

Stereoscopic Highlighting: 2D Graph Visualization on Stereo Displays

Basak Alper, Tobias Höllerer, JoAnn Kuchera-Morin, and Angus Forbes

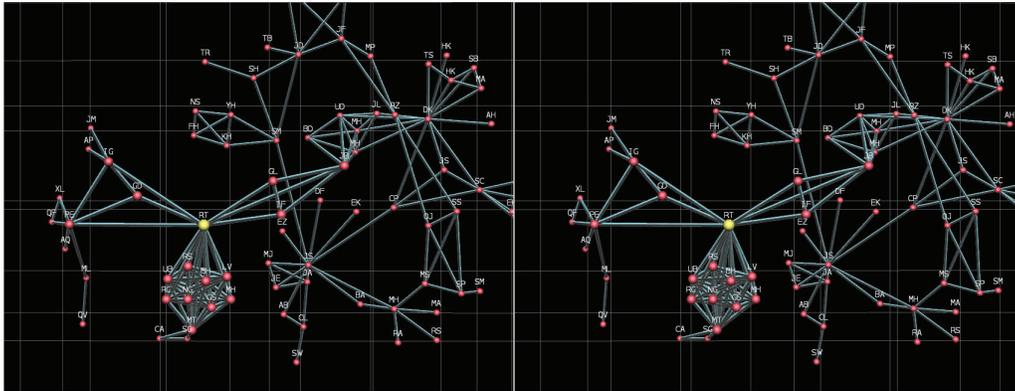


Fig. 1. The *stereoscopic highlighting* technique illustrated with separate images for the left and right eye. The highlighted yellow node and all of its immediate neighbors are rendered at a depth closer to the viewer. The two images above represent a stereo pair. By crossing eyes, the reader can perceive the 3rd dimension of the stereoscopic highlighting.

Abstract—In this paper we present a new technique and prototype graph visualization system, *stereoscopic highlighting*, to help answer accessibility and adjacency queries when interacting with a node-link diagram. Our technique utilizes stereoscopic depth to highlight regions of interest in a 2D graph by projecting these parts onto a plane closer to the viewpoint of the user. This technique aims to isolate and magnify specific portions of the graph that need to be explored in detail without resorting to other highlighting techniques like color or motion, which can then be reserved to encode other data attributes. This mechanism of *stereoscopic highlighting* also enables focus+context views by juxtaposing a detailed image of a region of interest with the overall graph, which is visualized at a further depth with correspondingly less detail. In order to validate our technique, we ran a controlled experiment with 16 subjects comparing static visual highlighting to stereoscopic highlighting on 2D and 3D graph layouts for a range of tasks. Our results show that while for most tasks the difference in performance between stereoscopic highlighting alone and static visual highlighting is not statistically significant, users performed better when both highlighting methods were used concurrently. In more complicated tasks, 3D layout with static visual highlighting outperformed 2D layouts with a single highlighting method. However, it did *not* outperform the 2D layout utilizing both highlighting techniques simultaneously. Based on these results, we conclude that *stereoscopic highlighting* is a promising technique that can significantly enhance graph visualizations for certain use cases.

Index Terms—Graph visualization, stereo displays, virtual reality.

1 INTRODUCTION

3D displays are becoming more ubiquitous. We witness increased acceptance and enthusiasm in the entertainment industry and an interconnected lowering of hardware costs. The technologies behind 3D displays are advancing at a rapid pace, creating higher quality stereoscopy. For instance, monitors with 140Hz refresh rate, capable of generating active polarized stereo images that do not require expensive LCD shutter glasses are easily accessible to end-users today. Additionally, 3D television sets featuring auto-stereoscopic

displays are making stereo more readily available to the general consumer. And the general public is becoming increasingly familiar with 3D displays and VR technologies through the increased accessibility of 3D movies and games. On the research side, volumetric displays are another promising technology, as recent studies have shown that users judge depth better using these displays than other stereo displays [12].

Recent scientific visualization techniques also take advantage of advances in VR and 3D display technologies. However, for information visualization tasks, utilizing the third dimension is a more complicated problem. Typically, data sets in information visualization have no inherent spatial encoding. Mapping data in 3D can complicate perception because issues related to occlusion and perspective and an ensuing necessity for viewpoint navigation can make it harder to glean relationships among data elements. However, 2D representations have their own limitations, such as issues related to the scalability of the data in terms of size and complexity. A primary example of a challenging problem for 2D visualization techniques is the design and layout of node-link representations for dense graphs [29].

- Basak Alper is with the Media Arts & Technology Program at University of California, Santa Barbara, E-Mail: basakalper@umail.ucsb.edu.
- Tobias Hollerer is with the Department of Computer Science at University of California, Santa Barbara, E-Mail: holl@cs.ucsb.edu.
- JoAnn Kuchera-Morin is with the Media Arts & Technology Program at University of California, Santa Barbara, E-Mail: jkm@create.ucsb.edu.
- Angus Forbes is with the Media Arts & Technology Program at University of California, Santa Barbara, E-Mail: angus.forbes@mat.ucsb.edu.

Manuscript received 31 March 2011; accepted 1 August 2011; posted online 23 October 2011; mailed on 14 October 2011.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

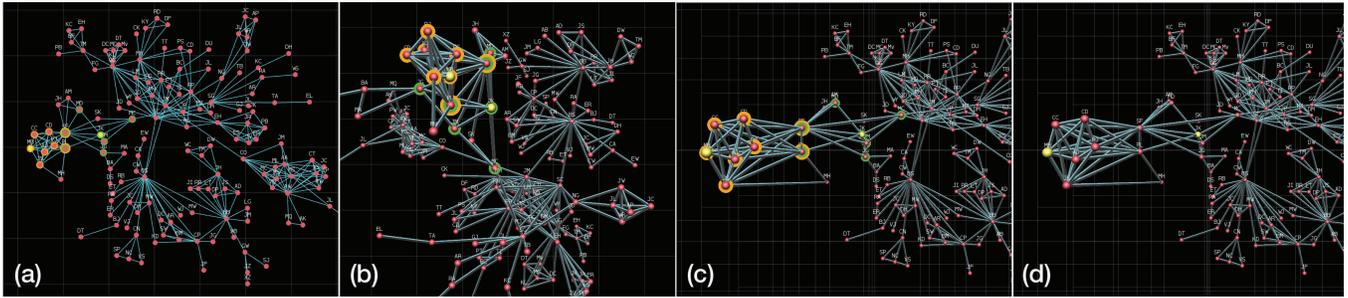


Fig. 2. Summary image of all 4 techniques: (a) 2D graph layout with static visual highlighting; (b) 3D graph layout with static visual highlighting; (c) 2.5DH graph layout utilizing both stereoscopic and static visual highlighting; (d) 2.5D layout utilizing only stereoscopic highlighting.

Node-link diagrams aim to provide an understanding of the overall structure of a graph while enabling users to identify individual links. For instance, presenting node-link graphs of small-world network structures, which are common for social networks, often show too many edge crossings, which makes it hard to discern adjacency without an on-demand highlighting mechanism [28]. Figure 2a shows the use of static visual cues for highlighting.

3D node-link visualizations can eliminate many edge crossings, and make it easier to identify the adjacency of specific elements [30]. Hence, certain tasks become easier, such as enumerating nodes that are accessible from a particular starting node. However, due to perspective, when the graph is rendered in 3D some nodes are rendered much closer to the virtual camera viewpoint than others (Figure 2b). With the addition of stereoscopic cues, much of the visual emphasis is placed on the nodes that are closer to the user’s viewpoint. But this emphasis is particular to the graph layout algorithm and the observer’s viewpoint, and independent of the data. When a graph is rendered as a 2D layout, the viewpoint remains relatively equidistant to each point on the graph. 3D layouts on the other hand often require increased viewpoint navigation. When a node of interest is rendered at a further depth in virtual space, the user may be required to rotate, zoom and adjust the view angle to get an optimal view. When the user is rotating the view angle of a 3D node and link diagram, however, the 2D projection on the screen is constantly changing. Thus, the graph appears different from different view angles, further complicating the task of building a mental map of the graph.

Our motivation for this study is to investigate alternative ways of making use of stereo displays for graph visualization. We propose a technique called *stereoscopic highlighting* that utilizes the visual emphasis provided by virtual depth to highlight points of interest on a 2D node and link diagram, as shown in Figure 2c-2d. Essentially, the technique involves bringing graph subsets of interest closer to the viewer in order to support tasks such as adjacency identification. By doing so, we aim to retain benefits of a 2D representation while leveraging the benefits of 3D display technologies.

We provide the results of a controlled experiment evaluating the effectiveness of *stereoscopic highlighting* with respect to static visual cues (colored concentric circles) on 2D and 3D graph layouts for five different tasks. Our results show that for certain tasks, the use of stereoscopic highlighting on a 2D graph layout, when used in conjunction with static visual highlighting, enhances user performance significantly. However, we did not find any significant difference in user performance when stereoscopic highlighting and static visualization highlighting were used independently. This is important, as it suggests that *stereoscopic highlighting* has the potential to supplement other techniques, but also to be used *in lieu of* static visual highlighting, freeing up visual variables for other uses.

Finally, we demonstrate how the stereoscopic highlighting technique can be utilized via a simple prototype implementation of a novel interactive graph visualization framework.

2 RELATED WORK

3D information visualization has been investigated for over two decades. Many examples of 3D information visualizations can be found in 0. *SemNet* [9], *Cone Trees* [26], *Data Mountain* [25], *Perspective Wall* [18] are well known examples. However, these visualizations did not make use of stereoscopy but instead used animated 3D visualizations that relied on kinetic depth, motion parallax, linear perspective and lighting for depth perception. A detailed description of these depth cues and their contribution to depth perception can be found in [7][27].

In all these visualizations, occlusion plays a detrimental role by hindering the visibility of elements or by making labels illegible. An evaluation of *Cone Trees* with respect to regular 2D tree visualizations reveals that users take longer to complete hierarchical navigation tasks, largely due to the problem of occlusion [4]. *Data Mountain* [25] devises a unique occlusion handling strategy by continually maintaining a minimum distance between all items, thereby mitigating the problem.

Occlusion is not the only problem. In a study evaluating the efficacy of 2D, 2.5D and 3D interfaces in physical and virtual environments, it has been shown that 3D (without stereoscopic depth cues) does not necessarily lead to a more effective use of spatial memory for high-level tasks such as building a mental map of document arrangements [5]. Other studies, however, point to appropriate use of 3D visualization.

In the case of graph visualization, work by Ware et al. demonstrates the advantages of 3D graph layouts with high-resolution stereo displays [29][30]. Their studies focus on general low-level graph analysis tasks such as identifying connections and path tracing between specified nodes. According to these studies, 3D visualizations are recommended only if real-time rotation and stereo cues are provided [29]. Their experiments reveal that for identifying adjacency tasks, for any given error rate, the amount of data that can be displayed is an order of magnitude greater when 3D depth cues of motion and stereo are both utilized. The improvement is about 60% when only stereo is used.

Conventionally, stereoscopy is used as a depth cue. However, it can also be used for segregating images that share the same texture with their background [13]. In military applications this is referred to as “camouflage breaking.” Work by Peterson et al. [22][23] uses stereoscopic displays to segregate overlapping labels. They concentrate on the issue of visual clutter associated with label overlap and propose utilizing stereo “label-layering” to address the problem.

In a similar manner to “camouflage breaking,” our technique incorporates stereoscopy to isolate subsets of a 2D graph as a separate layer, and thus create 2.5D layouts. There are many examples of 2.5D graph layouts where the third dimension is exploited to encode aspects of the data other than the relations among nodes [3]. For example, in [8] the third dimension is used to

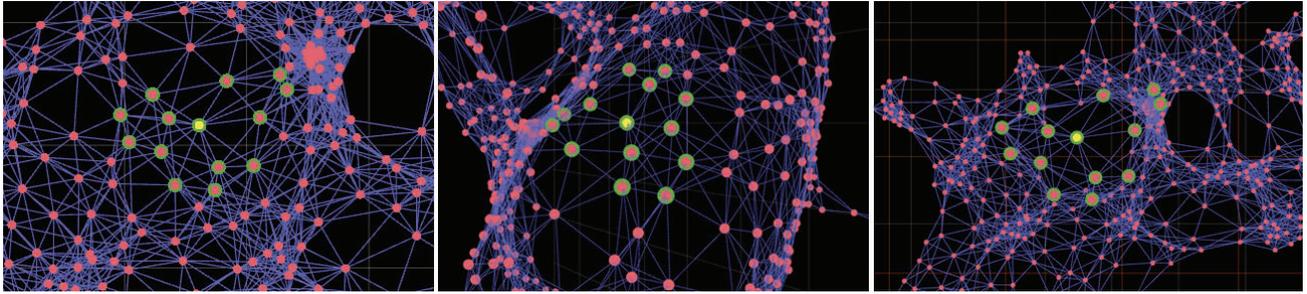


Fig. 3. Focus+context views for different visualizations illustrating how much of the overall graph is kept in view when highlighted nodes are scaled to the same size: (left) the 2D graph layout requires zooming in to the region; (middle) the 3D graph layout has a large amount of overlap when the highlighted nodes are viewed frontally; (right) our 2.5D graph layout can keep a larger portion of the overall graph visible for the same amount of highlighted nodes.

encode a hierarchical nesting of clusters in the graph via overlaid transparent layers. Another example utilizes the third dimension to encode evolution of the node-link graph over time [3]. This example uses each layer to encode a 2D layout at a given time step; nodes are then extruded between the layers, making structures that resemble “worms.” However, these works do not utilize stereoscopy nor do they provide quantitative measures of the usability of the techniques.

All of these prior 2.5D graph visualizations explore the use of the third dimension in the context of node-link diagrams. Our technique for 2.5D graph visualization exploits stereoscopy in order to support low-level graph analysis tasks. As will be discussed in the subsequent sections, we leverage the advantages of each of these various techniques—3D node and link diagrams using stereoscopic displays, camouflage breaking, and 2.5D graphs in general—and apply them in a novel way to node-link diagrams. Additionally, given a large node-link diagram, highlighting is another crucial aspect in supporting connectivity and adjacency tasks.

Much work has been done investigating highlighting methods for 2D node-link diagrams. Ware et al. have shown that highlighting is crucial for identifying adjacency tasks in node-link diagrams [28]. They also showed that motion cues can be used as effectively as static visual cues, such as color, for highlighting sub-graphs in medium-sized node-link diagrams. They noted that the combination of color and motion is more effective than the single highlighting mode [28]. Our study builds upon the result that combinations of visual features are more effective for highlighting by exploring the use of stereographic highlighting in conjunction with color.

Previous work by Munzner et al. [20] investigates the use of multiple static perceptual channels to highlight sub-graphs representing query results. These query results are positioned within 2D boxes and linked together with minimal edge crossings. Besides the obvious extension into three dimensions, our prototype application has a more extensive user interaction component, which allows users to manipulate the results as desired. That is, rather than placing the sub-graphs within 2D boxes, we place them into user-defined planes that can be extruded into 3D. This will be described in Section 5.

3 STEREO AUGMENTED 2D GRAPH VISUALIZATION

As discussed in the preceding section, 3D benefits node-link diagrams by segregating edges that would otherwise overlap on a 2D layout. However, a 3D graph layout complicates the task of building a mental map of the graph. We utilize stereoscopic depth to highlight regions of interest in a 2D graph by bringing these parts closer to the viewpoint of the user. Instead of using distinct geometries or colors to highlight multiple groups, depth variation is used as a means of differentiating between groups. The primary advantage of *stereoscopic highlighting* is the ability to bring relevant portions of the graph to the foreground without resorting to other highlighting

techniques such as color or motion. These other techniques then remain available to encode other data attributes.

A significant motivation for the *stereoscopic highlighting* technique is that it naturally lends itself to achieving focus+context views. In a standard 2D representation, when a large graph is in view, it is very hard to see details about an individual node. Zooming in to a portion of the graph is required to acquire detail. However, the global view is then lost (Figure 3, left). Similarly in a 3D graph layout, rotation and zooming in to the region of interest could result in the remainder of the graph going out of view (Figure 3, middle). However, when stereoscopic highlighting is used, regions of interest are projected onto a plane closer to the viewer, and hence appear magnified, while the rest of the graph is kept in the background at a further depth revealing less detail. In Figure 3, right, the highlighted node and its immediate neighbors are moved onto a plane closer to the viewer. Thus, for a particular node of interest, global view and details are overlaid in a single image. Moreover, this highlighting method lends itself to effective interaction metaphors.

By layering different portions of the graph at specific depth planes, a user can turn a 2D graph layout into a 3D visualization where depth is utilized in a discrete fashion. By doing so, the user can place any nodes that need to be analysed together at the same depth layer. These layers can be re-ordered interactively to vary focus+context during analysis. Stereoscopy also has the added benefit of allowing for visual search at different depth planes in parallel [21].

Nakayama and Silverman [21] have shown that stereoscopic disparity allows for parallel conjunctive search. If one of the visual features in a conjunctive search task is stereoscopic disparity, a second visual feature of either color or motion can be searched in parallel. Their experiments reveal that even in the presence of similar color and motion features on multiple planes, humans are still able to isolate a single depth plane while carrying out visual search tasks.

We implemented a prototype interactive graph visualization framework utilizing the novel *stereoscopic highlighting* technique. The prototype is a straightforward extension of the software we developed for our controlled experiment. The user starts with a standard 2D layout, but then, through a process of interactively extruding sub-graphs onto different 3D depth planes, creates a discrete 2.5D layout. The user is able to define both the number of planes as well as the number of nodes associated with each of the planes. Each plane can furthermore function as the resulting collection of a particular visual query. These collections can include, for example, something simple, such as all immediate neighbors of a selected node, or a more complex result, such as all nodes along the shortest path between two selected nodes. Giving users the ability to arrange these discrete planes in effect provides them the ability to create coordinated multiple views among sub-graphs of the larger data. We describe the empirical study in the next section and then provide further details of the prototype in Section 5.

4 EVALUATION

The purpose of our study is to evaluate the different approaches to utilizing stereo displays for representing node-link diagrams. Our goal is twofold: 1) To investigate the viability of stereoscopic depth as an on-demand highlighting method that could replace or augment highlighting with static visual cues in 2D graph layouts; and 2) Assuming stereoscopic capabilities, investigate how 3D layouts that take advantage of stereoscopy in the conventional way (as a depth cue) compare to 2D layouts with or without our new technique. In the case of 3D layouts we exclude the use of movement and rotation—in as fair a fashion as we could devise—in order to isolate the benefit of stereoscopic depth alone.

4.1 Experiment Design and Implementation

The visualization software for the experiment is implemented with LuaAV, a real-time audiovisual scripting environment [17]. The initial 2D and 3D graph layouts were calculated using the Fruchterman-Reingold heuristic [11]. The current implementation supports both 2D point and line rendering and 3D shaded sphere and tube rendering.

We designed a within-subjects experiment of four different visualization techniques with three data set sizes and five tasks (3x4x5). The four visualization conditions are:

1) *2D graph layout + static visual cues (2D-C)*: The graph is laid out in 2D (wherein the graph layout algorithm calculates distances for x, y coordinates only) and rendered in 2D with points and lines. A highlighted node is shown in yellow. When a node is highlighted, its immediate neighbors are also highlighted with the same color circles around them (Figure 2). When two nodes along with their immediate neighbors are highlighted, we use circles in two different colors as a differentiator. Nodes that neighbor both highlighted nodes are shown with two concentric circles of both colors.

2) *3D graph layout + static visual highlighting (3D-SC)*: The graph is laid out in 3D (the graph layout algorithm calculates distances for x, y, z coordinates, taking depth into account) and rendered using shaded tubes and spheres with stereoscopy. Highlighting is done with colored concentric circles in the same way as in (1). We used view-aligned circles to make sure that the view angle does not distort their appearance.

3) *2D graph layout + stereoscopic highlighting (2.5D-S)*: The graph is laid out in 2D on a single plane, and then highlighted nodes are pulled out on a different plane closer to the viewer. All neighbors of the node are also brought to the same depth as the highlighted node. If two nodes are highlighted, one is brought to the next closest plane to the viewpoint. Nodes shared between the two highlighted nodes are rendered on a plane in between the two highlighted nodes (Figure 2d). The graph is rendered with 3D geometries of shaded tubes and spheres with stereoscopy.

4) *2D graph layout + static visual cues + stereoscopic highlighting (2.5D-SC)*: This condition is the same as above, except we add the colored concentric circles for highlighting as described in (1) and (2) (See Figure 4).

4.1.1 Study Design Decisions

Our experiment focuses on investigating the viability of stereoscopy as a highlighting method and compares this particular use of it to the conventional use of stereoscopy as a depth cue (in the case of 3D graph layouts). Thus we controlled for other types of depth cues, e.g. motion parallax and kinetic depth, by excluding any interaction and animation, including rotation. The 3D graph layouts were at a disadvantage without rotation since certain nodes could be occluded from a static viewpoint. In order to overcome this, we optimized the 3D graph layout to prevent occlusion (for the static viewpoints that were used in the experiment). To do so, after calculating the layout in 3D, we projected the nodes' positions to the screen coordinate and calculated a second pass of repulsive forces between every other node. We pre-computed the graph layout for given viewpoints and

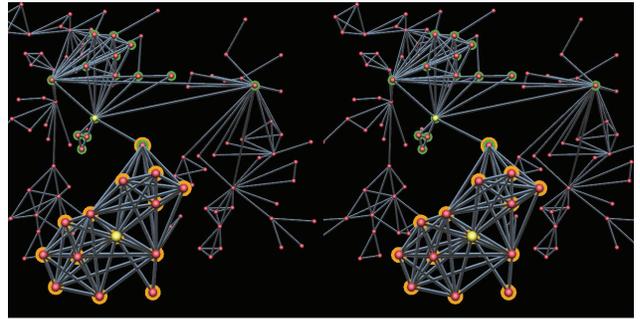


Fig. 4. A stereo pair image illustrating the 2.5D-SC condition where colored concentric circles and stereographic highlighting used simultaneously. By crossing eyes, the reader can perceive that the green nodes appear closer to the viewer than the background and orange nodes are even closer.

the modified graph algorithm eliminates occlusions for these specific viewpoints.

Note that with the exception of the first condition the size of the nodes and edges varied with distance to the viewpoint. This scaling with distance provided an extra depth cue for these cases. We preferred not to render the first condition with 3D geometries because in the 2D case edges rendered as tubes occupy significantly more screen space with no added benefit. We also displayed the label for each node directly on top of it. Labels were always rendered in the same pixel size in order to eliminate legibility differences.

4.2 Data Set and Tasks

We used 2004 InfoVis contest co-authorship data as the basis of our controlled experiment [10]. We identified three different sizes of connected components within the data. The smallest graph consisted of 90 nodes and 200 edges, while the medium size graph consisted of 120 nodes and 250 edges, and the largest graph consisted of 150 nodes and 350 edges (See Figure 5). We chose not to vary data size drastically because we did not intend to investigate how each technique scaled with different graph sizes. For smaller graphs this would not matter much but for the larger graphs it would require zooming into different regions of the data, thereby adding an extra dimension to our controlled study. We simplified the labeling of each node by using the authors' initials only in order to reduce clutter. We told our participants what the labels represent.

Since this was a preliminary study assessing the potential use of the stereoscopic highlighting technique, our goal was to focus on low-level topology based tasks that are common to the analysis of graph data [16]. Below we describe the four adjacency and common connection tasks used in our study. We believe that each of these tasks could possibly benefit from stereoscopic highlighting.

Task 1 (T1, adjacency): Find whether an element exists among highlighted nodes: *Is 'DK' a friend of highlighted yellow node 'CP'?* Note that all friends of DK are also highlighted either using same color concentric circles and/or by placement at a nearer depth plane. In order to ensure that the participant browsed all highlighted nodes carefully we intentionally omitted the target element 50% of the time.

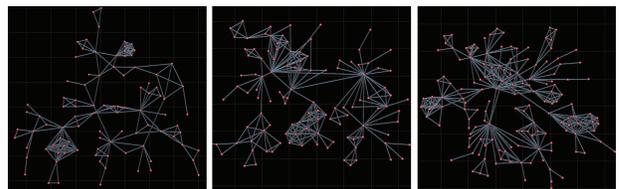


Fig. 5. 2D graph layouts of the data sets used in our experiment, in order of size.

Task 2 (T2, adjacency): Count all highlighted nodes: *How many friends does the highlighted yellow node 'SK' have?* This task requires the participant browse all the nodes and also to make sure that they count every element only once.

Task 3 (T3, adjacency, group intersection): Identify whether a node is a friend of both highlighted group of nodes: *Is 'SC' is a friend of highlighted yellow nodes 'BS' and 'MC'?* For this task, the participant had to browse for 'SC' among nodes that are highlighted with two concentric circles, and/or search the depth plane in between 'BS' and 'MC'.

Task 4 (T4, adjacency, group difference): Identify whether a node is a friend with one of the highlighted nodes but not with the other: *Is 'JP' a friend of 'RR' but not a friend of 'CK'?* This task required the participant to identify which color circles or depth indicate friendship with 'RR', and then browse for 'CK' among the nodes at that same depth or color, ignoring other highlighted nodes.

Task 5 (T5, common connection): Identify a node that has friends in both highlighted set of nodes: Friends of 'CP' and 'BB' are highlighted as two groups. Find a node outside of both communities but that has friends in both groups. For this task, we instructed the participants to identify critical elements in both groups, to identify nodes with many connections, to find nodes that are spatially closer to the cluster of the other highlighted group, and to search for a node among their friends.

4.3 Hypothesis

Based on our assessment of the various graph visualization techniques evaluated in this paper, we expected the following three main results: 1) For all tasks, 2.5D-SC would outperform 2D layouts that use only a single highlighting mode (i.e. 2D-C, 2.5D-S); 2) 2D layouts utilizing static visual highlighting (2D-C and 2.5D-SC) would be faster than 2.5D layout without static visual highlighting (2.5D-S). However, they wouldn't necessarily be more accurate; 3) When performing tasks where users are asked to search for an answer among nodes that are not highlighted, we expected 3D layouts with static visual highlighting to outperform all other techniques (3D-SC). By utilizing depth continuously 3D layouts segregate all nodes, not only the highlighted nodes, and thus it would seem justified to assume that this technique would outperform the others.

4.4 Participants and Apparatus

We recruited 17 (5 female) participants from an academic institution. The age of our participants ranged from 20 to 35, with a mean age of 28 years. Each participant completed 60 questions in about 45 minutes (excluding 10 minutes required for initial training).

Before the experiment, we asked participants how familiar they were with node and link representations. Only one participant indicated that he had used them professionally. 9 participants indicated that they had seen node-link diagrams previously, and knew what they were, but had never used them for anything. 7 participants had no experience whatsoever with node-link visualizations.

We ran the study in the Allosphere Virtual Reality environment [1][14] featuring a large spherical display with a 5m radius. For the experiment we used two Barco Galaxy NW-7 projectors providing a seamlessly stitched image of 3150 x 1050 pixels projected onto an approximately 9m x 3.5m patch on the sphere. The image is warp corrected for the curvature of the screen. The viewer stands in the centre on a platform with an approximate distance of 5m to each point on the screen. Stereo viewing is enabled via active stereo using StereoGraphics CrystalEyes 3D LCD shutter glasses. The projectors ran at 120Hz, providing a 60Hz image for each eye, synchronized with shutter glasses through an infrared transmitter.



Fig. 6. Summary views of the experimental setup in the Allosphere. The image on the bottom left shows what the participant sees on the stereoscopic display. The image on the bottom right shows the questionnaire displayed on the laptop screen.

4.5 Procedure

Each participant was made to stand at the centre of the sphere with a laptop computer placed in front of him or her at an approximate 45° angle. (This was to ensure that the laptop screen did not obstruct the stereo display.) The participants viewed the questions and provided their response via a mouse click on the laptop (See Figure 6). They were first shown a question with no data displayed and were instructed to click the 'view' button on the laptop screen once they fully understood the question. When the 'view' button was clicked, the corresponding visualization was displayed on the stereo display. The user then clicked on an answer and then clicked the 'next' button to confirm the answer. After clicking 'next', the participant was shown the next question on the laptop while the stereo display turned blank. This was done to ensure accurate measurement of the time taken for each question. We measured the time elapsed only between when the participant clicked on the 'view' and the 'next' button (after confirming his or her answer). We instructed them to take breaks as long as they wished when the stereo display turned blank.

Before the experiment we educated each participant about the data set, the types of visualizations and the tasks for about 10 minutes. We verified whether they were able to see stereoscopic depth by rendering identical pixel size points at different depths and asked them to verify how many different depth planes were visible to them. Only one participant failed, and we excluded him from the study.

We kept the order of the tasks the same for all users, advancing from easier to harder tasks. Participants answered 12 unique questions (the 3 data sets times the 4 visualization techniques) for each task (60 questions overall). We trained the participants separately for each task before they began answering the set of questions for that task. We instructed them about what strategies to utilize for that particular task in each visualization condition. We also kept the order of the data sets the same for all users, advancing from smaller to larger. In order to avoid learning effects regarding the dataset, we either rotated the graph (in the 2D cases) or changed the view angle (in the 3D cases). We also counterbalanced the order of visualization techniques across users so that each user answered the questions in a unique order, providing a different permutation of

the four visualization techniques. Finally, participants filled out satisfaction surveys after completing all tasks.

4.6 Results

4.6.1 Accuracy

On the average, 93.75% of the answers were accurate (SD=3.6) for all 16 participants. We analyzed the accuracy of our data using a within-subject design for Repeated Measures Analysis of Variance (RM-ANOVA). Our analysis showed a significant effect of accuracy for task, $F(4,12)=9.012$, $p=.001$. We did not observe any effect of accuracy for visualization technique or dataset. Post-hoc comparisons revealed significant effect of accuracy on interactions of visualization technique x task, $p=.002$.

Post-hoc comparisons show significant effects of accuracy between visualization techniques only for task T2 (counting). For this task 2.5D-SC was significantly more accurate than 2D-C and 3D-SC ($p=.009$ and $p=.002$ respectively). Although, ANOVA did not yield any significance for T5 (common connection), we observed that for this task using 3D-SC participants committed 50% less errors compared to 2.5D-S and 2D-C. A binomial test showed that 3D-SC was significantly more accurate than 2.5D-S and 2D-C (with $p = .01$ and $p=.04$ respectively). No significance was observed between 3D-SC and 2.5D-SC for this task.

4.6.2 Time

We analyzed completion times using the same within-subject design with RM-ANOVA. On average, each question took 15.64 seconds (SD=10.14) to complete.

RM-ANOVA showed a significant effect of time for task $F(4,12)=8.311$, $p<.02$. Unsurprisingly, participants required more time to answer for more complex tasks. Figure 7, top, summarizes

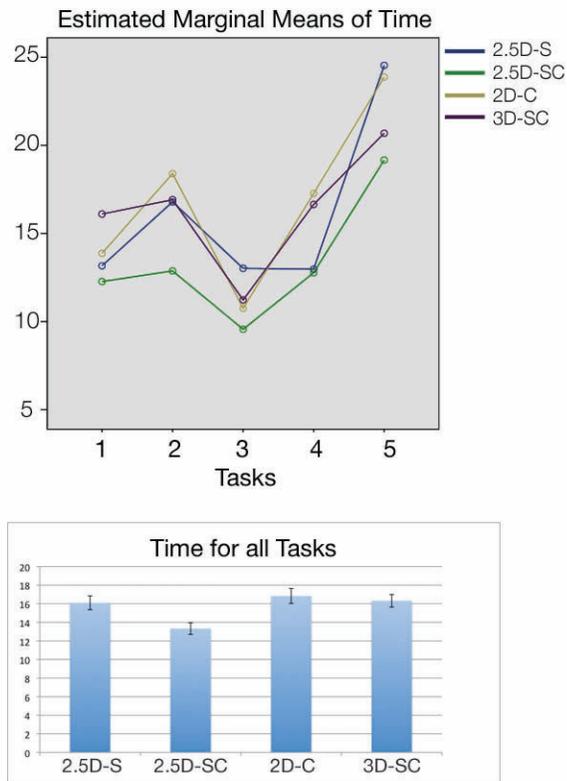


Fig. 7. (top) Plot of estimated marginal means of time (in seconds) for each visualization technique per task. (bottom) Average completion times (in seconds) for each visualization technique for all tasks. Error bars show +/- SE.

the estimated marginal means of time for each task per visualization technique. We also found a significant effect of time for visualization technique ($F(3,13)=10.46$, $p=.001$), as shown in Figure 7, bottom. We found that 2.5D-SC was about 15% faster than the average. Participants took an average of 13.3 seconds (SD= 8.7) with this technique.

RM-ANOVA also showed an interaction between visualization technique by task $F(12,4)=6.173$, $p<.04$. Post-hoc comparisons revealed no significant difference between visualization techniques for tasks T3 (group intersection) and T5 (common connection). However, they revealed a significant difference for T1, T2 (accessibility tasks) and T4 (group difference). Figure 8, top, illustrates the completion times for T1 and T2 per visualization technique. For T1, 2.5D-SC was significantly faster than 3D-SC, with $p=.002$. For T2, 2.5D-SC was significantly faster than 2.5D-S, 2D-C and 3D-SC with $p=.037$, $p=.004$, and $p=0.01$, respectively. For T4 both 2.5D-S and 2.5D-SC are significantly faster than both 2D-C ($p=.03$, $p=.025$ respectively) and 3D-SC ($p=.035$ and $p=.013$). Although we did not observe any statistical significance across visualization techniques for T5, Figure 8, bottom, illustrates a trend where 2.5D-SC is faster than 2D layouts with single highlighting method (2D-C and 2.5D-S), closely followed by 3D-SC.

4.6.3 Subjective User Ratings

At the end of the experiment we asked the participants to rate static visual cues (colored circles) and *stereoscopic highlighting* (depth) affecting their time and accuracy performance, using a Likert scale from 1 to 5. Participants rated depth and color very closely with a mean of 3.75 (SD=1.0) for the former, and 3.62 (SD=1.2) for the latter. 5 participants rated depth as more helpful, whereas 6 rated color as more helpful for the given tasks. Remaining subjects rated both equally. We also made them rate whether they performed faster or more accurately when both highlighting techniques were used.

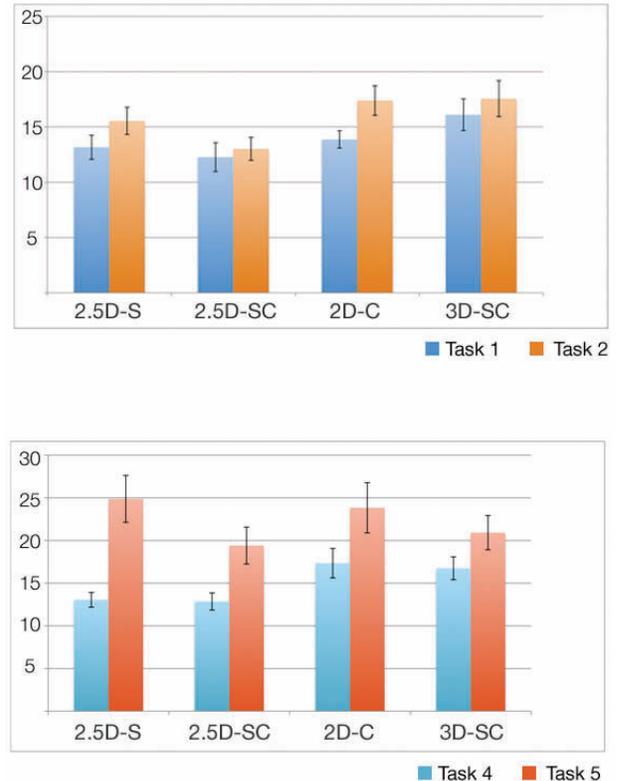


Fig. 8. (top) The chart of average completion times for tasks 1 & 2, (bottom) for tasks 4 & 5 per visualization technique (both measured in seconds). Error bars show +/- SE.

The subjects indicated that when both highlighting techniques were used they thought to be faster with a mean score of 4.81 (SD= 0.54), and more accurate with a mean score of 4.44 (SD=0.89).

Finally, we asked them to indicate their overall preference for all the tasks. All 16 participants picked 2.5D-SC. Two participants specifically indicated that they preferred 2.5D-SC overall but that 3D-SC was more useful for the last task.

4.6.4 Discussion

Our experiment revealed that *stereoscopic highlighting* augments 2D graph visualizations significantly for three of the five types of tasks we evaluated (T1, T2, T4). In fact, participants performed significantly better with 2.5D-SC for these three tasks. This strongly indicates that our first hypothesis—that 2.5D-SC would outperform the other 2D layouts that use only a single highlighting mode—is a valid assumption. Moreover, 2.5D-SC outperformed 3D-SC for these tasks.

For all the tasks, 2D-C was neither significantly more accurate nor faster than 2.5D-S; hence we could not verify the second hypothesis. We take this as a strong incentive to pursue further aspects of *stereoscopic highlighting* for a wider range of tasks.

We did not find any significant difference across techniques for T3, neither for accuracy nor time. This suggests that the questions for this task were too simple, since all participants performed it equally well using all techniques. Indeed, none of the questions for this task required the user to browse more than four nodes. In the future, using larger data sets may expose differences across the techniques.

For T4, 2.5D-S (only stereoscopic depth) performed almost as fast as 2.5D-SC, and significantly faster than 2D-C and 3D-SC. This result suggests that for this task, depth was a stronger cue than color. In the cases where multiple categories of highlighted nodes were

present, we found that depth cues became increasingly useful. In T4, nodes that are at the same depth appear spatially closer. We know from Gestalt theory that spatial proximity is a stronger association cue than color. As more categories are present, this becomes an increasingly significant advantage.

The fact that 3D layouts are more accurate than other layout techniques for T5 supports our final hypothesis. Although time performance for this task did not reveal any significant differences, Figure 8, bottom, suggests a trend where 3D-SC and 2.5D-SC are both faster than 2D layouts that utilize only a single highlighting method.

We will discuss the broader implications of these findings in the last section.

5 STEREO AUGMENTED 2D GRAPH VISUALIZATION: AN INTERACTION FRAMEWORK

In the previous section we have shown the benefits of stereoscopic highlighting for a variety of tasks. During the study, users carried out tasks given a graph, partitioned into predefined sub-graphs at fixed depth planes. In this section we are proposing the use of the *stereoscopic highlighting* technique within a novel interactive framework. We developed a prototype application directly based upon the results of the empirical studies. The application lets the user divide the graph into discrete layers, or *planes*, each positioned at distinct depths. The user interactively determines both the number of planes as well as what nodes are to be attached to each of the planes.

The main interaction affordance is the ability to extrude portions of the graph on to separate planes. To do so, a user first creates a plane by clicking a button on the GUI or keyboard. A new plane is immediately created in front of the existing ones. The user

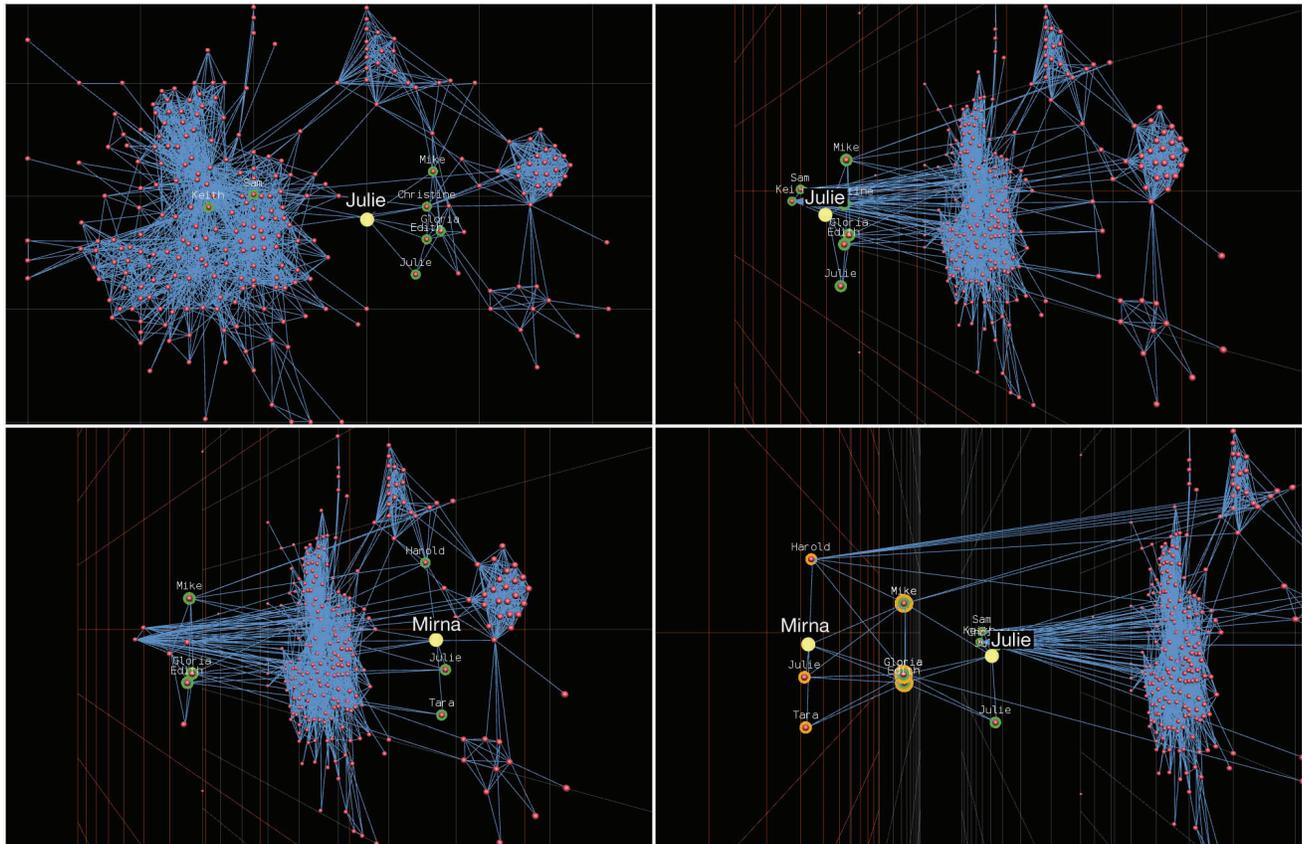


Fig. 9. Screenshots from the interactive prototype: (top, left) The starting 2D layout with a highlighted node; (top, right) The highlighted node, Julie, moved onto a different plane with her friends; (bottom, left) Another node, Mirna, is positioned on the selected base plane, and highlighting Mirna’s friends instantly reveals friends of Mirna and Julie on the first plane, (bottom, right) Moving Mirna’s friends onto a separate plane leaves Julie’s friends (who are not friends with Mirna) on the second plane. N.B., we rotated the view angle to make planes visible; however with stereoscopic cues it is possible to segregate nodes on different planes with a frontal view.

then highlights a node by clicking, and then drags the node towards the plane using the cursor. When released, the dragged node snaps to the layer that it is closest to. By clicking GUI or keyboard buttons, the user can bring first and/or second degree friends of the highlighted node to the same layer. This way it is relatively easy to populate elements on a newly created plane.

A similar interaction method can be applied to highlighting shortest path between two selected nodes. The user can drag multiple nodes on a plane, and then highlight both of them at the same time (with ctrl+click), and press a GUI button to highlight the shortest path between these two, which brings all the nodes on the shortest path to the same plane.

Following the steps above, the user can turn a 2D graph into a discrete 3D layout where nodes that need to be viewed together reside at the same depth. It is also possible to rearrange the planes interactively. The user can toggle cursor behavior between manipulating the position of a selected node or plane. The plane serves as a proxy for moving multiple elements at the same time, and thus preserves the spatial organization among the nodes that it contains.

This interaction framework lends itself to performing complex visual queries. For instance, it is possible to display the union of two visual queries by merging two planes in a single one. We illustrated the applicability of these visual queries on a Facebook network graph depicting all friends of one user. The data consisted of 267 nodes and 2041 edges between them. (We renamed the nodes in order to anonymize the dataset.) Figure 9, top left, shows the overall graph structure with a selected friend, Julie, and all her immediate neighbors. The user then carries the highlighted nodes on a separate plane (See Figure 9, top right). In Figure 9, bottom left, the user clicks on a different node, Mirna, on the base plane and highlights her immediate neighbors. This action clearly reveals that Mirna and Julie have two friends in common in this particular set. In Figure 9, bottom right, Mirna and his friends are carried onto a third plane, thus leaving on the second plane only the set of people who are friends with Julie but not with Mirna.

The user can then move the plane that holds Mirna's friends and place it anywhere in the 3D space. Multiple planes can be repositioned in the 3D space for optimal viewing. This creates coordinated multiple views among 2D sub-graphs with a visual style similar to *VisLink* [6]. Although this system is only in the form of a prototype, we believe that it points to promising applications using our *stereographic highlighting* technique.

6 CONCLUSION AND FUTURE WORK

The results of our experimental study validate the use of our novel *stereoscopic highlighting* technique in a wide range of graph visualization tasks. Our findings are further validated by the fact that users consistently preferred our technique to others. Moreover, our prototype illustrates the applicability of the technique to real-world use cases.

Our results do seem to contradict some findings in previous research [30] which have shown that 3D graph layouts more effective than 2D layouts for identifying adjacency tasks. We believe that this is due to differences in our evaluation methods, including the fact that we looked at much smaller data sets for a different range of tasks while using static visual highlighting for all techniques. It will be interesting to investigate how our technique scales with larger data sets when compared to 3D layouts. The results in this paper point out positive effects of stereographic highlighting, while keeping other depth cues out of consideration to the extent possible. In follow-up work, we will investigate the effect of stereographic highlighting in presence of other depth cues, such as motion parallax and kinetic depth.

One of the important details of our study that we did not investigate fully was the labeling of the nodes. As described above, we fixed the label size in order to reduce clutter and control legibility. However, this might have negatively affected the

perception of depth in the 2.5D-S, 2.5D-SC and 3D-SC cases since the label size did not scale in accordance with perspective. In a user application, scaling of the labels has to be done in a very sophisticated manner to ensure that scaling is performed in accordance with perspective while still securing legibility. We will investigate this in the future.

One important difference between *stereoscopic highlighting* and static visual cues is that the latter are nominal in character whereas the former is an ordinal property. That is, objects that are closer to the viewpoint have a greater visual emphasis, and a succinct order is defined. However, when utilized in an interactive framework, this property of *stereoscopic highlighting* can be useful in certain scenarios (as was discussed in Section 5).

A potential drawback with utilizing stereoscopic depth is the time that it takes for a viewer to get adjusted to stereo representations. We observed this in our controlled experiment, especially since we did not create head coupled stereo but instead computed stereo for a fixed location. On the other hand, color is observed instantly. Once the user is able to discern different depths clearly, however, depth has an advantage over color. Associating elements at the same depth is a type of spatial proximity. And as is known from Gestalt theory, spatial proximity is a stronger visual cue for association than is color [19]. And indeed, the results from our user study support this idea. We acknowledge that there is a certain percentage of the population that cannot correctly perceive stereo (around 8-10 percent). Given that there are similar blindness phenomena for other visual variables (such as color), we strongly believe that using different highlighting techniques in conjunction is advantageous.

We believe that this new technique, *stereoscopic highlighting*, and the associated interactive prototype implementation are very promising tools for visual graph analysis and interactive information visualization tasks. We will continue to investigate the applicability of the technique in more detail. In particular, applying *stereoscopic highlighting* to larger data sets should be a fruitful area of investigation, and we would like to apply our interactive framework with a large data set in the context of a real-world use case. The initial study leads us to believe that a wider range of interactive visual querying operations can benefit greatly from this technique. Another future area of research will investigate the option of duplicating nodes as the user interactively creates new sub-graphs on different planes.

We deployed the visualization in the Allosphere, which is a specialized VR environment. We are planning to investigate the hardware requirements for the *stereoscopic highlighting* technique to be effective on the more widely available consumer level stereoscopic displays.

Finally, in the future we will extend our evaluations to examine high-level sense making tasks, leading to a comparative evaluation of 3D and 2D graph layouts for such tasks.

ACKNOWLEDGMENTS

The authors would like to acknowledge the help of Bahar Koymen with some of the statistical analysis and Anand Ramaswamy for help with editing the manuscript. This work was partially supported by NSF grants IIS-1047678 and IIS-1058132, ONR grant N00014-09-1-1113, as well as an ARO MURI award for proposal #56142-CS-MUR.

REFERENCES

- [1] Allosphere. <http://www.allosphere.ucsb.edu/>, 2011.
- [2] U. Brandes and S.R. Corman, "Visual unrolling of network evolution and the analysis of dynamic discourse," *Information Visualization*, vol. 2:1, pp. 40-50, Palgrave Macmillan, 2003.
- [3] U. Brandes, T. Dwyer and F. Schreiber, "Visualizing related metabolic pathways in two and a half dimensions," *Graph Drawing*, pp. 111-122, Springer, 2004.

- [4] A. Cockburn and B. McKenzie. "An evaluation of cone trees," *PEOPLE AND COMPUTERS*, pp. 425-436, 2000.
- [5] A. Cockburn, A. and B. McKenzie, "Evaluating the effectiveness of spatial memory in 2D and 3D physical and virtual environments," *Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves*, ACM Press, pp. 203-210, 2002.
- [6] C. Collins and S. Carpendale, "VisLink: Revealing relationships amongst visualizations," *IEEE Transactions on Visualization and Computer Graphics*, pp. 1192-1199, 2007.
- [7] J. E. Cutting and P. M. Vishton, "Perceiving Layout and Knowing Distances: The Integration, Relative Potency, and Contextual Use of Different Information about Depth," *Handbook of Perception and Cognition*, edited by W. Epstein and S. Rogers, vol. 5, pp. 69-117, Academic Press, San Diego, CA, 1995.
- [8] P. Eades and Q.W. Feng, "Multilevel visualization of clustered graphs," *Graph Drawing*, pp. 101-112, Springer, 1997.
- [9] K.M. Fairchild, S.E. Poltrock, and G.W. Furnas, "Semnet: three-dimensional graphic representations of large knowledge bases," In *Cognitive science and its applications for human-computer interaction*, R. Guindon, Lawrence Erlbaum, 1988.
- [10] J.-D. Fekete, G. Grinstein and C. Plaisant, IEEE InfoVis 2004 Contest, the history of InfoVis, www.cs.umd.edu/hcil/iv04contest, 2004.
- [11] T.M.J. Fruchterman and E.M. Reingold, "Graph drawing by force-directed placement," *Practice and Experience* vol. 21:11, pp. 1129-1164, 1991.
- [12] T. Grossman and R. Balakrishnan, "An evaluation of depth perception on volumetric displays," Proceedings of the working conference on Advanced visual interfaces, pp. 193-200, ACM Press, 2006.
- [13] M. B. Holbrook, "Breaking camouflage: stereography as the cure for confusion, clutter, crowding, and complexity - three-dimensional photography," *Photographic Society of America Journal*, vol. 8, 1998.
- [14] T. Hollerer, J.A. Kuchera-Morin and X. Amatriain, "The allosphere: a large-scale immersive surround-view instrument," *Proceedings of the 2007 workshop on Emerging displays technologies: images and beyond: the future of displays and interaction*, ACM Press, p.3, 2007.
- [15] A. Katifori, C. Halatsis, G. Lepouras, C. Vassilakis, and E. Giannopoulou, E., "Ontology visualization methods—a survey," *ACM Computing Surveys (CSUR)*, vol. 39:4, p.10, ACM Press, 2007.
- [16] B. Lee, C. Plaisant, C.S. Parr, J.D. Fekete, and N. Henry, "Task taxonomy for graph visualization," *Proceedings of the 2006 AVI workshop on BEyond time and errors: novel evaluation methods for information visualization*, pp. 1-5, ACM Press, 2006.
- [17] LuaAV. <http://lua-av.mat.ucsb.edu/blog/>, 2011.
- [18] J.D. Mackinlay, G.G. Robertson and S.K. Card, "The perspective wall: Detail and context smoothly integrated," *Proc. SIGCHI conference on Human factors in computing systems: Reaching through technology*, pp. 173-176, ACM Press, 1991.
- [19] D. Marr, *Vision: A computational investigation into the human representation and processing of visual information*, Freeman, 1982.
- [20] T. Munzner, F. Guimbretiere, and G. Robertson, "Constellation: a visualization tool for linguistic queries from MindNet," In *Proceedings of the 199 IEEE Symposium on Information Visualization, (Info Vis'99)*, pp. 132-135, 1999.
- [21] K. Nakayama and G.H. Silverman, "Serial and parallel processing of visual feature conjunctions," *Nature*, vol. 320:6059, pp. 264-265, 1986.
- [22] S.D. Peterson, M. Axholt and S.R. Ellis, "Managing visual clutter: A generalized technique for label segregation using stereoscopic disparity," *Virtual Reality Conference, VR'08 IEEE*, pp. 169-176, 2008.
- [23] S.D. Peterson, M. Axholt and S.R. Ellis, "Objective and subjective assessment of stereoscopically separated labels in augmented reality," *Computers & Graphics*, vol. 33:1, pp. 23-33, Elsevier, 2009.
- [24] K. Risdén, M.P. Czerwinski, T. Munzner, D.B. Cook, "An initial examination of ease of use for 2D and 3D information visualizations of web content," *International Journal of Human-Computer Studies*, vol. 53:5, pp. 695-714, Elsevier, 2000.
- [25] G. Robertson, M. Czerwinski, K. Larson, D.C. Robbins, D. Thiel, and M. Van Dantzich, "Data mountain: using spatial memory for document management," *Proc. User interface software and technology*, p. 162, ACM Press, 1998.
- [26] G.G. Robertson, J.D. Mackinlay and S.K. Card, "Cone trees: animated 3D visualizations of hierarchical information," *Proc. SIGCHI conference on Human factors in computing systems: Reaching through technology*, pp. 189-194, ACM Press, 1991.
- [27] C. Ware, *Information visualization: perception for design*, Morgan Kaufmann, 2004.
- [28] C. Ware and R. Bobrow, "Supporting visual queries on medium-sized node-link diagrams," *Information Visualization*, vol. 4, pp. 49-58, 2005.
- [29] C. Ware and G. Franck, "Evaluating stereo and motion cues for visualizing information nets in three dimensions," *ACM Transactions on Graphics (TOG)*, vol. 15, pp. 121-140, 1996.
- [30] C. Ware and P. Mitchell, "Visualizing graphs in three dimensions," *ACM Transactions on Applied Perception (TAP)* vol.5, p.2, 2008.