The Mental Map and Memorability in Dynamic Graphs

Daniel Archambault*
Clique Strategic Research Cluster, University College Dublin

Helen C. Purchase†
School of Computing Science, University of Glasgow

ABSTRACT

In dynamic graph drawing, preserving the mental map, or ensuring that the location of nodes do not change significantly as the information evolves over time is considered an important property by algorithm designers. Many prior experiments have attempted to verify this principle, with surprisingly little success. These experiments have used several different algorithmic methods, a variety of graph interpretation questions on both real and fabricated data, and different presentation methods. However, none of the results have conclusively demonstrated the importance of mental map preservation on task performance. Our experiment measures the efficacy of the dynamic graph drawing in a different manner: we look at how memorable the evolving graph is, rather than how easy it is to interpret.

As observed in the previous studies, we found no significant difference in terms of response time or error rate when preserving the mental map. While preserving the mental map is a good idea in principle, we find that it may not always support performance. However, our qualitative data suggests that, in terms of the user's perception, preserving the mental map makes memorability tasks easier. Our qualitative data also suggests that there may be two features of the dynamic graph drawing that may assist in their memorability: interesting subgraphs that remain visible over time and interesting patterns in node movement. The former is supported by preserving the mental map while the latter is not.

Index Terms: H.1.2 [Information Systems]: User/Machine Systems—Human Factors; G.2.2 [Discrete Mathematics]: Graph Theory—Graph Algorithms

1 INTRODUCTION

The dynamic graph drawing problem concerns itself with ways of visually representing a graph as it evolves over time. Typically, several snapshots of the graph, or **timeslices**, are taken and these timeslices are placed in chronological order to depict how the graph evolves. Preserving the **mental map** [9, 18] in dynamic graph drawing is usually formalized as the principle that each node in a dynamic graph drawing algorithm should remain in the same area of the drawing as it evolves over time. In terms of **dynamic aesthetics**, this notion is expressed through the *minimal edit disruption* criteria where the *"placement of existing nodes and edges should change as little as possible when a change is made to the graph"* [7].

Many algorithms in the dynamic graph drawing community have focused on ways to preserve the mental map [6, 19, 8, 11, 15, 5, 4, 13] as, intuitively, this property should support user tasks. However, many human computer interaction experiments have been run on the topic [20, 21, 22, 2, 27] on a variety of tasks. In general dynamic graph drawing, no experiment has demonstrated a significant benefit for mental map preservation.

In all of the user studies to date on the mental map and dynamic graphs, the tasks selected by the experimental designers have been readability tasks. **Readability** tasks test the ability of participants to read a graph and extract structural information in order to answer a question. Mental map preservation can influence readability by restricting node movement. On general graphs, the readability of degrees [2, 21] and connections between sets of nodes [22] have been tested. Questions related to the dynamic nature of the graph series, the appearance of nodes or edges [2, 27], have also been tested with little to no significant effect when the mental map is preserved. However, to date, these experiments have not tested the effect of the mental map on memory in the context of dynamic graphs. An experiment which tests a **memorability** task would test the ability for participants to remember a dynamic graph series and reliably determine if it is the same or different. Memorability and readability tasks are related, but it is interesting to explicitly test the ability of users to remember the evolution of a dynamic graph sequence to see if it has changed. It may be that mental map preservation can better support memory tasks as users could potentially offload cognitive effort to the drawing if the nodes and edges stay in similar areas of the plane. Therefore, memorability tasks could be better supported by the mental map.

Memorability tasks are important for information discovery. Support for these tasks help with mental comparisons between graph series and with visualizing differences in multiple views of the data. Moreover, screen space is often at a premium for dynamic graph visualization techniques, and all the data often cannot be visible at once, requiring users to rely on visual memory.

The work presents a user study which investigates the following research question:

- Does preserving the mental map help users remember dynamic graph sequences and therefore assist in determining if changes have been made to them?

In our study, we found no significant difference in terms of response time or error rates on this task. This result is consistent with the many readability studies conducted with respect to mental map preservation on dynamic graph series. However, we found that preserving the mental map was significantly preferred for this task, and, qualitatively, certain features in the graph that would benefit from mental map preservation were used by the participants in order to perform the memorability tasks. We also found some evidence for two types of features used for memorization: static features, pertaining to the interconnections between nodes in the graph, and dynamic features, relating to the motion of certain subgraphs.

2 RELATED WORK

Currently, no experiment has evaluated the mental map in terms of the memorability of dynamic graph sequences. However, a number of works have looked at the readability of dynamic graph sequences and others have looked at visual memory or memorability tasks on different types of data in the information visualization community. In section 2.1, we look at the many experiments that have investigated the effectiveness of the mental map, and in section 2.2, we present experiments that have evaluated visual memory or memorability tasks in other areas of information visualization.

2.1 Mental Map Readability Experiments

A number of experiments have tested a variety of readability tasks on general and specific types of graphs in order to support the effec-

*e-mail: daniel.archambault@ucd.ie
†e-mail: Helen.Purchase@glasgow.ac.uk

tiveness of the mental map. However, currently, there is no study where the mental map has been shown to have a significant, positive effect on the performance of participants reading general dynamic graph drawings.

Purchase *et al.* [20] tested the preservation of the mental map in the context of dynamic, hierarchically drawn directed acyclic graphs. They tested the algorithm of Görg *et al.* [15], and found that maintaining the mental map was important for the comprehension of an evolving graph where nodes need to be identifiable by name, but that it was less important for tasks that focus on edges rather than nodes, or which do not require that nodes be nominally differentiated from each other. No other significant effect of the mental map was found in the study.

In Saffrey [22], participants were asked to read connections between tightly clustered groups of nodes in a dynamic graph. The algorithm used for this study was a novel algorithm that preserved the mental map based on geometric restriction. The user study found that preserving the user's mental map by restricting node movement does not always contribute positively to understanding, especially in the case of significant node overlap. In both of the studies presented in the paper, a layout based purely on a static drawing algorithm applied to each timeslice, with no restriction on node movement, produced the least errors.

Purchase *et al.* [21] examined the preservation of the mental map in the context of degree reading tasks. The experiment used the Graphael system of Ertin *et al.* [11], to draw the dynamic graph series. A high, medium, and no mental map preservation parameter was tested. The results of this experiment indicated that the extremes, no mental map preservation followed by high, produced the best performance, suggesting that individual preference could have been important. A compromise between the two performed significantly worse on this task.

Archambault *et al.* [2] tested mental map preservation in the context of animation and small multiples presentation methods for dynamic graphs. In their experiment, Graphael [11] was used and a no and high mental map preservation factor was tested in the experiment. The tasks used in the study were based on degree, appearance, growth, and path reading questions. The experiment found no significant difference of notable magnitude between no mental map preservation and high mental map preservation for either small multiples or animation.

Zaman *et al.* [27] had a mental map factor when evaluating the `DARLS` system for its effectiveness to determine the difference between two hierarchically drawn, directed acyclic graphs. In the first experiment of the paper, a low mental map preservation version of the algorithm was tested against a high mental map preservation version of the algorithm on tasks that involve finding the appearance of new nodes in the graph. For this task, no significant difference was found in terms of response time or error rates.

The main limitations of all of these experiments is that they test primarily readability or appearance questions. Also, for general, undirected dynamic graph drawing, the experiments generally find that for most tasks the mental map does not help readability, but does not significantly hurt readability either. In our experiment, we instead test the performance of a memorability task to see if preserving the mental map has an effect.

2.2 Visual Memory and Memorability Experiments

In the field of information visualization, a number of experiments have looked at visual memory or memorability tasks in the context of various types of information displays. As the work is related to our study, we summarize it in this section.

Lam *et al.* [16] looked at the effect of scaling, rotation, rectangular fisheye, and polar fisheye transformations on visual memory. The data sets used in this experiment were static graphs and with or without the superimposition of appropriate grids. The study maps out clear no-cost zones and demonstrates significant benefits when overlaying grids in all transformations except polar fisheye.

Tory *et al.* [24] compared coloured dot, contour, and landscape visualizations for non-spatial data at different levels of density in a visual memory context. Participants were asked to memorize a set of eight images of non-spatial data in a learning phase and then were presented with eight images in a testing phase and asked if the image appeared in the learning phase. The experiment found that dots and landscapes were significantly more accurate than contour images of the data. As the number of dots increased, dots were significantly more accurate than landscapes.

Two experiments have looked at *revisitation* tasks. Revisitation tasks require participants to remember where objects in a particular space are located and how they can be reached a second time whereas general *memorability* is a measure of the degree to which people remember things in the information space [23, 14]. Skopik *et al.* [23] examined the effect of fisheye distortions on revisitation tasks in a graph with and without landmarks. The first study required participants to revisit a node in a graph after the fisheye-distorted view had been shifted to a random location. The study found a significant effect on task completion time with increasing fisheye distortion. In the second study, they found that using visual landmarks significantly reduced task completion time under the same sorts of distortions. Ghani *et al.* [14] studied the effects of spatial features added to the background of the graph such as coloured patches, colour or texture encoding on the nodes, or combinations of the techniques had on revisitation tasks in static graph drawings. Their experiments found that coloured grids significantly improved response times on revisitation tasks and that landmarking is generally promising in this context.

3 The Experiment

To test the effect of mental map preservation on memorability, we performed a within subject experiment. We employed a 3 condition (No Mental Map, Medium Mental Map, High Mental Map) $\times$ 3 data set (`MIT`, `threads2`, `van de Bunt`) $\times$ 2 version $\times$ 2 repetition $\times$ 1 question design. Thus, each participant performed 36 tasks in the experiment.

3.1 Task

Our research question required participants to memorize a dynamic graph sequence and report if any changes have been made to it. Similar to many of the other studies in the information visualization community which have tested visual memory, we divided the blocks in our experiment into two phases: a memorization phase and a recall phase. In the **memorization phase**, participants were requested to memorize three graph series each six timeslices long. During this phase, participants were given the following instruction:

- Please memorize the following graph series.

All nodes and edges in the memorization graph were the same colour, a light grey, and participants were instructed to play the video as many times as desired to memorize the graph. After the memorization phase, participants entered a **recall phase**, where they were shown a series of graphs, some of which were identical to those in the memorization phase, and some which had a small percentage of nodes removed. During this phase, participants were asked if the graph was exactly the same as one presented in the memorization phase or if it had been modified. Participants were asked the following question:

- Is this graph **exactly the same** as one of the graphs presented in the memorization phase?

All nodes and edges in the recall phase were the same colour, a light grey, and participants were instructed to play the video as

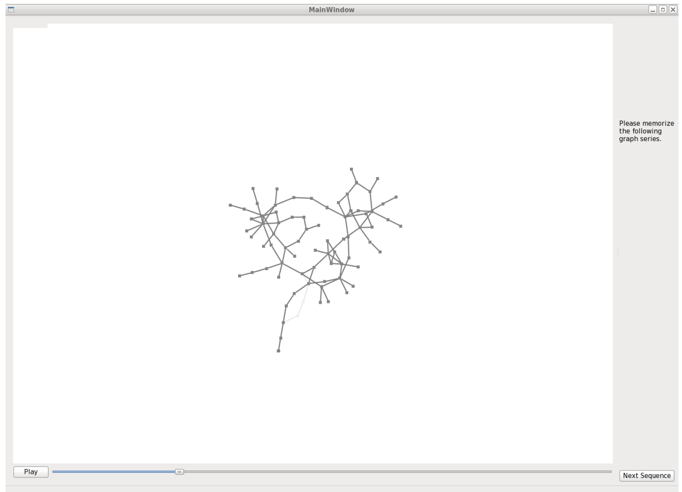

(a) Memorization Phase

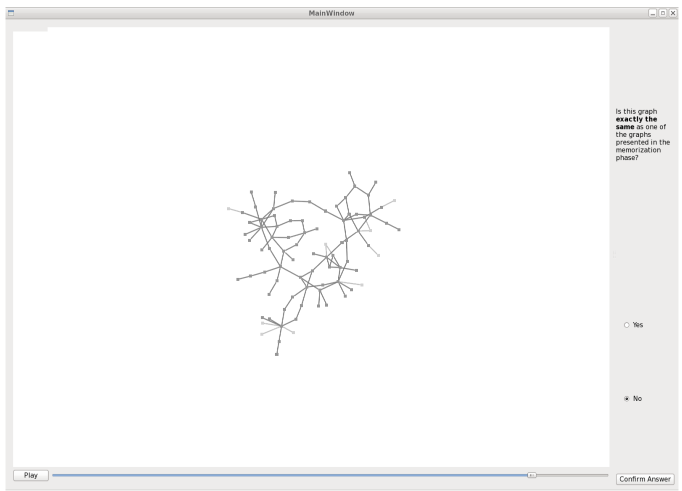

(b) Recall Phase

Figure 1: Interfaces for both phases of the experiment. **(a)** Memorization phase. Participants play the video and click on "next sequence" when they believe that they have memorized the graph. **(b)** Recall Phase. A question in the upper right corner asks if the graph was present in the memorization block. Participants are presented with two possible answers "yes" and "no". They select the appropriate radio button and click "submit answer" to respond. In both cases, `Threads2` is shown.

many times as desired to determine if the graph had been modified. Each recall phase consisted of twelve questions: four graphs associated with each graph in the memorization phase. Two of these graphs were identical to the graph presented in the memorization phase. The remaining two graphs were modified by selecting and removing 8% of the nodes in the graph and all their incident edges. We determined the appropriate percentage of nodes to remove from the graph through piloting – an amount of change that was discernible, but not too extreme. In order to ensure that participants watched the animation, the first and the last timeslice were always identical to the graph in the memorization phase and all changes were induced somewhere between timeslices two through five.

3.2 Interface

In this experiment, we use **animation**: the most common presentation method for dynamic graphs. In both the memorization and recall phases of the experiment, our animation widget was the same. In the experiment, we presented the animation of the dynamic graph series in the most common presentation method used by layout algorithm designers, similar to a movie player. In this presentation, the view of the graph takes up the entire screen and smooth transitions morph the graph from one timeslice to another. Nodes that are added to the data or removed from it are faded in or out of the display respectively. The positions of nodes and edges are linearly interpolated between frames.

Full control over the animation during the phases of the experiment was allowed to ensure that participants were able to memorize or determine the differences in the graph sequence to the best of their ability. The participants could play, pause, and rewind the video at any time. However, the animation started to play automatically at the beginning of the task so that participants could perform the task with little interaction, if desired. Participants could also drag the slider at their own rate. No other form of interaction, including zooming, was allowed. This interface was used previously in the experiments of Archambault *et al.* [1, 2].

Figure 1(a) shows the interface in the memorization phase. In this phase, the instructions in the upper right asked participants to try and memorize the graph shown in the widget. When participants

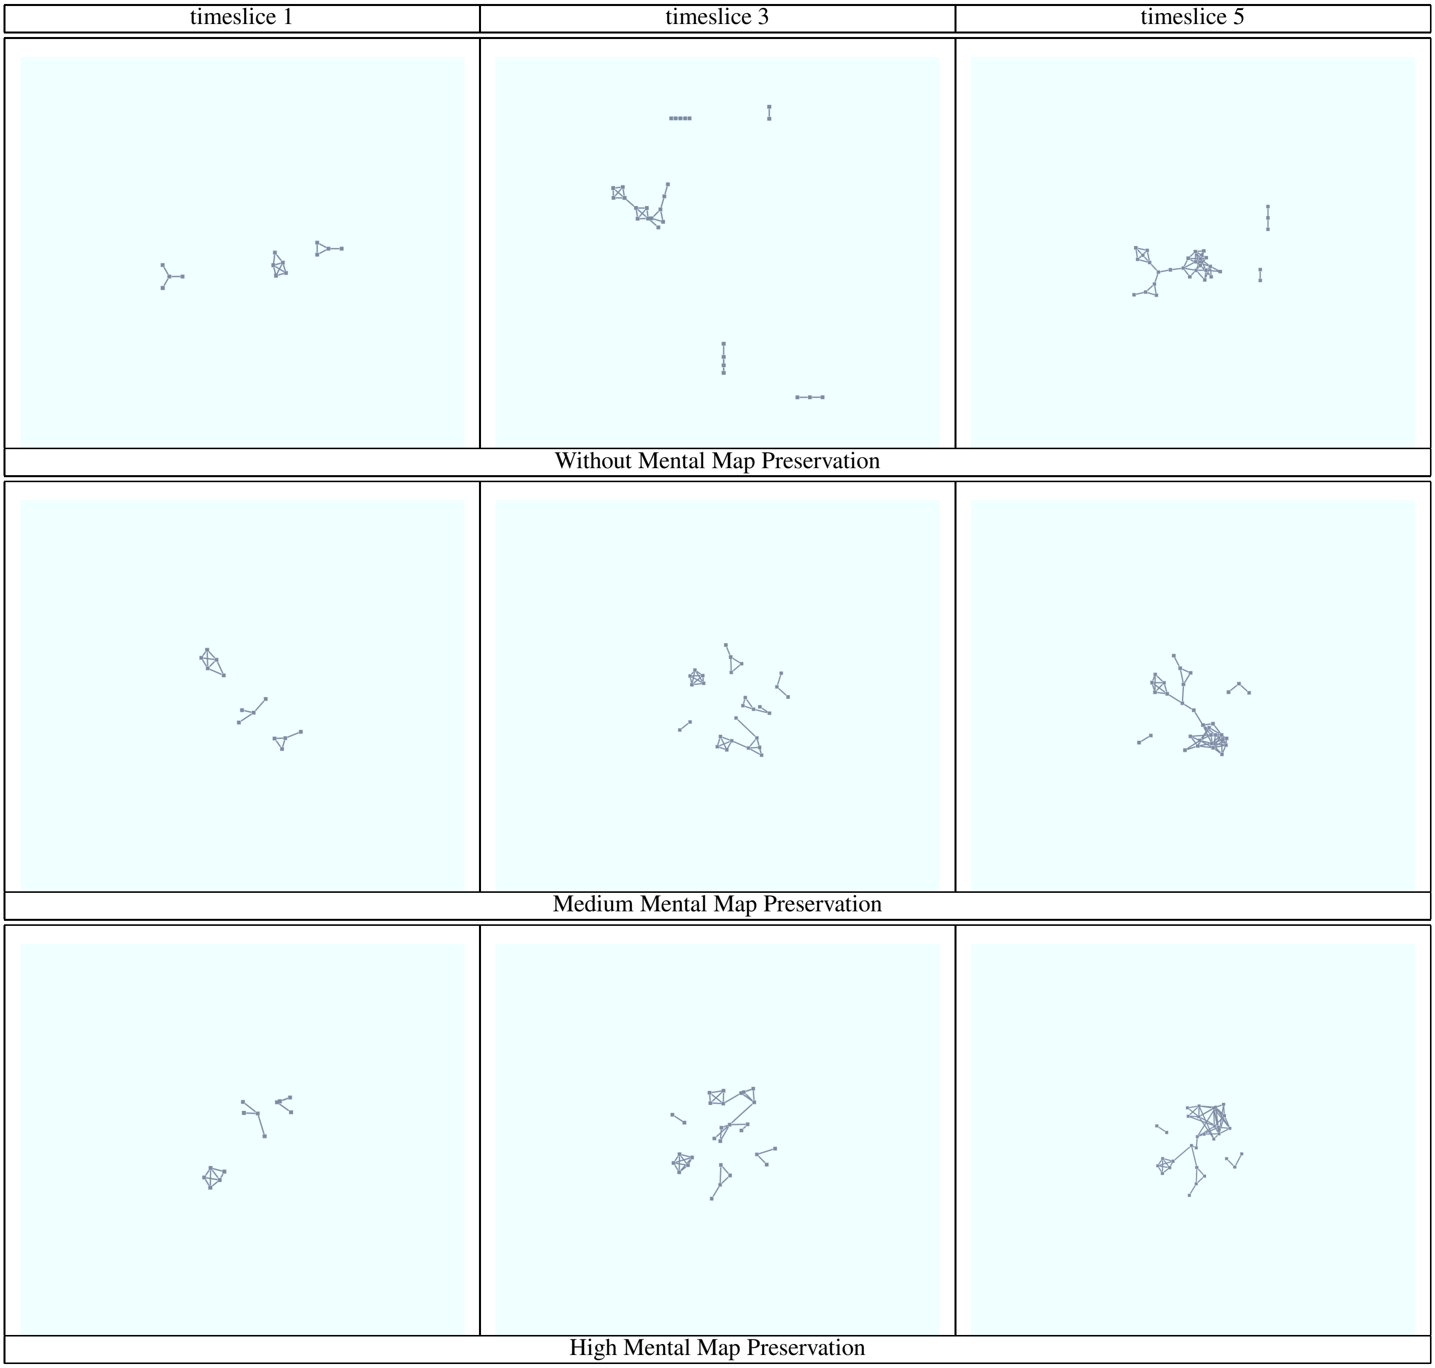

Figure 2: The three levels of mental map preservation used in this experiment on the `van de Bunt` data set. The "Without mental map preservation" level does not attempt to preserve the mental map at all. The medium level has inter-timeslice edges set to 60, and the high level has inter-timeslice edges set to 400 with nodes pinned to their average positions.

felt that they had memorized the graph, they clicked on "next sequence" to proceed to the next graph sequence to be memorized or the recall block. In Figure 1(b), the recall phase interface is shown. Participants played the video, decided if this graph had been present in the memorization block, and answered either "yes" or "no" using the radio buttons on the interface. In order to proceed to the next question, participants clicked on "confirm answer".

3.3 Data Sets

Three data sets were tested in this experiment: `MIT`, `Threads2`, and `van de Bunt`. All data sets were six timeslices long and drawn from real data sets to be of realistic size and structure. Graph size and length of the time series was informed through piloting and other experiments performed on dynamic graph drawing readability [21, 1, 2]. These data sets are described below.

`MIT` is six timeslices of the MIT Reality Mining project data set [10]. The data set consists the mobile telephone calls of a hundred students at MIT over the course of the 2004-2005 academic year. Nodes are mobile telephones and edges exist if a call was made between those two mobiles. The data set was divided into seven equal timeslices each consisting of about forty days of mobile telephone activity. Eight non-interacting two node components were removed from the data set. We selected the first six timeslices of this data for the experiment.

`Threads2`, used in the work of Frishman and Tal [13], is a graph series representing online newsgroup discussions. Nodes are authors of newsgroup articles, and an edge exists between two authors if one replied to the posting of another. We selected the last six timeslices of this data.

The `van de Bunt` data set [26] is a network that was used previously in dynamic graph readability experiments [1, 12]. The nodes in the graph are undergraduate students. An edge exists between two undergraduates if they self-reported in a survey that they had a relationship. In the original graph, there were a number of edge types which encoded if the relationship was best friends, friends, friendly neutral, or troubled. In our experiment, we used only the links that were best friends or friends and the last six timeslices of the data.

3.4 The Mental Map

Graphael [11] 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.

In this experiment, we used three levels of mental map preservation: no preservation, medium preservation, and high preservation. For the series of layouts that did not preserve the mental map, each frame is drawn independently with the force-directed algorithm used by Graphael. The medium preservation condition set inter-timeslice edges to a value of sixty. These relatively stiff inter-timeslice edges attempted to keep the same node in approximately the same area of the drawing in all the timeslices in which it was present. Thus, the heuristic would drive the layout to move as few nodes as possible as little as possible and attempt to preserve graph shape. For a high level of mental map preservation, we set the value of the inter-timeslice edges to 400 and took the average position of each node throughout all time, pinning them to a final position. When nodes were pinned they did not move at all through the entire animation sequence. An example depicting the evolution of `van de Bunt` under all three mental map conditions is shown in Figure 2. The approach for selecting inter-timeslice edge strength was based on visual inspection, piloting, and consistency with the values used in previous readability experiments [21, 2].

3.5 Experimental Design

Participants began the experiment with a session that demonstrated 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.

The experimental procedure required all participants to perform all tasks in the experiment. Thus, there were thirty-six experimental tasks in all: 3 mental map preservation conditions, 3 data sets, and 4 versions (2 unaltered, 2 modified). These thirty-six tasks were divided into three blocks of twelve questions each. Each of these memorization blocks had three graphs for which each data set and mental map preservation condition was asked exactly once. Memorization blocks for the participants were selected via all twelve solutions to the Latin squares problem and randomized for each participant. The recall phase consisted of the two variant graphs and the unaltered graph (asked twice) for each of the three data sets, giving a total of twelve questions. These questions were also independently randomized for each participant. Thus, after all three blocks were answered all three graphs under all three conditions were asked to each participant. These blocks were prefixed with a practice block of three memorization graphs and twelve recall graphs to help overcome any learning effects. Participants were not aware that this initial block of twelve questions did not form part of the experimental data collection.

The animations were rendered in real time using the Tulip framework [3]. No time limit was enforced per question or for the experiment overall. However, a warning label appeared on the screen after forty seconds had elapsed for each question, and participants were encouraged to finish their work quickly after that point. The animation started playing automatically after a delay of three seconds and took about ten seconds to play in its entirety.

Each experiment was conducted individually with the researcher and took approximately thirty minutes, including the pre-experiment training, practice tasks, experimental tasks for both experimental conditions, and post-experiment questionnaire. Overall, the data from twenty-five participants was used. Participants were drawn from members of the Complex and Adaptive Systems Laboratory of University College Dublin (UCD CASL).

4 Results

We present our overall results comparing the effects of no, medium, and high mental map preservation on memorability in this section: both overall and divided by data set. A Shapiro-Wilk test ($\alpha = 0.05$), was used to determine whether or not the data was normally distributed. For the normally distributed error data, a repeated measures ANOVA was used to determine significance. For the non-normally distributed response time and preference data, we used a Friedman non-parametric analysis. Significant differences between pairs of conditions were determined with adjusted pairwise comparison tests.

In all cases, a standard significance level of $\alpha = 0.05$ was used. After analyzing the data for the effect of mental map preservation, we analyzed the separate data sets for each of the three graphs: `MIT`, `Threads2`, and `van de Bunt`. In this case, we applied a Bonferroni correction, thus reducing the required significance level to $\alpha = 0.025$.

4.1 The Mental Map Condition

Figure 3 compares the three mental map preservation conditions with respect to both error rate and response time. We found that response time was not normally distributed while error rate was. No significant difference in terms of response time ($p = 0.19$) or error rate ($p = 0.27$) was found.

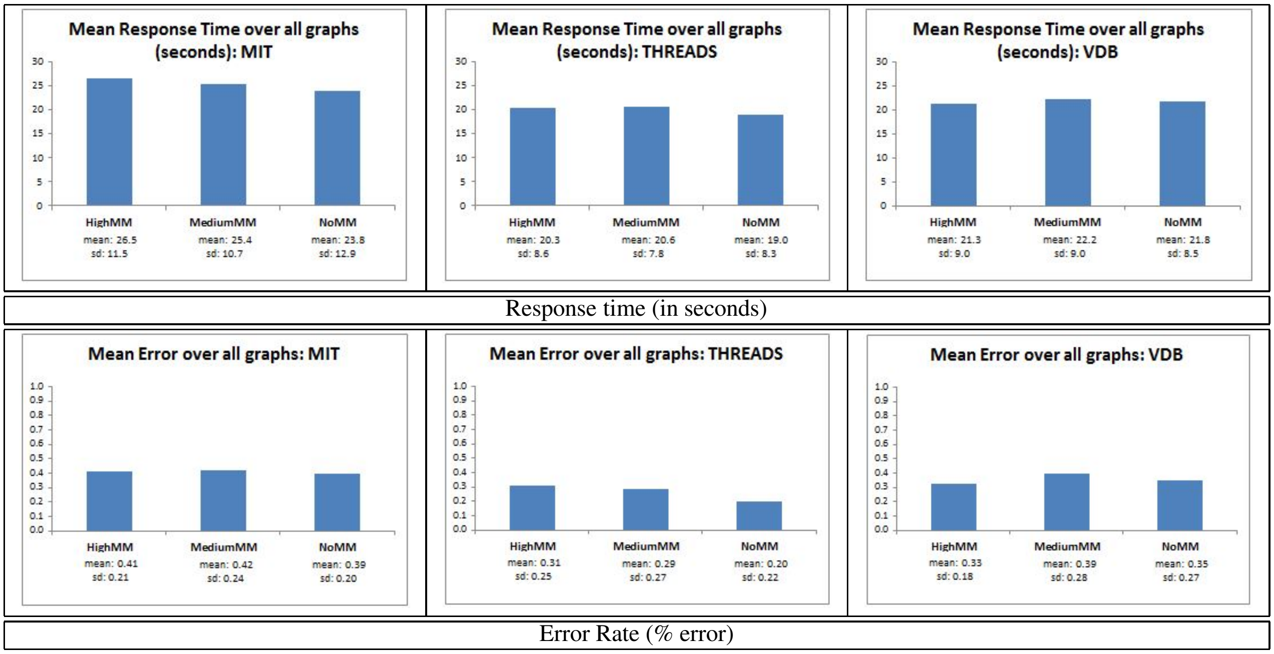

Figure 4: Response time and error rate divided by data set. No significant difference was found in either response time or error rate for any of the data sets individually.

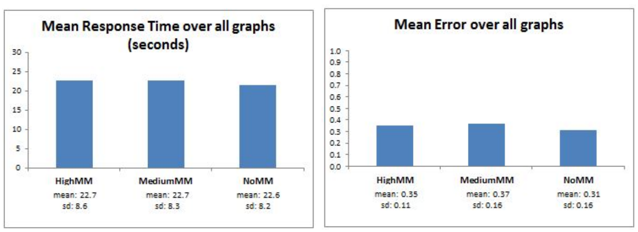

(a) (b)

Figure 3: Response time and error rate for high, medium, and no mental map preservation over all data sets. We found no significant difference in terms of response time ($p = 0.19$) or error rate ($p = 0.27$) in this data.

4.2 Mental Map Divided by Data Set

We subsequently divided the data by data set. Figure 4 shows the results of this analysis. No significant difference was found when any of the data sets were considered in isolation.

4.3 Preference Data

In terms of the quantitative preference data, we found significant results for both ranking questions of our survey. For these questions, numbers were a ranking score provided by the participants with 1 corresponding to most preferred or useful and 3 corresponding to least preferred or useful. Figures 5 and 6 show the results of this analysis. In question 1, we asked:

1. In some of the dynamic graphs, nodes moved large distances across the screen. In some, nodes moved short distances. In others, nodes did not move at all. Please rank these three methods according to how **easy** they made it to remember the graphs (1 is easiest and 3 is most difficult):

    - nodes move large distances
    - nodes move short distances
    - nodes do not move at all

For this question, the data was not normally distributed and we found that no movement was significantly preferred to large movement ($p = 0.011$) and small movement ($p = 0.021$).

In our second question, we asked:

2. Please rank how important the following features were to identifying the graphs (1 is most important and 3 is the least):

    - large node movement
    - constant node position
    - certain groups of nodes that were unique/constant

This data was not normally distributed. We found that certain groups of nodes ($p = 0.003$) and constant position ($p = 0.011$) were considered significantly more important than large node movement.

5 Discussion

We found no significant difference in terms of response time or error rate overall or on any of the data sets individually. This result is consistent with a number of experiments that have looked at mental map preservation in terms of graph readability [20, 21, 22, 2, 27], even though our experiment looks at the different task of memorability. Our results, when taken in context with the other five studies on the topic, imply that preserving the mental map by keeping the same node in the same area of the plane does not seem to have a quantitative significant effect on task performance. This result is surprising given the importance of preserving the mental map in the dynamic graph drawing community.

In the survey data, participants found it easier to memorize diagrams drawn with a high degree of mental map preservation when compared to medium or no mental map preservation. Thus, even though quantitatively the mental map condition did not make a difference, subjectively participants felt that it did. Thus, it may be

	Ease of Method Memorability		
	Large Distances	Short Distances	No Movement
Mean	2.27	2.23	1.46
Median	2.5	2.0	1.0
Std. Deviation	0.83	0.71	0.71
	Importance of Feature		
	Node Movement	Const Node Position	Subgraphs
Mean	2.54	1.73	1.62
Median	3.0	2.0	2.0
Std. Deviation	0.76	0.78	0.64

Figure 6: Table of preference data. The top part of the table compares the three mental map preservation conditions while the bottom part compares the importance of certain layout features. We found that no movement in node position was significantly preferred to large movement and small movement. We also found that subgraphs and constant position were significantly more important than large node movement.

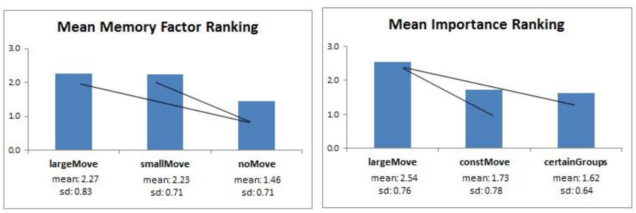

(a) (b)

Figure 5: Results of our ranked survey data. Lines between bars indicate where significance was found. **(a)** Question 1: We found that no node movement was significantly easier than short movement and large movement. **(b)** Question 2: We found certain groups of nodes (subgraphs) and constant node position significantly more important than large node movement.

important for user confidence and perceived ease of memorization that the mental map is preserved as much as possible. Secondly, participants considered identifying certain groups of nodes significantly more important for memorizing the graph sequences than large or small node movement. As preserving the relative position of certain groups of nodes benefit from mental map preservation, qualitatively the mental map seems important in this case.

From our observations, and from the survey responses of the participants, we suggest that there is evidence of two features at work that are important in terms of the memorability of the graph:

- **Static Features**: static patterns in the interconnections between nodes in certain subgraphs
- **Motion Features**: patterns of movement in the graph.

The no mental map condition of `Threads2`, in particular, had a characteristic motion and participants noticed when it brought in newly appearing nodes from the periphery. Below, we provide some examples from our qualitative feedback regarding the features participants used to memorize the graph:

- Static Features:
  - shapes such as a kite, star, or ring
  - forms that I can recognize: house or star
  - interesting subgraphs such as cliques
  - relate graph shapes to geometric shapes and real world components
- Motion Features:
  - look and feel of the design and movement rather than the small details
  - specific key movements of certain parts of the graph
  - order of changes, inversions of subgraphs
  - used unique movement of graph to remember the sequence
  - how the graph expands or shrinks on the computer screen
- Combination
  - if I could relate part or all of the graph to a real world object and describe the changing graph in these terms. A crab with three legs grows a tail; this tail fattens and moves left and then a 3-star grows on the tail

As it seems that these two features aid in the memorability of dynamic graphs for animation, perhaps the mental map is more useful for other presentation methods, such as small multiples. In a **small multiples** [25] interface, all timeslices are presented on the screen at once with each timeslice placed inside its own window. The user scans the matrix of windows to see how the graph evolves. As there are no transitions between frames, there is no motion and the preservation of the mental map may be more important in this context. As small multiples has been shown to be more effective than animation on a number of tasks [2, 12], confirming this conjecture would be interesting.

To date, a number of experiments have not witnessed a quantitative benefit, either in terms of response time or error rate, when preserving the mental map on dynamic graph drawings. Thus, the principle in general may not be as important as originally assumed. However, we should continue to try and identify a task that may benefit from the mental map. Tasks, such as revisitation tasks [23, 14], would be a promising direction for this research.

6 Conclusions and Future Work

In this paper, we present a human computer interaction experiment to test the mental map in the context of dynamic graph memorability. As observed in the previous studies, we found no significant difference in terms of response time or error rate when preserving the mental map. It is becoming increasingly clear that preservation of the mental map, while a good idea in principle, and while valued by the graph drawing community, may not always support performance. However, our qualitative data suggests that preserving the mental map appears to make memorability tasks easier in terms of the user's perception. While including an algorithmic feature that attempts to minimize the movement of nodes may not actually improve performance, it could improve perception of its usefulness.

Further studies are required to fully understand how this definition of mental map preservation and others support user tasks.

We have tried to collect data that allows us to make generalizations about the mental map and memorability by including more than one data set and three preservation levels. It would have been infeasible to test a much wider range of data sets and preservation levels. The generalization of these results are therefore limited by these parameters. It was necessary for our participants to have some knowledge of graphs, meaning that our results only hold for this particular population. Running the experiment in a laboratory situation, on context-free graphs (even if based on real data sets) means that these results may not extrapolate to the visualization of graphs within an application context.

It is important to note that we are using a definition of the mental map whereby the same node should stay in the same area of the plane. Other definitions of the mental map exist [17], which preserve qualities of a dynamic graph through its evolution such as orthogonal ordering of components, average relative distance, edge orientation, and nearest neighbours. It would be interesting to test these definitions in readability and memorability task contexts.

Finally, in certain areas of the data mining literature, there has been some criticism of timeslicing as a technique for dealing with dynamic graphs in general. An interesting idea would be to look at more event-based techniques for dynamic graph drawing and visualization. In an event driven approach, time is not divided into equal length intervals containing all edges present during that interval. Instead, these techniques visualize the appearance of each node and each edge as its own proper event, achieving a higher resolution in time. Further work is required to develop systems that visualize dynamic graphs in this way and further evaluation is required to determine if these visualizations are effective.

Acknowledgements

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/I140. We would also like to thank John Hamer for writing the dynamic graph layout algorithm based on the Graphael model. Finally, we would like to thank all of the participants who took part in this experiment.

References

[1] 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.

[2] 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.

[3] 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.

[4] K. Boitmanis, U. Brandes, and C. Pich. Visualizing internet evolution on the autonomous systems level. In *Proc. of Graph Drawing (GD '07)*, volume 4875 of *LNCS*, pages 365–376, 2007.

[5] U. Brandes, D. Fleischer, and T. Puppe. Dynamic spectral layout of small worlds. In *Proc. of Graph Drawing (GD '05)*, number 3843 in LNCS, pages 25–36, 2005.

[6] R. F. Cohen, G. D. Battista, I. G. Tollis, R. F. Cohen, R. Tamassia, and I. G. Tollis. A framework for dynamic graph drawing. *ACM Symp. on Computational Geometry*, pages 261–270, 1992.

[7] M. K. Coleman and D. S. Parker. Aesthetics-based graph layout for human consumption. *Software - Practice and Experience*, 26(12):1415–1438, 1996.

[8] S. Diehl and C. Görg. Graphs, they are a changing – dynamic graph drawing for a sequence of graphs. In *Proc. of Graph Drawing (GD '02)*, volume 2528 of *LNCS*, pages 23–31, 2002.

[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] N. Eagle, A. Pentland, and D. Lazer. Inferring social network structure using mobile phone data. *Proc. of the National Academy of Sciences*, 106(36):15274–15278, 2009.

[11] 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 '03)*, volume 2912 of *LNCS*, pages 98–110, 2003.

[12] 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.

[13] Y. Frishman and A. Tal. Online dynamic graph drawing. *IEEE Trans. on Visualization and Computer Graphics*, 14(4):727–740, 2008.

[14] S. Ghani and N. Elmqvist. Improving revisitation in graphs through static spatial features. In *Proc. of Graphics Interface (GI 2011)*, pages 175–182, 2011.

[15] C. Görg, P. Birke, M. Pohl, and S. Diehl. Dynamic graph drawing of sequences of orthogonal and hierarchical graphs. In *Proc. of Graph Drawing (GD ' 04)*, volume 3383 of *LNCS*, pages 228–238, 2004.

[16] H. Lam, R. A. Rensink, and T. Munzner. Effects of 2-d geometric transformations on visual memory. In *Proc. 3rd Symp. Applied Perception in Graphics and Visualization (APVG 06)*, pages 119–126, 2006.

[17] Y. Lee, C. Lin, and H. Yen. Mental map preserving graph drawing using simulated annealing. In *Proc. of the Asia-Pacific Symp. on Information Visualisation*, pages 179–188, 2006.

[18] 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.

[19] S. North. Incremental layout in dynaDAG. In *Proc. of Graph Drawing (GD '95)*, volume 1027 of *LNCS*, pages 409–418, 1995.

[20] H. C. Purchase, E. Hoggan, and C. Görg. How important is the "mental map"? – an empirical investigation of a dynamic graph layout algorithm. In *Proc. Graph Drawing (GD '06)*, volume 4372 of *LNCS*, pages 184–195, 2006.

[21] 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.

[22] 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.

[23] 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.

[24] M. Tory, C. Swindells, and R. Dreezer. Comparing dot and landscape spatializations for visual memory differences. *IEEE Trans. Visualization and Computer Graphics (Proc. InfoVis 2009)*, 16(6):1033–1040, 2009.

[25] E. Tufte. *Envisioning Information*. Graphics Press, 1990.

[26] M. van Duijn, E. Zeggelink, M. Huisman, F. Stokman, and F. Wasseur. Evolution of sociology freshmen into a frendship network. *The Journal of Math. Sociology*, 27(2):153–191, 2003. `http://www.stats.ox.ac.uk/~snijders/siena/vdBunt_data.htm`.

[27] L. Zaman, A. Kalra, and W. Stuerzlinger. The effect of animation, dual view, difference layers, and relative re-layout in hierarchical diagram differencing. In *Proc. of Graphics Interface (GI 2011)*, pages 183–190, 2011.