the *linking* approach [14], we added inter-time links in the graph in order to ensure layout stability. The *Slider* controls the amount of stability by interactively changing the relaxed lengths of inter-time springs. In the implementation of the *Highlighting* technique used for our experiment, when the user moves the mouse pointer over a node, a different combination of fill and stroke colors is used to highlight each different type of “adjacent” graph item, as shown in Fig. 2.

For the experiment we used real-world datasets: the dynamic graphs of social relationship between university freshmen collected by van Duijn, consisting of 38 nodes and 5 time points, and the one collected by van de Bunt, consisting of 49 nodes and 7 time points [40]. Through a threshold mechanism we derived two dynamic unweighted (i.e., binary) graphs from the original dynamic weighted graphs. Each task involved only subsets of 3 time slices.

Tasks. We selected six different types of tasks (Table 1). As a criterion for the selection of the tasks, we considered existing studies on the importance




Fig. 2. A: a dynamic graph over three time slices; B: the same graph highlighted on a mouse-over event.

of layout stability and tried to elicit a set of similar tasks, in order to have comparable results. Furthermore, we considered the task taxonomy by Ahn et al. [1] in order to have a meaningful and representative set of tasks along its three axes. As for the graph entities, we included tasks referring to all the levels: the *entity* level (nodes, links), the *group* level (paths, components), and the *graph* level (the entire graph). As for the properties, we disregarded tasks referring to *domain properties* and only considered tasks referring to *structural properties*, which are specific aspects for graphs. As for the temporal features, we included tasks referring to *individual events* and *contraction & growth*, scoping out more complex tasks, which can be investigated in a follow-up study. In order to better describe the nature of our tasks and to enable a better interpretation of results, we also categorized our tasks according to other existing taxonomies [6, 8], as shown in Table 1.

Table 1. Task types, examples, and classifications

Task	Description	by [1]	[6]	[8]
1.NO	Node occurrence e.g., <i>When is the first appearance of node 27?</i>	Event/Node	Local/Distinguishable	When
2.LO	Link occurrence e.g., <i>When is the last appearance of link 6-9?</i>	Event/Link	Local/Distinguishable	When
3.ND	Node degree e.g., <i>When does the smallest degree (number of connections) of node 10 occur?</i>	Event/Group	Local/Indistinguishable	When
4.SP	Shortest path e.g., <i>When does the largest geodesic distance between node 7 and node 9 occur?</i>	Event/Group	Local/Distinguishable	When
5.CC	Connected components e.g., <i>Is the number of connected components increasing, decreasing, or stable?</i>	Growth/Graph	Global/Indistinguishable	What
6.AL	All links e.g., <i>Is the total number of edges increasing, decreasing, or stable?</i>	Growth/Graph	Global/Indistinguishable	What

Subjects’ Pool and Study Settings. We conducted the experiment remotely by using the Evalbench toolkit [2] (Fig. 1). In order to assess the technical setup, the estimated overall length of the evaluation session, and the understandability of textual descriptions of our tasks, we performed two pilot tests with direct

observation of subjects, and then we implemented small adjustments before the main remote study. For the main study we recruited 64 volunteer subjects among undergraduate students at the fifth semester of a bachelor programme in Visual Computing. All the subjects had normal or corrected vision. Right after the recruiting, we instructed the subjects with a 15 minute briefing, describing the visualization and the interactions to be evaluated, and recalling the necessary concepts from graph theory (e.g., the notion of geodesic distance as shortest path, or the notion of connected components). The subjects were instructed to be fast and accurate in solving the tasks, without assigning any priority between speed and accuracy. The evaluation software included a training session for each of the four interfaces. During the training sessions, the software does not collect data; it shows the correct answer after completion of each task and allows repetitions until the subject feels confident of having understood the task types and the interface. The test, including the training sessions, had an average duration of 20 min.

Hypotheses. We designed our experiment to test three hypotheses: **(A)** the *Slider* reduces error rates at the cost of longer completion times, in comparison with the non-interactive interface; **(B)** the *Highlighting* reduces error rates at the cost of longer completion times, in comparison with the non-interactive interface; **(C)** the *Highlighting* outperforms the *Slider*.

We hypothesize that each interaction reduces error rates in comparison with the non-interactive interface, because both interactions comply with the *rule of self-evidence* and address the *adjacency task*. The *rule of self-evidence* for multiple views prescribes the use of “perceptual cues to make relationships among multiple views more apparent to the user” [9]. The *Highlighting* complies with this rule, by drawing attention to different instances of the same node across different time slices; the *Slider* also complies with this rule, by allowing the user to select the maximum stability and fix node positions across different time slices. The *adjacency task* (i.e., “Given a node, find its adjacent nodes”) has been identified as the only graph-specific task [28]. The *Highlighting* obviously addresses this task, as well as the *Slider* does, by allowing the user to select the minimum stability and exploit the proximity *Gestalt* principle [12]. Conversely, we hypothesize that both interactions increase the task completion time in comparison with the non-interactive interface. We make this hypothesis in analogy with the existing comparative evaluations between animation and (static) timeline views [4, 17], while we consider interactive timeline views as a middle way. More specifically, in terms of interaction costs [27], the *Highlighting* might increase the completion time because of the physical-motion cost of tracking elements with the mouse, while the *Slider* might imply view-change costs of reinterpreting the perception when the layout rearranges. For both techniques, there might be the decision costs of forming goals, such as deciding whether the available interaction is useful to solve the given task, and how. Moreover, the simple fact that the GUI provides an interactive option might lead users to explore its use, in order to form a solving strategy before solving a task, or to possibly increase the confidence about the solution afterwards. Furthermore, we hypothesize that the

Highlighting will have better performance than the *Slider*. We derive this hypothesis from the observation that the *Highlighting* is a common and relatively simple interaction, which at least partially exploits pre-attentive processing, while the *Slider* is based on a novel and complex concept. In other words, while the *Highlighting* directly addresses the issue of connecting entities along two dimensions (time and graph structure), the *Slider* implicitly introduces another dimension, since the stability lies in the parameter space of the layout algorithm.

4 Analysis

We preprocessed data collected from 64 subjects in order to assess whether they were eligible for analysis and we had to discard one subject whose logs were corrupted. The analysis was then performed on data from 63 subjects, consisting of 3024 samples in total. We checked the completion times for normality with the Shapiro-Wilk goodness-of-fit test but the check failed. We then applied a logarithmic transformation to the completion times and checked again the normality with a positive result. The verification of the Gaussian condition assured the applicability of parametric tests; we could perform the analysis of variance through an ANOVA with the subject as a random variable. When the ANOVA found a factor to have a statistically significant effect, we compared the two levels of that factor with a pairwise post-hoc Student's t test; when the ANOVA found the interaction between factors to be statistically significant, we performed an all-pairs Tukey's honestly significant difference (HSD) post-hoc test. The error can be understood as a dichotomous (i.e. binary) variable, since there are only two possible outcomes for each data sample (correct, not correct). Hence, we analysed the error by logistic regression as a generalized linear model (GLM) with a binomial distribution and a logit transformation as the link function, computing likelihood ratio statistics. When a factor was found to be a significant effect, we analysed the contrast between its levels in terms of pairwise comparisons between estimated marginal means.

5 Results and Discussion

Figure 3 shows time and error by *Highlighting* and *Slider*, grouped by *Task*; time is represented by box-plots with first, second (median), and third quartile, while error is represented by bars (mean) and error bars (standard error). Figure 4 shows statistically significant differences. In light of these results, we can verify our hypotheses.

Hypothesis A is partially confirmed. The *Slider* decreases the error rate for all tasks but the easiest one (1.NO), and it increases the completion time for tasks 3.ND and 4.SP only.

Hypothesis B is partially confirmed. The *Highlighting* decreases the error rate for all tasks but the most difficult one (6.AL); it increases completion times for




Fig. 3. Time (left-hand side, as box plots) and error (right-hand side, as bars representing means and error bars representing standard error) by *Highlighting* and *Slider*, grouped by *Task*. $\color{green}\blacksquare$ H^0S^0 $\color{red}\blacksquare$ H^0S^1 $\color{blue}\blacksquare$ H^1S^0 $\color{purple}\blacksquare$ H^1S^1

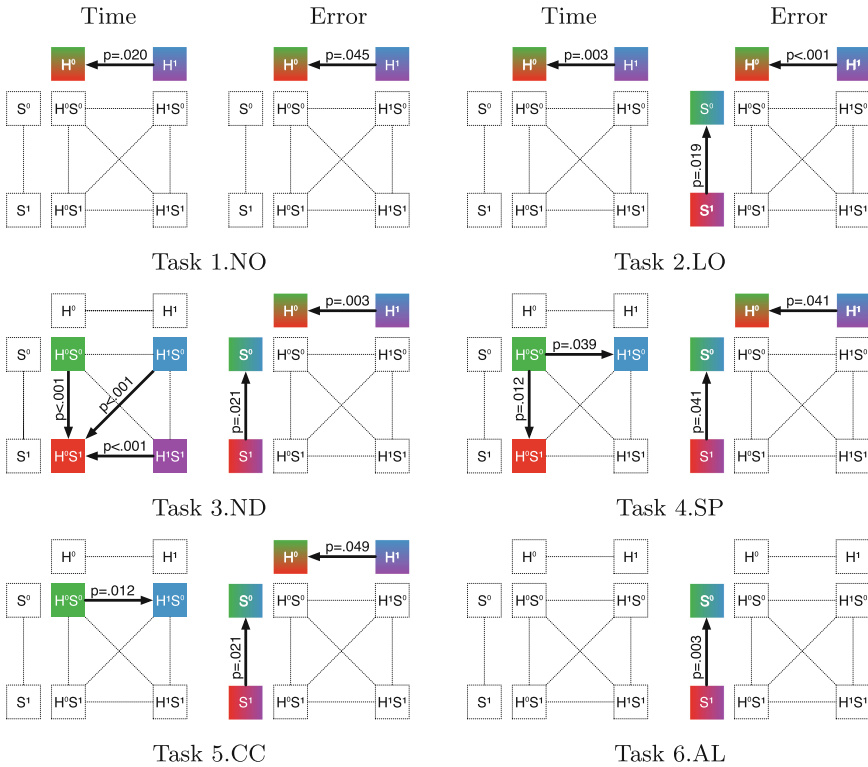


Fig. 4. Statistically significant differences for time and error, by *Task*. An arrow means that the source is faster or, respectively, more accurate than the destination, with the reported probability. Lines represent all-pairs comparisons between factor combinations ($\color{green}\blacksquare$ H^0S^0 , $\color{red}\blacksquare$ H^0S^1 , $\color{blue}\blacksquare$ H^1S^0 , and $\color{purple}\blacksquare$ H^1S^1), as well as pairwise comparisons by *Highlighting* (H^0-H^1 , top) and *Slider* (S^0-S^1 , left).

tasks 4.SP and 5.CC, but it reduces it for task 2.LO, and does not affect the remaining tasks.

Hypothesis C is partially confirmed. The *Highlighting* outperforms the *Slider* for tasks 1.NO, 2.LO, and 3.ND. For task 4.SP, the *Highlighting* and the *Slider* score equally: each of them decreases the error rate (by the same amount) and also increases the completion time if used alone, but when used together they do not increase the completion time, showing a desirable effect interaction. For task 5.CC, both factors reduce the error rate, but the *Highlighting* also increases the completion time when used alone. For task 6.AL, the only significant effect is that the *Slider* reduces the error rate.

Besides the verification of our hypothesis, which are mostly confirmed, our user study provides interesting insights about the differences between tasks. First of all, we observe that the differences in error rate and completion time among the tasks are significant, hence we can confirm that in general our tasks have different levels of difficulty. Secondly, we observe that the effectiveness of the tested interaction techniques varies with the levels of difficulty of the tasks. In a very brief but accurate summary we can say that, for easier tasks, the *Highlighting* decreases error rates and in some cases even decreases completion times; conversely, for more difficult tasks, it is the *Slider* that decreases error rates. Moreover, for tasks 3.ND, 4.SP and 5.CC, one technique increases completion times if used alone, but it does not if used in combination with the other one. Looking back at the classification of our tasks (Table 1), we can also identify the relevant aspects. We can observe that, for those tasks involving simpler temporal features of distinguishable single entities (1.NO and 2.LO), or indistinguishable groups (3.ND), the *Highlighting* is more effective. For those tasks that refer to more complex temporal features at the graph level, even if indistinguishable (5.CC and 6.AL), the *Slider* is more effective. Task 4.SP is about the changes of the geodesic distance between two nodes and requires the distinct identification of several nodes and links along the shortest path. In this case both techniques are equally accurate; if (and only if) they are used together, they do not even slow down the analysis. We can conjecture (by also considering our observations during the pilot experiments) that during the completion of a such complex task, the *Slider* can be used to switch back and forth between the minimum stability (to guess geodesic distances and shortest paths based on Euclidean distances) and the maximum stability (to identify instances of nodes and links across different time slices), while the *Highlighting* helps with tracking objects. As for task 6.AL, the *Slider* resulted to be effective; we hypothesize (by also considering our pilot observations) that subjects simply set the minimum stability and looked at the total graph area as an estimator of the density. We would have expected the *Highlighting* to be also effective, since the analysis of the degree a few central nodes might provide a good estimator of the graph density, given the power-law distribution of real-world networks. The results show that the test subjects did not exploit this expert strategy.

Design Implications. Both the *Slider* and the *Highlighting* are effective interaction techniques for dynamic graph visualization, and their use generally

improves user performances. In those circumstances where it might be not possible to include them both (for example, if the color channel is employed to encode attributes of multivariate graphs, or if the GUI is already overloaded with many controls), our evaluation provides an indication to designers according to the user tasks to be supported. Our results suggest that the *Highlighting* is indicated for tasks involving temporal features of distinguishable single entities or indistinguishable groups, while the *Slider* is indicated for tasks involving complex temporal behaviours at the graph level. The joint use of both interactions is beneficial for the most complex task involving temporal behaviours of connectivity paths.

6 Conclusion

We have presented an evaluation of two interaction techniques for dynamic graph visualization, namely the interactive control of the layout stability and the interactive highlighting of adjacent nodes and links. The results mostly confirm our hypotheses: both interactions decrease the error rate, in some cases at the cost of a longer completion time. We observed significant differences between tasks, with the highlighting performing better for some tasks, and the stability control performing better for others. We acknowledge the limitations of our experiment, whose findings might not be directly generalizable to large-scale datasets. The highlighting interaction for dynamic graphs is much more complex than the standard connectivity highlighting for static graphs, and may require training to be understood and used effectively. The stability control might have a different effect when combined with 3D visualization and interaction techniques (e.g., the vertigo zoom [18]). However, our study provides preliminary clues for visualization designers who need to choose the most appropriate interaction technique for their users' tasks. Further studies are needed to obtain a comprehensive understanding of the role of interaction in visualization of dynamic graphs.

Acknowledgements. The authors wish to thank the anonymous study subjects for their participation, as well as Theresia Gschwandtner, Simone Kriglstein, and Margit Pohl for their feedback on the manuscript.

This work was partially supported by the Austrian Research Promotion Agency (FFG), project *Expand*, grant 835937.

References

1. Ahn, J.W., Plaisant, C., Shneiderman, B.: A task taxonomy for network evolution analysis. *IEEE Trans. Visual. Comput. Graph.* **20**(3), 365–376 (2014)
2. Aigner, W., Hoffmann, S., Rind, A.: EvalBench: a software library for visualization evaluation. *Comput. Graph. Forum* **32**(3pt1), 41–50 (2013)
3. Archambault, D., Abello, J., Kennedy, J., Kobourov, S., Ma, K.L., Miksch, S., Muelder, C., Telea, A.: Temporal multivariate networks. In: Kerren et al. [26], pp. 151–174

4. Archambault, D., Purchase, H., Pinaud, B.: Animation, small multiples, and the effect of mental map preservation in dynamic graphs. *IEEE Trans. Visual. Comput. Graph.* **17**(4), 539–552 (2011)
5. Archambault, D., Purchase, H.C.: The mental map and memorability in dynamic graphs. In: *Proceedings of the Pacific Visualization Symposium, PacificVis 2012*, pp. 89–96. IEEE, Washington, DC (2012)
6. Archambault, D., Purchase, H.C.: The map in the mental map: experimental results in dynamic graph drawing. *Int. J. Hum.-Comput. Studies* **71**(11), 1044–1055 (2013)
7. Archambault, D., Purchase, H.C.: On the application of experimental results in dynamic graph drawing. In: *Proceedings of the 1st International Workshop on Graph Visualization in Practice, GViP 2014*, vol. 1244, pp. 73–77. CEUR-WS (2014)
8. Bach, B., Pietriga, E., Fekete, J.D.: GraphDiaries: animated transitions and temporal navigation for dynamic networks. *IEEE Trans. Visual. Comput. Graph.* **20**(5), 740–754 (2014)
9. Baldonado, M.Q.W., Woodruff, A., Kuchinsky, A.: Guidelines for using multiple views in information visualization. In: *Proceedings of the Working Conference on Advanced Visual Interfaces, AVI 2000*, pp. 110–119. ACM, New York (2000)
10. Beck, F., Burch, M., Diehl, S., Weiskopf, D.: The state of the art in visualizing dynamic graphs. In: Borgo, R., Maciejewski, R., Viola, I. (eds.) *EuroVis - STAR*, pp. 83–103. The Eurographics Association (2014)
11. Becker, R.A., Cleveland, W.S.: Brushing scatterplots. *Technometrics* **29**(2), 127–142 (1987)
12. Bennett, C., Ryall, J., Spalteholz, L., Gooch, A.: The aesthetics of graph visualization. In: Cunningham, D.W., Meyer, G., Neumann, L. (eds.) *Computational Aesthetics in Graphics, Visualization, and Imaging*. The Eurographics Association (2007)
13. Boyandin, I., Bertini, E., Lalanne, D.: A qualitative study on the exploration of temporal changes in flow maps with animation and small-multiples. *Comput. Graph. Forum* **31**(3pt2), 1005–1014 (2012)
14. Brandes, U., Mader, M.: A quantitative comparison of stress-minimization approaches for offline dynamic graph drawing. In: Kreveld, M., Speckmann, B. (eds.) *GD 2011*. LNCS, vol. 7034, pp. 99–110. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-25878-7_11](https://doi.org/10.1007/978-3-642-25878-7_11)
15. Card, S.K., Mackinlay, J.D., Shneiderman, B. (eds.): *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, San Francisco (1999)
16. Coleman, M.K., Parker, D.S.: Aesthetics-based graph layout for human consumption. *Softw.: Pract. Exp.* **26**(12), 1415–1438 (1996)
17. Farrugia, M., Quigley, A.: Effective temporal graph layout: a comparative study of animation versus static display methods. *Inf. Visual.* **10**(1), 47–64 (2011)
18. Federico, P., Aigner, W., Miksch, S., Windhager, F., Smuc, M.: Vertigo zoom: combining relational and temporal perspectives on dynamic networks. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces, AVI 2012*, pp. 437–440. ACM, New York (2012)
19. Federico, P., Aigner, W., Miksch, S., Windhager, F., Zenk, L.: A visual analytics approach to dynamic social networks. In: *Proceedings of the International Conference on Knowledge Management and Knowledge Technologies, i-KNOW 2011*, pp. 47:1–47:8. ACM, New York (2011)
20. Ghani, S., Elmqvist, N., Yi, J.S.: Perception of animated node-link diagrams for dynamic graphs. *Comput. Graph. Forum* **31**(3pt3), 1205–1214 (2012)

21. Heer, J., Boyd, D.: Vizster: visualizing online social networks. In: IEEE Symposium on Information Visualization, INFOVIS 2005, pp. 32–39, October 2005
22. Heer, J., Card, S.K., Landay, J.A.: Prefuse: a toolkit for interactive information visualization. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2005, pp. 421–430. ACM, New York (2005)
23. Herman, I., Melançon, G., Marshall, M.S.: Graph visualization and navigation in information visualization: a survey. *IEEE Trans. Visual. Comput. Graph.* **6**(1), 24–43 (2000)
24. Keim, D.A.: Information visualization and visual data mining. *IEEE Trans. Visual. Comput. Graphics* **8**(1), 1–8 (2002)
25. Kerracher, N., Kennedy, J., Chalmers, K.: The design space of temporal graph visualisation. In: Elmqvist, N., Hlawitschka, M., Kennedy, J. (eds.) *EuroVis - Short Papers*, pp. 7–11. The Eurographics Association (2014)
26. Kerren, A., Purchase, H.C., Ward, M.O. (eds.): *Multivariate Network Visualization*. LNCS, vol. 8380. Springer, Heidelberg (2014)
27. Lam, H.: A framework of interaction costs in information visualization. *IEEE Trans. Visual. Comput. Graph.* **14**(6), 1149–1156 (2008)
28. Lee, B., Plaisant, C., Parr, C.S., Fekete, J.D., Henry, N.: Task taxonomy for graph visualization. In: Proceedings of the AVI Workshop on Beyond Time and Errors: Novel Evaluation Methods for Information Visualization, BELIV 2006, pp. 1–5. ACM, New York (2006)
29. McGuffin, M., Jurisica, I.: Interaction techniques for selecting and manipulating subgraphs in network visualizations. *IEEE Trans. Visual. Comput. Graph.* **15**(6), 937–944 (2009)
30. Misue, K., Eades, P., Lai, W., Sugiyama, K.: Layout adjustment and the mental map. *J. Vis. Lang. Comput.* **6**(2), 183–210 (1995)
31. Moscovich, T., Chevalier, F., Henry, N., Pietriga, E., Fekete, J.D.: Topology-aware navigation in large networks. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2009, pp. 2319–2328. ACM, New York (2009)
32. Pike, W.A., Stasko, J., Chang, R., O’Connell, T.A.: The science of interaction. *Inf. Visual.* **8**(4), 263–274 (2009)
33. Pretorius, A., Purchase, H., Stasko, J.: Tasks for multivariate network analysis. In: Kerren et al. [26], pp. 77–95
34. Purchase, H.C., Hoggan, E., Görg, C.: How important is the “mental map”? – an empirical investigation of a dynamic graph layout algorithm. In: Kaufmann, M., Wagner, D. (eds.) *GD 2006*. LNCS, vol. 4372, pp. 184–195. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-70904-6_19](https://doi.org/10.1007/978-3-540-70904-6_19)
35. Purchase, H.C., Samra, A.: Extremes are better: investigating mental map preservation in dynamic graphs. In: Stapleton, G., Howse, J., Lee, J. (eds.) *Diagrams 2008*. LNCS (LNAI), vol. 5223, pp. 60–73. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-87730-1_9](https://doi.org/10.1007/978-3-540-87730-1_9)
36. Rey, G.D., Diehl, S.: Controlling presentation speed, labels, and tooltips in interactive animations. *J. Media Psychol.: Theor. Methods Appl.* **22**(4), 160 (2010)
37. Saffrey, P., Purchase, H.: The mental map versus static aesthetic compromise in dynamic graphs: a user study. In: Proceedings of the Conference on Australasian User Interface, AUIC 2008, pp. 85–93. Australian Comp. Soc. (2008)
38. Smuc, M., Federico, P., Windhager, F., Aigner, W., Zenk, L., Miksch, S.: How do you connect moving dots? Insights from user studies on dynamic network visualizations. In: Huang, W. (ed.) *Handbook of Human Centric Visualization*, pp. 623–650. Springer, New York (2014)

39. Spence, R.: *Information Visualization: Design for Interaction*. Addison Wesley, Harlow, New York (2007)
40. Van De Bunt, G.G., Van Duijn, M.A.J., Snijders, T.A.B.: Friendship networks through time: an actor-oriented dynamic statistical network model. *Comput. Math. Organ. Theory* **5**(2), 167–192 (1999)
41. von Landesberger, T., Kuijper, A., Schreck, T., Kohlhammer, J., van Wijk, J., Fekete, J.D., Fellner, D.: Visual analysis of large graphs: state-of-the-art and future research challenges. *Comput. Graph. Forum* **30**(6), 1719–1749 (2011)
42. Ware, C.: *Information Visualization: Perception for Design*. Morgan Kaufman, San Francisco (2004)
43. Ware, C., Bobrow, R.: Supporting visual queries on medium-sized nodelink diagrams. *Inf. Visual.* **4**(1), 49–58 (2005)
44. Wybrow, M., Elmqvist, N., Fekete, J.D., von Landesberger, T., van Wijk, J., Zimmer, B.: Interaction in the visualization of multivariate networks. In: Kerren et al. [26], pp. 97–125
45. Yi, J.S., ah Kang, Y., Stasko, Y., Jacko, J.: Toward a deeper understanding of the role of interaction in information visualization. *IEEE Trans. Visual. Comput. Graph.* **13**(6), 1224–1231 (2007)
46. Zaman, L., Kalra, A., Stuerzlinger, W.: The effect of animation, dual view, difference layers, and relative re-layout in hierarchical diagram differencing. In: *Proceedings of the Conference Graphics Interface, GI 2011*, pp. 183–190. Canadian Human-Computer Comm. Soc. (2011)