

Mental Map Preservation Helps User Orientation in Dynamic Graphs

Daniel Archambault¹ and Helen C. Purchase²

¹ Clique Strategic Research Cluster, University College Dublin
daniel.archambault@ucd.ie

² School of Computing Science, University of Glasgow
helen.purchase@glasgow.ac.uk

Abstract. We present the results of a formal experiment that tests the ability of a participant to orient themselves in a dynamically evolving graph. Examples of these tasks include finding a specific location or route between two locations. We find that preserving the mental map for the tasks tested is significantly faster and produces fewer errors. As the number of targets increase, this result holds.

1 Introduction

Dynamic graph drawing approaches depict time-evolving graphs where nodes and edges can be added and removed. Algorithms to draw dynamic graphs typically represent the evolving graph as a series of timeslices. A **timeslice** encodes the structure of the graph at a given time. The timeslices, or the sequence of graphs, are often placed in chronological order, depicting graph evolution.

An often cited aesthetic property of dynamic graph drawings is drawing stability or keeping parts of the drawing in the same area of the plane. This **preservation of the mental map** [9, 17, 8] is often cited as a desired property and Coleman and Parker [8] defined this concept as *dynamic aesthetics*:

- *The placement of existing nodes and edges should change as little as possible when a change is made to the graph.*

Formal human computer interaction experiments have been run to test the effectiveness of the mental map with users on a variety of tasks. A common and repeated conclusion of these experiments [18, 20, 3, 1] is that no significant positive effect was found between conditions that preserved the mental map and conditions that did not. In other words, choosing to preserve the mental map did not improve the performance of participants in the experiment, either in terms of response time or error rate, when compared to a representation of the dynamic graph whereby each and every frame of the series was drawn independently.

Results in the psychology literature [19, 23, 16] may suggest a partial explanation for these unintuitive results. In this body of work, it has been confirmed that participants can track up to five randomly moving objects, or **targets**, against a background of distractor objects with high accuracy [19]. As many of

the experiments in dynamic graph drawing [18, 20, 3, 1] involved a small number of independently moving nodes, it may be the case that participants relied little on the mental map. Secondly, nodes of interest in these tasks were often preattentively highlighted using colour [20, 3]. This highlighting would ease tracking the nodes involved in the task, alleviating some of the need for the mental map. Finally, the tasks did not require participants to use the dynamic drawing as a *map*. Map-based orientation tasks would require the participant to orient themselves in the data, such as finding specific locations or paths.

Animation and small multiples are two representations for dynamic data which have been used to visualize dynamic graphs. **Animation** presents the dynamic graph as a movie and uses linear interpolation to morph the position of nodes as they evolve through time. A **small multiples** [22] representation of a dynamic graph places each timeslice in a matrix of static images in chronological order. Recent experiments [3, 11] have indicated that small multiples may be a more effective presentation method for readability tasks, but neither presentation method has been evaluated on orientation tasks.

This paper presents a user study on the following research question:

1. Does mental map preservation [8] help with user orientation in dynamic graphs?

We also examine two secondary research questions:

1. Does the number of targets influence performance?
2. Does animation or small multiples influence performance?

Our results indicate that the mental map significantly helps users orient themselves in the dynamic data, regardless of the number of targets. Our results also suggest that animation could be potentially be beneficial for these tasks.

2 Related Work

We summarize related work in terms of performed readability and memorability experiments and results in psychology.

2.1 Readability and Memorability Experiments

Several experiments have been run in order to try and demonstrate the benefits of preserving the mental map on a variety of tasks. Purchase *et al.* [18] tested three levels of mental map preservation (no, medium, high) on degree reading tasks in an animation context. The experiment found that no mental map followed by high mental map preservation performed significantly better than a compromise between the two. Saffrey *et al.* [20] tested a low and high level of mental map preservation in a shortest path reading context. Endpoints of the path were preattentively highlighted using colour. The authors found that mental map preservation can be harmful to readability, especially when significant node

overlap occurs. Archambault *et al.* [3] tested a no and high level of mental map preservation on readability tasks in a small multiples and animation context. Focus nodes involved in the tasks were preattentively highlighted using colour. The authors found that preserving the mental map did not have a significant effect for either small multiples or animation. Archambault *et al.* [1] tested memorability of dynamic graph drawing sequences in an animation context. Although they found that drawings which preserve the mental map were significantly preferred, the experiment found no significant performance difference in preserving the mental map.

Previous experiments [20,3] have used colours to preattentively highlight focus nodes in experimental tasks and have only tested readability and memorability tasks. In this experiment, we test the ability of the participant to remain oriented in the information space.

2.2 Results in Psychology

Three results in the field of psychology inspire the design of this experiment [19, 23, 16]. Pylyshyn and Storm [19] demonstrated that participants could track up to five randomly moving targets on a screen simultaneously with relatively high accuracy against a field of distractors. Yantis [23] replicated these results and found that if the targets had coordinated movement, the accuracy of tracking increased. Finally, Liu *et al.* [16] studied a similar problem in an air traffic control scenario. The experiment found that accuracy reduces with increased object speed, but seemed to be unaffected in terms of scene motion (panning, rotating, and zooming). Our experiment, to determine the effect of mental map preservation in a dynamic graph drawing context, is inspired by this work.

3 Experiment

To test the map properties of a dynamic graph presentation, we performed a within-participant experiment. We employed a 2 mental map preservation levels (No Mental Map Preservation (NMM) and Mental Map Preservation (MM)) \times 2 presentation method (Animation vs Small Multiples) \times 3 target levels (Low, Med, High) \times 2 data set (two subsequences of `threads2`) \times 2 question design. Thus, each participant performed 48 measured tasks in the experiment.

Target, in this experiment, refers to the number of nodes of interest in the question posed to the participant. The targets are either a set of independent nodes or a path in the graph. The number of targets was determined using the results of the experiments in psychology [19, 23, 16].

3.1 Mental Map Preservation

Graphael [10] is a dynamic graph drawing algorithm where the degree of mental map preservation can be adjusted relatively easily. In the algorithm, inter-timeslice edges exist between nodes which represent the same node across timeslices. This graph, containing all real edges and inter-timeslice edges, is drawn

once with a force-directed algorithm. If these inter-timeslice edges are stiff, nodes stay in the same area of the plane, while if they are loose, nodes move more freely. A number of algorithms [6] could have been used. However, this algorithm is classed as a linking strategy and there is evidence [7] that these strategies best conform to the Colman and Parker definition while maintaining layout quality.

This experiment tests two levels of mental map preservation. For the *no mental map preservation condition*, all timeslices of the graph series are drawn independently using the force-directed algorithm used by Graphael. For the *mental map preservation condition*, inter-timeslice edge strength is set to a value of sixty. The approach for selecting inter-timeslice edge strength was based on visual inspection, piloting, and consistency with the values used in previous readability experiments [18, 3, 2, 1].

3.2 Interfaces

In this experiment, we test both animation and small multiples as methods for dynamic graph presentation (Fig. 1). In the **animation** condition, the dynamic graph is presented in an interface similar to a movie player. Smooth, linear interpolation morphs the graph from one timeslice to another and nodes that are added and removed from the graph are faded into and out of the display. The depiction of the dynamic graph takes up the entire screen. Participants have complete control over the animation through a play/pause button and a slider which can be manipulated to evolve the graph at a desired rate. No other form of interaction, including zooming, is allowed.

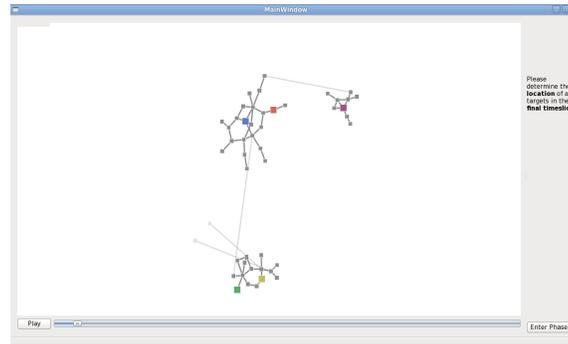
The **small multiples** interface presents all timeslices in a matrix. Timeslices are placed in chronological order from left to right and top to bottom. Participants were instructed to read the small multiples matrix in this way during training. The participant scans the windows in order to determine graph evolution. No interaction with the timeslices, including zooming, is allowed.

3.3 Data Set

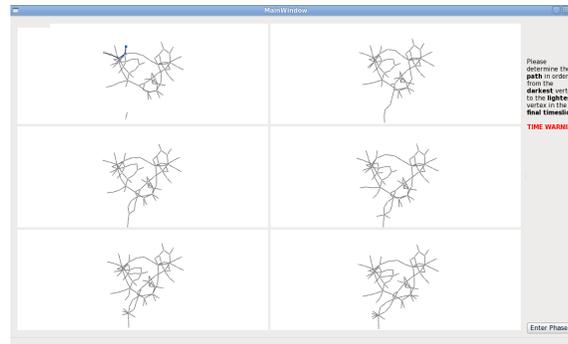
Two six timeslice long subsequences of a real world data set are used in this experiment. **Threads2** [12] is a graph series representing online newsgroup discussions and how they evolve over time. Nodes are authors of newsgroup articles, and an edge exists between two authors if one replied to the posting of another. Two subsequences of **Threads2** (timeslices 5 - 10 and timeslices 14 - 19) are used. The first subsequence ranges from 46 nodes and 52 edges to 62 nodes and 72 edges. The second subsequence ranges from 72 nodes and 87 edges to 85 nodes and 104 edges. The maximum of changes between timeslices is five for nodes and four for edges. The graph size and the length of the time series was informed through piloting and earlier experiments [18, 3, 2, 1].

3.4 Tasks

The tasks in this experiment test the map properties of dynamic graph evolution directly or the ability of the participants to orient themselves in the dynamic



(a) Animation



(b) Small Multiples

Fig. 1. Animation and small multiples interfaces. **(a)** Animation and NMM. **(b)** Small multiples and MM. The participant uses the interface to determine the location of the coloured nodes or the path in the final timeslice. Animation is like a movie player and small multiples is read left to right and top to bottom. Once the participant clicks on *Enter Phase 2*, the last timeslice of the graph sequence is shown and the participant clicks on nodes to indicate their answer. The *time warning* label is shown in the small multiples interface

data. The first task is essentially a revisitation [21, 13] task in a dynamic graph drawing setting while the second task tests the ability of participants to follow long paths.

Each task has two phases. In phase 1, the participant uses the visualization with the aim of determining target locations (highlighted in the first timeslice) in the final timeslice. Once the participant is satisfied that they have located all targets, the participant clicks on *Enter Second Phase*. In phase 2, the last timeslice of the graph series is shown full screen. During this phase, the participant cannot return to the dynamic visualization of the data and answers the question by clicking on nodes in the final timeslice. Response time is the sum of the times spent in phase 1 and 2.

The first task tests the ability of participants to revisit a particular node in the dynamic graph sequence. If the graph were interpreted like a map, this question is similar to locating landmarks in a city (eg. *Where is St. Stephen's Green in Dublin?*). The targets are independently moving, uniquely coloured nodes spread through the graph (isoluminant shades of blue, red, yellow, green, or purple). The target level was defined by the number of nodes coloured: 1 for low, 3 for medium, and 5 for high – these values were drawn from previous work in psychology [19, 23, 16]. Nodes are highlighted in the first timeslice but are not highlighted in subsequent timeslices. The participant uses the visualization to determine the location of all nodes in the final timeslice.

In phase 1 of this question, participants were asked:

- Please determine the **location** of all targets in the **final timeslice**.

Subsequently, in phase 2, participants were asked:

- Where is the **<colour>** vertex?

One of the above indicated colours is the value for **<colour>**. The answer to the question is correct if the selected node is the node of that colour in the first timeslice.

The second task tests the ability of participants to follow a path through the evolving graph. If the graph were viewed as a map, this question would be similar to finding a route through a city (eg. *How do I get from St. Stephen's Green to O'Connell Bridge?*). The targets are nodes along a path in the graph with decreasing saturation indicating direction. As the path is connected, the motion of these targets is more coordinated than in the previous task. Informed by Yantis [23], the number of targets (the length of the path) is 3, 5, and 7 nodes for low, medium, and high target levels. The nodes in the path are only highlighted in the first timeslice, and the participant uses the visualization to find the path in the final timeslice.

In phase 1, participants were asked:

- Please determine the **path** in order from the **darkest** vertex to the **lightest** vertex in the **final timeslice**.

In phase 2, participants were asked:

- Click on the **vertices** of the path **in order (X vertices)**.

The value of X is the number of vertices on the path. Correctness is determined based on $(X - d)/X$ where X is the length of the path in the answer and d is the Levenshtein distance (edit distance) [15] between the entered answer and the correct answer. Edit distance measures the number of insertions, deletions, and changes between the entered answer and the correct answer and is at most X (replacement of the entire string).

3.5 Experiment Design

The experiment was divided into two, counterbalanced blocks (animation and small multiples) of 2 mental map preservation levels (No Mental Map and Mental Map Preservation) \times 3 target levels \times 2 data sets (subsequences of `Threads2`) \times 2 questions or 24 tasks per presentation method factor. These twenty-four tasks were randomized individually per participant. To overcome the learning effect, the tasks were prefixed with 8 practice tasks. The results of the practice tasks were discarded and participants were not made aware that these tasks did not form part of the experiment. The practice tasks presented every mental map preservation level once, each data set once, each question four times, and each target level at least twice. Thus, per presentation method factor, the participant was required to complete 32 tasks. These tasks were divided into two blocks of 16 tasks between which participants could take a short break.

The experiment required participants to answer all tasks under one presentation method (animation or small multiples) first, followed by all tasks on the other presentation method. Therefore, any cognitive shift required to move from one interface to another only occurred once. We counterbalanced between participants by presenting animation first to even-numbered participants and small multiples first to odd-numbered participants. The participant could take a short break between experimental conditions.

Before the start of each presentation method block, participants had a demonstration session which introduced them to the experimental interface. During this session, the participants could ask questions, find out about the experimental tasks and learn how to find the answers to the questions.

All three interfaces were rendered in real time using the Tulip framework [4]. No time limit was enforced per task or for the experiment overall. However, a warning label appeared on the screen after fifty seconds had elapsed for each task, and participants were encouraged to finish the task after that point.

Overall, there were twenty-eight participants used in the final results. Participants were drawn from members of the Complex and Adaptive Systems Laboratory of University College Dublin (UCD CASL).

4 Results

We present the results for our experiment divided according to the two questions tested. As location and path tracing tasks are very different in nature, we treat the data for each separately. We compare No Mental Map preservation (NMM) to Mental Map preservation (MM) in the context of Animation (Anim) and Small Multiples (SM). A Shapiro-Wilk test, with a significance level of $\alpha = 0.05$, was used to determine whether or not the data was normally distributed. We found that neither error rate nor response time was usually normally distributed. As a consequence, we used an exact Wilcoxon signed rank test when comparing two data sets at a significance level of $\alpha = 0.05$. When we divided the data by target level, we applied a Bonferroni correction, thus reducing the significance

level to $\alpha = 0.017$. When comparing more than two distributions, a Friedman test was used as usually one distribution was not normal. Post-hoc analysis was conducted with Nemenyi test. In all bar charts, black lines connect pairs of bars with significant differences. Mean and median values, separated with a hyphen, are indicated below each bar. The standard error is indicated on each bar.

4.1 NMM vs MM: Location Task

This task asked participants to locate vertices in the final timeslice of a dynamic graph, the vertices having been indicated in the first timeslice. The task is similar to locating a region on an evolving map.

Overall: We found significant difference both in terms of error rate and response time for this task (Figs. 2(a) and 2(b)). MM produced both significantly faster response times ($p < 0.001$) and fewer errors ($p < 0.001$) than NMM.

When comparing NMM_Anim, NMM_SM, MM_Anim, and MM_SM, we found significant differences in terms of response time ($\chi^2 = 61.12$, $df = 3$, $p < 0.001$) and error rate ($\chi^2 = 65.74$, $df = 3$, $p < 0.001$) with an order of MM_Anim, MM_SM, NMM_Anim, NMM_SM. Post-hoc analysis revealed pairwise significant differences between any MM and NMM pair (Figs. 3(a) and 3(b)) but not between Anim and SM pairs.

Divided by Target Level: When comparing NMM_Anim, NMM_SM, MM_Anim, and MM_SM, we found that the order of performance in terms of both response time and error rate was: MM_Anim, MM_SM, NMM_Anim, NMM_SM for a low and medium number of targets and MM_SM, MM_Anim, NMM_Anim, NMM_SM for error rate on a high number of targets. Post-hoc analysis revealed pairwise differences between any MM and NMM pair (Fig. 4) but not between Anim and SM pairs.

4.2 NMM vs MM: Path Task

This task asked participants to identify a path in the final timeslice of a dynamic graph, the path having been indicated in the first timeslice. The task is similar to finding a route on a dynamic map.

Overall: We found significant difference both in terms of error rate and response time for this task. MM produced both significantly faster response times ($p < 0.001$) and fewer errors ($p < 0.001$) than NMM (Figs. 2(c) and 2(d)).

When comparing NMM_Anim, NMM_SM, MM_Anim, and MM_SM, we found significant differences in terms of response time ($\chi^2 = 68.10$, $df = 3$, $p < 0.001$) and error rate ($\chi^2 = 75.02$, $df = 3$, $p < 0.001$) with an order of MM_Anim, MM_SM, NMM_Anim, NMM_SM. Post-hoc analysis revealed pairwise differences between any MM and NMM pair (Figs. 3(c) and 3(d)) but not between Anim and SM pairs.

Divided by Target Level: When comparing NMM_Anim, NMM_SM, MM_Anim, and MM_SM, we found that the order of performance in terms of both response time and error rate was: MM_Anim, MM_SM, NMM_Anim,

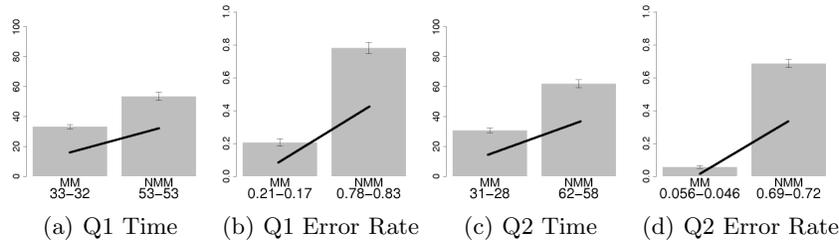


Fig. 2. Response time and error rate for NMM and MM on questions 1 and 2. Response time in seconds and error rate as defined in Section 3.4.

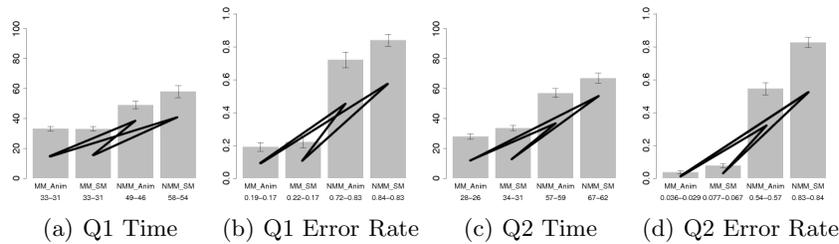


Fig. 3. Response time and error rate across MM_Anim, MM_SM, NMM_Anim, NMM_SM. Response time in seconds and error rate as defined in Section 3.4.

NMM_SM for all three target levels. Post-hoc analysis revealed pairwise differences between any MM and NMM pair (Fig. 4) but not between Anim and SM pairs.

Table 1. Preference data for features of the dynamic graph layout. **Large move** is large node movement. **Pattern move** is patterns in node movement. **Const pos** is constant node position. **Const patterns** patterns in connection that remain constant.

	large move	pattern move	const pos	const patterns
Overall	3.83	2.79	1.48	1.90
Location	3.83	2.79	1.38	2.00
Path	3.83	2.72	1.66	1.79

4.3 Preference Data

Table 1 shows the results for characteristics of the graph layout that participants felt helped them the most overall and in both questions. Participants were asked to rank these features from 1 through 4 with 1 corresponding to the most important feature. Overall and on both questions, participants found that constant node position and static patterns in the graph helped the most while patterns in node movement and large node movement helped the least.

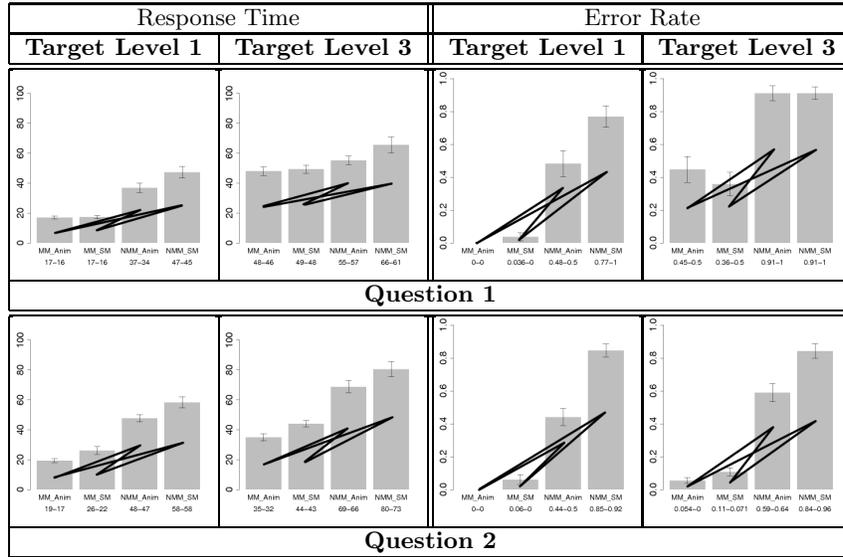


Fig. 4. Significant differences in response time and error rate across MM_Anim, MM_SM, NMM_Anim, NMM_SM when divided by target level. Response time in seconds and error rate as defined in Section 3.4. All three target levels were significant. Low and high target levels only shown for space reasons.

5 Discussion

For both the location and the path questions, preserving the mental map was found to be both significantly faster in terms of response time and to produce significantly fewer errors. The probable reason is that stable drawings help participants orient themselves in the dynamic data, allowing them to re-locate areas of the graph or to follow long paths through it.

We noticed some similar effects to works in the psychology literature [19, 23, 16] when dividing the data by target level. When the mental map was preserved, objects did not collide and moved slowly. Therefore, participants were able to perform the tasks with low error rates even for a large number of targets. For the path task, where motion of the nodes was coordinated [23], error rate was further reduced. When the mental map was not preserved, error rate was rather high. This result was probably due to collisions between graph elements and the higher velocity of nodes and edges when not preserving the mental map. Additionally, a dynamically evolving graph is a more complex visual entity.

The only significant differences we found were between the MM and NMM conditions. Within these conditions we did not find any significant difference between animation and small multiples in terms of response time or error rate. However, animation generally seemed to perform better than small multiples for these tasks. Therefore, animated transitions could help, consistent with previous experiments [5].

Ghani *et al.* [14] recently investigated the effect of the mental map on determining the order of addition/deletion of nodes when animating a dynamic graph. In this experiment, each addition/deletion defined a new timeslice. To preserve the mental map, vertices were pinned using an aggregation of all timeslices. The no mental map condition drew each timeslice independently. Participants were presented a set of randomly coloured targets inserted/deleted at different times. Answers were an entered sequence of colours indicating the perceived order of removal/addition. The study found that the pinning strategy outperformed laying out each timeslice independently for this task. Where this work investigated temporal causality, our experiment tested navigation in dynamic data. It would be interesting to further investigate the relationship between these two types of tasks in future work.

6 Conclusion

We have presented an experiment which demonstrates that preserving the mental map, or diagram stability, is important in a Coleman and Parker [8] sense. The tasks rely on user orientation in the dynamically evolving graph data. Our tasks required participants to locate nodes in the data and to follow long paths. Although this result indicates that the mental map can support certain tasks, many previous experiments have not revealed a benefit [18, 20, 3, 1], suggesting that the benefits of preserving the mental map are highly dependent on the nature of the user's task.

Acknowledgments

The first author would like to acknowledge the support of the Clique Strategic Research Cluster funded by Science Foundation Ireland (SFI) Grant No. 08/SRC/I1407. We would also like to thank John Hamer for writing the dynamic graph layout algorithm based on the GraphAEL model and Bruno Pinaud. We would like to thank all of the participants who took part in this experiment. Finally, we would like to thank the anonymous reviewers of GD 2012 for their insightful comments.

References

1. D. Archambault and H. C. Purchase. The mental map and memorability in dynamic graphs. In *Proc. of the IEEE Pacific Visualization Symposium (PacificVis 2012)*, pages 89–96, 2012.
2. D. Archambault, H. C. Purchase, and B. Pinaud. Difference map readability for dynamic graphs. In *Proc. of Graph Drawing (GD 2010)*, volume 6502 of *LNCS*, pages 50–61, 2010.
3. D. Archambault, H. C. Purchase, and B. Pinaud. Animation, small multiples, and the effect of mental map preservation in dynamic graphs. *IEEE Trans. on Visualization and Computer Graphics*, 17(4):539–552, 2011.

4. D. Auber. Tulip : A huge graph visualization framework. In P. Mutzel and M. Jünger, editors, *Graph Drawing Software*, Mathematics and Visualization, pages 105–126. Springer-Verlag, 2003.
5. B. B. Bederson and A. Boltman. Does animation help users build mental maps of spatial information? In *Proc. of the 1999 IEEE Symposium on Information Visualization*, pages 28–35, 1999.
6. U. Brandes, N. Indlekofer, and M. Mader. Visualization methods for longitudinal social networks and stochastic actor-oriented modeling. *Social Networks*, 34(3):291–308, 2011.
7. U. Brandes and M. Mader. A quantitative comparison of stress-minimization approaches for offline dynamic graph drawing. In *Proc. of Graph Drawing (GD 2011)*, volume 7034 of *LNCS*, pages 99–113, 2011.
8. M. K. Coleman and D. S. Parker. Aesthetics-based graph layout for human consumption. *Software - Practice and Experience*, 26(12):1415–1438, 1996.
9. P. Eades, W. Lai, K. Misue, and K. Sugiyama. Preserving the mental map of a diagram. In *Proc. of Compugraphics*, pages 24–33, 1991.
10. C. Erten, P. J. Harding, S. Kobourov, K. Wampler, and G. V. Yee. GraphAEL: Graph animations with evolving layouts. In *Proc. of Graph Drawing (GD 2003)*, volume 2912 of *LNCS*, pages 98–110, 2003.
11. M. Farrugia and A. Quigley. Effective temporal graph layout: A comparative study of animation versus static display methods. *Journal of Information Visualization*, 10(1):47–64, 2011.
12. Y. Frishman and A. Tal. Online dynamic graph drawing. *IEEE Trans. on Visualization and Computer Graphics*, 14(4):727–740, 2008.
13. S. Ghani and N. Elmquist. Improving revisitation in graphs through static spatial features. In *Proc. of Graphics Interface (GI 2011)*, pages 175–182, 2011.
14. S. Ghani, N. Elmquist, and J. S. Yi. Perception of animated node-link diagrams for dynamic graphs. *Computer Graphics Forum (Proc. EuroVis 2012)*, 31(3):1205–1214, 2012.
15. V. I. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics-Doklady*, 10(8):707–710, 1966.
16. G. Liu, E. L. Austen, K. S. Booth, B. D. Fisher, R. Argue, M. I. Rempel, and J. T. Enns. Multiple-object tracking is based on scene, not retinal, coordinates. *Journal of Experimental Psychology: Human Perception and Performance*, 31(2):235–247, 2005.
17. K. Misue, P. Eades, W. Lai, and K. Sugiyama. Layout adjustment and the mental map. *Journal of Visual Languages and Computing*, 6:183–210, 1995.
18. H. C. Purchase and A. Samra. Extremes are better: Investigating mental map preservation in dynamic graphs. In *Diagrams 2008*, volume 5223 of *LNAI*, pages 60–73, 2008.
19. Z. W. Pylyshyn and R. W. Storm. Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, 3(3):179–197, 1988.
20. P. Saffrey and H. C. Purchase. The “mental map” versus “static aesthetic” compromise in dynamic graphs: A user study. In *Proc. of the 9th Australasian User Interface Conference*, pages 85–93, 2008.
21. A. Skopik and C. Gutwin. Improving revisitation in fisheye views with visit wear. In *Proc. of SIGCHI conference on Human Factors in Computing Systems (CHI 2005)*, pages 771–780, 2005.
22. E. Tufte. *Envisioning Information*. Graphics Press, 1990.
23. S. Yantis. Multielement visual tracking: Attention and perceptual organization. *Cognitive Psychology*, 24:295–340, 1992.