

Augmented Textual Data Viewing in 3D Visualizations Using Tablets

Charles Roberts*

Basak Alper†

JoAnn-Kuchera Morin‡

Tobias Höllerer§

Media Arts and Technology Program
UC Santa Barbara

ABSTRACT

Many data sets contain structural 3D components along with associated textual data. While structural data is often best visualized on large stereoscopic displays, such displays can pose problems presenting accompanying textual information. Chief among these problems is a tradeoff between size dependent legibility of text and the occlusion of structural data if text is presented at larger sizes. Our solution to this problem integrates structural data shown on a large display while users select features of the structure and view the associated textual data on personal tablets.

Our solution also lends itself to collaborative browsing tasks. In our initial implementation each user can individually select structural components and view the associated text on their own tablet; thus, everyone becomes an active participant in data mining. When individual users find textual data they deem of interest to the group they can share it with collaborators by temporarily pushing the text box to the large display. The current status and each individual's browsing history are shown on the large display in order to provide awareness of other team members' actions.

Index Terms: H.5.2 [Information Interfaces And Presentation]: User Interfaces—Input devices and strategies ; H.5.3 [Information Interfaces And Presentation]: Group and Organization Interfaces—Computer-supported cooperative work;

1 INTRODUCTION

Researchers often need to analyze data sets involving 3D structural models along with textual data associated with specific parts of the model. Examples include, but are by no means limited to, protein folding structures accompanied by statistical analysis results, 3D MRI models accompanied by medical data pertaining to specific features, and 3D social network data where each node is associated with user specific information. The textual information can be another dimension in the data or metadata in the form of annotations or higher level statistical summaries.

3D structural data is better visualized on large, stereoscopic displays where scale and a higher number of depth cues aid users' understanding of the model's spatial structure. Such displays also create a greater sense of immersion and realism [16] [17]. However, these displays are not necessarily optimal for visualizing text. Rendering text in 3D space poses problems related to scaling and display orientation [13] [10]. Although previous work [13] has shown that view aligned text increases readability significantly, legibility issues related to text size persist. Readable text requires either very high resolution or a large rendering size. Textual information associated with different parts of the 3D model at different depths cannot be freely scaled along with the model, as linear scaling based

on depth might yield unnecessarily large text that occludes parts of the scene, or, alternatively, illegibly small text.

One potential solution to this problem involves dividing the display into two areas: one for textual information and the other for the 3D model. In this solution, text is commonly rendered at a specified area of the screen with a uniform font size. Linkage between text and its associated part of the 3D model is provided using a static visual cue such as a connecting line. There are two major drawbacks of this approach. First, a large amount of space is required for displaying lengthy text passages (Figure 1); this space is no longer available for the 3D visualization. Second, the separation of the text display area and the 3D visualization necessitates the use of external visual cues to help users map specific pieces of text to their associated location in the visualization. These cues often clutter the visualization.

In order to solve problems related to rendering 3D structural data and text together, we propose displaying textual data on hand-held tablet devices instead of the shared display. With their high screen resolutions and flexible viewing distance, tablets maintain legibility at small font sizes allowing a greater amount of text to be displayed. Constraining text to personal tablets also allows multiple users to select, view and author the information without distracting collaborators (Figure 2). Finally, software that is running on the tablet is disassociated from rendering complex 3D data; thus, we create an interaction framework between tablet devices and a shared display that is generalizable to a variety of visualizations.

The main issue to address in such an implementation is enabling contextual awareness that lets users effortlessly perceive the connections between textual data on their personal tablet display and associated parts of the 3D structural data on the shared display. Additionally, this contextual awareness must be achieved using minimal interactive visual correlations between the users tablet and the shared display in order to minimize clutter.

In this paper we present the paradigm of separating 3D data and textual data to different displays where each is viewed optimally. We provide details of the interaction model between the shared display and the tablet. We conclude by describing a specific implementation of our interaction model using 3D visualization of social network data.




Figure 1: Image illustrates how much occlusion text can create if rendered in place in 3D on the left. On the right, the same information is shown on the tablet display.

*e-mail: charlie@charlie-roberts.com

†e-mail: basakalper@umail.ucsb.edu

‡e-mail: jkm@create.ucsb.edu

§e-mail: holl@cs.ucsb.edu




Figure 2: A 3D social network visualization is shown on the shared display, featuring three different colored cursors for three users. Network nodes that are being viewed by each user share the same color with their respective cursor. Nodes that have been previously visited by a user are indicated using colored concentric circles around the nodes. The middle and right images show the tablet displays of two users. Note that relative positioning of the nodes in the bottom half of the tablet screens roughly corresponds to the relative positioning of these users on the shared display surface.

2 RELATED WORK

The idea of using a tablet / mobile display to interact with physical environments or large public displays has long been explored [19] [6] [14]. Fitzmaurice et al. [6] investigated using spatially aware hand held displays to provide access to more information by tracking the position of the display so that it shows different parts of a large workspace when physically moved. Yee [19] extended the work on spatially aware displays by incorporating pen input and various other interaction techniques. One drawback of this technique is problematic ergonomics due to the amount of movement that is needed to access information while interacting with large displays.

Rekimoto [14] uses a video see-through hand held device with a magnifying glass like approach wherein objects in a physical environment are augmented by overlaying metadata associated with the object on the see-through display. Boring et al. [3] turned a camera-equipped mobile device into a mouse-like touch device that allows users to manipulate content on displays lacking touch input or displays that are at a distance. Later, they extended their interface to accommodate multiple users by showing individual content on mobile displays that would have otherwise cluttered the large display and distracted other users [4]. This is similar to our approach of collaborative browsing, where multiple users analyze textual data on their personal display. However, we also provide the added ability for users to push text they deem of interest to the group to the large display for discussion.

Multi-display groupware systems, which often mix multiple networked personal and public displays have been explored [2] [18]. These systems allow collaborators to switch between individual tasks and tasks involving the collective group. Typically, the individual tasks are performed on a personal display (e.g., laptop displays) while the collective workspace is a shared display (e.g. a wall display or a tabletop) [2]. A common approach is to replicate the visualizations on both the shared and personal displays. This approach is costly and it is often not possible for tablet devices to replicate computationally intensive 3D renderings. This also breaks a useful abstraction between the personal displays and the data being visualized that allows software running on the personal display to be generic.

Another role of the public workspace is to help individuals maintain a sense of the whereabouts and activities of collaborators. There is a conflict between supporting individual tasks and supporting coordination, communication and situational awareness in the group workspace. Studies have shown that users perform individual tasks better on personal displays [9]. Although having a personal workspace is desirable when users need to focus on a specific task,

it deteriorates shared workspace awareness as users do not make frequent use of the shared display where information about collaborators' actions is shown.

Our system encourages the use of shared display as the visualization is not replicated on personal devices. Reading text, which has a high cognitive load, is done on the personal device while the 3D structural data is only available on the shared display. The user also needs to use the shared display to access textual data.

3 OVERVIEW

3.1 Goals and Contributions

We project structural data onto a shared display while allowing individual users to selectively view associated textual data on tablets. In separating the display of text from the display of structural data we improve user experience in a number of ways:

Use the right tool for the right job. By directing resources towards a single task, we provide affordances to both the text viewer and the structural visualization that would be difficult to achieve in an integrated application. The difficulties of creating legible text overlaying 3D scenes are well documented [13] [10]. Having a dedicated reading interface improves text legibility while also enabling the comparison of textual data from different nodes without causing occlusion of the visualization. Tablets have been shown to be better reading interfaces than mounted horizontal or vertical displays [12].

Promote collaboration. Our approach fosters increased situational awareness in collaborative data mining. Since structural data is not presented on personal displays, users will regularly reference the shared display. In addition to structural data the shared display also presents information on other users' past and present activity. This helps provide individuals better awareness of their collaborators actions and allows them to focus (if they choose) on regions of interest in the structural data that have not yet been explored.

Unencumbered interaction. By specifying that the text viewer will not perform costly 3D rendering tasks, we ensure that it is lightweight and able to run on tablets instead of laptops. Tablets provide better mobility and ergonomic comfort when compared to laptops. This improved mobility is especially important in large immersive environments that enable users to move about freely.

3.2 Interaction Model

We propose an interaction model that can be generalized to a broad range of 3D visualizations. This model is defined as follows:

1. Each tablet handshakes with the server running the visualization engine. The server then sends a unique ID number to the




Figure 3: Tablet device interface consists of a top area where textual data is displayed and a bottom area where users control a cursor using touch. Aliases of the selected nodes are shown in the bottom area.

tablet; this ID number will be included in all future messages that the tablet sends to the server.

2. Each tablet device controls a cursor on the shared display using 2D coordinates derived from single touch input. Cursor positioning uses absolute position instead of relative incremental movement, thus providing a direct mapping between the control area on the tablet and the large shared display.

3. A simple raycasting technique is used to perform hit testing on the three dimensional structure using the cursor's position. When a region of interest is hit, a rollover message is sent to the tablet controlling the cursor. This message provides both the location (in screen coordinates) and identification of the targeted area.

4. When tablets receive a rollover message, they use the reported location to place an alias of the rolled-over item on the control area of the tablet. Relative positions of the rolled-over nodes on the tablet correlate to their positions on the shared display surface (Figure 2).

5. Users can select items that appear on their tablet by double tapping in the control area. Upon selection, tablets send a notification to the visualization giving the ID of the selected structural element. The visualization software responds by sending the text associated with the element to the tablet to be displayed in a vertical text column.

6. Upon user request, the text associated with a particular structural element is presented in the shared visualization. This lets users push interesting textual information to the large display for everyone to view simultaneously and encourages group discussion.

7. The visualization provides indicators showing the history and current actions of users by visually encoding viewed or previously viewed areas of the visualization.

4 TEST BED IMPLEMENTATION

In our test bed implementation tablets (either Android or iOS devices) display text associated with a 3D node-link diagram of social network data. The data visualized is the 2004 InfoVis contest co-authorship data [5] where each node represents an author and edges indicate a co-authorship relationship. The text data for each node is a list of publications written by the author. The 3D graph layout is generated with Fructerman-Reingold force-directed algorithm [7], and rendered with 3d shaded spheres and tubes. The 3D graph visualization is currently deployed in the AlloSphere environment (Figure 4), which features a spherical display measuring five-meter radius [1]. Users stand on a bridge that spans the cen-

ter of the structure. The distance between users and the surface of the sphere necessitates the usage of relatively large text in order to maintain legibility. This necessity creates occlusion problems (mentioned throughout this paper) when simultaneously trying to view structural data.

The text data viewer software runs in a web environment that is part of the iOS and Android application Control [15]. Control is a generic solution for realizing touchscreen interfaces to control musical, artistic and virtual reality projects. Interfaces are designed using HTML/CSS/JavaScript. The JavaScript runtime has access to hardware features of the devices Control runs on; this enables low-latency UDP networking that would not normally be available to a web interface. Control was designed and developed by the first author.

Figure 3 shows the interface Control presents to users for navigating the graph visualization and viewing text data. At the top is a collection of text data associated with nodes that the user has selected. This collection scrolls horizontally to display more nodes, and each individual node scrolls vertically if there is not enough height to display all publications of an author. The bottom portion of the interface tracks users finger position and sends this information to the visualization to control a virtual cursor. When the visualization notifies the tablet that this cursor has highlighted a node, the tablet displays a square labeled with the authors initials at a location mimicking the node's position on the shared screen surface. Double-tapping the control area sends a request to the visualization for the text data associated with tapped node. When the text is displayed, a line connects the text box to the node it is associated with. This creates a visual link that helps users place text box within the 3D structure. Each expanded node presents buttons allowing users to push the text box to the AlloSphere surface or remove the text box from the tablet display.

We recruited three participants to test the visualization and tablet interface and collect formative feedback. After providing a brief overview and short demo of the interface, participants were asked to complete two tasks; each task was performed twice. The first consisted of identifying an author with more than ten publications while the second consisted of identifying common publications between two authors. Users expressed satisfaction with the tablet interface and all were able to complete the assigned tasks accurately in a short amount of time.

Without any explanation, users quickly adopted the unconventional cursor movement of the tablet, where the cursor's position is directly mapped onto the screen in absolute terms, unlike a mouse cursor which moves relative to its last position. Two participants indicated that they lost the cursor initially, but understood the behavior quickly after. Although they were given the option, none of the subjects chose to push text to the large screen. When asked why they had not used this feature, two participants indicated that they strongly preferred the tablet display to read the text. One participant did indicate that if the same contrast and resolution was possible she would have used the large display for reading. We asked participants whether they understood and made use of the visualizations appearing on the bottom half of the tablet display. Two of them indicated that it grabbed their attention towards the end and all three were able to accurately describe the relation between the abstracted visualizations on the tablet and the large screen. No users expressed any discomfort related to holding the tablet, however, one user said a stand to place the tablet on might have increased comfort.

5 CONCLUSION AND FUTURE WORK

We are encouraged by the initial evaluations of our prototype implementation. Our system removes obtrusive occlusion associated with textual data display in visualizations, fosters collaboration as well as individual tasks, and increases ergonomic comfort compared to interaction techniques involving laptop-based personal dis-




Figure 4: Three users standing on the bridge of the AlloSphere manipulating a graph visualization.

plays.

We will move forward by testing our technique with different datasets in order to prove it is generalizable across a variety of data. We are particularly interested in working with graph visualizations of Wikipedia links where each node represents an article and each edge represents an internal link. Unlike our current static dataset, this graph will be generated dynamically based upon an initial article selected by users. As links between citations are revealed, we believe that articles featuring important concepts and people related to the initial submission will quickly emerge visually. One advantage of having a dedicated textual data reader is that we can make the article associated with each node directly readable from within the viewing application. In addition to internal Wikipedia article citations, we also plan to implement our technique with biological and molecular structural data.

We currently use a simple interaction model in order to navigate users cursors in 3D space that relies on raytracing to determine selection in the 3D environment from absolute 2D positions generated on the tablet. Although this model is sufficient to navigate and select nodes in our test bed implementation, we believe that more precise control of navigation will be necessary for 3D structures with intricate details. We plan to explore using multi-touch gestures to enhance navigation and selection in future implementations by building off of previously described techniques [8] [11].

As we iterate through different implementations and refine the system outlined in this paper, we will also conduct user studies in order to evaluate the effectiveness of our technique for collaborative knowledge discovery. These studies will isolate the effects of moving textual information onto a dedicated reading device and quantify the improvements that our system offers over other text and metadata display techniques in 3D environments.

6 ACKNOWLEDGMENTS

This research was partially supported by NSF grants IIS-0855279, IIS-1047678, and IIS-1058132; by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant No. W911NF-09-1-0553 and Cooperative Agreement No. W911NF-09-2-0053; and by a graduate fellowship from the Robert W. Deutsch Foundation.

REFERENCES

[1] X. Amatriain, J. Kuchera-Morin, T. Höllerer, and S. Pope. The allosphere: Immersive multimedia for scientific discovery and artistic exploration. — *IEEE MultiMedia*, pages 64–75, 2009.

[2] J. Biehl, W. Baker, B. Bailey, D. Tan, K. Inkpen, and M. Czerwinski. Impromptu: a new interaction framework for supporting collaboration in multiple display environments and its field evaluation for co-located software development. In *Proceeding of the twenty-sixth*

annual SIGCHI conference on Human factors in computing systems, pages 939–948. ACM, 2008.

[3] S. Boring, D. Baur, A. Butz, S. Gustafson, and P. Baudisch. Touch projector: mobile interaction through video. In *Proceedings of the 28th international conference on Human factors in computing systems*, pages 2287–2296. ACM, 2010.

[4] S. Boring, S. Gehring, A. Wiethoff, A. Blöckner, J. Schöning, and A. Butz. Multi-user interaction on media facades through live video on mobile devices. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pages 2721–2724. ACM, 2011.

[5] J. Fekete, G. Grinstein, and C. Plaisant. Ieee infovis 2004 contest, the history of infovis, 2004.

[6] G. Fitzmaurice, S. Zhai, and M. Chignell. Virtual reality for palm-top computers. *ACM Transactions on Information Systems (TOIS)*, 11(3):197–218, 1993.

[7] T. Fruchterman and E. Reingold. Graph drawing by force-directed placement. *Software: Practice and experience*, 21(11):1129–1164, 1991.

[8] C. Fu, W. Goh, and J. Ng. Multi-touch techniques for exploring large-scale 3d astrophysical simulations. In *Proceedings of the 28th international conference on Human factors in computing systems*, pages 2213–2222. ACM, 2010.

[9] C. Gutwin and S. Greenberg. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In *Proceedings of the 1998 ACM conference on Computer supported cooperative work*, pages 207–216. ACM, 1998.

[10] J. Jankowski, K. Samp, I. Irzynska, M. Jozwicz, and S. Decker. Integrating text with video and 3d graphics: The effects of text drawing styles on text readability. In *Proceedings of the 28th international conference on Human factors in computing systems*, pages 1321–1330. ACM, 2010.

[11] A. Martinet, G. Casiez, and L. Grisoni. The design and evaluation of 3d positioning techniques for multi-touch displays. In *3D User Interfaces (3DUI), 2010 IEEE Symposium on*, pages 115–118. IEEE, 2010.

[12] M. Morris, A. Brush, and B. Meyers. Reading revisited: Evaluating the usability of digital display surfaces for active reading tasks. In *Horizontal Interactive Human-Computer Systems, 2007. TABLETOP’07. Second Annual IEEE International Workshop on*, pages 79–86. IEEE, 2007.

[13] N. Polys, S. Kim, and D. Bowman. Effects of information layout, screen size, and field of view on user performance in information-rich virtual environments. *Computer Animation and Virtual Worlds*, 18(1):19–38, 2007.

[14] J. Rekimoto. The magnifying glass approach to augmented reality systems. In *International Conference on Artificial Reality and Tele-Existence*, volume 95, pages 123–132, 1995.

[15] C. Roberts. Control: Software for end-user interface programming and interactive performance. In *Proceedings of the International Computer Music Conference (ICMC), 2011*. ICMC, 2011.

[16] D. Tan, D. Gergle, P. Scupelli, and R. Pausch. Physically large displays improve performance on spatial tasks. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 13(1):71–99, 2006.

[17] A. Van Dam, A. Forsberg, D. Laidlaw, J. LaViola Jr, and R. Simpson. Immersive vr for scientific visualization: A progress report. *Computer Graphics and Applications, IEEE*, 20(6):26–52, 2000.

[18] J. Wallace, S. Scott, T. Stutz, T. Enns, and K. Inkpen. Investigating teamwork and taskwork in single- and multi-display groupware systems. *Personal and Ubiquitous Computing*, 13(8):569–581, 2009.

[19] K. Yee. Peephole displays: pen interaction on spatially aware handheld computers. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1–8. ACM, 2003.