Comparing Graph Layouts for Vertex Selection Tasks

Roman Klapaukh
Victoria University of Wellington
Wellington
New Zealand
roma@ecs.vuw.ac.nz

David J. Pearce
Victoria University of Wellington
Wellington
New Zealand
djp@ecs.vuw.ac.nz

Stuart Marshall
Victoria University of Wellington
Wellington
New Zealand
stuart@ecs.vuw.ac.nz

ABSTRACT

Different graph layouts can affect a user's ability to complete both passive understanding and active interaction tasks. While most research exploring the effects of graph layout looks at a user's ability to accomplish a passive understanding task, this paper's novel contribution is looking at their ability to complete a selection task. Specifically we compare two graph layout algorithms with respect to their suitability for free-form multi-selection.

The two layout algorithms are drawn from our previous work which established that they have significantly different understandability metric scores. A motivation for this choice was to explore whether graphs with significantly different metric scores will also have significantly different performance for selection tasks. We carried out our comparison by means of a user experiment that followed a within-subjects design, where 74 users were given a PlayStation Move controller to select vertices in 20 pairs of graphs. We found that while there was no difference in the speed of interaction there was a difference in the number of errors users made between the two layout algorithms.

Keywords

Graph Interaction, Graph Layout, NUI Controllers

Categories and Subject Descriptors

H.5.2 [**User Interfaces**]: Evaluation

General Terms

User Evaluation

1. INTRODUCTION

Different graph layout properties have been shown to optimise different sorts of user tasks (Blythe et al., 1996; Purchase et al., 2002). These properties are often expressed in terms of metrics, allowing them to be explicitly measured. The evaluations of these metrics generally focus on how easy it is to passively read and understand a graph. While reading and understanding are important, they are both passive. Optimising a graph layout for passive tasks does not necessarily optimise it for interaction. Interaction itself is a broad category of tasks which can be divided up into five: Navigation, Selection, Manipulation, System Control and Symbolic Input (Bowman et al., 2004).

We explore selection as it is relevant to all graphs, irrespective of size or structure. Furthermore, interactions for system control or manipulation often require a selection subtask to determine the scope of the interaction. Selection's inherent sensitivity to the relative locations of the vertices suggests that selection is sensitive to the graph layout.

We chose two variants of the force-directed graph layout algorithm to use for this study. Force-directed graph layout algorithms are widely used in real systems (including D3 and GraphViz) to lay out graphs (AT&T; Bostock et al., 2011). We draw our variants from our previous work which is the largest metric-based evaluation of force-directed layout algorithms conducted to date (Klapaukh et al., 2014). In that work, experiments were run across 13,720 graphs to compare several variants of the force-directed graph layout algorithm. The experiments measured the number of overlapping pixels, and considered algorithms which minimised overlaps to be better than those which did not. We identified several variants of the standard force-directed layout algorithm that reduced the number of overlaps with statistical significance. These metrics focused on how easy the graph is to read, as overlaps hide information from the user. The experiments showed that the control algorithm (H), and the HWED algorithm (one of the top performers) had significantly different metric scores.

In this paper we extend the aforementioned work to examine whether the differences in metrics we recorded also result in a difference in user performance at a free-form selection task. We conducted a user study where graphs laid out by two different graph layout algorithms (HWED and the control (H)) were compared with free-form selection tasks using a PlayStation Move controller.

We chose to use the PlayStation Move because entertainment consoles have become a standard household good, owned by over 50% of households in North America alone (Entertainment Software Association, 2013; Entertainment Software Association of Canada, 2012) many of which now come with gesture controllers (Gallagher, 2010; Microsoft, 2010; Nintendo of America Inc.). These entertainment consoles can be used to display entertainment media (e.g. video

OzCHI '15, December 07-10, 2015, Melbourne, Australia

games, movies, music) which have natural graph representations based on similarity and purchasing recommendations (Bogers, 2010), and so are a prime target for graph interaction.

Contributions. This paper describes a user experiment where 74 participants select vertices in graphs laid out by two different algorithms to see if there was a difference in user performance where a difference in metric values has already been shown. The experiment was carried out in a controlled lab environment, with a within-subjects experimental design. Overall we find that:

- There is no significant difference in time taken to complete a selection between graphs laid out using the two layout algorithms.

- Participants made fewer errors when using graphs laid out by HWED (the modified algorithm) than when using the control algorithm (H).

In summary, there is a difference in user performance for a selection task between two graph layouts previously shown to be different with respect to passive understanding metrics.

2. RELATED WORK

We start by looking briefly at graph layout, and then focus on human computer interaction research.

2.1 Graph Layout

A graph consists of a set of vertices and edges, where each edge connects exactly two vertices. A wide variety of data can be represented in this way, such as social networks (Brandes and Wagner, 2004) and scene graphs (Autodesk, Inc., 2013).

A common approach to graph layout is force-directed graph layout, which originated with Eades (1984) Spring Embedder. Much work has been done on modifying the force-directed algorithm, but has largely been evaluated using only small numbers of graphs (e.g. (Frishman and Tal, 2009; Huang et al., 2006; Li et al., 2005; Lin et al., 2009)). More recently we published a large experiment comparing the effects of adding additional forces to Eades' algorithm (Klapaukh et al., 2014). The experiments showed that the standard algorithm, using Hooke's Law for attraction and Coulomb's Law for repulsion (called H) performs worse with respect to overlaps in the resulting layout than a variant which adds charged walls, charged edges, and varying repulsion for nodes based on their degree (referred to as HWED).

However, we only evaluated the algorithms with respect to three metrics: numbers of edge crossings, overlaps, and occluded pixels. While these are widely known and standardised metrics (Purchase, 2002), these metrics have been primarily evaluated in the context of passive reading and understanding tasks (Blythe et al., 1996; Purchase et al., 2002). This paper extends the information provided by these metrics by comparing these graphs in the context of an interactive selection task.

2.2 Human Computer Interaction

Selection is one of the core interaction tasks (Bowman et al., 2004; LaViola and Marks, 2010). As such, we use game controllers to select vertices in small graphs. As the graph is small, navigation is assumed to be unnecessary, and the whole graph is expected to fit on screen.

In terms of types of gestures, this paper looks exclusively at deictic (pointing) gestures (Karam and Schraefel, 2005). Deictic gestures are used under the assumption that they will be familiar given the wide spread of touch devices which use direct touch selection.

Research on the speed of interaction tasks has resulted in rules like the Steering Law (Accot and Zhai, 1997) which predicts that the narrower and longer the path that a user needs to take for an action is, the longer the action will take. The steering law has been validated for gesture devices like the Nintendo Wii Remote (Wrzesien et al., 2011). More recently this has been disputed for screen pointing tasks similar to those used in this research using the PlayStation Move (Kopper et al., 2010). Recent literature suggests that two part models which also consider gain (the ratio of the distance moved in the real world to the distance moved on the screen) are better predictors (Shoemaker et al., 2012). For large displays such as televisions with gesture controllers, values of gain may vary dramatically from values when using a computer mouse with a standard monitor. The conclusion that narrower and longer courses are harder to navigate still holds, but with the added consideration that it is affected by how far a user must move their hand to move the cursor across the display.

2.3 Gesture Controllers

Rather than using typical VR controllers, this paper uses a controller that was already on the mass market and in common circulation to ensure that the work had a concrete application in the current state of hardware. Three gesture controllers have come out with games consoles on the mass market that were considered for this work: the Nintendo Wii remote (Nintendo of America Inc.), the Microsoft Kinect (Microsoft, 2010), and the Sony PlayStation Move (Gallagher, 2010). Due to their earlier release dates more work has been done with the Nintendo Wii remote and Kinect (e.g. (Kamel Boulos et al., 2011; Williamson et al., 2010; Wrzesien et al., 2011)). We use the PlayStation Move controller for these experiments for its modern API.

Using a gesture controller as opposed to a computer mouse can cause additional difficulties. Bowman et al. (2002) found that pushing buttons on a hand-held controller causes the pointer to move, reducing accuracy. They called this the Heisenberg effect. Additionally, continuous use of interfaces where the user keeps their arm up can cause fatigue (Pfeil et al., 2013). We attempt to address both these points in our experiment.

3. EXPERIMENT

This section describes an experiment to evaluate the difference in user performance and preference when performing vertex selection tasks using a PlayStation Move controller, on graphs laid out by a control force directed layout algorithm and the HWED algorithm, as described in our previous work (Klapaukh et al., 2014). The experiment was approved by the Victoria University of Wellington Human Ethics Committee. Each participant had to select a marked set of vertices in a pair of graphs (actually the same graph laid out by a different layout algorithm). They then answered three questions about their preferences and perceived speed. The system recorded all user interactions to determine the true time spent on each graph.

3.1 Graph Layout Algorithms

The two graph layout algorithms compared are a standard force-directed graph layout algorithm (H), and the HWED algorithm from our previous work (Klapaukh et al., 2014). We showed that graphs produced by HWED have a reduced probability of overlaps, and potentially increased space between the vertices. The purpose of this experiment is to explore the effect of differences in metrics on human performance. While the metrics reported in our previous work showed significant differences in the layout, it is unclear whether those differences will affect a user interacting with the graph.

3.2 Graph Interaction

Choosing to work in the context of an entertainment console limited the common devices to either a gesture controller or D-pad (including the analogue stick). However, the D-Pad is not ideal for this task. Using the D-pad to jump from vertex to vertex by moving along edges may be difficult as there may be more than one edge going in a single direction. Additionally, it is hard to draw free form polylines around an area using the D-pad (directional or analogue), due to both awkwardness of interaction (using the D-pad to control a cursor) and the slow speed of interaction that is required to keep control of the cursor.

Our experiments use the PlayStation Move rather than the Nintendo Wii remote as the Move is more modern and has an accessible API to work with provided by Sony's Move.Me software (Sony Computer Entertainment America LLC, 2013). This took care of tracking the controller using a PlayStation 3, ensuring that the controller behaved as standard. The Microsoft Kinect system was not considered as it provided no clear means either to disengage from the system, or to easily "click" at a given coordinate without pausing, due to its lack of buttons. Pausing would significantly affect the time taken to complete each selection.

3.3 Data Set

This paper uses undirected simple graphs from the IEEE (2009) VAST challenge. The VAST challenge dataset contains a fake social network. It contains 6,000 individuals, their connections and cities of origin. All data is fabricated, and the place names are fictitious. It is represented by individual plain text files that specify the people (vertices), and their connections (edges).

Each graph was made by taking the friend network of each user out to two degrees. It is hypothesised that these graphs are more likely to be structured like real social networks than randomly generated data, and they are publicly available. A random subset of 20 of these graphs having no more than 50 vertices was used. There was no limitation on the number of edges.

The clusters of vertices to select were manually assigned, to control for the layout possibly not making the same groupings of vertices in the layout. This allowed us to make the tasks in each pair of graphs as similar as possible.

Each cluster was assigned such that the entire cluster could be selected in a single selection and contained up to 5 vertices. This was done to reduce the scope of the experiment and is based on the hypothesis that multiple selections will behave the same as repeated single selections.

3.4 Test Participants

Participants were recruited by asking students attending Victoria University of Wellington to participate and to invite their friends. Participants were informed the study would take 15 – 30 minutes, and they would be put into a draw to win gift vouchers for participation, as well as a prize for the fastest performance. Participants had to be at least 14 years of age, with any below 18 requiring parental consent.

3.5 Setup

The display was a 40" Panasonic TV with a resolution of 1920x1080 pixels. The user was seated on a computer chair which they could adjust as they chose. The chair was positioned opposite the camera which was at the bottom centre of the screen. The camera was raised on some books, as the table was both too wide, and the user too high relative to it. A sticky note was hung on each side of the screen with one of the buttons on the PlayStation controller. The entire set-up can be seen in Figure 1.

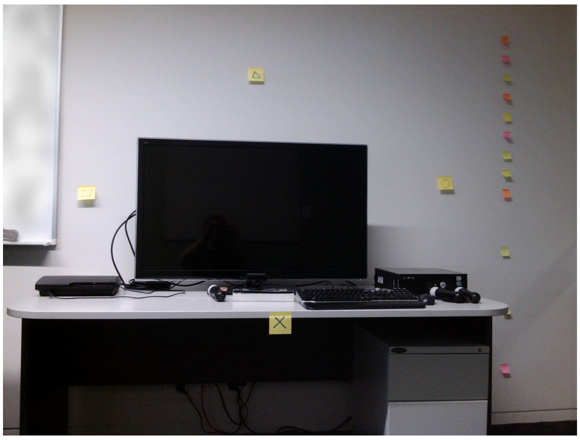

Figure 1: Experimental setup. The PlayStation eye camera is propped up on a book as the desk is too low. The vertical column of Post-it notes contains measurements that participants could use to estimate their height.

As the room being used was quite narrow, and the table quite wide, participants could not move around much without the camera losing track of the controller. For some participants, their arm's reach could extend out of the camera viewing angle. Fortunately this did not prove to be an issue with only one participant performing large enough gestures for this to manifest, and only in the practice graphs. The participants were all initially seated in a fixed location in order to minimise differences in gain.

Both the software used to interface to the PlayStation Move[1] and track and display the user interaction with graphs[2] were written for this study and are open source.

3.5.1 *PlayStation Move*

The PlayStation Move system tracks the location and orientation of a special controller in 3D space using a camera. The Move controller (shown in Figure 2) has a sphere on the top which is uniformly illuminated by internal lights. Direct tracking of this glowing sphere combined with information from sensors within the controller allows it to be used as a

[1] https://github.com/klapaukh/JMoveMe

[2] https://github.com/klapaukh/GraphLayout

mouse replacement, provided that the sphere can be kept within view of the camera.

Figure 2: A PlayStation Move controller. The sphere (on the left) can be lit up in a range of solid colours

There are two practical difficulties with the system. The first is that the camera must be set up with a clear view of the participant. This requires the participant be sufficiently far from the screen that they fit in the camera's view. The second is that the camera has to be connected to a PlayStation 3 console in order to do the processing.

3.6 Tasks

The experiment was divided up into four parts, taking 15 – 30 minutes in total. Questionnaires in the study used a seven point Likert scale. Initially each participant filled out a pre-questionnaire about demographics and previous experience. They then calibrated the system, and performed four practice tasks. After the practice tasks the proper experiment began.

Participants were presented with two graphs, one after another, where they had to perform a similar selection task. They were supervised by the author. Unknown to them the two graphs were logically the same, but had been laid out using different algorithms. Once the marked vertices were selected the application automatically took the participant to the next graph. When both graphs had been completed they were presented with a question screen about the two previous graphs as in Figure 3. To leave this screen they had to answer all the questions and press the `x` button (so participants could change their answers). Each participant was presented with the same 20 pairs. The order of the graph pairs and the algorithms within each pair were shuffled. Upon completion of the selection tasks participants answered a post-questionnaire.

3.7 Interaction

The PlayStation Move was set up in laser mode, which treats the controller as a laser pointer (similar to ray casting). This takes care of the issue of reachability as it becomes an issue of angle rather than reach. To calibrate the controller, the participant pointed the controller at each of four Post-it$^{\text{TM}}$ notes around the TV (Figure 1) and pushed the button shown. These could be done in any order.

A green dot with a black outline and 20 pixel diameter was shown as the cursor. This was the raw position as read from the PlayStation Move Software. If the camera lost sight of the controller an error message was displayed on the screen, but the system would continue to attempt tracking.

In order to enter selection mode, participants had to hold down the trigger (under their index finger). This started drawing a self-closing pink polyline that showed in real-time which vertices would be selected when the button was released. The polygon obeyed the inside-outside rule, which meant that items could be deselected partway through a selection (Figure 4). Selected vertices were shown with a large green tick on them. Similarly, holding the Move button (under their thumb) drew a blue polyline with yellow fill that deselected vertices.

Allowing the selection to begin at an arbitrary point avoids Bowman et al.'s Heisenberg effect, as all precise movements can be done while the button is being held down.

4. RESULTS

The experiment had 74 participants. R version 2.14.2 (2012-02-29) was used for all the data analysis (R Development Core Team, 2010). 13 participants were female, the mean age was 24 years old, and the mean height was 176 cm. Charts showing the distribution of participant age and height can be seen in Figure 5. 75% of participants reported that their exposure to the PlayStation Move controller or similar controllers such as the Nintendo Wii remote was "rarely" (3 out of 7) or less. However, all participants had previously used such a device.

The binomial test was used to check that the ordering of which layout came first for each pair of graphs was random (to verify the implementation). It gave a 95% confidence interval of $[0.48, 0.53]$ suggesting that it was in fact random.

The participants answered three questions after each pair of graphs. The binomial test was used to see if the order mattered. For the two questions relating to performance - speed and accuracy - there was no bias based on position at 95% confidence. However, user preference slightly favoured the second layout shown with a 95% confidence interval of $[0.445, 0.498]$. The binomial test was used to determine if users correctly identified which layout they performed faster on. The 95% confidence interval is $[0.69, 0.74]$, suggesting that participants could usually identify which layout they had performed faster on.

Performing the binomial test to compare the algorithms gave the same results. Speed and accuracy have no significant difference based on algorithm ($p = 0.2983$ and $p = 0.5116$ respectively). Preferences show a significant difference of very small magnitude in favour of the modified algorithm ($p = 0.0363$, confidence interval = $[0.5018, 0.5550]$).

The next tests looked at speed differences by layout. The Shapiro-Wilk normality test showed that the distribution for times to complete a graph is not normal ($p < 2.2 \times 10^{-16}$ in both cases). Hence the Wilcoxon rank sum test was used to compare the times per graph based on different layouts, finding no significant difference ($p = 0.8039$). The system also determined the total path length drawn to complete each task. Again, the Shapiro-Wilk normality test found the distribution to be not normal ($p < 2.2 \times 10^{-16}$ in both cases), so the Wilcoxon rank sum test was used to show that the values were significantly different at 95% ($p = 4.3 \times 10^{-4}$). The difference in path lengths can be seen in Figure 6a, and the similarity in times in Figure 6b.

As a measure of how often participants made mistakes, the system recorded the number of times participants used the deselection tool per graph. The Wilcoxon rank sum test showed a significant difference at 95% ($p < 2.2 \times 10^{-16}$) in the use of deselection. Looking at the histogram of the log of deselection tool usage frequencies (Figure 7), shows that participants using the modified algorithm were less likely to deselect. As participants could not advance to the next graph without completing the task correctly, this means that they were less likely to make an error in their selection.

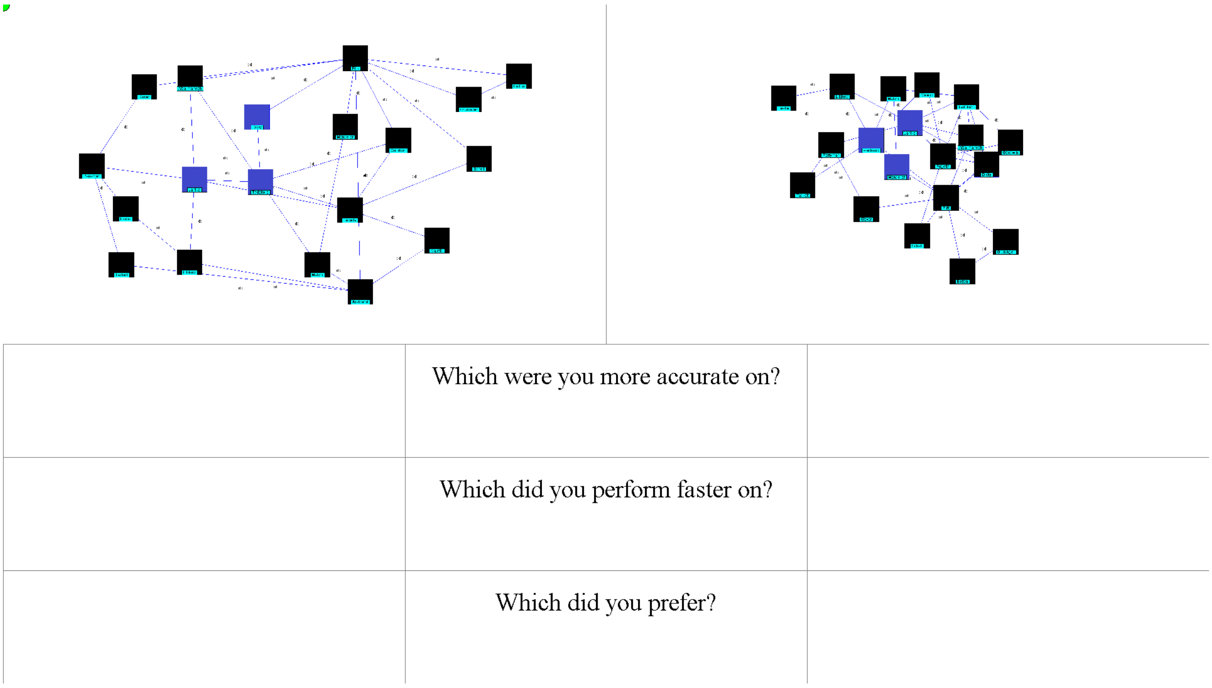

Figure 3: A screen asking the three questions about the two tasks that have been completed immediately prior. In this example the graph on the left has been laid out with HWED, and the graph on the right with H.

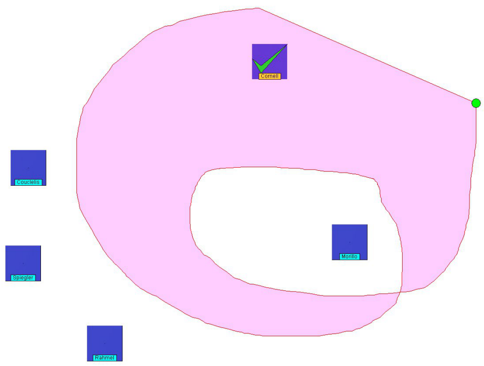

Figure 4: The polygon can have holes in it. The inside-outside rule states that a point is inside the polygon if and only if a ray going out to infinity to crosses an odd number of walls.

Participant feedback was positive. Each participant answered four main questions, on a Likert scale from 1 to 7, asking their agreement with the following statements: the system was fun to use, the system was novel to use, I would use this system again, the system did what I wanted.

Bar-charts of feedback for each of the questions can be seen in Figure 8. Values of 1–3 indicate disagreement, 4 neutrality and 5–7 agreement. These clearly show that participants largely chose ratings which agreed with all of the sentences, indicating positive feeling.

5. DISCUSSION

The feedback on the system was positive. The majority of participants indicated both that the system was fun and that they would use it again (Figure 8). This is good support for the concept of integrating such controls into the standard user interfaces of consoles. Moreover, only a small proportion of participants particularly disliked the system.

A number of participants attempted to use the controller with their arm fully extended at the beginning of the experiment. As discussed in the related work, this can cause excess fatigue and render the system unusable for long term interaction. While the participants who continued to hold their arm extended for the entire duration (6 minutes on average) did report feeling tired at the end, the majority either started off holding the controller comfortably or moved into a more comfortable position as the experiment progressed.

A number of participants started out doing repeated single item selection. All participants who started out doing selections in this way changed to using only a single free-form selection by the end of the study. No participants changed from multi-item selection to repeated single-item selection, suggesting that free-form multi-item selection is something participants would use, and so is desirable in a real system.

Participants had clear difficulty remembering choices that they had made previously. A number commented that they felt like they always preferred the first graph. This is contrary to the analysis which shows that the second graph was preferred, and was also not noticed by the experimenter.

The tendency to prefer the second graph, while not very strong, goes against the assumption that participants prefer graphs based on the layout algorithm rather than their ordering. Statistical tests suggest that the graphs were in fact shuffled, so this effect cannot be explained by layout algorithm. A possible reason is the similarity between each pair of graphs. For each pair, the vertices that need to be selected are in the same area on the screen (as much as possible), and

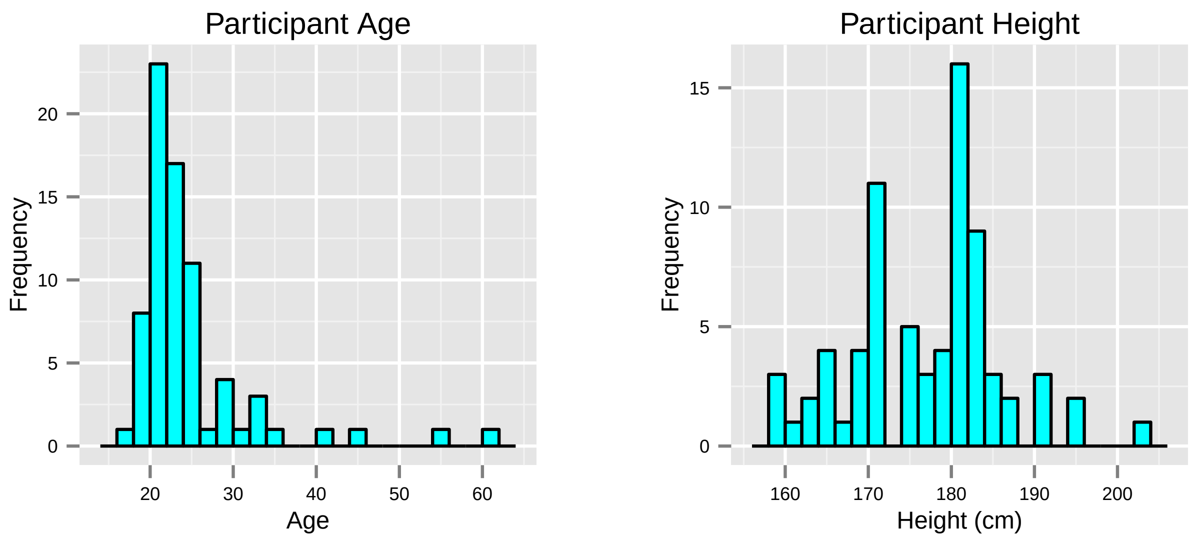

(a) Distribution of participant ages.

(b) Distribution of participant heights

Figure 5: Distributions of participant ages and heights. Most participants were young university students.

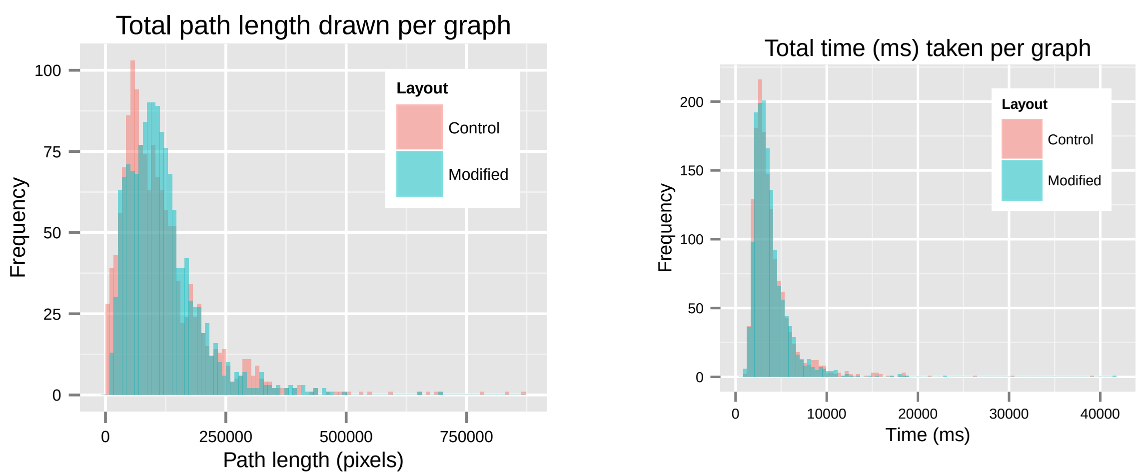

(a) The total path length used to make a selection.

(b) The total time taken to make a selection.

Figure 6: The frequencies of path length for a selection, and time taken for a selection. There is no significant difference between the times, but the lengths for the modified algorithm are longer suggesting that the user's arm is moving faster.

the participants may have learnt the location of the selection, and therefore considered the second task slightly easier than the first. Alternatively the first graph may not have been remembered clearly, and the second picked for that reason. However, the magnitude of this difference is small enough to be of little practical interest.

The experimental analysis showed that the modified layout (HWED) made selection tasks easier for the participants. There are two main factors that contribute to this conclusion. The first is that there was no difference in timing between the different layouts, but the paths drawn by the participants were longer in the modified case. Results on the steering law (Accot and Zhai, 1997) and the modifications due to gain (Kopper et al., 2010; Shoemaker et al., 2012), imply that, at a constant gain, if the difficulty of the paths is the same, the longer path should take more time. In this case, the longer path took less time suggesting that the user's arm was moving faster. As the gain is constant, the longer paths in the modified layout must have a lower index of difficulty. This also provides additional evidence that the modified layout (HWED) algorithm results in sparser layouts than the control.

The second is that participants used the deselect functionality less when they were performing selection tasks on

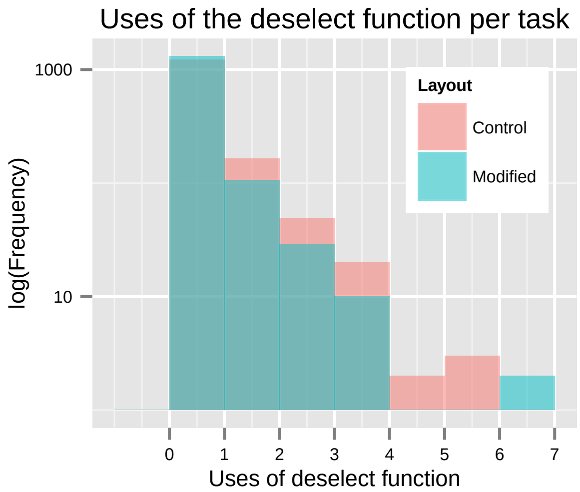

Figure 7: Number of times the deselection functionality was used. Note that the frequencies are shown as logarithms, and that for non-zero uses the control algorithm generally has higher incidence. This difference statistically significant.

graphs laid out with the modified algorithm. This means that they generally made fewer mistakes, while taking no more time to do it, reinforcing the results from the metric evaluation. Furthermore, this creates a safety net for novices so they can make fewer mistakes without having to go slower.

It may be that these results are strongly biased towards novices. As the PlayStation Move controller is not used widely for graph selection tasks, it is unlikely that the participants would have had a large amount of expertise before the experiment. While a number of them had previously used gesture controllers, they were unlikely to have been performing similar tasks. This is further supported by participant comments to the effect that they felt that they liked the more spacious layouts at the beginning, but started to prefer denser layouts as they became more confident and familiar with the system.

5.1 Limitations and Future Work

The sample of participants is drawn from quite a narrow range. It is largely young men studying computer science or engineering at university. This is only a single demographic amongst a variety that use entertainment consoles. While some participants are drawn from other areas, there are not enough of them to say that they would definitely exhibit the same behaviours, although in this experiment no differences were observed. However, we hypothesise that these results provide a strong indication about the sorts of people who did participate, and believe that further research on other demographic groups would find similar results.

This work limited itself to selection tasks that could be accomplished with a single free-form selection. It would be interesting to see how this extended to more complicated selections, where a single contiguous lasso selection would not be a practical mechanism. It would also be interesting to see if factors like decreasing the pointer sensitivity (by filtering the console's tracking) or having previous experience with the system changed participants' perceptions of the system.

6. CONCLUSION

This paper investigates the difference in user performance for free-form selection tasks between two graph layout algorithms that have already been shown to have a difference with respect to passive reading and understanding metrics.We performed a user experiment where 74 participants interacted with 20 graphs per algorithm using a PlayStation Move controller. Overall we find that there is a difference in performance between these two different layout algorithms. Users made less errors and moved faster when using the graphs laid out using the algorithm with better metric scores.

References

Johnny Accot and Shumin Zhai. Beyond Fitts' law: models for trajectory-based HCI tasks. CHI, pages 295–302. ACM, 1997.

AT&T. Graphviz. `www.graphviz.org/` accessed: 20 Nov 2012.

Autodesk, Inc. Autodesk Maya 2013 API Documentation: Querying the Scene Graph, 2013. `http://docs.autodesk.com/MAYAUL/2013/ENU/Maya-API-Documentation/index.html?url=files/GUID-0B85C721-C3C6-47D7-9D85-4F27B787ABB6.htm,topicNumber=d30e5360` accessed: 03/04/2013.

Jim Blythe, Cathleen McGrath, and David Krackhardt. The effect of graph layout on inference from social network data. In *Graph Drawing*, volume 1027 of *LNCS*, pages 40–51. Springer Berlin Heidelberg, 1996.

Toine Bogers. Movie Recommendation using Random Walks over the Contextual Graph. In *CARS*, 2010.

Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. D3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, December 2011. ISSN 1077-2626. doi: 10.1109/TVCG.2011.185. URL `http://dx.doi.org/10.1109/TVCG.2011.185`.

D. A. Bowman, C. A. Wingrave, J. M. Campbell, V. Q. Ly, and C. J. Rhoton. Novel Uses of Pinch Gloves for Virtual Environment Interaction Techniques. *VR*, 6(3):122–129, 2002.

Doug A. Bowman, Ernst Kruijff, Joseph J. LaViola, and Ivan Poupyrev. *3D User Interfaces: Theory and Practice*. Addison Wesley Longman, 2004.

Ulrik Brandes and Dorothea Wagner. Analysis and Visualization of Social Networks. In Michael Jünger and Petra Mutzel, editors, *Graph Drawing Software*, Mathematics and Visualization, pages 321–340. Springer Berlin Heidelberg, 2004. ISBN 978-3-642-62214-4. doi: 10.1007/978-3-642-18638-7_15. URL `http://dx.doi.org/10.1007/978-3-642-18638-7_15`.

Peter Eades. A Heuristic for Graph Drawing. *Congressus Numeratum*, 42:149–160, 1984.

Entertainment Software Association. 2013 essential facts about the computer and video game industry, 2013.

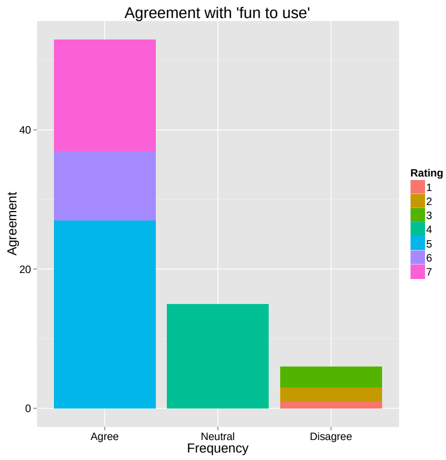

(a) Participant ratings for "fun".

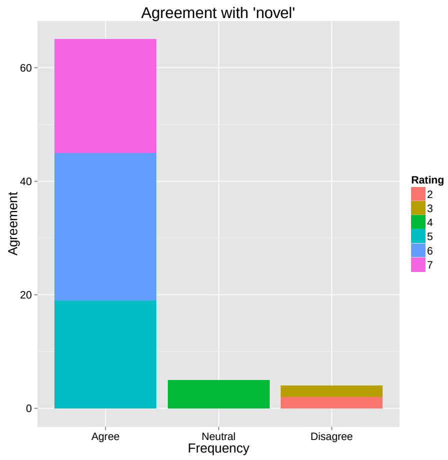

(b) Participant ratings for "novel".

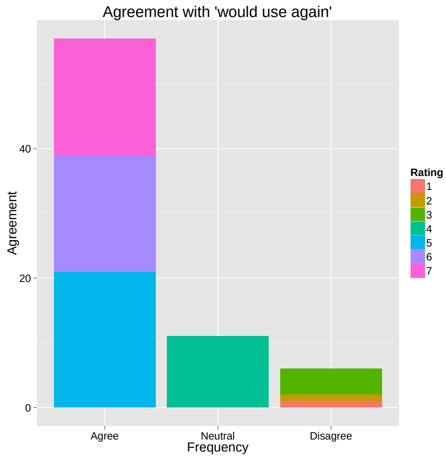

(c) Participant ratings for "use again".

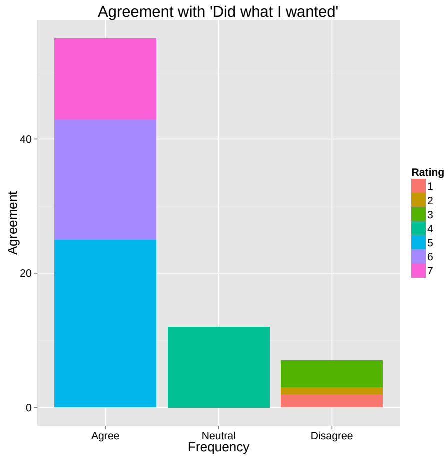

(d) Participant ratings for "did what I wanted."

Figure 8: Participant ratings for questions. 1 is disagree, 4 is neutral, 7 is agree. For all questions the majority of results are at 4 and above suggesting that participants had a favourable reaction to vertex selection using the PlayStation Move Controller.

Entertainment Software Association of Canada. Essential Facts 2012, 2012.

Yaniv Frishman and Ayellet Tal. Uncluttering Graph Layouts Using Anisotropic Diffusion and Mass Transport. *IEEE TVCG*, 15(5):777–788, 2009.

James Gallagher. Everything You Need To Know About PlayStation Move, 2010. `http://blog.eu.playstation.com/2010/09/07/everything-you-need-to-know-about-playstation-move/` accessed: 18 Feb 2013.

Xiaodi Huang, A. S. M. Sajeev, and Wei Lai. A Scalable Algorithm for Adjusting Node-Node Overlaps. In *CGIV*, pages 43–48. IEEE Computer Society, 2006. ISBN 0-7695-2606-3.

IEEE. VAST Challenge, 2009. `www.cs.umd.edu/hcil/VASTchallenge09` accessed: 24 April 2013.

Maged N Kamel Boulos, Bryan J Blanchard, Cory Walker, Julio Montero, Aalap Tripathy, and Ricardo Gutierrez-Osuna. Web GIS in practice X: a Microsoft Kinect natural user interface for Google Earth navigation. *International Journal of Health Geographics*, 10(1):45, 2011. URL `http://www.ncbi.nlm.nih.gov/pubmed/21791054`.

Maria Karam and Monica Mc Schraefel. A taxonomy of gestures in human computer interactions. Technical report, University of Southampton, 2005. URL `http://eprints.soton.ac.uk/261149/`.

Roman Klapaukh, David J. Pearce, and Stuart Marshall. Towards a vertex and edge label aware force directed layout algorithm. In *ACSC*, volume 147 of *CRPIT*, pages 29–37. ACS, 2014.

Regis Kopper, Doug A. Bowman, Mara G. Silva, and Ryan P. McMahan. A human motor behaviour model for distal pointing tasks. *International Journal of Human-Computer Studies*, 68(10):603–615, 2010.

Joseph J. LaViola and Richard L. Marks. An introduction to 3D spatial interaction with video game motion controllers. In *ACM SIGGRAPH Courses*, pages 2:1–2:78. ACM, 2010.

Wanchun Li, Peter Eades, and Nikola S Nikolov. Using Spring Algorithms to Remove Node Overlapping. In *APVIS*, volume 45 of *CRPIT*, pages 131–140, 2005. ISBN 1-920682-27-9.

Chun-Cheng Lin, Hsu-Chun Yen, and Jen-Hui Chuang. Drawing graphs with nonuniform nodes using potential fields. *JVLC*, 20(6):385–402, 2009.

Microsoft. Kinect for Xbox 360: Science Fiction Comes to Your Living Room, 2010. `http://www.microsoft.com/en-us/news/features/2010/nov10/11-03Kinect.aspx` accessed: 18 Feb 2013.

Nintendo of America Inc. |Nintendo - Corporate Information | Company History. `http://www.nintendo.com/corp/history.jsp` accessed: 18 Feb 2013.

Kevin Pfeil, Seng Lee Koh, and Joseph LaViola. Exploring 3D gesture metaphors for interaction with unmanned aerial vehicles. In *IUI*, pages 257–266. ACM, 2013.

Helen C. Purchase. Metrics for Graph Drawing Aesthetics. *JVLC*, 13(5):501–516, 2002.

Helen C. Purchase, David Carrington, and Jo-Anne Allder. Empirical evaluation of aesthetics-based graph layout. *Empirical Software Engineering*, 7(3):233–255, 2002.

R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2010. URL `http://www.r-project.org`.

Garth Shoemaker, Takayuki Tsukitani, Yoshifumi Kitamura, and Kellogg S. Booth. Two-part models capture the impact of gain on pointing performance. *ACM TOCHI*, 19(4):28:1–28:34, 2012.

Sony Computer Entertainment America LLC. Move.me - Software Tool for PS3 System that uses PlayStation Move Technology, 2013. `https://us.playstation.com/ps3/playstation-move/move-me/` accessed: 04/04/2013.

Brian Williamson, Chadwick A. Wingrave, and Joseph J. LaViola. RealNav: Exploring natural user interfaces for locomotion in video games. In *3DUI*, pages 3–10. IEEE, 2010. ISBN 978-1-4244-6846-1.

Maja Wrzesien, María José Rupérez, and Mariano Alcañiz Raya. Input Devices in Mental Health Applications: Steering Performance in a Virtual Reality Paths with WiiMote. In *INTERACT (2)*, volume 6947 of *LNCS*, pages 65–72. Springer, 2011.