



Proceedings of the Workshop on Data Exploration for Interactive Surfaces DEXIS 2011

Petra Isenberg, Sheelagh Carpendale, Tobias Hesselmann, Tobias
Isenberg, Bongshin Lee (editors)

**RESEARCH
REPORT**

N° 0421

May 2012

Project-Teams AVIZ



Proceedings of the Workshop on Data Exploration for Interactive Surfaces DEXIS 2011

Petra Isenberg*, Sheelagh Carpendale†, Tobias Hesselmann‡,
Tobias Isenberg§, Bongshin Lee¶ (editors)

Research Report n° 0421 — May 2012 — 47 pages

Abstract: By design, interactive tabletops and surfaces provide numerous opportunities for data visualization and analysis. In information visualization, scientific visualization, and visual analytics, useful insights primarily emerge from *interactive* data exploration. Nevertheless, interaction research in these domains has largely focused on mouse-based interactions in the past, with little research on how interactive data exploration can benefit from interactive surfaces. These proceedings represent the results of the DEXIS 2011 Workshop on Data Exploration for Interactive Surfaces. It was held in conjunction with the ACM International Conference on Tabletops and Interactive Surfaces (ITS) in Kobe, Japan on November 13, 2011. The introduction summarizes the published papers of the workshop and points to results from workshop discussions. The remainder of the proceedings is made up of the position papers submitted to the workshop.

Key-words: interactive tabletops and surfaces, visualization, data exploration

* petra.isenberg@inria.fr

† sheelagh@ucalgary.ca

‡ tobias.hesselmann@offis.de

§ isenberg@cs.rug.nl

¶ bongshin@microsoft.com

**RESEARCH CENTRE
SACLAY – ÎLE-DE-FRANCE**

Parc Orsay Université
4 rue Jacques Monod
91893 Orsay Cedex

Résumé : Les tables et surfaces interactives sont conçues pour offrir de nombreuses possibilités en termes de visualisation des données et leur analyse. Dans la visualisation d'information, la visualisation scientifique et la visualisation analytique, une bonne compréhension émerge principalement d'une exploration *interactive* des données. Néanmoins, dans le passé, la recherche en interaction dans ces domaines a surtout porté sur des interactions basées sur la souris, avec peu de recherches sur les avantages des surfaces interactives. Ce rapport de recherche comprend les résultats du DEXIS 2011, un atelier de travail portant sur l'exploration de données avec des surfaces interactives. Il a été tenu en conjonction avec la Conférence Internationale de l'ACM sur Tabletops and Interactive Surfaces (ITS) à Kobe, au Japon le 13 Novembre 2011. L'introduction résume les articles publiés dans cet atelier de travail et les résultats de nos discussions. Le reste du rapport se compose des articles présentés à l'atelier.

Mots-clés : tables et surfaces interactives, visualisation, exploration de données

Contents

1	Preface by the Organizers	4
2	Interstitial Space in MDEs for Data Analysis (<i>Anthony Tang and Pourang Irani</i>)	9
3	Position Paper: Design and Evaluation of Visual Learning and Data Exploration Applications in Multiple-Surface Environments (<i>Chia Shen</i>)	12
4	GeoLens: Multi-User Interaction with Rich Geographic Information (<i>Ulrich von Zadow, Florian Daiber, Johannes Schöning, and Antonio Krüger</i>)	16
5	Interactive Multitouch Surfaces for Software Visualization (<i>Craig Anslow, Stuart Marshall, James Noble, and Robert Biddle</i>)	20
6	Position Paper: Touch Interaction in Scientific Visualization (<i>Tobias Isenberg</i>)	24
7	Mobile Interactive Displays for Medical Visualization (<i>Steven Birr, Raimund Dachsel, and Bernhard Preim</i>)	28
8	Branch-Explore-Merge: Real-time Revision Control for Conflict Resolution in Collaborative Visual Exploration (<i>Will McGrath, Brian Bowman, Niklas Elmqvist, and Pourang Irani</i>)	32
9	CoSpaces: Workspaces to Support Co-located Collaborative Visual Analytics (<i>Narges Mahyar, Ali Sarvghad, Tyler Weeres, and Melanie Tory</i>)	36
10	Multitouch Magic Fisheye: Precise Interaction with Dense Data on Tabletop (<i>Frédéric Vernier and Chia Shen</i>)	40
11	Towards Utilizing Novel Interactive Displays for Information Visualization (<i>Christian Tomin-ski, Heidrun Schumann, Martin Spindler, and Raimund Dachsel</i>)	44

1 Preface by the Organizers

By design, interactive tabletops and surfaces (ITS) provide numerous opportunities for data visualization and analysis. In information visualization, scientific visualization, and visual analytics, useful insights primarily emerge from an interactive data exploration. Nevertheless, interaction research in these domains has largely focused on mouse-based interactions in the past, with little research on how interactive data exploration can benefit from interactive surfaces. We assert five apparent benefits of interactive surfaces for visualization systems:

1. As interactive surfaces become part of our everyday environments, they provide new ubiquitous data analysis platforms in which data can be accessed and analyzed anywhere and at any time (e. g., on mobile phones and tablets, in meeting rooms, or on public surfaces);
2. Interactive surfaces offer research opportunities on novel interaction paradigms that can improve data exploration experiences or encourage alternative forms of data exploration;
3. Novel visualization designs and interactions promote the use of visualizations for a broad range of people;
4. In particular, large interactive surfaces offer the possibility of depicting and interacting with much larger visualization spaces than possible previously; and
5. As data analysis is increasingly turning into a collaborative process, interactive surfaces offer novel research possibilities on dedicated collaborative visualization platforms.

While the combination of interactive surface technology and visualization research promises rich benefits, much remains to be learned about the effects of supporting a visual data exploration on interactive surfaces. For instance, we need to learn more about (a) how to re-design desktop- and mouse-based systems for alternative forms of input, (b) what motivates people to explore data using novel vs. traditional interfaces, and (c) how novel input modalities change the ability of people to understand data and draw insights from it. In addition, interactive surfaces often come in the forms of larger or screens, more screens, higher resolutions, sometimes less accurate inputs, and multiple simultaneous inputs, all of which create additional challenges for visualization designers.

At DEXIS 2011, we brought together researchers and practitioners from scientific visualization (VIS), information visualization (InfoVis), visual analytics (VAST), and human-computer-interaction (HCI) to discuss and shape the field of visualization and analysis on interactive surfaces. We discussed ongoing research, exchanged experiences about challenges and best practices, identified open research questions, and developed a research agenda. In these proceedings we collate the knowledge gathered during and after the workshop in order to contribute to the future research in the field.

1.1 Workshop Outline

We organized discussions at DEXIS along three main topics: Data Analysis Environments, Data-Specific Challenges, and Interaction Techniques. The authors of each paper were asked to set their work in relation to the topics in a brief position statement. After these presentations, we discussed the topic in break-out groups and then summarized the discussion in the plenary. We outline the topics we discussed according to the papers presented in the corresponding session.

1.2 Workshop Sessions

1.2.1 Topic 1: Data Analysis Environments

Chair: Tobias Isenberg

The first workshop session was dedicated to the role of the data analysis environment and its influence on the configuration of visualizations and ITS systems. In particular, we focused on the challenges that arise from the specific data analysis environments, insights on how these environments should be designed for interactive surfaces, and experiences with ITS in different types of environments (e. g., work settings and exhibitions). The following papers were discussed in this session:

1. Anthony Tang and Pourang Irani.
Interstitial Space in MDEs for Data Analysis (page 9).
2. Chia Shen.
Position Paper: Design and Evaluation of Visual Learning and Data Exploration Applications in Multiple-Surface Environments (page 12).
3. Ulrich von Zadow, Florian Daiber, Johannes Schöning, and Antonio Krüger.
GeoLens: Multi-User Interaction with Rich Geographic Information (page 16).

In the discussions following these papers we identified a number of challenges for research at the intersection of ITS and visualization. Social aspects characterize the first group of challenges. These include different types of users and their domain-specific tasks, interactions, and visualization requirements as well as different types of collaboration settings. These challenges had been specifically mentioned in the presented papers. Tang and Irani discussed challenges of surface connectivity and how the space between surfaces could be used to help people make a transition of data and results between collaborative and individual work (tasks) as well as between different types of devices. Shen discussed several design guidelines as well as evaluation measures for data analysis environments with a strong social focus on collaboration, engagement, and cognition. The GeoLens interaction techniques presented by von Zadow et al. were observed in use by over 150 visitors of an interactive science exhibition in Germany. The observations point to interesting challenges designing data exploration tools for walk-up-and-use scenarios. We further discussed hardware challenges that strongly influence data analysis. These begin with the physicality of the surface itself, its affordances, and the type of fluidity and interaction modalities it supports but also include finding the right location for an interactive surface.

Both social and hardware challenges of data analysis environments together point to the problem that one needs to design these spaces to be highly personalizable and adjustable to different needs. These needs may depend on a variety of factors and research needs to gain a better understanding of these individual needs and how they can be transformed into more generalizable design considerations for data analysis environments.

1.2.2 Topic 2: Data-Specific Challenges

Chair: Tobias Hesselmann

In this session we discussed issues related to specific types of data (e. g., software structure, medical data) and their influence on the design of visualization on interactive surfaces. The session consisted of the following papers:

1. Craig Anslow, Stuart Marshall, James Noble, and Robert Biddle.
Interactive Multi-touch Surfaces for Software Visualization (page 20).
2. Tobias Isenberg.
Position Paper: Touch Interaction in Scientific Visualization (page 24).
3. Steven Birr, Raimund Dachzelt, and Bernhard Preim.
Mobile Interactive Displays for Medical Visualization (page 28).

The three papers discussed different aspects of the session topic. Anslow et al. presented their challenges in moving from single-user desktop-based software visualizations to collaborative tabletop visualization tools for software development. These included those of synchronizing different representation types and providing dedicated tabletop navigation techniques in the data space. Isenberg discussed interaction challenges for typically three-dimensional-scientific visualizations. He pointed to specific underexplored future research directions including the development of 3D interaction toolkits, better support for precise exploration control, the interaction with stereoscopic displays, and the interaction with 3D data in collaboration. Birr et al. presented several challenges and ideas for introducing mobile surfaces in medical environments. The presented techniques were put in close context of the medical domain and its requirements among which the most important ones were scalability and platform independence as well as interaction and communication.

The discussions following these papers were centered on unanswered questions about designing data interactions and data representations for interactive surfaces. These included: How do we need to redesign visualizations and data exploration tools for collaboration? How specifically do we have to augment data representations for collaboration? What are data-specific interaction challenges, for example, 2D vs. 3D vs. text-data, etc.? How can solutions from a particular domain be generalized or transferred to another domain?

1.2.3 Topic 3: Interaction Techniques

Chair: Bongshin Lee

The last session focused on issues related to the interaction with visualizations on interactive surfaces. We examined the role of interaction in ITS vs. Desktop PC computing, and discussed the validity of interactions across different devices and types of data. The following papers were presented:

1. Will McGrath, Brian Bowman, Niklas Elmqvist, and Pourang Irani.
Branch-Explore-Merge: Real-time Revision Control for Conflict Resolution in Collaborative Visual Exploration (page 32).
2. Narges Mahyar, Ali Sarvghad, Tyler Weeres, and Melanie Tory.
CoSpaces: Workspaces to Support Co-located Collaborative Visual Analytics (page 36).
3. Frédéric Vernier and Chia Shen.
Multitouch Magic Fisheye: Precise Interaction with Dense Data on Tabletop (page 40).
4. Christian Tominski, Heidrun Schumann, Martin Spindler, and Raimund Dachzelt.
Towards Utilizing Novel Interactive Displays for Information Visualization (page 44).

McGrath et al. discussed the challenge of allowing parallel and diverging interaction and revision control for collaborative data analysis tools. They also discussed social challenges of public and private views during collaboration and how they can be brought into context. Similarly, Mahyar et al. discussed interaction challenges in the social context of collaboration. Their system, CoSpaces, uses tab-based portals for

interaction across workspaces and for maintaining awareness of others. Vernier and Shen presented a specific interaction technique, the multitouch magic fisheye that showed some of the possibilities of utilizing multiple fingers for controlling data exploration interactions for interactive surfaces. Finally, Tominski et al. discussed three gaps in our current research at a higher level. These include technology, integration, and guidelines gaps and point to challenges that still need to be addressed as we develop visualization systems for interactive surfaces. In the plenary discussions we concentrated on the following questions:

- How does touch interaction fundamentally differ from mouse settings and how does the new input modality impact visualization?
- How can we encourage collaboration through interaction?
- Can we develop a general data exploration interaction vocabulary?
- What are special requirements when display space and interaction space overlap?

1.3 Future Research Directions

The workshop ended with a brainstorming session on future research directions and challenges. We identified 10 main topics which require further attention:

1. *Multi-display environments*: Multi-device interaction, multi-device integration, surface ecologies for data analysis, etc.
2. *Collaboration*: Supporting different collaboration styles, transitioning between styles, merging group work, etc.
3. *Interaction techniques*: Develop a vocabulary of touch interaction for data exploration, support different touch proficiency, improve touch literacy, support touch with different data types, and develop gestures for data exploration.
4. *Representations*: Understand how visual encoding needs to change depending on screen size and input modality and what is the role of perception?
5. *Evaluation*: Develop dedicated study methodologies for understanding the role of touch for data exploration, understand how to measure efficiency and effectiveness of interaction in data exploration, etc.
6. *Multi-modality*: How can we enhance data exploration with input modalities other than touch?
7. *Privacy*: How can we support exploration of private data and what are privacy sensitive actions in data exploration?
8. *Role of surface type for data analysis*: Understand what different types of surfaces and surface settings are most suited for and what are the most suited tasks, data types, representations, and interactions for different devices?
9. *Guidelines*: Can we find generalizable guidelines or design considerations for data analysis environments on interactive surfaces?
10. *Promote the research direction*: We need to more strongly promote the value of research on interaction in novel environments, we need to get a cross-disciplinary dialog started between communities on the topic, and we need to design and deploy more systems that promote and demonstrate the value of interactive surfaces for data analysis.

1.4 Acknowledgements

We would like to thank the participants of DEXIS 2011 for their excellent contributions and lively discussions. The exchange showed that data exploration on interactive surfaces is an exciting area of research that offers great potential.

Petra Isenberg, Sheelagh Carpendale, Tobias Hesselmann, Tobias Isenberg, Bongshin Lee
DEXIS 2011 Organizers

1.5 Copyright Notice

The copyright for all of the following papers including all material presented therein is retained by the respective authors of the workshop papers.

Interstitial Space in MDEs for Data Analysis

Anthony Tang

Department of Computer Science
University of Calgary
tonyt@ucalgary.ca

Pourang Irani

Department of Computer Science
University of Manitoba
irani@cs.umanitoba.ca

ABSTRACT

Multi-display environments comprise large shared displays as well as personal devices. In this work, we discuss how the interstitial space—the space between displays and devices—can also be made into an interactive space. Such a space can support collaborative data analysis by providing a focus+context workspace; providing a means to transition between collaborative and individual work, and by providing a means to transition tasks between devices.

INTRODUCTION

Multi-display environments are rich digital workrooms comprising wall-displays, tabletop displays and personal devices. These display ecologies represent great opportunities for collaborative data analysis: shared displays can facilitate group tasks, while personal devices can support independent tasks. Figure 1 illustrates an imaginary digital workroom with the typical large, high-resolution, interactive displays to which the research community has devoted much of its recent efforts. It also shows two spaces, the interstitial space (the space between displays) that we argue represents an interesting design opportunity for the community. While considerable effort has gone into designing interaction techniques and visualizations for what we typically consider as “interactive large displays,” very little work has considered the interstitial space in these rooms, and in particular, the role that the space between the displays can play. In this position paper, we consider how this space can be used to support collaborative data exploration and analysis, and present several design factors for interstitial spaces that we seek to explore.

For explanatory purposes only, we describe here an imaginary instantiation of such spaces in a digital workroom. Interstitial space is comprised of surface space between displays. In this imaginary scenario, this interstitial space is made visible through a low-resolution projection onto both the walls and the ground in this digital workroom. These projected surfaces are not touch-interactive; instead, people interact with it by using their mobile devices as a proxy. If on the floor, users can interact with such spaces with their feet or through shadows. We can already see that conceptually, the MDE becomes a Focus+Context environment, where the interactive displays are the Focus, while the interstitial space provides Context. Yet, how can this space be leveraged to support the collaborative analysis process? Let us further examine the collaborative visual analysis process.




Figure 1. (1) and (2) are typical large shared displays in an MDE. (3) the wall, and (4) the ground are interstitial spaces that can be used to support auxiliary tasks or work as a scratch-pad for collaborative data analysis.

SUPPORTING COLLABORATIVE ANALYSIS

We outline three ways in which interstitial space might support collaborative analysis in MDEs. First, it may support specific “tasks” or “sub-processes” in a collaborative visual analysis process. Second, it might support transitions between shared and independent work. Third, it could be used to support transitions between different device types.

To begin, we focus on Isenberg and colleagues’ work on the visual analysis process (2008). Here, the authors outlined several processes that individuals in a group engage in when analyzing visual data together: Browse, Parse, Discussion Collaboration Style, Establish Task Strategy, Clarify, Select, Operate, and Validate. Several of these processes intuitively map to how interstitial space could be used. During the “Browse” process, for example, people look through the data, implicitly sorting it into several piles based on how they might expect to use the data. Interstitial space could be used here for groups of items that may not seem as important (i.e. saving valuable “interaction” space). In so doing, it can also support a faster “Select” process, as items can remain visible without having to be “piled” into groups.

Another common design requirement arising out of studies of collaboration is to support fluid transitions between shared and independent tasks (i.e. collaborative work and independent work). Interstitial space can support this transition by providing a visible storage or transport medium for different workspaces. Moving different workspaces across

to different displays through the interstitial space supports collaborators' awareness of what is happening in the workspace, and how work and tasks are being divided among individuals and sub-groups. For example, if data is to be examined by a sub-group or an individual, the interstitial space can be used as temporary "ether", a "scratch-pad" or as a space where content is placed temporarily for these or other individuals to retrieve at a later time. We hypothesize that allowing users to off-load their internal "scratch-pads" onto interstitial space will facilitate analysis on only core components to the task.

Finally, because devices and displays in MDEs are physically (and visually) detached from one another, interstitial space can actually function as a visual bridge between the devices. It can be used for visualizing the movement of content or information across devices and displays in such an ecology.

DESIGN FACTORS FOR INTERSTITIAL SPACES IN MDES

In these early stages, we have considered several factors that influence the design of interstitial space—factors that influence how it is realized, how it is interacted with, and the affordances it provides.

Conceptual model: transient vs. ambient. Prior work that has considered this interstitial space (Xiao et al., 2011) has primarily viewed this space for transient interaction. That is, content in this space is only intended to be here for a short period of time. This relates to our concept of it as being a visual transport medium for content—something like a "clipboard" for the MDE. It can also be leveraged to support collaborator awareness of our interactions in the workspace (e.g. Shoemaker et al., 2006). Yet, we can also consider it as a low-resolution ambient space that either: (a) exposes functionality to manipulate the high-resolution interactive space (e.g. controls for visualizations could be placed in interstitial space to save room from the actual workspace), (b) provides low-resolutions visualizations that react to the interactions taking place on the shared displays.

Organization: semantic vs. spatial. One thing to consider is how content is to be organized in interstitial space: will it be organized semantically, or spatially. Recently, researchers have considered reinventing the interface to exploit users' spatial memory, through semantic association of interface components with spatial layout. Interstitial space provides a larger repository that can further enable more associations. However, separating the organization either semantically or spatially can be left to designers based on the affordances they wish to embed.

Objects: dynamic vs. static. How should content in interstitial space appear?

Content: artefacts that relate to work / artefacts that relate to the people in the MDE. To this point, we have mainly considered that artefacts in interstitial space would be data elements related to the analysis. An alternative conceptuali-

zation of this would be to place content as it relates to individuals in the interstitial space in such a way that it tracks or follows individuals. This content could relate to their independent tasks, or be tools that relate to those individuals. Having this information track and follow along with an individual would provide easy access to it.

BASIC TASKS WITH INTERSTITIAL SPACE

There are some basic tasks that people will need to accomplish with interstitial space: placing and retrieving content from interstitial space, querying data in this space and making the results visible, deciding how to eliminate or erase content from this space. While these problems have received less consideration, several basic mobile device interaction techniques could be used to facilitate these tasks. For example, a mobile device could act as a peephole for shared displays (as in Boring et al., 2010). Users can simply 'scan' interstitial space with their mobile device to make 'invisible' content appear on it. Other mechanisms might include displaying interstitial content in a minimalistic ways, using methods of world-in-miniature (e.g. Biehl & Bailey, 2006), or even with mechanisms that provide accurate cues to off-screen content (e.g. Gustafson et al., 2008).

Yet, we envision that designing appropriate interaction techniques will rely on an understanding of how these spaces will be used in domains such as collaborative analysis. One method for developing this understanding is to study how the interstitial space is organized (i.e. study how people organize content within the space); second, to develop methods to provide people representations of this organization and content, and then finally to use this understanding to iteratively design interaction techniques.

A CASE FOR IMPROVED ANALYSIS

Our primary argument infers that interstitial space will augment traditional methods of analysis and data inquiry. We elaborate on this aspect by walking through a simple case.

Let's consider a group of two analysts (for simplicity) inquiring into a recent case of a hit-and-run incident in a city. The police inform them that potential witnesses implied that a cab driver was involved and provide the analysts with GPS data from that cab's company, based on the time interval of the incident. The analysts now have to prod the provided data to assist the police in determining the suspect.

Effect of interstitial space on on-the-fly queries. As suggested earlier, instead of employing object piles, analysts can off-load immediately unnecessary content into interstitial space. For example, the analysts may 'store' multiple forms of space-time views of the data and instead of having it clutter the usable space can have it placed on the 'side lines.' A view showing all cab movements (from the GPS) can be placed into this space, while the analysts explore city wide camera recordings. On demand, the analysts may query the route data, which can then be presented onto the primary display for analysis. Based on our hypothesis that

such space can better facilitate semantic-spatial associations, retrieving objects of interest on demand will be highly efficient. We plan on studying the effectiveness of such forms of object placement/retrieval methods in comparison to more traditional methods for analytic outcomes.

Attributing relevance. In the analysis process, some information is more relevant than other. To avoid completely erasing that knowledge (as it may have taken the analysts some time to produce it), organization in interstitial space can attribute relevance to the derived information. For example, if the analysts have now attributed the incident on a handful of cab drivers, the history information in how these were obtained and the relevance associated to each item in history can be organized in IS space for later presentation. Such forms of history tracking typically require large trees or lists, which can be avoided if properly partitioned in this additional space.

Linking between alternative forms of analysis. Another feature of interstitial space that we leverage for this analytic task could be for linking between various steps taken by the two analysts. Analysts may at times be working separately on the same source of data and such spaces can provide a common ground among their individual approaches. For example, both analysts may have refined their inquiry such that only one piece of the information may be left to solve the problem. Data from interstitial space could be fused, by auxiliary routines (metaphorically that run in the background) to suggest a refined solution. By removing the back-and-forth comparisons away from their primary displays, analysts may be able to find better solutions to their problem.

CONCLUSIONS

We are in the early days of exploring the design space of multi-display environments. Whereas most researchers have focused their efforts (rightfully so) on the main interactive shared displays in these spaces, our focus in this position paper is to consider how the low-resolution interstitial space can be used to support collaboration. We have discussed how an interstitial space might be used to support analysis, as well as described a set of design factors that can guide exploration into this space.

REFERENCES

1. Biehl, J.T., and Bailey, B.P. 2006. Improving interfaces for managing applications in multiple-device environments. In *Proc AVI '06*, 35-42.
2. Boring, S., Baur, D., Butz, A., Gustafson, S., and Baudisch, P. 2010. Touch projector: mobile interaction through video. In *Proc CHI '10*, 2287-2296.
3. Gustafson, S., Baudisch, P., Gutwin, C., and Irani, P. 2008. Wedge: clutter-free visualization of off-screen locations. In *Proc CHI '08*, 787-796.
4. Isenberg, P., Tang, A., and Carpendale, S. 2008. An exploratory study of visual information analysis. In *Proc CHI '08*, 1217-1226.
5. Shoemaker, G., Tang, A., and Booth, K.S. 2007. Shadow reaching: a new perspective on interaction for large displays. In *Proc UIST '07*, 53-56.
6. Xiao, R.B., Nacenta, M.A., Mandryk, R.L., Cockburn, A., Gutwin, C. 2011. Ubiquitous Cursor: A Comparison of Direct and Indirect Pointing Feedback in Multi-Display Environments. In *Proc GI '11*, 135-142.

POSITION PAPER

Design and Evaluation of Visual Learning and Data Exploration Applications in Multiple-Surface Environments

Chia Shen

School of Engineering and Applied Sciences
Harvard University
chia_shen@harvard.edu

ABSTRACT

In this position paper, I will describe a set of interactive visual prototype systems that have been developed in an interactive environment with a multi-touch tabletop and a large display data wall. I will summarize findings from our initial pilot studies of these experimental systems, offer design guidelines, and then propose a set of evaluation criteria to the interactive tabletops and surface research community.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Design, Human Factors

Keywords: Multi-touch, Visualization

INTRODUCTION

What makes an interactive surface better suited for a particular data visualization application than other form factors? Another way to ask this question is “How can a particular visual data exploration application benefit from the potentially higher cost of being developed on one or multiple interactive surfaces?”. This cost can be in the form of the actual price of the display device, in the form of development cost as touch input and multi-user applications require non-standard interfaces, or in the form of user cost such as learning to interact with non-standard UI as well as the possibility of enduring multi-user interference.

In the past few years, we have been addressing the above questions by continuously investigating the design and development of a variety of visual data exploration systems on a heterogeneous multiple interactive surface environment. In this position paper, I summarize our experience, findings, and propose three measurement criteria as evaluation metrics for visual data explorations.

The types of visual data that we have studied include:

- Scientific data from multiple independent source
- Coordinated data with multiple representations
- Simulation data
- Animation and video
- 2D and 3D visualization
- Large public scientific databases

In the following, I will briefly describe each of the multi-surface data exploration prototypes and identify the key findings that will help us to formulate design guidelines and evaluation metrics.



Figure 1: CThru, an interactive video, 2D and 3D system for learning the inner life of a cell. (Reproduced from Figure 1 in Jiang et al CHI2009.)

FOUR MULTIPLE SURFACE DATA EXPLORATION AND ANALYSIS PROTOTYPES

Cthru

CThru [3] is a video-centered information space for self-guided learning. In CThru, core multimedia education contents of different formats are used to construct a cross-referenced information space running along with the main video, conveying a variety of domain specific knowledge and information. Users watching the video can freely jump to or out of the information space for in-depth knowledge, thus a sequential video playback is replaced by the experience of immersing and exploring in a multi-dimensional space. CThru runs in a multi-display environment with a large display wall and an interactive tabletop (Figure 1).

Copyright is held by the author.

DEXIS 2011 Workshop on Visual Data Exploration on Interactive Surfaces. Workshop in conjunction with ACM ITS 2011, Kobe, Japan. November 13, 2011.

During our evaluation of CThru, one interesting issue arose during the interaction – the unequal distribution of user attention between the wall and tabletop. Our observation showed that the users' visual foci stayed on the tabletop most of time, even though the data wall provides three times as much resolution as the tabletop and is much larger. Most users only shifted their foci to the wall when reading text.

INVOLV

INVOLV [1] allows users to explore and navigate over one million species of life on earth catalogued in the Encyclopedia of Life (www.eol.org) database using a voronoi tree map interface on a multi-touch tabletop, in conjunction with a data wall (Figure 2). On the multi-touch table, users can use the treemap to browse and navigate to each group of species, while on the data wall, the associated webpages from www.eol.org and www.tolweb.org on the current species are opened and displayed.

Given the findings from CThru, we added a pair of 'audio' and animation cues to bring users' attention to the data wall. Whenever a user's interaction on the tabletop goes to a new family of species, a pair of animated visual "nuggets" will flow from the table to the two webpage windows on the datawall. This multi-modal sound and animation sequence have brought users' attention across the display space between the table and the display datawall.




Figure 2: INVOLV, a voronoi treemap based system for exploring the www.eol.org database. (Horn et al ISVD2009.)

WALDEN

The WALDEN simulation [5] shown in Figure 3 seeks to model changes in abundance of a selected group of plants found in Concord Massachusetts, as a response to climate change and human intervention in the local ecosystem. On the multi-touch table, a user is able to manipulate the average and the standard deviation of annual temperatures and watch the effects of this simulated climate change on the set of flowering plants on the datawall. A user can in addition select a particular plant on the tabletop to see the geo-



Figure 3: WALDEN, a dynamic simulation modelling changes in abundance of a select group of flowering plants found in Concord, Massachusetts in response of climate change. (Reproduced from Figure 1 in Schneider et al 2011.)

graphic distribution and illustrations of this plant on the datawall.

Our case study reveals that cognitive overload occurs when individuals encounter a large amount of information on separate displays, while pairs of users may benefit from collaborative learning in such a setting.

From our observations of both individual and pairs of users, having a large amount of display area is a double-edged sword on a cognitive level: visualizations are not limited by space, but individuals find the amount of information overwhelming. They complained that there was "too many things to look at" or ignored a graph containing important information. Dyads, on the other hands, used the amount of information as a way to make unique contributions to their conceptual understanding of a domain. They used the multiple views from the simulation as a way to propose different perspective on phylogenetics and confront their hypotheses with the results displayed.

Figure 6 provides a more detailed view of two sessions, one individual session and the other a dyad session: every 10 seconds is coded according to the quality of the comments made by the users. A score of 0 means silence or off-task behavior while 5 reflects a conceptual understanding of the system. From this graph we can observe that the individual abundantly commented the interface but almost never went beyond making simple observations. The dyad, on the other hand, produced richer and more varied comments. For instance, they made more complex hypotheses by considering branches of the phylogenetic tree. This allowed them to make finer distinction in terms of phylogenetic relationships (e.g. "this common ancestor seems to be responsible for the decrease in abundance"). They also questioned the structure of the tree (e.g. "is the length of a

branch significant?", "are some part of the tree hidden/not displayed?", "if species are in a close spatial proximity on the radial tree, are they really more closely related?"), whereas the individual did not. More importantly, members of the dyad made more hypotheses and confronted their predications more often than the single user.

WESPACE

The WeSpace [8] is a multi-surface collaboration space. It is a general-use tool for workplaces in which simultaneous visual exploration rendered from multiple data sources by multiple people is a crucial part of the work-flow process. WeSpace is the outcome of close collaboration between our research team and a population of scientists – astrophysicists from the Harvard Smithsonian Center for Astrophysics (the CfA). It is designed to enable modern day-to-day spontaneous collaborative sessions that are mediated and augmented with computational display devices.

During our observation sessions in which three scientists carried out working meetings examining different types of telescopic data, we recorded inputs on the multi-touch table from all users. While the distribution of input was fairly equal when the meeting was viewed as a whole, a different story emerged when we examined input from each of the three users over time. Figure 5 shows the relative number of input events performed by each of the three users for each 5 minute time period. While the overall distribution of input was relatively even, the distribution during the 5 minute samples was not. The logs suggest that the control of the system passed from user to user at different times during the meeting as the participants took turns directing the conversation. While the majority of input was normally made by one participant, it was rare to see any one user monopolizing the table. These patterns match our observations of the meeting that the scientists took turns introducing new data and hypotheses, with their colleagues reacting to these additions.

GUIDELINES AND EVALUATION METRICS

In this section, I will summarize the findings from our own experience and evaluation outcomes from our past and current research systems into design guidelines. I will then propose a set of three measures that can be used as evaluation metrics for these multiple surface interactive visual data exploration applications.

Design Guidelines

1. *"Let there be sound and motion"*: Non speech cues and motion for visual foci shift are useful design tools. It is well-known in the visual design arena that motion is a good visibility enhancement and can be used effectively for directing a viewer's orienting response [7]. Vision (high resolution acuity but a small of area of a window) and hearing (3D surround but low resolution) are interdependent and work well together [2]. In a multiple surface environment where some display surfaces are out of the immediate visual field of the user,

utilizing a combination of sound and animation cues should be considered.

2. *"Let there be collaboration"*: The large visual space lends well to collaboration. Two or more people seem to gain better cognitively a multi-surface environment for learning, exploration and analysis. In choosing the types of visual data exploration applications, we should take this into consideration.
3. *"Let there be equal input"*: In order to allow people to collaborate, we need to provide mechanisms for each user to conveniently direct her personal input. This may include inputting data, redirecting attention, and exploring and navigating shared visualization.
4. *"Engage them!"*: We have also observed that interactive surfaces may offer higher levels of engagement from the users. To attract and engage the users deeply into the visual representation is an important criterion.

Evaluation Measures

Here I list three levels of evaluation measures that we have found effective and useful:

1. Cooperation metric: Can the interactive visual design support two or more people work and play together? How fluid can people input, converse through visualization, and explore without interference? This can be measured using the collaboration profile developed by Shaer et al in [6].
2. Engagement metric: Does the system engage its users? This can be measure through a few different ways, such as self-reporting as in [4], or in comparison to other systems or other form factors.
3. Cognitive metric: Does the design meet the goals of the application? If the application domain is education, how well do the learners achieve the projected learning gains? If the application do-




Figure 4: WeSpace, a systems that allows astrophysicists to dynamically interactive with their own data from individuals' laptops on a multi-touch tabletop and a large datawall. (Reproduced from Figure 1 in Wigdor et al 2009.)

main is science discovery, does the system allow scientists to effectively make inferences and predictions? If the application domain is visual analytics, does the system enable the users to arrive at analytical insights?

CONCLUSION

As interactive surfaces become more and more prevalent, we need to develop well-studied design guidelines and have a formal set of evaluation measures that can be shared within the research and design community. In this position paper, I have described our past and ongoing work in designing and developing multiple-surface interactive visual data exploration systems. I have also proposed a set of initial design guidelines and evaluation measures.

REFERENCES

1. Horn, M.S., Tobiasz, M., Shen, C., Visualizing Biodiversity with Voronoi Treemaps, *Proceedings of Sixth annual International Symposium on Voronoi Diagrams in Science and Engineering*, Copenhagen, Denmark. June 23 - 26, 2009.
2. Jacko, J., *The Human-Computer Interaction Handbook: fundamentals, evolving technologies, and emerging applications*. CRC Press, 2008. ISBN-10: 0805858709
3. Jiang, H., Viel, A., Bajaj M., Lue R.A. and Shen C., CThru: Exploration in a Video-Centered Information Space for Educational Purposes. *Proceedings of CHI 2009*, Boston, MA, April 4-9, 2009, 1237-1246.
4. O'Brien, H. L., Toms, E. G., Kelloway, E. K., and Kelley, E. Developing and evaluating a reliable measure of user engagement. *Proceedings of the American Society for Information Science and Technology* 45, 1 (Jan. 2008), 1–10.
5. Schneider, B., Tobiasz, M., Willis, C., Shen, C., When Two is Too Many for One: Cognitive Overload in a Multi-Surface Multi-Touch Environment. *Submission in review*.
6. Shaer, O., Strait, M., Valdes, C., Feng, T., Lintz, M., and Wang, H. Enhancing genomic learning through tabletop interaction. In *Proc. of CHI '11*, ACM Press (2011)
7. Ware, C., *Visual Thinking for Design*. Morgan Kaufmann, First Edition (April 2008). ISBN-10: 0123708966.
8. Wigdor, D., Jiang, H., Forlines, C., Borkin, M., Shen, C. The WeSpace: The Design, Development and Deployment of a Walk-Up and Share Multi-Surface Visual Collaboration System. *Proceedings of CHI 2009*, Boston, MA, April 4-9, 2009, 1237-1246.




Figure 5: Relative input contribution from each of the three scientists. (Reproduced from Figure 4 in Wigdor et al 2009.)




Figure 6: Comparing quality of users' verbal reflection between a dyad and an individuals. (Reproduced from Figure 3 in Schneider et al 2011.)

GeoLens: Multi-User Interaction with Rich Geographic Information

Ulrich von Zadow

Archimedes Solutions GmbH
Berlin, Germany
uz@archimedes-exhibitions.de

Florian Daiber

DFKI GmbH
Saarbrücken, Germany
florian.daiber@dfki.de

Johannes Schöning

People Interactive GmbH
Köln, Germany
schoening@people-interactive.de

Antonio Krüger

DFKI GmbH
Saarbrücken, Germany
krueger@dfki.de

ABSTRACT

Zoomable 3D map viewers have been used many times as excellent demos of the potential of multi-touch systems. While they showcase the possibilities of multi-touch very well, they are inherently not capable of supporting multi-user interaction. We present enhanced magic lenses - GeoLenses - as user interface concept that expands on prior magic lens work. GeoLenses are fully multi-user capable while still being intuitive to use. We built a complete end-user application based on this concept and deployed it in an exhibition. Our observations of application usage by about 150 visitors suggest that the concept is viable and easy to understand for the general public.

Author Keywords

Geographical information systems, GIS, multi-touch, multi-user, tabletop computing, in situ, magic lens

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces

INTRODUCTION

Demos of 3D multi-touch map viewers abound¹. The interaction paradigms of these demos are immediately understandable and showcase the possibilities of multi-touch very well. However, they are inherently not multi-user capable, since the interaction of one user will always disrupt the interaction of all other users. A good example for this can be seen in a video² of two children trying to simultaneously interact with a map application (see Figure 2): After a short while one child blocks interaction by simply touching the screen.

Magic Lenses support multi-user interaction. In the context of GIS (Geographical Information Systems), they consist of screen areas that modify what is displayed inside - for instance, by magnifying the map view or by displaying different layers of data. Since several lenses can be active at one time, interaction by multiple users does not present a problem. There is a significant amount of research on this



Figure 1. Users simultaneously interacting with the GlobalData application on our multi-touch table.

topic. However, no data on use outside of research laboratories exists and usage in a real application context has not been demonstrated.

We are interested in the use of multi-touch, multi-user GIS systems for the general public. To this end, we have expanded existing Magic Lens concepts to support additional geolocated information and additional methods for configuring the data that is displayed inside the lens. We constructed a full end-user application, deployed it in an exhibition and report on observations of visitors using the system.

RELATED WORK

Tabletops and wall-sized interactive displays are well suited for collaborative work. Co-located collaborative work at tabletops has been an active research field for many years. In the early 90s Tang [16] did observational studies and derived requirements for collaborative work at tabletops and walls. Scott et al. [15] proposed system guidelines for collaborative work at tabletop displays. The DiamondTouch [6], with its capability to distinguish between different users, was often used to explore collaborative work at tabletops (e.g. [7, 8, 17, 19]).

Installations in museums and other public spaces present unique challenges. Hinrichs et al. [10] sum these up very well. Hornecker and Stifter [12] report that time pressure and the number of exhibits in an average exhibition result in many exhibits being abandoned quickly: 'the first ten seconds need to provide an incentive to continue.' In addition, several studies of interactive surfaces in the wild (e.g. [5, 11,

¹For examples, see <http://www.youtube.com/watch?v=3gxVjzmTnS4> and <http://vimeo.com/4973440>.

²<http://www.evl.uic.edu/cavern/tactile/kid.mov>




Figure 2. Children blocking each others input (screenshots from video²): (a) Children interacting with the table (panning and zooming the map); (b) - (c) The child in the foreground is accidentally blocking the interaction with the table and then consciously blocking the other's input by pushing permanently on the screen; (d) Intervention of instructor to correct the childrens' interaction.

14], to some extent [10]) determine that many users were interested more in the technology or visual effects than the content.

The notion of the “magic lens” was first introduced by Bier et al. in 1993 [1]. Bier et al.'s original lenses are transparent or semi-transparent user interface elements which can be placed over objects to change their appearance and/or facilitate interaction. Since then, the concept has been applied frequently to geographic visualization. Carpendale et al. made significant contributions, introducing virtual magnifying lenses in GIS applications [3, 4]. Here, they are used as a focus+context technique showing detailed information without losing the context. The goal of the technique is to allow zooming in while retaining a sense of the global position of the zoom region. In the context of multi-user environments, these lenses allow multiple users to simultaneously inspect details in different foci in a global context. Forlines et al. added user interface tools to the lenses [7]. Furuchi et al. [8] used a magic lens concept to provide several layers of geographic information. However, beyond layers and zoom, configuration of the data set or its visualization is not possible in prior work. Beyond geographical information layers, no additional data is visualized in these approaches. All work on magic lenses described above was done in lab or demo settings; real-world data sets were seldom used, and no study with end users has been performed.

With GlobalData we present a ready-to-use end-user application with a significant quantity of real-world data that can be explored interactively.

THE GLOBALDATA MULTI-TOUCH APPLICATION

The GlobalData application was commissioned as part of the Science Express Germany Exhibition³. The Science Express Germany was a mobile exhibition train focusing on research in Germany. It was open for six months in 2009 and was visited by more than 250.000 people. GlobalData was also presented as a demo at ITS 2010 [18].

We used an Archimedes SESSION DESK⁴ – a 125 × 80 cm multi-touch device based on DI – for the exhibit. A horizontal tabletop was a natural choice given that this shows the maps in the correct perspective and is open to collaboration. This was aided in part by the narrowness of the

train wagons: the table was only accessible from three sides.

GlobalData was developed using the open-source media development platform libavg⁵.

Content

As content we used the National Geographic Special Edition on the Planet Earth 2008 [13], including many maps, statistical data, images, additional videos and documents. From this, the GeoLens concept was derived: Multi-user interaction is facilitated using personal “magic lenses” on a 2D map of the earth. To integrate the content into the application concept, we made an early decision to have several different map views, per-map geolocated images and videos, per-map geolocated text snippets and a legend for each map. In the end, we had 4 maps and a total of 63 geolocated info points, with most of the points containing textual as well as image or video data.

In accordance with the topic of this railcar, the map views selected were: Population (per-country population density data), Habitat (density of farm and pasture land), Overfishing (density of ocean fishing) and Water (density of irrigated areas). The population view deserves special mention: It shows population data for the years 1950-2008 and projected data for 2030. This data was not part of the National Geographic Magazine, but compiled from GapMinder⁶ and mapped to a world map using standard GIS software.

Interaction

In idle mode, the device shows a stylized map of the earth. No usage instructions are provided. However, users can open circular GeoLenses by simply touching the surface (see also Figure 3). These circles show the same map segment as the underlying base map and superimpose different data layers on it. GeoLenses can be dragged and resized using the standard pinching motions. GeoLenses are novel enough to attract attention to the exhibit. Still, they are immediately understandable as a user interface concept and thus easy to use. All interaction takes place in the context of a lens.

While the device has no concept of lens ownership, personal space issues that come into play in such a collaborative environment [9] are dealt with smoothly. The initial size of a magic lens roughly corresponds to the intimate space of the

³http://de.expedition-zukunft.org/alias/healthy+_productive/994242.html

⁴http://www.archimedes-exhibitions.de/%23/exhibits/in-serie/_/sessiondesk.html

⁵<http://www.libavg.de>

⁶<http://www.gapminder.org/data/documentation/gd003>



Figure 3. Screenshot of GeoLens showing the user interface components: (1) Geolocated point of interest; (2) Map view menu; (3) Geolocated text; (4) Open/close button for geolocated images and videos; (5) Geolocated image; (6) Open/close button for map legend; (7) Time slider for population data; (8) Language toggle.

person interacting with it. Dragging and resizing operations allow collective ownership and ownership transfer. All additional content and user interface elements specific to a lens are grouped around it. The area inside the lens is intentionally kept clear of user interface elements and contains map data and location markers for georeferenced items. The goal was to keep the lens itself visually clean and avoid user interface conflicts between the elements and dragging or sizing operations.

A menu near the lens switches between different map views. Buttons on the lens perimeter open and close the map legend and geolocated images and videos. An additional button is used to switch between different languages. The geolocated text snippets are also displayed close to the lens. These change automatically without user interaction as the lens is moved around the tabletop. When the population view is active, an additional slider allows the user to choose the year that is visible in the GeoLens. Note that this represents an additional degree of freedom not present in prior work.

OBSERVATIONS

We observed about 150 exhibition visitors in early June 2009. Group size ranged from 1–5 people, with a few larger groups (school classes) and a wide age distribution, including a significant number of senior citizens. As is typical for museum exhibits, mean interaction time was about 3 minutes. Most of the groups recognized the collaborative potential of the application and interacted together on the table. As with many public displays, getting people to interact with this exhibit proved to be a critical point (see [2]). Some visitors passed by the GlobalData exhibit without noticing its potential at all (see Figure 5c: visitor watching world population clock instead). As soon as one person started interacting, others invariably followed.

As hoped and anticipated, collaboration and communication between visitors was widespread. In many cases, people would open lenses by themselves but communicate with each other at the same time, teaching each other how to do things or showing each other interesting info on the map. In



Figure 4. Future work: The GeoLens concept transferred to a different domain.

contrast to prior studies [10], it was common for more than one group to interact at one time.

We observed that about one third of the people interacting used only a single finger. One simple reason was that a lot of people were carrying things (backpacks, jackets), but others simply did not expect multi-touch functionality. A significant number of visitors belonging to the second group did not attempt to drag the lenses either, but simply opened new ones as necessary. Luckily, the application supported this type of interaction as well. Age-wise, we found that older users were just as capable of interacting with the application as younger ones. As was to be expected, younger ones experimented more and were mostly interested in the visual effects, while older people were more thoughtful in their actions and spent more time browsing the actual content.

One of the most important observations we made was that most people were actually able to access the content after a short period of exploration. Many visitors were genuinely interested in the content and spent some time moving lenses around, reading the text and generally learning about the actual subject matter. At first, this may not seem like a significant achievement. However, it is in contrast to several prior studies (e.g. [5, 11, 14]) of interactive surfaces in the wild, where the focus of the interaction was the technology and not the content.

CONCLUSIONS AND OUTLOOK

Our observations suggest the principal viability of the GeoLens approach for data exploration by untrained users - the principles of interaction seem to be immediately apparent to the majority, and access to a large amount of content is possible in this way. However, a full in-situ user study would be necessary to verify these findings and prompt further improvements.

The concept of GeoLenses would appear to generalize very well. Sliders, menus and buttons can be used to configure the content inside magic lenses in arbitrary ways. Geolocated data can be viewed, GIS layers selected and configured using controls arranged at the perimeter of the lens. The time slider in particular could be used in very different application settings as well - see Figure 4.




Figure 5. Observations in the train: (a) - (b) Visitors interacting with the table. (c) One person ignores the interactive table, watching the world population clock instead.

Further expansion of the GeoLens concept seems possible. While we enable more fine-grained control of the visualization in this particular application, further use cases as well as usability studies in a controlled environment would allow statements regarding the general feasibility of the approach. The combination of GeoLenses with zooming lenses is another promising research direction. Also, the current solution simply prevents lenses from overlapping; the ability to merge lenses could aid cooperation.

REFERENCES

1. E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose. Toolglass and magic lenses: the see-through interface. In *Proc. of SIGGRAPH '93*, pages 73–80. ACM, 1993.
2. H. Brignull and Y. Rogers. Enticing people to interact with large public displays in public spaces. In *Proc. of INTERACT '03*, pages 17–24. IOS Press, 2003.
3. M. S. T. Carpendale, D. J. Cowperthwaite, and F. D. Fracchia. 3-dimensional pliable surfaces: for the effective presentation of visual information. In *Proc. of UIST '95*, pages 217–226. ACM, 1995.
4. S. Carpendale, J. Ligh, and E. Pattison. Achieving higher magnification in context. In *Proc. of UIST '04*, pages 71–80. ACM, 2004.
5. F. Daiber, J. Schöning, and A. Krüger. Whole body interaction with geospatial data. In *Proc. of Smart Graphics '09*, pages 81–92. Springer, 2009.
6. P. Dietz and D. Leigh. DiamondTouch: A multi-user touch technology. In *Proc. of UIST '01*, pages 219–226. ACM, 2001.
7. C. Forlines and C. Shen. Dtlens: Multi-user tabletop spatial data exploration. In *Proc. of UIST '05*, pages 119–122. ACM, 2005.
8. M. Furuichi, Y. Mihori, F. Muraoka, A. Estenther, and K. Ryall. Dtmapp demo: Interactive tabletop maps for ubiquitous computing. In *UbiComp*, 2005.
9. E. Hall. *Distances in Man: The Hidden Dimension*. Double Day, 1966.
10. U. Hinrichs, H. Schmidt, and S. Carpendale. Emdialog: Bringing information visualization into the museum. *IEEE Transactions on Visualization and Computer Graphics*, 14:1181–1188, November 2008.
11. E. Hornecker. “I don’t understand it either, but it is cool” - visitor interactions with a multi-touch table in a museum. In *Proc. of Tabletop*, pages 113–120. IEEE, 2008.
12. E. Hornecker and M. Stifter. Learning from interactive museum installations about interaction design for public settings. In *Proceedings of the 18th Australia conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments, OZCHI '06*, pages 135–142, New York, NY, USA, 2006. ACM.
13. National Geographic Deutschland. *NATIONAL GEOGRAPHIC Collector's Edition No. 9: Planet Erde 2008*. 2008.
14. P. Peltonen, E. Kurvinen, A. Salovaara, G. Jacucci, T. Ilmonen, J. Evans, A. Oulasvirta, and P. Saarikko. It’s mine, don’t touch!: interactions at a large multi-touch display in a city centre. In *Proc. of CHI '08*, pages 1285–1294. ACM, 2008.
15. S. D. Scott, K. D. Grant, and R. L. Mandryk. System guidelines for co-located, collaborative work on a tabletop display. In *Proc. of ECSCW'03*, pages 159–178. Kluwer Academic Publishers, 2003.
16. J. C. Tang. Findings from observational studies of collaborative work. *Int. J. Man-Mach. Stud.*, 34(2):143–160, 1991.
17. E. Tse, C. Shen, S. Greenberg, and C. Forlines. Enabling interaction with single user applications through speech and gestures on a multi-user tabletop. In *Proc. of AVI '06*, pages 336–343. ACM, 2006.
18. U. von Zadow, F. Daiber, J. Schöning, and A. Krüger. Globaldata: multi-user interaction with geographic information systems on interactive surfaces. In *Proc. of ITS '10*, pages 318–318. ACM, 2010.
19. M. Wu and R. Balakrishnan. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. In *Proc. of UIST '03*, pages 193–202. ACM, 2003.

Interactive Multi-touch Surfaces for Software Visualization

Craig Anslow, Stuart Marshall, James Noble
School of Engineering and Computer Science
Victoria University of Wellington
Wellington, New Zealand
{craig, stuart, kjx}@ecs.vuw.ac.nz

Robert Biddle
School of Computer Science
Carleton University
Ottawa, Canada
robert_biddle@carleton.ca

ABSTRACT

Most software systems are developed by teams of people. The tools used to develop and maintain these systems are primarily designed from a single-user perspective and are bound to Integrated Development Environments (IDEs). These design decisions do not allow users to collaboratively navigate through software visualizations or to analyse software easily. We are investigating whether multi-touch table interaction techniques are more effective for co-located collaborative software visualization than existing single-user desktop interaction techniques. The implications of our research will help inform developers how to design better visualization applications for interactive multi-touch surfaces.

ACM Classification: H1.2 [User/Machine Systems]: Human Factors; H5.2 [Information interfaces and presentation]: User Interfaces. - Multi-touch user interfaces.

General terms: Experimentation, Human Factors.

Keywords: Multi-touch, software visualization, user study.

1. INTRODUCTION

Maintaining large software systems requires understanding the underlying structure, which is a hard task. Understanding software is often a social activity and involves teams of software developers. Software visualization aims to help with techniques to visualize the structure, behaviour, and evolution of software [4]. Visualization tools designed for a single user perspective make it hard for developers to analyse software when working together in a co-located environment (within the same room) using the same interface.

Multi-touch table user interfaces are an example of a co-located collaborative tool which could be used for software analysis. We are investigating whether multi-touch table interaction techniques are more effective for co-located collaborative software visualization than existing single user desktop interaction techniques. In this paper we discuss past work, our multi-touch software visualization prototype, a qualitative user study, challenges, and future work.

2. PAST

An approach to understand how participants use interactive multi-touch surfaces follows a qualitative research method. This method has been successfully adopted to provide insight into the design of tabletop displays for information visualization [9], visual analytics [10], and collaborative design [16].

Ko et al. have explored how software development tools support collaborative software understanding [11]. Storey et al. have explored collaborative software visualization [18]. These studies and very few tools have explored collaborative interactive multi-touch surfaces for software development or software visualization.

Parnin et al. suggest the addition of peripheral interactive spaces to programming environments for supporting developers in maintaining their concentration using touch based devices ranging from portable to tabletops [15]. Anslow et al. conducted a qualitative user study to find out the effectiveness of an established software visualization technique using a large visualization wall [1]. Boccuzzo et al. have made an extension to their 3D tool for software exploration with multi-touch features [2]. Hardy et al. created a visual system that uses digital pens and Wii Remotes for interaction to assist with software development processes such as what developers are working on, a summary of the architecture, and work flow activities [7]. Bott et al. created an interactive CRC card system that uses Wii Remotes for collaborative requirements engineering [3].

3. PRESENT

We are developing a prototype called SourceVis for large multi-touch tables to analyse the structure, evolution, and vocabulary of software systems, see Figure 1. The visualizations adapt existing information and software visualization techniques and modify them to support multi-touch and multi-user interaction. SourceVis is built upon MT4j [13].

3.1 Interaction

We envision developers working in groups (2–3) with our prototype to explore at a high level what parts of a software system are large, problematic, and need to be refactored.

Users first load a system by tapping on the menu to select a name of a system before any visualizations can be displayed. Users start a visualization by tapping on the icon of the visualization. Each visualization is displayed in a rotatable and scalable window with options to display at full screen or close. Multiple visualizations can be displayed at once.

Users can interact with individual elements in the visualizations. To select an element users tap, drag with one finger, and rotate and resize with two fingers. Double tapping an element displays properties about that element. Multiple elements can be grouped by drawing a shape around them using a lasso gesture and then subsequently move them around the




Figure 1: SourceVis. Users interacting with visualizations: A. Metrics Explorer, B. Class Blueprint, C. System Hotspots View, D. Wordle Vocabulary.

display. A tap and hold gesture displays a new visualization type. Zooming uses a pinch gesture with one or two hands and navigation by scrolling with two fingers.

For a multi-user scenario one developer might be looking at the overview of a system while another developer is looking at the details of a class. This allows users to orient visualizations to where they are standing and make some visualizations larger than others depending on the size of the system.

3.2 Visualizations

The visualizations are grouped into three categories: exploration, structure, and evolution. The exploration category contains visualizations that show metrics about a system, vocabulary employed in entities, and a touch enabled web browser for documentation. The structure category adapts Polymetric Views [12] to multi-touch. The evolution category shows how a system has evolved over time focusing on structural changes and developer revision histories. We now describe some exploration and structure visualizations.

Metrics Explorer (Annotated A. in Figure 1). Shows metrics about the different modules in a system such as the number of packages, classes, methods, and variables [6]. All the pack-

ages in a system are first displayed alphabetically. Tapping a package displays the metrics about the package and the classes it contains. Likewise tapping a class displays metrics, methods, and variables. In the figure a user has done a tap and hold gesture on one of the classes which has displayed the associated Class Blueprint which has partially obscured the Metrics Explorer.

Class Blueprint (B. in Figure 1). Shows the dependencies and references between methods and attributes within a class and adapted from Lanza et al. [12]. The visualization is broken into five layers. The first four layers relate to methods and the final layer to attributes. The methods and attributes have a different fill colour depending on which layer they belong to. Likewise the edges for the dependencies and references. The weight of an edge can be adjusted by moving a slider up (for thickest) or down (for invisible). In the figure a user has selected one of the interface methods by tapping and holding which has highlighted in green one accessor method that it calls and one attribute it accesses. With their other hand they have selected one of the attributes which has highlighted a method that references that attribute which is the same accessor method the interface method calls.

System Hotspots View (C. in Figure 1). Shows large packages and classes in a system and adapted from Lanza et al. [12]. Packages are displayed down the Y axis and classes from each package along the X axis. In the figure a user has double tapped on the package label which has displayed properties about the package in a linked yellow box. The properties include the total metrics, and options for visually sorting the classes in the package alphabetically, ascending, or descending by individual metrics or groups of the metrics. Each class is represented as a rectangle where the width indicates the number of variables and height number of methods. The colour of a class is represented as the number of lines of code. The darker the rectangle the more lines of code the class contains. Different border colours represent the type of class (e.g. red is interface, blue abstract class, no border a concrete class). In the figure a user has also double tapped on the large black class `java.awt.Component` which has displayed the class properties including metrics. Classes can be moved around the visualization to be compared with other classes and can be grouped together. A tap and hold gesture on a class displays the associated Class Blueprint.

Vocabulary (D. in Figure 1). The modified Word Cloud and Wordle provide a quick overview of the vocabulary used in the entities (e.g. packages, classes, methods) of a software system to understand the coding standards employed. In the figure a user has moved some words around, grouped some of the larger words together, and filtered out some of the smaller words. We intend to link this visualization with others so that selecting or grouping words highlights entities in other displayed visualizations that use these words.

3.3 User Study

We conducted a user study with 10 participants using SourceVis. The aim of the study was to collect data about how effective our software visualization techniques are for program comprehension in order to validate our interactive and visualization design decisions following a qualitative approach [9].

Procedure. Participants were given an information sheet, consent form, and a pre-study questionnaire. The questionnaire asked about their demographics and background experience. With each participant's consent we recorded their actions and asked them to think aloud. The study was conducted with a 48 inch low cost rear diffuse illuminated multi-touch table that we built, based on some existing designs [5]. Following the pre-study questionnaire, participants were given a warm up exercise by experimenting with the example applications from MT4j for five minutes. For the user tasks participants were asked 14 program comprehension questions similar to the types of questions software developers ask within industry [17]. The questions asked participants to identify, count, and find information within the same set of visualizations. The sample data set was the Java Standard API. Participants recorded their answers to the questions on a sheet attached to a clipboard. We recorded the time it took participants to complete the user tasks. Participants completed a post-study questionnaire which asked for their opinion on the effectiveness, strengths, and weaknesses of the interaction capabilities and the visualizations.

Participants. There were eight males and two females, who worked in pairs. One group of participants had known each other for 18 months while the other pairs knew each other for 12 months, 6 months, and 2 months. The other pair did not know each other. The age of participants was in the range of 25-29. All participants had a bachelors degree in computer science and three had a masters degree. Of the participants; one was currently an honours student (4th year undergraduate), three masters students, and six PhD students. Four participants had used some software visualization tools before but not on a frequent basis. Seven of the participants had used desktop touch screens or touch tables before. All had experience in programming using the Java Standard API.

Limitations. The small number of participants, who were a convenience sample, of graduate computer science students. The warm up exercise was some example applications not our visualizations. This was the first time participants had used our prototype before. The questions we asked were not numbered on the sheet provided nor did we vary the order, but all participants answered the questions in the order they were listed on the sheet. This may have led to a learning bias. When measuring how long it took for participants to answer the questions we had to account for the time participants spent thinking aloud and recording their answers on the sheet attached to the clipboard.

3.4 Results

Perceived Effectiveness. In the post-survey the Word Cloud ranked as the most effective technique followed by Metrics Explorer and Wordle. The System Hotspot Views and Class Blueprint ranked the same and slightly below the others.

Time and Errors. The first pair took 20 minutes to complete the user tasks, second pair 28 minutes, third pair 22 minutes, fourth pair 24 minutes, and fifth pair 21 minutes, for a mean average of 23 minutes. Pairs one, two, and five answered all the questions correctly, for a total of 36 (100%). Pair three received 34 (94%), and pair four 33 (92%).

Metrics Explorer. This visualization provided an overview and gave participants a clear summary of the metrics about a system. The white background made it easier to read the names of packages and classes, and colour to highlight the selected entities made them stand out. One participant would have liked to have seen inheritance information about classes. A couple of participants were not sure what to expect when they tapped on the name of an entity such as where information was going to be displayed. Nor did some of them know the best location for the metrics information, some suggested putting it next to the name of the entity instead of the left hand side of the visualization. Some packages contain many classes which required lots of scrolling to find a class, and some participants suggested adding a text-based search.

Class Blueprint. Participants liked how this visualization showed what methods and attributes were connected to each other. All participants commented that the highlighting of edges made it an effective way to answer questions when using this visualization. They also liked how multiple users could highlight more than one edge at a time. The slider that adjusted the weight of the edges was a welcomed addition, but some participants were not aware of it. Since the edges cross each other and methods this made it confusing for some participants to be able to read the names of methods. One participant suggested that if different methods were selected then only show the intersection if one exists.

System Hotspots View. This visualization made it easy for participants to compare the different entities in the system as packages were initially laid out in alphabetical order and classes grouped in packages, plus the ability to move entities around the visualization. Once the participants remembered the information cues (metrics and colour encoding) it was easy to identify certain aspects of a system such as the types of classes and large classes. The properties windows made it easy to determine precise information about a package or a class. The sorting provided a quick way to answer some of the identify and count questions. The example system used was large, participants found that it was hard to get an overview of all the information because when the visualization started it was half zoomed in. This meant lots of scrolling to find what they were looking for. Some classes in the visualization were small and required zooming in to validate their colour.

Vocabulary. All participants found the word font size and background colours made it easy to understand these visualizations. Long words made it slightly harder to interpret. It was hard to compare words if they were not next to each other as the absolute size is not easy to see and nor is it clear what metric was being used. Separating grouped words with a gesture was not easy to do. Adding colours to the different words helped to distinguish between them. Words can be overlapped which helped when comparing the size of two words together, but sometimes there was too much overlapping. A slider was added to filter out words, but it confused some participants as they suggested it was not intuitive.

3.5 Lessons Learnt

Support Collaboration. Seven of the participants stated that it was easier to work in groups than individuals because you can each look at several parts of the same visualization at once, divide the tasks up, and discuss the answers.

Display Multiple Visualizations. Most of the participants did not display multiple visualizations at once, even though they knew they could. They only did this when they navigated to a Class Blueprint. Some commented having different ways of looking at a system with multiple visualizations would be beneficial and suggested adding synchronization features.

Provide Visualization Help. Participants were not familiar with our visualizations, so providing help documentation such as how to interact with the visualization and what the encoding means would help improve usability.

Use Higher Resolution. The resolution of our touch screen was only 1280x800 pixels so having a higher resolution would make it easier to display more visualizations at once.

Display Radar View. It is important for users not to lose navigation context when visualizations are large and require lots of scrolling. Displaying a small radar view of the current context would help with navigation, but would come as an expense by taking up some screen real estate.

User Study. Make questions more subjective which will involve more exploration with the visualizations, use industry professionals as participants, conduct a study over a longer period of time and in a real world setting, and compare our prototype with a control visualization tool.

4. FUTURE

Challenges. Our visualizations are primarily viewed in isolation, we need some way to synchronize them together especially when some of them use the same underlying data source. Since our visualizations of the software are not displayed within an IDE it is important to be able to link the visualizations somehow with the underlying source code. Our focus has been to visualize the source code through exploration and navigation techniques, but some participants expressed that they wanted to enter text to do this and to program using the touch table. In order to support text entry we could adopt an existing text-entry method [8]; or to support programming we could scale a programming model approach for tablets that uses tiles and behaviour constructs [14].

Future Work. We intend to create a more comprehensive prototype and are currently working on evolution visualizations. Once our prototype is more mature we plan to conduct a large collaborative quantitative between subjects user experiment similar to Wettel et. al. involving industry professionals [19].

ACKNOWLEDGMENTS

This work is supported by the New Zealand Ministry of Science and Innovation for the Software Process and Product Improvement project, a Telstra Clear Postgraduate Scholarship, and the Canadian NSERC SurfNet Network.

REFERENCES

1. C. Anslow, J. Noble, S. Marshall, E. Tempero, and R. Biddle. User evaluation of polymetric views using a large visualization wall. In *Proc. of SoftVis*. ACM, 2010.
2. S. Boccuzzo and H. C. Gall. Multi-touch collaboration for software exploration. In *Proc. of International Conference on Program Comprehension (ICPC)*. IEEE, 2010.
3. F. Bott, S. Diehl, and R. Lutz. CREWW: collaborative requirements engineering with Wii-remotes (NIER track). In *Proc. of ICSE*. ACM, 2011.
4. S. Diehl. *Software Visualization: Visualizing the Structure, Behaviour, and Evolution of Software*. Springer Verlag, 2007.
5. J. Schöning et al. *Tabletops - Horizontal Interactive Displays*, chapter Building Interactive Multi-touch Surfaces, pages 27–49. Springer Verlag, 2010.
6. N. E. Fenton and S. L. Pfleeger. *Software Metrics: A Rigorous and Practical Approach*. PWS Publishing, 1998.
7. J. Hardy, C. Bull, G. Kotonya, and J. Whittle. Digitally annexing desk space for software development (NIER track). In *Proc. of ICSE*. ACM, 2011.
8. U. Hinrichs, M. Hancock, C. Collins, and S. Carpendale. Examination of text-entry methods for tabletop displays. In *Proc. of the Workshop on Horizontal Interactive Human Computer Systems (TABLETOP)*. IEEE, 2007.
9. P. Isenberg. *Collaborative Information Visualization in Co-located Environments*. PhD thesis, Uni of Calgary, 2009.
10. P. Isenberg, D. Fisher, M. Morris, K. Inkpen, and M. Czerwinski. An exploratory study of co-located collaborative visual analytics around a tabletop display. In *Proc. of VAST*. IEEE, 2010.
11. A. Ko, R. DeLine, and G. Venolia. Information needs in collocated software development teams. In *Proc. of ICSE*. IEEE, 2007.
12. M. Lanza and S. Ducasse. Polymetric views—a lightweight visual approach to reverse engineering. *IEEE Trans. on Soft. Eng.*, 29(9):782–795, 2003.
13. U. Laufs, C. Ruff, and J. Zibuschka. Mt4j: a cross-platform multi-touch development framework. In *Proc. of the Workshop on Engineering Patterns for Multi-Touch Interfaces at Symposium Engineering Interactive Computing Systems (EICS)*. ACM, 2010.
14. S. McDirmid. Coding at the speed of touch. In *Proc. of Symposium on New Ideas in Programming and Reflections on Software (Onward!)*. ACM, 2011.
15. C. Parnin, C. Görg, and S. Rugaber. Codepad: interactive spaces for maintaining concentration in programming environments. In *Proc. of SoftVis*. ACM, 2010.
16. S. Scott, S., Carpendale, and K. Inkpen. Territoriality in collaborative tabletop workspaces. In *Proc. of CSCW*. ACM, 2004.
17. J. Sillito, G. Murphy, and K. De Volder. Questions programmers ask during software evolution tasks. In *Proc. of FSE*. ACM, 2006.
18. M-A. Storey, C. Bennett, I. Bull, and D. German. Remixing visualization to support collaboration in software maintenance. In *Proc. of ICSM*. IEEE, 2008.
19. R. Wettel, M. Lanza, and R. Robbes. Software systems as cities: A controlled experiment. In *Proc. of ICSE*. ACM, 2011.

Position Paper: Touch Interaction in Scientific Visualization

Tobias Isenberg

DIGITEO/CNRS/INRIA, France, and University of Groningen, the Netherlands; isenberg@cs.rug.nl

ABSTRACT

Direct-touch interaction is receiving an increasing amount of attention in the HCI domain and has recently been applied to some problems in information visualization. However, the field of scientific visualization has not seen much work on how direct-touch interaction could benefit the interactive visualization process. One reason may be that scientific visualization typically deals with datasets that are defined in 3D space, while touch input is two-dimensional. Therefore, to control scientific datasets in interactive visualizations one has to define intuitive mappings from 2D input to 3D manipulations for a variety of data types and exploration techniques. In this position paper we discuss some of the challenges of using direct-touch interaction in the field of scientific visualization that arise from the specific constraints of scientific visualization and argue for the development of integrated techniques to overcome these problems.

Keywords: Direct-touch interaction, scientific visualization.

INTRODUCTION

Scientific visualization is the science of creating graphical representations of scientific data that is the basis of research in virtually all domains of science. Scientific visualization (as opposed to information visualization) deals predominantly with data that is spatially explicit in 3D, i.e., for each data point a precise 3D location is known. Such data includes CT/MRI scans in medicine; particle simulations in physics, astronomy, or swarm behavior; molecular models in genetics, biology, chemistry, or material sciences, or shapes of functions or sets in mathematics, to name but a few examples.

It is a recognized fact that a good visualization needs to support user activities beyond viewing the data [35]. Scientists need to be able to drill down and find details about what seems important, relate information on many levels of granularity, and gain an encompassing picture about relationships and correlations present in their data to form hypotheses and plan next step actions. To support interactivity in scientific visualization two main aspects require dedicated research attention: (a) *interaction models for control and exploration* of the visualization to support problem solving and (b) *interactive rendering speeds* to achieve real-time refresh rates for minimal disruption of higher-level tasks. Many new techniques have been and are being developed based on GPU processing which address the second challenge successfully. The challenge of providing novel solutions for interactive exploration, navigation, and problem solving with three-dimensional scientific visualizations [21] is one that has received less attention in the field of scientific visualization to date. We therefore argue in this paper that we need to embrace emerging interactive display technology and inves-

tigate its use for creating engaging and intuitive interactive next-generation scientific visualization work environments.

DIRECT-TOUCH INTERACTION

In particular, touch-sensitive displays have numerous important but unexplored benefits for scientific visualization: they provide *enough area and resolution* to explore visualizations, they facilitate *awareness* in *collaborative settings* [17], and they offer *natural direct-touch interaction* which provides *somesthetic information and feedback* beneficial for effective interaction both in real and virtual environments [28], and the direct-manipulation metaphor of touch interaction allows people feel *in control of their data* [19] and has shown to *outperform mouse input* for specific tasks [22]. Despite these advantages, however, to date there is little support for interactive scientific visualization on these large displays. While touch interaction has previously been explored in a general visualization context (e.g., [9, 10, 18, 26]) much remains to be learned about *employing direct-touch input specifically for three-dimensional scientific visualization*.

CONSTRAINTS OF SCIENTIFIC VISUALIZATION

Traditionally, PC-based environments or dedicated hardware setups [2] such as virtual reality environments (e.g., the Responsive Workbench [23] or the CAVE [3]) have often been employed for scientific visualization. While such settings have numerous advantages for creating and viewing visualizations, both have disadvantages for the interaction with 3D data. In PC settings, for instance, one typically interacts indirectly through a mouse. In VR environments one can interact directly by means of tracked objects in physical space (e.g., wands) but this type of control often leaves viewers with the feeling of interaction in “empty space” without adequate haptic feedback or rest positions. Touch interaction, in contrast, provides direct control (in 2D) with somesthetic feedback [28] which can alleviate this problem by allowing users to feel “in control of the data.” Unlike manipulating 2D data, however, touch interaction with 3D data or within 3D environments is a challenging task [30] because it requires an under-constrained mapping from 2D input parameters (touch) to transformations in 3D space (mouse-based interaction faces the same problem). To address this issue of 2D-to-3D mapping, several interaction techniques have been described in the past [5, 6, 7, 8, 11, 13, 14, 15, 20, 24, 25, 27, 36, 37, 38] and from which we can take inspiration. Yet, scientific visualization has a number of unique characteristics which make dedicated research necessary.

Only few of the previously named approaches [5, 11, 38] deal with the specific *needs and requirements of scientific visualization*: First, scientific visualization has broad applicability and developed techniques cannot be specific to one

type of digital object or a specific geometric projection. Virtually all of the above referenced methods address the 3D touch interaction problem only for a specific type of object (e.g., medium sized closed shape, planar representation, or particle cloud) in a specific environment (e.g., top-down projected space with ground plane or object(s) freely floating in space). Our goal is to develop and encompassing interaction landscape for scientific visualization which supports multiple different interaction methodologies in an environment that can change depending on the type of data being visualized. For example, when interacting with a volumetric dataset it is often necessary to independently control one or more cutting planes (i.e., planar objects) that are used to reveal the inside of the dataset together with an interaction of particle layer data (i.e., many points too small to use their surface to constrain an interaction). Secondly, for scientific visualization it is essential that interaction can be controlled precisely in 3D for space-accurate exploration. Many existing techniques lack the necessary precise control (e.g., [37]) for scientific exploration and even the ones with good control are still subject to the inherent imprecision of direct-touch input. Lastly, scientific visualization interaction includes many dedicated actions beyond general 3D navigation and object manipulation. These additional interactions are essential to scientific data exploration and include selection, object and parameter manipulation, interaction with the time axis, etc.

FUTURE RESEARCH DIRECTIONS

As a consequence of these unique constraints, we suggest a number of research directions for facilitating the exploratory interaction with three-dimensional scientific datasets using touch-sensitive displays.

Integrated Direct-Touch Interaction Toolkit

First, we plan to develop a toolkit for three-dimensional direct-touch interaction with scientific data. This toolkit needs to comprise a set of integrated techniques and methods to support two main visualization feedback loops: the *data manipulation loop* and the *exploration and navigation loop* [35]. The toolkit's purpose, therefore, is to make the 3D interaction techniques readily accessible for a variety of scientific visualization applications and also to outside researchers. The *data manipulation loop* for visualization application includes basic data interactions such as selection and positioning of objects in space. While these are basic operations, they are fundamental to many follow-up interactions in the navigation and exploration loop. For example, data representations need to be found, selected, and possibly positioned before they can be effectively compared or correlated. Interactions for the *navigation and exploration loop* are complex as they need to encompass theories of pathfinding, map use, spatial metaphors, awareness, and feedback [35].

The development of techniques dedicated to scientific visualization for both loops needs to start by developing data manipulation methods for general view changes (i.e., camera or projection manipulations) not only for the visualized data but also for dedicated data exploration elements such as cutting planes or drilling tools. Therefore, it is necessary to support a catalog of interaction needs for scientific visualization—abstracted across datasets and tasks. For this purpose it will

be necessary to analyze several existing scientific visualization tools from various application domains and data types (e.g., brain visualization, astronomic particle visualization, fluid flow simulation) and identify their most fundamental and common interaction requirements. The resulting toolkit may not need to completely re-invent new 3D interaction techniques but may incorporate some of the previously developed approaches for direct-touch interaction with 3D objects [5, 6, 7, 8, 11, 13, 14, 15, 20, 24, 25, 27, 36, 37, 38]. However, care must be taken to make the different modes of manipulation compatible with each other.

In a second stage, it will be necessary to integrate methods for the exploration and navigation loop. For this purpose one first needs to add selection strategies that are compatible with the view selection techniques. The selection techniques also need to go beyond the common tap-to-select because scientific datasets can comprise a variety of different data types (e.g., volumetric data, particle data, line data, or surface data). Based on the ability to select data and/or subspaces, mechanisms for the manipulation of selected objects (relocate, reorient, or resize), for specification of parameters (e.g., transfer function manipulation, placing seed particles, etc.), for interaction with the scale of the displayed dataset (potentially across several magnitudes of scale), and many others need to be integrated. Moreover, domain-specific interaction-techniques need to be supported, e.g., specific ones for geological data [31, 32]. Similar to the constraint for the selection techniques, also the data manipulation techniques need to be compatible with the remaining techniques of the toolkit. One of the major challenges, therefore, will be to provide the set of interaction techniques in an integrated manner such that they do not negatively affect each other.

Precise Control Issues

An important additional challenge that arises when employing direct-touch interaction is that touch input is inherently imprecise due to the size of our fingers as interaction tools—while scientific visualization often comes with a requirement of precise location and control of 3D data. Here, two aspects of precise control play an important role. The first aspect is the translation of imprecise touch input into control of similar precision as the mouse. Here, we can learn from HCI research which in the past has developed several strategies to provide such precision (e.g., [1]). The second aspect of precise control arises from scientific visualization's need to single out specific parameters and to control them without affecting others. This aspect implies that, in addition to fully integrated interactions, we need to also support partial interactions. For example, instead of only the known pinching interaction (RST, [16, 27]), scientific visualization needs support for navigation along or rotation around a single axis. Therefore, an Interaction Toolkit needs to support techniques that allow users to single-out certain parameters.

Stereoscopic Displays

Additional challenges arise when one wants to retain the benefits of direct-touch control but, at the same time, wants to take advantage of the improved depth perception provided by traditional dedicated visualization environments with stereoscopic displays. Research has shown that touch interaction

with stereoscopic displays is challenging because it is strongly affected by parallax between the images displayed for both eyes [33]. Only when a virtual object is within a small distance from the screen can people perceive their touches as affecting the virtual objects [34]. Moreover, when viewing an object displayed as reaching out of a display people often attempt to touch “in thin air.” In contrast—if an object is displayed “below the surface”—people may not even perceive the display surface as being present and hit the (to them invisible) display in an attempt to touch the object behind it. While solutions such as transparent props [29], separating the touch surface from the stereoscopic display [12], or tilted setups in connection with shallow-depth data [4] can be used to alleviate this problem somewhat, these solutions lack support for dynamic visualization elements or the diverse character of scientific data in general.

Large Displays and Multi-User Settings

The large size of traditional visualization displays prompts the question of multi-user visualization settings, in particular if interaction techniques based on multi-touch input are used. Not only do certain 3D touch interaction techniques not scale well to large settings (e.g., [38]) or are not compatible with vertical display setups (e.g., [37]), the possibility for several people using the same visualization environment simultaneously raises additional questions. For example, one could envision interactive discussions between colleagues, the use of touch-controlled visualization in group discussion, and the interactive touch-based presentation of visualizations to a larger audience. Yet, it is not simple to extend single-user interaction techniques to multi-user ones. For example, a single-user 3D exploration method being applied on a large wall display may not be suitable for a presentation setting since the interacting person is largely occluding the interaction and visualization space. In such situations it may be necessary to separate the interaction from the visualization space and to employ dedicated awareness techniques. An additional important challenge is that current multi-touch display technology does not support tracking of user identity. Without user identity one has to develop heuristics to determine which user issued a certain interface command and react to synchronous input accordingly. Multiple concurrent changes to an object require addressing computational challenges and, more importantly, how conflicts are handled.

CONCLUSION

We believe that—despite the discussed issues/challenges—touch interaction can have a tremendous impact on how visualization is being used by domain scientists (and beyond). Direct-touch interaction has the potential to facilitate the use of scientific visualization on a much larger variety of display and user settings, instead of being restricted to largely single-user, mouse/keyboard-based interaction in PC environments or specialized 3D visualization hardware. Thus, instead of only being the end product of a scientific exploration process, intuitive touch-based interactive visualization technology can be tightly integrated into the scientific exploration process, and could actively be used for gaining an understanding of the analyzed data. This means that scientists may be able to use 3D interaction techniques to not only discuss ideas but instead to collaboratively create and manipu-

late visuals that illustrate their data, resulting in fundamental insights and an easy way to communicate these to others.

REFERENCES

1. H. Benko, A. D. Wilson, and P. Baudisch. Precise Selection Techniques for Multi-Touch Screens. In *Proc. CHI*, pp. 1263–1272, New York, 2006. ACM Press. doi> 10.1145/1124772.1124963
2. D. A. Bowman, E. Kruijff, J. J. LaViola, Jr., and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley, Boston, 2005.
3. S. Bryson. Virtual Reality in Scientific Visualization. *Communications of the ACM*, 39(5):62–71, May 1996. doi> 10.1145/229459.229467
4. T. Butkiewicz and C. Ware. Exploratory Analysis of Ocean Flow Models with Stereoscopic Multi-Touch. In *IEEE Visualization Posters*, 2011.
5. D. Coffey, N. Malbraaten, T. Le, I. Borazjani, F. Sotiropoulos, and D. F. Keefe. Slice WIM: A Multi-Surface, Multi-Touch Interface for Overview+Detail Exploration of Volume Datasets in Virtual Reality. In *Proc. I3D*, pp. 191–198, New York, 2011. ACM. doi> 10.1145/1944745.1944777
6. A. Cohé, F. Dècle, and M. Hachet. tBox: A 3D Transformation Widget Designed for Touch-Screens. In *Proc. CHI*, pp. 3005–3008, New York, 2011. ACM. doi> 10.1145/1978942.1979387
7. J. Edelmann, S. Fleck, and A. Schilling. The DabR – A Multitouch System for Intuitive 3D Scene Navigation. In *Proc. 3DTV*, Piscataway, NJ, USA, 2009. IEEE. doi> 10.1109/3DTV.2009.5069671
8. C. Forlines and R. Lilien. Adapting a Single-User, Single-Display Molecular Visualization Application for Use in a Multi-User, Multi-Display Environment. In *Proc. AVI*, pp. 367–371, New York, 2008. ACM. doi> 10.1145/1385569.1385635
9. C. Forlines and C. Shen. DTLens: Multi-User Tabletop Spatial Data Exploration. In *Proc. UIST*, pp. 119–122, New York, 2005. ACM. doi> 10.1145/1095034.1095055
10. M. Frisch, J. Heydekorn, and R. Dachsel. Investigating Multi-Touch and Pen Gestures for Diagram Editing on Interactive Surfaces. In *Proc. ITS*, pp. 149–156, New York, 2009. ACM. doi> 10.1145/1731903.1731933
11. C.-W. Fu, W.-B. Goh, and J. A. Ng. Multi-Touch Techniques for Exploring Large-Scale 3D Astrophysical Simulations. In *Proc. CHI*, pp. 2213–2222, New York, 2010. ACM. doi> 10.1145/1753326.1753661
12. M. Hachet, B. Bossavit, A. Cohé, and J.-B. de la Rivière. Toucheo: Multitouch and Stereo Combined in a Seamless Workspace. In *Proc. UIST*, pp. 587–592, New York, 2011. ACM. doi> 10.1145/2047196.2047273
13. M. Hachet, F. Decle, S. Knödel, and P. Guitton. Navidget for Easy 3D Camera Positioning from 2D Inputs. In *Proc. 3DUI*, pp. 83–89, Los Alamitos, 2008. IEEE Computer Society. doi> 10.1109/3DUI.2008.4476596

14. M. Hancock, S. Carpendale, and A. Cockburn. Shallow-Depth 3D Interaction: Design and Evaluation of One-, Two- and Three-Touch Techniques. In *Proc. CHI*, pp. 1147–1156, New York, 2007. ACM. doi> 10.1145/1240624.1240798
15. M. Hancock, T. ten Cate, and S. Carpendale. Sticky Tools: Full 6DOF Force-Based Interaction for Multi-Touch Tables. In *Proc. ITS*, pp. 145–152, New York, 2009. ACM. doi> 10.1145/1731903.1731930
16. M. S. Hancock, S. Carpendale, F. D. Vernier, D. Wigdor, and C. Shen. Rotation and Translation Mechanisms for Tabletop Interaction. In *Proc. Tabletop*, pp. 79–88, Los Alamitos, 2006. IEEE Computer Society. doi> 10.1109/TABLETOP.2006.26
17. E. Hornecker, P. Marshall, N. S. Dalton, and Y. Rogers. Collaboration and Interference: Awareness with Mice or Touch Input. In *Proc. CSCW*, pp. 167–176, New York, 2008. ACM. doi> 10.1145/1460563.1460589
18. T. Isenberg, M. Everts, J. Grubert, and S. Carpendale. Interactive Exploratory Visualization of 2D Vector Fields. *Computer Graphics Forum*, 27(3):983–990, May 2008. doi> 10.1111/j.1467-8659.2008.01233.x
19. T. Isenberg, U. Hinrichs, and S. Carpendale. Studying Direct-Touch Interaction for 2D Flow Visualization. In *Proc. CoVIS*, pp. 17–20. Univ. of Munich, Germany, Depart. of Computer Science, 2009. Published in April 2010 in technical report # LMU-MI-2010-2.
20. Y. Jung, J. Keil, J. Behr, S. Webel, M. Zöllner, T. Engelke, H. Wuest, and M. Becker. Adapting X3D for Multi-Touch Environments. In *Proc. Web3D*, pp. 27–30, New York, 2008. ACM. doi> 10.1145/1394209.1394218
21. D. F. Keefe. Integrating Visualization and Interaction Research to Improve Scientific Workflows. *IEEE Computer Graphics and Applications*, 30(2):8–13, Mar./Apr. 2010. doi> 10.1109/MCG.2010.30
22. K. Kin, M. Agrawala, and T. DeRose. Determining the Benefits of Direct-Touch, Bimanual, and Multifinger Input on a Multitouch Workstation. In *Proc. Graphics Interface*, pp. 119–124, Toronto, 2009. CIPS.
23. W. Krüger and B. Fröhlich. The Responsive Workbench. *IEEE Computer Graphics and Applications*, 14(3):12–15, May 1994. doi> 10.1109/38.279036
24. A. Martinet, G. Casiez, and L. Grisoni. 3D Positioning Techniques for Multi-Touch Displays. In *Proc. VRST*, pp. 227–228, New York, 2009. ACM. doi> 10.1145/1643928.1643978
25. A. Martinet, G. Casiez, and L. Grisoni. Design and Evaluation of 3D Positioning Techniques for Multi-touch Displays. Technical Report RR-7015, INRIA, France, 2009.
26. C. North, T. Dwyer, B. Lee, D. Fisher, P. Isenberg, K. Inkpen, and G. Robertson. Understanding Multi-touch Manipulation for Surface Computing. In *Proc. Interact*, pp. 236–249, Berlin/Heidelberg, 2009. Springer Verlag. doi> 10.1007/978-3-642-03658-3_31
27. J. L. Reisman, P. L. Davidson, and J. Y. Han. A Screen-Space Formulation for 2D and 3D Direct Manipulation. In *Proc. UIST*, pp. 69–78, New York, 2009. ACM. doi> 10.1145/1622176.1622190
28. G. Robles-De-La-Torre. The Importance of the Sense of Touch in Virtual and Real Environments. *IEEE MultiMedia*, 13(3):24–30, July–Sept. 2006. doi> 10.1109/MMUL.2006.69
29. D. Schmalstieg, L. M. Encarnação, and Z. Szalavári. Using Transparent Props for Interaction with the Virtual table. In *Proc. I3D*, pp. 147–153, New York, 1999. ACM. doi> 10.1145/300523.300542
30. F. Steinicke, K. H. Hinrichs, J. Schöning, and A. Krüger. Multi-Touching 3D Data: Towards Direct Interaction in Stereoscopic Display Environments coupled with Mobile Devices. In *Proc. AVI Workshop on Designing Multi-Touch Interaction Techniques for Coupled Public and Private Displays*, pp. 46–49, 2008.
31. N. Sultanum, E. Sharlin, M. C. Sousa, D. N. Miranda-Filho, and R. Eastick. Touching the Depths: Introducing Tabletop Interaction to Reservoir Engineering. In *Proc. ITS*, pp. 105–108, New York, 2010. ACM. doi> 10.1145/1936652.1936671
32. N. Sultanum, S. Somanath, E. Sharlin, and M. C. Sousa. “Point it, Split it, Peel it, View it”: Techniques for Interactive Reservoir Visualization on Tabletops. In *Proc. ITS*, pp. 192–201, New York, 2011. ACM. doi> 10.1145/2076354.2076390
33. D. Valkov, F. Steinicke, G. Bruder, and K. Hinrichs. 2D Touching of 3D Stereoscopic Objects. In *Proc. CHI*, pp. 1353–1362, New York, 2011. ACM. doi> 10.1145/1978942.1979142
34. D. Valkov, F. Steinicke, G. Bruder, K. H. Hinrichs, J. Schöning, F. Daiber, and A. Krüger. Touching Floating Objects in Projection-based Virtual Reality Environments. In *Proc. EGVE/EuroVR/VEC*, pp. 17–24, Goslar, Germany, 2010. Eurographics Association. doi> 10.2312/EGVE/JVRC10/017-024
35. C. Ware. *Information Visualization – Perception for Design*. Morgan Kaufmann Series in Interactive Technologies. Morgan Kaufmann Publishers, Amsterdam, Boston, Heidelberg, and others, 2nd edition, 2000.
36. A. D. Wilson. Simulating Grasping Behavior on an Imaging Interactive Surface. In *Proc. ITS*, pp. 137–144, New York, 2009. ACM. doi> 10.1145/1731903.1731929
37. A. D. Wilson, S. Izadi, O. Hilliges, A. Garcia-Mendoza, and D. Kirk. Bringing Physics to the Surface. In *Proc. UIST*, pp. 67–76, New York, 2008. ACM. doi> 10.1145/1449715.1449728
38. L. Yu, P. Svetachov, P. Isenberg, M. H. Everts, and T. Isenberg. FI3D: Direct-Touch Interaction for the Exploration of 3D Scientific Visualization Spaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1613–1622, Nov./Dec. 2010. doi> 10.1109/TVCG.2010.157

Mobile Interactive Displays for Medical Visualization

Steven Birr¹, Raimund Dachzelt², Bernhard Preim¹

¹ Visualization Group, ² User Interface & Software Engineering Group
Otto-von-Guericke-University Magdeburg,
Germany

steven.birr@ovgu.de, dachzelt@ovgu.de, preim@ovgu.de

ABSTRACT

Medical visualization proves to be essential in modern health care systems. Collaborative diagnosis and treatment decisions as well as medical education and training can be enhanced using interactive 2D and 3D graphics. However, besides developing effective visualizations, the challenge remains on how to support optimal interaction especially on mobile devices. This paper presents possible future scenarios of using mobile displays by different medical experts in collaborative environments. Typical tasks, technical requirements, open research issues and challenges on interacting with mobile medical visualization are discussed.

General terms: Design, Human Factors

Keywords: Medical visualization, mobile displays, interaction techniques, multitouch, collaboration

MOTIVATION

Medical visualization techniques play an essential role in current and future health care systems. Main application areas are diagnosis, treatment planning, intraoperative support and education [1]. Amongst others, interactive 2D and 3D visualizations are used for the diagnosis of diseases. Such visualizations might include measurements (extend of pathologic structures) and annotations like encircled regions or arrows to enhance their interpretation. Visualizations are generated to be discussed among colleagues and to enhance the planning of surgical interventions. Therefore virtual 3D models are very useful since they provide an overview of the morphology. Spatial relations between pathologic lesions and structures at risk may be evaluated better with 3D visualizations (Figure 1). For example, Krüger et al. presented a medical visualization method for neck dissection planning [2]. Training systems like the *LiverSurgeryTrainer* [3] (Figure 2) were developed for computer-assisted, preoperative planning of liver tumor resections based on individual patient anatomy and pathology.

VISION

A challenging task, not only for medical visualization, is to provide efficient input devices and easy-to-learn, yet effective interaction techniques for exploring and editing medical visualizations. Nowadays, a multitude of heterogeneous information is being generated from several medical doctors in different situations and usage contexts.




Figure 1: 3D surface visualization of relevant anatomical structures for neck surgery planning. Left: The desired object, an enlarged lymph node, is highlighted. Right: The distance between an enlarged and potentially malignant lymph node to a muscle is color-coded to support the decision whether the muscle should be removed or not [2].

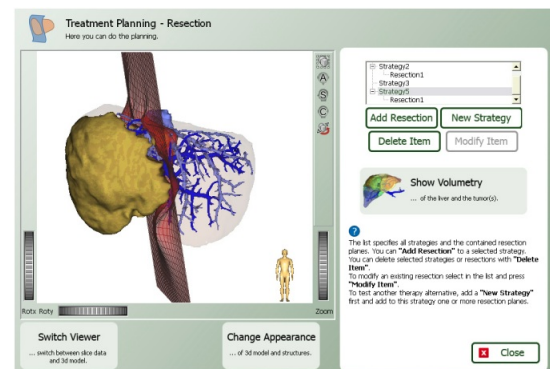


Figure 2: Liver surgery treatment planning with the training software *LiverSurgeryTrainer* [3]. A deformable resection plane can be controlled and modified interactively by the user. The challenging goal is the exact resection of the tumor and saving the functionality of the remnant liver parenchyma.

Radiologists generate digital reports during their diagnosis; surgeons use these reports and additional image data to prepare surgical interventions. Moreover, surgeons generate new information, such as intraoperative findings, that need to be documented as fast and as comfortable as possible. Nevertheless, personal conversations between medical experts are very essential in order to clarify ambiguous medi-

cal reports or image details. Current multi-touch smartphones and tablets are very promising for these cooperative clinical workflows. Together with wireless connectivity they enable access to relevant patient information at different places, e.g. at the bedside of the patient. Portable medical systems enable the location-independent selection of laboratory data and radiological images, zooming in selected data, specifying measurements or entering digital notes into a hospital information system. Often, information is conveyed from doctor to patient in a verbal manner which is sometimes misunderstood because of the jargons and medical terms used. Hence, portable devices are highly welcome by medical doctors since they enable explaining surgical strategies using movies, animations or 3D renderings. Moreover, gesture- and touch-based interfaces are considered very attractive by a large majority of medical experts. The problem however remains on how to design effective and intuitive interaction techniques for daily use.

Modern smartphones and tablets are fast enough to visualize large medical datasets and interactive 3D elements directly on the device. Single and multitouch gestures are the typical mode of interaction for these devices. Several gestures already exist for the exploration of data. For example, pinch gestures are already established for zooming images. Nevertheless, the design of appropriate gestures for interacting with more complex visualizations is a challenging task. A gesture set should be carefully limited and the gestures should be unique enough to avoid misinterpretation with similar gestures (cf. [4]).

In this article, several future scenarios and research topics concerning medical visualization on mobile displays are sketched and discussed. Typical tasks and requirements as well as possible solutions for exploring medical data on mobile devices are presented in the following section. Furthermore, we want to raise open research topics in this special and important application domain. In our research agenda we want to reach three main goals:

- The enhancement of patient interviews with mobile multimedia material as a visual support on mobile devices.
- The support of medical education and training with interactive 3D visualizations on mobile devices.
- The support of collaborative workflows using multiple mobile and large displays.

EXAMPLE SCENARIOS

Scenario 1: Patient Interview

A typical task in clinical routine is the doctor-patient-conversation. During that ward round, clinical values (e.g. blood pressure, pulse, temperature etc.) are monitored and documented, often by handwritten notes, in paper-based medical records. Furthermore, the patient has the chance to ask questions, e.g. concerning surgical procedure and possible complications. These tasks can be enhanced by using mobile devices. Additional information could be entered in a digital patient record on a smartphone or tablet.




Figure 3: Virtual 3D model of a liver surface with annotation labels for medical education purpose. Labels indicate anatomical elements of the liver, like metastasis and the parts of the liver portal vein.

Thus, misunderstandings due to unreadable handwritten notes are avoided. Further, the digital input enables automatically generated diagrams based on the stored values. These diagrams are useful to visually gain insight into the abstract data and to recognize vital trends. Moreover, multimedia elements and animations can be used to explain certain activities planned for the surgery, e.g. resecting a tumor or stitching vessels. Furthermore, the talk can be enhanced by using links to websites providing additional textual information, high-quality movies and animated 3D content.

An appropriately designed multimedia presentation can be very useful for visually explaining the steps during surgery and for improving the patient's trust. An important requirement for that scenario is a lightweight mobile device and simple interaction techniques supporting a clear and comprehensible patient interview using imagery and schematic graphics without too much or realistic detail.

Scenario 2: Medical Education and Training

Nowadays, human anatomy teaching is often based on static schoolbooks or on lectures where anatomical facts and concepts are transferred from one teacher to many students. But lectures and books poorly convey the three-dimensional nature of anatomical structures [5]. New e-learning portals are trying to fill the gap by providing Web-based 3D multimedia contents combined with social Web 2.0 techniques.

We are trying to envision the following scenario: A medical student wants to test his/her knowledge about human anatomy and different pathologies. The student opens a website, which allows access to a case-based medical education system. It provides different patient-individual datasets with various kinds of vessel anomalies and tumors/metastasis. The Web application can be used to get and train knowledge about vessel systems and territories, different diseases and their effects as well as various resection methods. The mobile application enables exploring medical 2D and 3D data by using single- and multi-touch gestures. For example, a single-touch gesture could be used for scrolling up/down in the image stack. Multi-touch input can be used

to interactively zoom and pan the imagery. Pan is controlled by dragging two fingers in the multi-touch interface, and the zoom level is adjusted with two-finger pinching. Furthermore, medical 3D models can be rotated, translated and zoomed freely by the user, which allows for more realistic and detailed representation of anatomical and pathological structures. However, free exploration of 3D scenes like rotating the scene, zooming in/out and enabling/disabling different structures can be a complex and tedious task for unfamiliar users. Therefore, easy-to-learn interaction modes have to be considered to ease the exploration of the 3D models and to reduce the learning effort. We suggest deploying certain widgets, like sliders or thumbwheels, to enable the user to rotate and zoom the 3D objects around fixed axes. Thus, unwanted viewpoints might be avoided. A simple drag touch input on a thumbwheel widget meets the requirements to interactively explore the 3D scene. We also suggest using the mobile device *itself* for direct interaction with virtual 3D visualizations. For example, tilting the device to a certain orientation might rotate or zoom the 3D model around fixed axes in that direction.

The educational application is enhanced by an anatomical quiz where multiple choice questions can be answered by touch input. After finishing the task, the user gets a visual feedback on his/her answer. If needed, anatomical labels can be activated in order to help answering the questions (Figure 3). Additionally, Web 2.0 elements, like forums, blogs and chats are desirable in order to ask questions to the tutor or for communication with fellow students.

This mobile scenario has several advantages: It allows medical e-learning everywhere *just in time* and multi-touch input for interactive selection and manipulation of virtual anatomical structures. Compared to the first scenario, the interaction with medical 2D images and 3D structures has to be focused and designed carefully. Among the requirements for that scenario are suitable constraints for rotating, zooming, translating and editing 3D objects. Furthermore, appropriate visualization and animation techniques are important to highlight certain anatomical elements and thereby supporting the learning process of the user.

Scenario 3: Multi-Disciplinary Collaboration

Multi-disciplinary team meetings are very essential in healthcare. Based on their domain-specific expertise, medical specialists (e.g. radiologists, surgeons and internists) present and discuss relevant information about patient cases. The overall goal is to deliberate on a careful therapy concept. Now, imagine an enhanced cooperative future scenario using tablets.

Instead of bringing a bulk of paper reports into the discussion, every doctor has his/her own mobile device equipped with wireless connection and multi-touch input. The person's own "private view" on his/her data is displayed on the tablet. For example, a *radiologist* can access and explore 2D medical image data as well as segmented structures and manual annotations. A *surgeon* benefits from prepared and

interactive 3D models of anatomical and pathological structures. The overall goal is a computer-supported and efficient communication between all experts. The individual view of every person should be synchronized in a collaborative way to gain insight into the specific patient data.

The *radiologist* might send medical examinations (X-Ray images or volume data from computed tomography) to an interactive whiteboard turning his/her private view into a public one. Using e.g. a flick gesture on the device causes sliding through the 2D image stack. Since whiteboard and tablets are synchronized, the displays are updated continuously. One of the primary motivations is to enable the participants to point on certain medical structures from any location in the room. This could be achieved by touching and dragging with a single finger on the touch-screen device. Different users are distinguished by colored cursors.

Another desirable feature would be to highlight important findings, e.g. tumors, to discuss treatment options with colleagues and to review these annotated areas in future meetings. To fulfill these needs, several easy-to-use controls have to be considered. For example, a toolbar with large buttons could be deployed to enable different annotation modes like pointing, freehand sketching and placing arrows or text labels on the imagery. These annotations should of course be linked to the medical record after the meeting and physicians should at any time be able to access them. The annotated areas might be collected in a "clip gallery" that enables users to directly navigate to important areas of interest, e.g. to review proposed surgical structures. It is also conceivable to match annotated lesions over several imagery examinations to identify progress trends of tumors.

The *surgeon's* aim is to present interactive medical 3D renderings to the colleagues and to discuss an optimal access path during surgery. In order to let all people have a look on the 3D model, the surgeon might shake the mobile device. Thus, the private view is sent to all connected tablets and can be explored by every person individually. The surgeon could tilt the device itself to rotate the virtual 3D model around the fixed rotation axis or to zoom in. Every person can follow the demonstration since the view is displayed on each device simultaneously. If one wants to get a detailed view on the data, it is possible to interact with it on the personal tablet for making measurements or leaving annotations. One could raise concerns about the issue of multiple users trying to interact simultaneously with the radiology images. In order to avoid chaos during the meeting, technical policies should be implemented: Each physician is able to interact with the imagery one after another, but every person should be able to adapt annotations of his/her colleagues.

Compared to both scenarios described before, supporting collaborative processes is the most important issue here. Images, movies, annotations or 3D objects can be shared easily by synchronizing the interactive devices. This multi-

media-based procedure yields a better understanding of medical reports instead of just “talking about” the findings.

ISSUES & CHALLENGES

Scalability and Platform Independence. Several devices with different hardware and platform systems as well as different display sizes and interaction capabilities have to be considered. The visualization should be scalable, which enables the rendering on small and large screens. A platform-independent visualization on smartphones, tablets as well as on interactive tabletop surfaces is desirable to address a variety of users. Ideally, visualizations are accessible in real-time and sharable between different devices and operating systems. Special care has to be taken for scaling interaction techniques. For example, while measurement tools by touch can be easily performed on a tabletop, a smartphone’s size prohibits fine adjustments.




Figure 4: Various layers of a 3D information space can be revealed by lifting and lowering a *tangible magic lens* [6] on top of a projected human body. Right: Directly interacting with volumetric data of a human head by tilting a lens in 3D space.

Interaction and Communication. Demanding issues are appropriate methods for interacting and communicating with and between mobile devices. Intuitive gesture sets have to be implemented to allow smooth exploration of 3D visualizations and effective collaboration on heterogeneous data. User interface elements should be scalable, customizable and easily operated by the user to reach a high user experience. Single-touch and multi-touch interfaces have to be carefully tested and evaluated with end users in order to verify if users understand certain gestures and can use them for their daily workflow. In particular, in medical systems exact input is desirable, e.g. when making measurements. It is a challenging task to find out if touch input is accurate enough, compared to mouse input, and provides a sufficient level of trust during diagnosis. In contrast, during multi-disciplinary team meetings it is not important to exactly outline a lesion but it is sufficient to sketch important areas quickly. Therefore, well-arranged user interfaces and intuitive widgets, e.g. large sized touch buttons, should be deployed. Furthermore, besides interacting *on* portable devices via touch or pen, novel interaction techniques such as tangible magic lenses proposed by Spindler et al. [6] can be

used as natural metaphors to explore large information spaces *with* the spatially-aware device (Figure 4). A conceivable idea would be to adapt this metaphor to off-the-shelf mobile devices. A future use case could be to use a tracked, lightweight mobile display to explore large medical volumetric data, thus serving as a window into virtuality (c.f. Figure 4, right).

CONCLUSION

We have discussed several possible scenarios on using mobile devices in collaborative medical environments. The presented scenarios demonstrate the advantages of using interactive devices in patient consultations, medical education and multi-disciplinary team meetings. The most important benefits are portability, social collaboration and context awareness. Appropriate technologies have to be implemented to support the communication between devices and users. We have envisioned several single- and multi-touch interaction techniques for mobile 2D and 3D visualizations, which take well-suited constraints (e.g. fixed rotation axes) and widgets (e.g. interactive thumbwheels) into considerations. An ongoing question is, how already approved high-level visualization techniques can be ported to smartphones or tablets. Thereby, several technical restrictions like smaller displays, limited size of memory, rendering speed and bandwidth have to be considered.

REFERENCES

1. Preim, B. and Bartz, D. *Visualization in Medicine*. Morgan-Kaufmann-Verlag, 2007
2. Krüger, A., Tietjen, C., Hintze J., Preim, B., Hertel, I. and Strauß, G. Interactive Visualization for Neck Dissection Planning. *EuroVis*, 2005, pp. 295–302
3. Logge, C., Cordes, J., Mühler, K., Preim, B. and Hillert, C. Training in Living Donor Liver Transplantation - The LiverSurgeryTrainer. *ILTS 13th Annual International Congress*, 2007
4. Frisch, M., Heydekorn, J., Dachsel, R.: Diagram Editing on Interactive Displays Using Multi-Touch and Pen Gestures, *Proc. Diagrams 2010*, Springer, 182–196.
5. Brenton, H., Hernandez, J., Bello, F., Strutton, P., Purkayastha, S., Firth, T. and Darzi A. Using multimedia and Web3D to enhance anatomy teaching. *Computers & Education* 49,1 (2007), 32–53
6. Spindler, M., Tominski, C., Schumann, H., Dachsel, R.: Tangible Views for Information Visualization. *Proc. ITS 2010*, ACM Press, 157–166.

Branch-Explore-Merge: Real-time Revision Control for Conflict Resolution in Collaborative Visual Exploration

Will McGrath, Brian Bowman, and Niklas Elmqvist
Purdue University
West Lafayette, IN, USA
{wbmcgrat, bowmanb, elm}@purdue.edu

Pourang Irani
University of Manitoba
Winnipeg, MB, Canada
irani@cs.umanitoba.ca

ABSTRACT

Collaborative work is characterized by fluid transitions of coupling between individual participants, and the degree of coupling varies smoothly throughout the whole collaborative session. To facilitate these seamless transitions for collaborative information visualization while avoiding conflicts between users, we propose a collaboration protocol called Branch-Explore-Merge (BEM) based on real-time revision control of the shared state being explored during the visual analysis process. The protocol allows participants to diverge from the shared analysis path, make their own decisions, and then potentially merge back their findings into the shared display. We apply this general concept to collaborative search in multidimensional datasets, and propose an implementation where the shared view is a large tabletop display and the private views are embedded in handheld tablets.

Keywords: Exploration history, public views, private views, collaborative coupling, tabletop displays, tablets.

ACM Classification: H.5.2 [Information Interfaces and Presentation]: User Interfaces—*Interaction styles*; I.3.6 [Computer Graphics]: Methodology and Techniques—*Interaction techniques*

INTRODUCTION

Collaborative coupling is defined as the closeness of collaboration between two participants working on a shared task [2, 18], and ranges from *coupled* work, where the participants are working in lockstep on the task, all the way to *uncoupled* (or decoupled) work, where they are solving the task independently. Ample research (e.g. [10, 11, 20, 21, 23]) has shown that collaborative work consists of an endless sequence of changes in coupling, where participants smoothly transition between working coupled or uncoupled with other participants at different points in time. As a result, designers of groupware [1], regardless of problem domain, have long strived for explicitly supporting seamless coupling transitions in their systems while avoiding conflicts arising from interfering actions performed by different collaborators.

Several collaborative visualization systems have adopted this recommendation by making distinctions between private and public views [7], arguing that this will explicitly support both ends of the coupling spectrum. For example, Lark [23] adapts and generalizes the concept of coordinated multiple views (CMV) [16] to multi-user and collaborative settings on




Figure 1: The BEMVIEWER prototype being used for collaborative search in a multidimensional dataset.

the basis that private views of data will better enable uncoupled work. The Hugin [12] system incorporates interaction workspaces that are private to a participant, and provides an elaborate access control mechanism for releasing and acquiring views to and from a public area. However, while these mechanisms certainly avoid conflict by explicitly separating the workspaces, they provide little support for combining the results of different workspaces to reach a shared result. Furthermore, a proliferation of private workspaces will consume valuable screen real estate on the tabletop display, yielding less available space for the shared workspaces. In a sense, adding private views to a tabletop display almost goes against the collaborative nature of this display medium.

In this paper, we explore an alternative approach to supporting varying degrees of coupling for collaborative visualization tasks on shared displays. Our approach is called Branch-Explore-Merge (BEM) and is essentially an adaptation of source revision control—such as RCS, CVS, and Subversion—to visual exploration on a collaborative surface. The idea is to allow participants to branch the current state of the visualization, explore the data separately and independently from other participants, and then merge back any new results into the shared display. Branching is a *local* operation and simply means that the participant may now deviate from the shared state in their private views; merging, on the other hand, is a *global* operation that changes the shared state and thus requires consensus among participants (enacted by a voting protocol). To address the space consumption aspect of private views, we take the BEM concept a step further by suggesting that all private views be offloaded on separate physical displays, such as smartphones or tablets, leaving the collaborative surface to be used exclusively for shared state.

We have built a prototype implementation of the BEM model that we call BEMVIEWER (Figure 1) where the visual exploration is conducted in a simple 2D scatterplot visualization of a multidimensional dataset. The scatterplot supports progressive filtering through *query sculpting* [6] where the collaborators can iteratively filter out items in the dataset using a combination of range and lasso set selections. The resulting visual query consists of a conjunction of such set selections and form the shared state that the BEM model operates on. BEMViewer renders a shared scatterplot of the data and the visual query on a tabletop display. Each participant is also given a multitouch Android tablet (Samsung GalaxyTab in our implementation) through which they can connect to the BEMViewer server. The tablet applet provides a local scatterplot visualization and an interface to branch, explore, and merge results back to the tabletop display.

BACKGROUND

Our work lies at the intersection of computer-supported cooperative work [1], visualization, and novel computing platforms. Below we present the background for all these topics.

Collaborative Visualization

Collaborative visualization [8] is defined as the shared use of visual representations by multiple users, and has been named one of the grand challenges for visualization research [22]. Studies have shown that involving multiple participants in sensemaking generally improves the results (in quality or in time, or both). For example, Mark et al. [13, 14] found significant benefits for collaborative visualization in both distributed and co-located settings. Balakrishnan et al. [3] show that a shared visualization significantly improves performance compared to separate visualizations or none at all.

Several existing frameworks and toolkits for collaborative visualization exist. Scientific visualization has devoted much effort towards collaboration, mainly for distributed settings (see the survey by Brodlie et al. [5]). For information visualization, the Command Post of the Future [22] was one of the earliest examples. The emerging field of visual analytics is also branching out into collaborative settings, the multi-analyst framework by Brennan et al. [4] being one example.

Novel Interaction for Visualization

Novel input and output surface technologies are poised to make a significant impact on visualization research. Digital tabletop displays supporting direct (often multi-) touch interaction have been shown to be particularly well-suited to collaborative information visualization [17]. Examples include Isenberg and Carpendale's tabletop visualization system for tree comparison [7], the Cambiera system [9] for face-to-face collaborative analysis of documents, and the Hugin toolkit [12] for mixed-presence visualization on tabletops. The Lark system [23] extends the concept of coordinated multiple views (CMV) [16] to multi-user collaboration.

Collaborative Browsing, Filtering and Search

People working together on a common problem often have shared information needs, and so *collaborative browsing and search* is a common task in many collaborative sessions [15, 24]. With the exception of the Cambiera [9] system reviewed above, very few collaborative visualization systems

are designed for collaborative search. The system perhaps most relevant to our work is Facet-Streams [11], where multiple participants use physical widgets on a tabletop display to construct shared visual queries. However, while the Facet-Streams system does support independent work as well as mechanisms for combining queries constructed by different participants, the private workspace of each participant is limited by the overall size of the tabletop display.

BRANCH-EXPLORE-MERGE

The motivation for the Branch-Explore-Merge protocol is to embrace the continually changing coupling in collaborative work to allow for participants to seamlessly move from closely to loosely coupled work. This is achieved by adopting a revision control mechanism for interactive visual exploration. Below we discuss the three protocol components.

Model

Branch-Explore-Merge (BEM) assumes a collaborative application with a shared state S , a visual representation $V(S)$, and a set of interaction techniques i that operate on the state to produce a modified state ($i(S) \rightarrow S'$). In an actual implementation of the BEM protocol, the visual representation $V(S)$ would be rendered on a shared display that all participants can see (and typically interact with). Meanwhile, all participants would also have one or more private displays that also allow for rendering the visual representation as well as interacting with the state. Although not an intrinsic aspect of the BEM protocol, we recommend that private displays be physically separate devices from the shared display.

Upon starting a BEM session, all private displays are *synchronized*, meaning that they are using the same shared state S , and the private displays will update as that state is changed.

Branch Operation

Branching in the BEM protocol is a local (i.e., non-conflicting) operation that only affects the particular participant who initiated the branch. The result of a branch operation is simply to *desynchronize* the participant's private display from the shared display; in other words, the global shared state S is copied to a local state S_i for that participant.

Branching is explicitly invoked by, e.g., clicking on a button, but an implementation may also allow implicit branching; for example, automatically branching (and desynchronizing) a participant who starts to modify the visual representation on their private view instead of on the shared view.

Explore Operation

After having desynchronized from the shared state using a branch operation, the participant's private view will henceforth render the visual representation $V(S_i)$ instead of $V(S)$. Furthermore, any interaction performed on S_i is specific to that view only. In other words, the participant is now free to independently explore the state in uncoupled mode.

Merge Operation

Merging is invoked by a participant with a desynchronized private view when the participant wants to add his or her results back to the shared state (alternatively, the participant can always revert back to the shared state if they decide that

a branched exploration is not worth pursuing further). During the merge operation, the shared display is transformed into a *visual merge* which shows the proposed changes to the shared state (designing the visual merge is an important aspect, but is specific to the particular visual representation). Unlike branching, merging is a global operation that may potentially cause conflicts with other participants as well as the shared state. For this reason, we need to introduce a conflict resolution mechanism to handle this situation.

In the BEM protocol, we use a voting mechanism where participants cast a vote on whether to allow a particular change to the shared display. The voting policy can vary on the application; simple majority is the most straightforward and useful one. Of course, even if a proposed visual merge is voted down, the participant who initiated the merge will retain that proposed state on his or her own private view. The participant can then choose to refine the state or persuade other participants to accept the change, or simply discard that state. Furthermore, for situations where a visual merge is accepted by simple majority, any naysayers may choose to automatically desynchronize from the new shared state and receive the old state on their private views to continue their work.

IMPLEMENTATION

We have implemented the BEM protocol in a prototype collaborative visualization system called BEMVIEWER (Figure 1 and 2) for collaborative visual search in multidimensional datasets using 2D scatterplot displays. Below we describe the details of the BEMViewer system.




Figure 2: Tablet and tabletop displays in our prototype.

Visual and Interaction Design

The BEMViewer prototype uses a 2D scatterplot visualization that fills the center of the tabletop surface (Figure 3). The shared state consists of a query string that the participants are iteratively building using *query sculpting* [6]: successively adding terms of range or lasso selections on one or two dimensions to a conjunction of such terms. This allows for searching in a dataset by sequentially adding constraints. The private tablet interface also consists of a scatterplot and the same interface controls and interactions as the tabletop.

The prototype currently allows up to four participants—one per table side—to connect and join the collaborative search process. Each connected participant gets an individual control panel (also visible on Figure 3) on their side of the table

that contains a visual representation of the query as well as interface controls to branch, explore, and merge.




Figure 3: Screenshot of the tabletop interface.

Implementation Notes

BEMVIEWER has two components: the tabletop application, and the tablet app. Both are implemented in Java. The tabletop implementation uses Piccolo2D for vector graphics and the TUIO interface for capturing touch interaction. Network communication between the apps and the tabletop is implemented using a standard TCP-based protocol.

Beyond Scatterplots

There is nothing in the branch-explore-merge protocol that limits the complexity or nature of the visual mapping. The only requirement is the capability to, given two visualization states, render their difference as a visual difference, and then supporting the merge operation. This leaves the space open for using virtually any visualization design in the protocol. For example, a collaborative social network visualization tool may easily support expanding and collapsing the node-link diagram as filtering operations, and a visual difference would highlight the topological changes to the network that a merge would cause. An investigative analysis tool (such as Jigsaw [19]) could be adapted for collaborative work by placing a report of findings on the shared displays and allowing participants to add and remove to this report.

CONCLUSIONS AND FUTURE WORK

We have proposed a collaborative protocol for explicitly supporting branching and merging of visual exploration in co-located collaborative visualization. Our approach incorporates both public and private views and allows participants to branch their private views from the main visualization state to perform independent work. Merging the new state to the shared and public display consists of a visual merge operation and a common vote among all participants. We have also presented an implementation of this collaborative protocol called BEMViewer, which lets a group of participants collaboratively search in a multidimensional dataset on a tabletop display with handheld tablets for private views.

Our future work will consist of evaluating BEMViewer and the BEM protocol through controlled user experiments. We also want to explore other visual representations, interactions, and datasets. In general, we think that there is tremendous potential in pairing mobile devices with large displays, and we look forward to studying this further in the future.

REFERENCES

1. R. M. Baecker. *Readings in Groupware and Computer-Supported Cooperative Work*. Morgan Kaufmann Publishers, San Francisco, 1993.
2. K. Baker, S. Greenberg, and C. Gutwin. Empirical development of a heuristic evaluation methodology for shared workspace groupware. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, pages 96–105, 2002.
3. A. D. Balakrishnan, S. R. Fussell, and S. B. Kiesler. Do visualizations improve synchronous remote collaboration? In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 1227–1236, 2008.
4. S. E. Brennan, K. Mueller, G. Zelinsky, I. Ramakrishnan, D. S. Warren, and A. Kaufman. Toward a multi-analyst, collaborative framework for visual analytics. In *Proceedings of the IEEE Symposium on Visual Analytics Science & Technology*, pages 129–136, 2006.
5. K. W. Brodlie, D. A. Duce, J. R. Gallop, J. P. R. B. Walton, and J. D. Wood. Distributed and collaborative visualization. *Computer Graphics Forum*, 23(2):223–251, 2004.
6. N. Elmqvist, P. Dragicevic, and J.-D. Fekete. Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1141–1148, 2008.
7. P. Isenberg and S. Carpendale. Interactive tree comparison for co-located collaborative information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1232–1239, 2007.
8. P. Isenberg, N. Elmqvist, J. Scholtz, D. Cernea, K.-L. Ma, and H. Hagen. Collaborative visualization: definition, challenges, and research agenda. *Information Visualization*, 10(4):310–326, 2011.
9. P. Isenberg and D. Fisher. Collaborative brushing and linking for co-located visual analytics of document collections. *Computer Graphics Forum*, 28(3):1031–1038, 2009.
10. P. Isenberg, A. Tang, and M. S. T. Carpendale. An exploratory study of visual information analysis. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 1217–1226, 2008.
11. H.-C. Jetter, J. Gerken, M. Zöllner, H. Reiterer, and N. Milic-Frayling. Materializing the query with facet-streams – a hybrid surface for collaborative search on tabletops. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 3013–3022, 2011.
12. K. Kim, W. Javed, C. Williams, N. Elmqvist, and P. Irani. Hugin: A framework for awareness and coordination in mixed-presence collaborative information visualization. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*, pages 231–240, 2010.
13. G. Mark, K. Carpenter, and A. Kobsa. A model of synchronous collaborative information visualization. In *Proceedings of the International Conference on Information Visualization*, pages 373–383, 2003.
14. G. Mark, A. Kobsa, and V. M. González. Do four eyes see better than two? Collaborative versus individual discovery in data visualization systems. In *Proceedings of the International Conference on Information Visualization*, pages 249–255, 2002.
15. M. R. Morris, J. Lombardo, and D. Wigdor. WeSearch: supporting collaborative search and sensemaking on a tabletop display. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, pages 401–410, 2010.
16. J. C. Roberts. State of the art: Coordinated & multiple views in exploratory visualization. In *Proceedings of the International Conference on Coordinated Multiple Views in Exploratory Visualization*, pages 61–71, 2007.
17. Y. Rogers and S. E. Lindley. Collaborating around vertical and horizontal large interactive displays: which way is best? *Interacting with Computers*, 16(6):1133–1152, 2004.
18. K. Schmidt and C. Simone. Coordination mechanisms: Towards a conceptual foundation of CSCW systems design. *Computer Supported Cooperative Work*, 5(2/3):155–200, 1996.
19. J. T. Stasko, C. Görg, and Z. Liu. Jigsaw: supporting investigative analysis through interactive visualization. *Information Visualization*, 7(2):118–132, 2008.
20. A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative coupling over tabletop displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pages 1181–1190, 2006.
21. J. C. Tang. Findings from observational studies of collaborative work. *International Journal of Man-Machine Studies*, 34(2):143–160, 1991.
22. J. J. Thomas and K. A. Cook, editors. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Society Press, 2005.
23. M. Tobiasz, P. Isenberg, and S. Carpendale. Lark: Coordinating co-located collaboration with information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1065–1072, 2009.
24. M. Twidale, D. M. Nichols, and C. D. Paice. Browsing is a collaborative process. *Information Processing and Management*, 33(6):761–783, 1997.

CoSpaces: Workspaces to Support Co-located Collaborative Visual Analytics

Narges Mahyar

University of Victoria
PO Box 3055, STN CSC
Victoria, BC, Canada
V8W 3P6
nmahyar@cs.uvic.ca

Ali Sarvghad

University of Victoria
PO Box 3055, STN CSC
Victoria, BC, Canada
V8W 3P6
asarv@cs.uvic.ca

Tyler Weeres

University of Victoria
PO Box 3055, STN CSC
Victoria, BC, Canada
V8W 3P6
tweeres04@gmail.com

Melanie Tory

University of Victoria
PO Box 3055, STN CSC
Victoria, BC, Canada
V8W 3P6
mtory@cs.uvic.ca

ABSTRACT

We introduce CoSpaces, a system designed for co-located collaborative Visual Analytics on large interactive surfaces. A core design idea within CoSpaces is the use of tab-based portals to access to other work areas, supporting awareness. Underlying the tabs is a record-keeping mechanism that enables tracking of analysis history and note taking; such records are useful not only for managing individual analysis activities, but also for maintaining awareness of other users' activities. A usability study of CoSpaces suggests that these new design ideas can effectively support group analysis tasks.

Keywords: Collaboration, Visual Analytics, Interactive Surfaces, Record-Keeping, Workspace, Awareness, Portal.

INTRODUCTION

We introduce CoSpaces (collaborative workspaces), a system designed for co-located collaborative Visual Analytics on interactive tabletops. Tabletop displays are well suited to collaborative work since they allow users to interact and explore a dataset simultaneously. Previous research has shown that while working together, collaborators tend to move back and forth between loosely and tightly coupled work [7, 13]. However, when working independently on a complex analytics task, users may lose track of progress made by others. When work becomes loosely coupled, users need to maintain awareness of each other's activities [3]. It has been suggested that awareness can be increased by recording and presenting a visualization history, which is also believed to facilitate insight generation and reduce the need to redo earlier work [1, 9]. With CoSpaces, we explore how a history mechanism combined with a tab metaphor can enable users to review the work of others without disruption.

In CoSpaces, partitioning of work is accomplished via Worksheets (Figure 1), or flexible work areas that accommodate changes in the collaboration style. Our primary design contribution is a tab-based portal view to other worksheets, which enables users to see and reuse each other's work without interference.

BACKGROUND

Many collaborative tasks require changes in collaboration style, where people move back and forth between individual and group work [2]. According to Tang et al. [13], collaborators tend to frequently switch between loosely and closely coupled work styles when working over a tabletop. Another study [11] demonstrated that users preferred to work individually on some parts of a problem when the system used

was capable of supporting such individual activities. Yet research shows that even during loosely coupled work, maintaining awareness (understanding who you are working with, what is being worked on, and how your actions will affect others) is critical to ensure efficient and effective team coordination and decision-making [8].

In co-located collaboration, people are able to gather implicit information about team members' activities from body language, alouds, and other consequential communications [3]. Nonetheless, awareness becomes a challenge when group members are working in a loosely coupled fashion since conversation may be disruptive [5]. This is particularly true in complicated visual analytics tasks, where users can easily duplicate each other's work (e.g., by creating the same charts of a data set). Therefore, there should be channels for providing awareness with minimal interruption and cost.

Few visualization tools for co-located collaboration provide explicit mechanisms for awareness. With Lark [14], users can create several copies of data views at various points along a visually presented "Visualization Pipeline". Changes at upstream locations (i.e. closer to the dataset on the pipeline) are propagated into all the downstream data views. Though it reveals the downstream changes, in line with the authors we believe that this approach works better for coordinating work rather than providing awareness of what others have done.

Colour-coding has been used as a mechanism for providing awareness in co-located collaborative tools. An example is Cambiera [6], a system designed for collaborative visual analysis of text documents. Each user's searches are marked with varying shades of one colour. This enables collaborators to recognize and track their own and each other's work. The implementation of colour-coding in Cambiera is more suitable for providing collaborators with information about each other's search interest (i.e. searched keywords). In CoSpaces, colour-coding is used slightly differently. Instead of assigning distinctive colours to users, they are assigned to Worksheets. Using this approach, "analytic activities" rather than analysts are marked and identified by colours.

COSPACES SYSTEM DESCRIPTION

Each Worksheet defines a work territory, either personal or shared. Worksheets can be freely resized and positioned and users may create as many Worksheets as they need. Personal versus shared Worksheets are identical as far as the system is concerned; ownership is defined by the way in which they are used. Worksheets can be used to organize work categorically,

and also to create personal and shared territories. A Worksheet has five main sections, as shown in Figure 2.

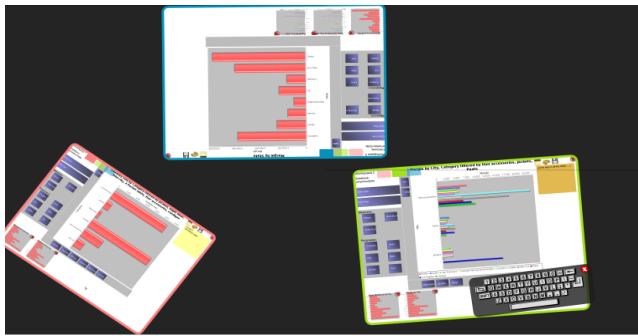


Figure 1: A snapshot of CoSpaces' user interface. Dark background is the common work area (tabletop's surface). There are three open Worksheets where collaborators can simultaneously analyze data.

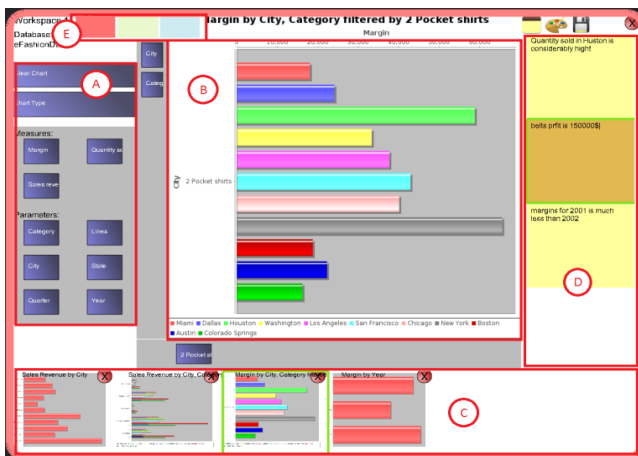


Figure 2: Detail of a Worksheet: Analysis pane (A) that gives users control over the charts, Visualization pane (B) that shows the current chart, scrollable History pane (C) where thumbnails of previous charts are shown, Notes pane (D), and Tabs (E) that provide a portal view to another worksheet.

The main design contribution of CoSpaces is the use of a tab metaphor to address the awareness problem. Coloured tabs at the top of each Worksheet (Figure 2E) are associated with other existing Worksheets. Each tab is colour-coded to match the border colour of the Worksheet that it links to. Tabs act as portals to view other Worksheets. Tapping on a tab replaces the local worksheet content with a *view* of another Worksheet. Tapping on the local tab switches the view back. An example remote view is shown in Figure 4.

When another Worksheet's tab is selected, the contents of all worksheet panes reflect the remote information, including the current visualization of data as well as recorded items in the history pane and notes taken in the notes pane. This provides the viewer with complete and up to date information about the remote Worksheet. The user may browse charts in the history pane to learn about another user's past analytical activities and interests in the data space.



Figure 3: CoSpaces in use.




Figure 4: A remote view of a red Worksheet shown within a blue Worksheet. While remotely viewing, widgets on the Analysis pane are deactivated (grayed-out) to avoid unintentional interruption.

Reading notes in the note pane can notify the viewer about all the externalized insights, findings and hypotheses generated by another user. To prevent unintentional changes and interruption, a Worksheet's remote view is read-only. Widgets on the Analysis pane of the remote view are grayed-out to visually imply their deactivation. To avoid causing disruption to another user, navigation in a remote view does not propagate to the other Worksheet's local view. Although manipulation of remote content is prohibited, items from a remote Worksheet's history pane can be copied to the local Worksheet.

A Worksheet automatically captures and saves a copy of the current analysis-state right before a change has been applied. An example of an analysis-state change is when a user changes the quantitative measure on a bar chart to another data attribute. As part of the analysis-state, we also capture a thumbnail picture of the current chart, which is placed in the history pane in chronological order (Figure 2C). To avoid an overwhelming number of saved items, we use a simple heuristic inspired by the chunking rules devised by Heer et al. [4]. An analysis-state is saved only when a change in the current mapping of data takes place. In other words, adding or removing filters will not result in a save. We have also provided the ability to save a desired analysis-state explicitly through a save button. Moreover, users can delete any undesired item from the list of recorded analysis-states.

An analyst working with data on a Worksheet can externalize her cognitive products such as findings, hypotheses and so on

using the notes pane. Tapping a note button on the top of a Worksheet opens an on-screen keyboard. Also a yellow post-it like text area is created in the notes-pane. Considering the importance of connecting externalized material to the visual representation [1, 9, 12], a new note is automatically linked to the current visualization of data (i.e. the chart) and the note.

IMPLEMENTATION

CoSpaces is multi-touch application written in JAVA. Multitouch for Java (MT4J) provides multitouch functionality, and communicates with the touch detection library Community Core Vision (CCV) using the TUIO protocol. JFreeChart is used to create the graphical charts.

EVALUATION

We conducted a usability study to gain initial feedback about CoSpaces. At this early stage of research, we did not attempt to validate utility of CoSpaces for long term tasks by real analysts. The main objective of the study was to test CoSpaces's capability to support awareness under conditions of changing collaboration styles.

We recruited 20 computer science students (16 graduates, 4 undergraduates) in the form of 10 pairs. They were all familiar with basic data analysis activities and concepts such as creating a simple statistical chart of tabular data. We used a rear-projected 70-inch (diagonal) tabletop with a resolution of 3840 x 2160 provided by combining four 1080p projectors. The tabletop uses infrared light and a rear mount infrared camera to detect a virtually unlimited number of touches.

Tasks and Procedure

Participants performed two tasks. After receiving a 20-minute introduction to the system features, participants started Task 1, which took about 30 minutes. Task 1 enabled participants to learn how to use CoSpaces. They could stop and ask either of the two co-present observers if they had any questions about either the system or the data. After completing Task 1, groups were given a 5-minute break. Task 2, which took almost 40 minutes, was an open-ended analytical question that required both loosely and closely coupled work. These two tasks were followed by a questionnaire and a follow up interview that took almost 20 minutes.

RESULTS

All the reported observations are based on Task 2, since Task 1 was only intended as practice. Though we gathered both quantitative and qualitative data, here we focus on qualitative observations and participants' comments from the interviews.

Our observations corroborated our speculated benefits of using tabs for providing awareness. Many participants gave us positive feedback about being able to view each other's work progress via tabs. For instance, in the follow up interview a participant expressed that "...real time update of [the] other's view was interesting, because [I] could keep [myself] updated all the time...". Another participant mentioned "...being able to see others' workspaces, [and] keep track of them in own workspace" was one of the most useful features of the system. Participants' quantitative assessments of the usefulness of the tabs were also positive. Out of 20 participants, 17 assessed Tabs as useful in their evaluation. The average score given to Tabs was 4.95 out of 6.0 with a STDEV of 1.07.

Participants used tabs to investigate another user's current chart (17 times), review their collaborator's work history (7 times), copy an item to their own Worksheet (12 times) and review the other's notes (3 times). On average, tabs were used 2 times per group during task 2. Participants spent between 20 seconds to 2 minutes using tabs each time.

Because the groups contained only 2 people, and they often positioned themselves side by side (8 out of 10 groups), participants could easily look over at each other's Worksheets to see what the other person was working on. Participants frequently did this for an update on current work. However, we observed that while working individually, none of the participants attempted to get close enough to the other's Worksheet to have a detailed review of his/her work history; for this purpose, they used tabs instead. This behaviour avoids unnecessary interruption and imposition on the personal territory of another person. Thus, while tabs are useful for observing another user's current work, they may be even more beneficial for reviewing a collaborator's past activities.

All the pairs engaged in both loosely and closely coupled collaboration, as anticipated based on the design of task 2. The important observation here was that the design of the Worksheet effectively supported both collaboration styles, as well as the transition between the two.

Additionally, we observed that the flexible nature of the Worksheet and its fluid design not only supported changes in collaboration style but also facilitated the analytical reasoning process. The ability to create multiple new Worksheets as well as creating a Worksheet from an item in the history pane facilitated exploratory analysis. We also observed that users stacked or placed Worksheets side by side to compare visualizations and/or have discussions (Figure 3). This was enabled by the tabletop's substantial screen real-estate.

We observed that participants frequently saved, reused and manipulated recorded items. In total, participants manually saved charts to the history 90 times. They also regularly reloaded items from the history pane (146 times). Reuse happened both during the analysis, when participants often worked individually, and towards the end of the analysis session, when participants typically engaged in a closer collaboration to share their findings. Less often (8 times), participants used saved items to create new Worksheets. Participants often used this feature when they wanted to compare two previously created charts side by side.

We also noticed that many participants deleted unwanted charts (81 times) to keep a clean history pane. It seems that our simple heuristic for reducing the number of automatically saved charts was insufficient. One participant remarked, "I think it is not overwhelming to save charts explicitly, what is overwhelming is having too many charts automatically saved!"

Since our notetaking mechanism was simplistic, we did not expect it to work perfectly. Nonetheless, we observed that many groups used the on screen keyboard to take notes (total of 71 times, used by 8 out of 10 groups). Participants mostly took notes while working individually. When they wanted to share their results, they read through notes or reloaded charts that were linked to notes to discuss their findings. In other

words, notes were the primary mechanism for recording important material. The fact that notes were linked to the charts was extremely important to participants. At least 3 participants explicitly mentioned the linking as important.

DISCUSSION

We found that tabs were a useful way to maintain awareness, but the frequency of their use was somewhat less than we anticipated. One possible explanation is the relatively short length of the analysis session (40 minutes). With a longer, more complicated task, or more users, we suspect that users might need to review each other's work more often to avoid duplication of effort. Our initial observations suggest that tabs will be a suitable way to accomplish this. In addition, some participants mentioned that the colour-coding of the tabs was not quite sufficient when there were many Worksheets. This problem could be overcome by adding additional visual cues, such as labeling a tab with the name of the corresponding Worksheet or the owner's name and photo.

We note that in our design, we did not address data privacy. All information in a Worksheet is accessible when viewed remotely from another Worksheet. Results of an earlier study [10] indicated that users needed space to work independently without interference, but that the space did not necessarily need to be private. Accordingly, we designed CoSpaces for small closely-knit teams who share a common objective and therefore benefit from sharing all of their information and findings. Such teams are by no means universal, however. Collaborative teams may involve individuals with competing interests or from different organizations; these people may wish to keep some information and findings private. CoSpaces would need to be substantially extended to support this scenario. Possibly individual laptops containing personal data, linked to a large display for shared information, would be useful in such situations. Alternately, privacy mechanisms could be added so that users could prevent parts of a Worksheet or entire Worksheets from being viewed remotely.

CONCLUSION AND FUTURE WORK

We introduced CoSpaces, a prototype designed for co-located collaborative visual analysis of tabular data on interactive tabletops. CoSpaces introduced the concept of tab portal views to address the challenge of awareness, especially during periods of loosely coupled work. Our user study indicated that tab views are a promising design direction for supporting awareness in collaborative visual analytics, when combined with flexible workspaces and record-keeping tools such as linked notes and thumbnails of past analysis states.

Future work could improve interface design details and extend the functionality of CoSpaces. We plan to expand the record-keeping module by incorporating more efficient note taking mechanisms, rich text editing, and improved heuristics for automated analysis-state capturing. The observation that users manually saved and deleted many recorded items in our study suggests that they need greater control over the recorded history. Further studies are required to assess our design ideas for long-term use over multiple sessions.

ACKNOWLEDGMENTS

We thank SAP and the Natural Sciences and Research Council of Canada (NSERC) for funding this research.

REFERENCES

1. Chen, Y. and Barlowe, S. Click2Annotate: Automated Insight Externalization with rich semantics. *Proc. IEEE VAST*, (2010), 155-162.
2. Gutwin, C. and Greenberg, S. Design for individuals, design for groups: tradeoffs between power and workspace awareness. *Proc. ACM CSCW*, (1998), 207-216.
3. Heer, J., Ham, F. van, Carpendale, S., Weaver, C., and Isenberg, P. Creation and collaboration: Engaging new audiences for information visualization. Kerren et al. eds, *Information Visualization*, (2008), 92-133.
4. Heer, J., Mackinlay, J., Stolte, C., and Agrawala, M. Graphical histories for visualization: supporting analysis, communication, and evaluation. *IEEE Trans. Visualization and Computer Graphics*, 14, 6, (2008), 1189-96.
5. Isenberg, P. and Carpendale, S. Interactive tree comparison for co-located collaborative information visualization. *IEEE Trans. Visualization and Computer Graphics*, 13, 6 (2007), 1232-9.
6. Isenberg, P. and Fisher, D. Collaborative Brushing and Linking for Co-located Visual Analytics of Document Collections. *Computer Graphics Forum*, 28, 3 (2009), 1031-1038.
7. Isenberg, P., Tang, A., and Carpendale, S. An exploratory study of visual information analysis. *Proc. ACM CHI*, (2008), 1217-1226.
8. Kulyk, O. and Veer, G. van der. Situational awareness support to enhance teamwork in collaborative environments. *Proc. 15th European Conf. Cognitive ergonomics*, (2008), 1-5.
9. Lipford, H.R., Stukes, F., Dou, W., Hawkins, M.E., and Chang, R. Helping users recall their reasoning process. *Proc. IEEE VAST* (2010), 187-194.
10. Mahyar, N., Sarvghad, A., and Tory, M. A closer look at note taking in the co-located collaborative visual analytics process. *Proc. IEEE VAST* (2010), 171-178.
11. Morris, M.R., Ryall, K., Shen, C., Forlines, C., and Vernier, F. Beyond social protocols: Multi-user coordination policies for co-located groupware. *Proc. ACM CSCW*, (2004), 262-265.
12. Shrinivasan, Y.B., Gotzy, D., and Lu, J. Connecting the dots in visual analysis. *Proc. IEEE VAST*, (2009), 123-130.
13. Tang, A., Tory, M., Po, B., Neumann, P., and Carpendale, S. Collaborative coupling over tabletop displays. *Proc. ACM CHI*, (2006), 1181-1190.
14. Tobiasz, M., Isenberg, P., and Carpendale, S. Lark: coordinating co-located collaboration with information visualization. *IEEE Trans. Visualization and Computer Graphics*, 15, 6 (2009), 1065-1072.

Multitouch Magic Fisheye : Precise interaction with dense data on tabletop

Fred Vernier

LIMSI-CNRS Labs. Univ. Paris-Sud
BP 133, 91403 Orsay cedex, France
Frederic.vernier@limsi.fr

Chia Shen

SDR Lab, SEAS, Harvard University
33 Oxford Street, Cambridge, MA 02138 USA
Chia_shen@harvard.edu

ABSTRACT

Fisheye lens can be useful on multi-touch devices where dense usemap based applications need rich interaction to zoom, pan, rotate, select, annotate, etc. Direct input precision is often not critical to pan, rotate or zoom a map, but selection of items remains critical on dense maps. Surprisingly, only a very recent work on mouse based fisheye lens revealed how fisheye lens can dramatically improve precision. Unfortunately the disclosed techniques heavily rely on mouse pointer and cannot be transposed on tabletops. In this paper, we present a multi-touch interactive fisheye lens called MMF – Multitouch Magic Fisheye. MMF decouples the lens definition phase and the interaction phase on different fingers, enabling flexible input gestures with higher precision. We then present design issues (activation and lens offset) and discuss user strategies (choice of fingers and hand) to achieve a smooth integrated gesture with activation, lens adjustments and precise selection. We finally describe two concrete realizations of MMF through the implementation of a 1D Combobox menu and a 2D lens.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Design. **Keywords:** Multi-touch, fisheye lens, precision input.

INTRODUCTION

Fisheye lenses are now well known techniques but very few popular software use it. Attempts to explain it often refer inputs problems (poor input precision and overshooting). In fact, when fisheye lens is attached to the mouse cursor it does not magnify the motor space (see [6,8,1] for more details), resulting in the user overshooting magnified targets. Surprisingly, only very recent research work [1] suggested to improve fisheye efficiency by increasing the input precision in magnified areas. They achieve measured improvement by releasing the constraint of mapping lens and cursor's position to the same fluid movement of user's input device. The solutions proposed by [1] decouple fisheye movements and interaction with magnified content (like target acquisition) at the cost of an additional articulating task (modifier key, ring manipulation) or use cursors speed to introduce fast and precise modes [6,1]. Another lens-based touch-tabletop interaction technique that offers a solution for precise data selection is the focus+context technique of the iLoupe [9]. iLoupe was designed for tabletop and the focus area provides a precise access to data.

The quest for precision has also taken place on direct touch systems (tabletops, touch screens, etc.) but for another reason: direct touch surfaces do not offer the same precision as its mouse counterpart and legacy applications barely work on touch screens. Benko et al [3] introduced five techniques based on Dual Finger: Offset, Midpoint, Stretch, X-Menu and Slider. These techniques all rely on a first gesture to gain precision and then on a SimPress click (harder press of the finger on the surface). Such pressure-based interaction is not supported by every touch surface. "Stretch" is the closest to a fisheye lens and outperform the other 4 techniques but suffers from occlusion.

Fluid DTMouse [4] is a multitouch mouse emulation to support legacy applications on the DiamondTouch table. It displays a mouse cursor in between two fingers (like Midpoint of [3]) but use a third finger to trigger mouse button avoiding SimPress. Fluid DTMouse alleviates occlusion and increases precision. Precision is in fact doubled when moving only one finger because the cursor moves precisely at half speed (see Midpoint discussion of [3]). The third finger (DTMouse) offers better precision than SimPress (Dual Finger Midpoint), yet neither technique benefits from a fisheye which potentially provides greater than a 2x gain in resolution and precision if appropriate mechanisms can be designed to enable motor space resolution increase.

The idea of fisheye lens on direct touch surface has been explored with DTLens [5] and a technique for "Rubbing the fisheye" [8]. DTLens uses two fingers to set both lens location and size. DTLens supports only limited interaction with the underlying content and the annotation task discussed in [5] does not require high precision. "Rubbing the fisheye" addresses the precision issue with a compound gesture (rubbing the finger on the surface to zoom-in or zoom-out the fisheye then a 1sec timer let the user release her finger and hit one of the zoomed target). The second half of the gesture introduces a timer which severely limits the efficiency of the technique, and focus location is barely predictable since the user may not accurately set the center of the magnification. One of the biggest problems of the technique is the lack of movement of the fisheye lens once it has been invoked.

This paper presents Multitouch Magic Fisheye (MMF), a fisheye technique efficiently mapping multitouch inputs to fisheye lenses in order to maximally improve their precisions. We present two implementations of the technique: a 1D Fast Fisheye Combobox (FFC) to show how it works

on the selection of a country among a list of 200 entries and a 2D map fisheye to show how it works on spatial data.

MULTITOUCH MAGIC FISHEYE PRINCIPLE

The Multitouch Magic Fisheye (MMF) starts with applying the Midpoint input technique to a fisheye output. By using the 2+1 fingers posture to manipulate the fisheye lens we first leverage the multitouch capabilities to avoid occlusion and improve precision and predictability of focus location. Unlike Fluid DTMouse where the third finger serves only as a trigger, MMF empowers the third finger to operate as a direct touch selector in the middle of the fisheye lens. As this finger remains independent of the first two finger ones (defining the lens) it can become a powerful input device. This precise (but still direct-touch) third finger can:

- Tap (select) magnified items
- drag local items (may be dropped later in non fisheye mode when first two finger have been released)
- drop items on local areas (may have been grabbed earlier in non fisheye mode)
- adjust fisheye properties with a related gesture (e.g., center offset or magnification factor)
- pan the fisheye (will add small offset between first 2 fingers and fisheye boundaries)

In essence, once the fisheye magnified motor space is defined by the first two-finger touch, MMF treats this space as a new separate and independent touch zone, allowing external touch input gestures to operate inside this movable secondary input zone. As with Fluid DT Mouse, this strategy of using 2+1 fingers remains independent from how single touches are managed (i.e. normal direct touch).

ACTIVATION AND DESIGN ISSUES

MMF can be triggered by detecting 2 simultaneous contacts (separated by a plausible distance) on full multitouch tabletops (e.g. optical tabletop systems) or by detecting plausible bounding box (on DiamondTouch like systems). This activation we call direct, can tolerate more than one finger touch area, e.g., the ring finger and/or pinky finger to touch the table. Such tolerance is important because user observation reveals that it is more comfortable to sometimes rest two or three fingers on the table surface (i.e., more fingers imply a hand's weight is divided among more contact points and reduce friction during movements). The pinky and the ring fingers can also be used to distinguish thumb from the middle finger, as they are closer together. These additional fingers also indicate which hand (left or right) is used as they are always on the same side of the thumb to middle finger axis. The other way to release the finger ambiguity is to wait for index finger to tap the surface one time before activating the fisheye lens. Lens is made more explicit and because the index finger always falls on the other side of the thumb-middle finger axis closer to the index finger it tells the system where the hand is resting on the surface. This activation strategy we call "lazy" as it waits for the third finger to appear. It is also very well suit-

ed to fight the overshoot problem. When only 2 fingers touch the table a temporary cursor can help the user to position the center of the lens before it appears.

MMF can then use the information on the hand location to best set and adjust the lens position. The fisheye's focus should be located with an offset from the two touch points of the hand to have the index finger comfortably ready over the focus area. In lazy activation, this offset can be set precisely under the 3rd finger tap.




Figure 1: Expected location of fisheye's focus according to relative position of the fingers involved in the Magic Multitouch Fisheye.

As in Fluid DTMouse user strategy to touch with thumb+middle fingers, leaving the index finger in the air, seems convenient while a bi-manual equivalent strategy was also reported being natural by users [3]. The goal of using a finger of the second hand as the third input finger is twofold:

- Improve stability (third finger is not connected to the same kinesthetic joints of the skeleton thus will involuntarily interfere)
- Increase reach among magnified items (moving index finger in between thumb and middle finger has a small comfortable area)

We highlight how the hand posture of Fluid DTMouse mimics hand posture over a traditional mouse (thumb and middle finger holding the mouse and index finger triggering the button). This comparison can be extended to MMF's similarity to the Apple Magic mouse [7] (thumb and middle finger holding the mouse while index sliding on the multitouch top surface of the mouse). This latter comparison explains the term "magic" used in the name of our technique as we expect magic mouse users to be easily become familiar with MMF.

MAGIC FISHEYE COMBOBOX

In this section, we present a concrete implementation of MMF to illustrate its utility and benefits. Magic Fisheye Combobox (MFC) is a Combobox widget displaying a Fisheye Menu [2] on demand using the Multitouch Magic Fisheye interaction described above. It is illustrated in

Figure 2. It is implemented as a QuartzComposer (QC) patch listening TUIO events, processing multitouch events in a JavaScript engine and displaying efficiently a list of textures generated with QC components. MFC allows title items to be inserted in the list (e.g., the Alphabetic letters in *Figure 2*) with different levels of magnification while not in focus. Title miniatures are limited to a readable level, text is aligned alternatively left and right and then shifted to appear slightly on the side of the regular items). Regular items are either right or left align according to where the hand is.

MFC operates in single finger mode to open/close the Combobox. Single finger can also roughly set the fisheye position and preselect the closest item under the finger. In the 2-finger mode MFC sets the fisheye position and the size at the same time, it then preselects the item exactly in the middle of the 2 fingers. If finger(s) is(are) released on top of the combo list in one of these two modes, the pre-selected item triggers a “selection” event. Finally, the user can also tap with the 3rd finger while in the 2-finger mode in order to select any visible item in focus (i.e. “Colombia” instead of “China” in *Figure 1*). When the first two lens defining fingers are released, the Combobox is immediately closed to avoid a possible conflicting event. The precision of MFC using 2+1 fingers is enabled by 3 factors:

- As the defining fingers do not occlude the pre-selected item and because the pre-selection occurs on a sliding motion (not on tap) the base precision depends on surface technology and not on fat finger effect. The precision of the focus area placement can be sub-pixels in many tabletop technologies.
- As discussed in [3] for Dual Finger Midpoint, by moving only one of the 2 fingers in the 2-finger mode, the pre-selected item changes along with fisheye’s center at a doubled precision compared to the moving finger.
- As the third finger tap or slides on top of the magnified items (in magnified motor space), the precision is multiplied by the magnification factor of the fisheye lens.

Following MMF principles, MFC is an accurate technique to select items from a long list of items. This precision is achieved without sacrificing usability:

- Lens’s size AND position are mapped directly to a human naturally integrated feature (location and gap between two fingers)
- Magnification factor can be automatically adjusted so items in focus are always large enough for the upcoming third finger (fisheye magnification strength can be automatically adjusted to magnify items at the size of thumb’s contact that is usually bigger than index finger).
- The index finger can tap directly on visible items like anywhere else on the direct touch surface.

We decided to not display a cursor in between the first two defining fingers of MFC because the feedback of preselect-

ed item is sufficiently visible. However MMF and MFC are not incompatible with other fisheye usability improvements such as transparency [8,1] or speed coupling flattening [6]. A visible cursor could be useful in such cases.

MFC uses the direct activation strategy to speed up the selection. We also combine the first 2 touches with the




Figure 2: Fast Fisheye Combobox with the country “Cambodia” preselected in light blue among a list of 200 (end of the list omitted). The thumb and middle finger both define fisheye position and size while index finger moves down to select one magnified item

opening of the Combobox. If any one of the first 2 fingers touches the Combobox button, it opens the list and set the fisheye in between the 2 fingers. The list can also be open with a simple touch first. Two fingers are then put down to control the fisheye.

MAGIC 2D FISHEYE LENS

In the previous section, we applied the MMF principle to a 1D fisheye menu because it is similar to well understood widgets (menu or Combobox) and thus relatively easy to test and compare during the development process. FisheyeMenu [2] in the literature has undergone a challenging experiment showing fisheye is not the fastest technique. In our case, the top challenges for direct-touch surfaces are precision and screen footprint [3,4]. Handheld or laptop ouch screens are still limited in size and the touch surface is shared in tabletop conditions. Both contribute to the needs for better support for precise input and best utilization of screen space. In this section, we apply MMF lenses on 2D content as in all the other previous work we reviewed [1,5,6,8]. As illustrated in *Figure 3*, we faced a challenging trade-off between fisheye size, distance of the focus area to the index finger and occlusion

- The hand can occlude focus (in *Figure 3* the selection “Vert-Galant” is partially occluded)
- The index finger may be too far from focus (in *Figure 3* the index finger reaches the center but not the left of the focus)
- The lens should be large enough to leave undistorted objects in the focus (in *Figure 3*, the train station north of “Vert-Galant” has a label too long to fit in the fisheye).

Unlike 1D fisheye we couldn’t find a satisfactory trade-off for single hand usage. Many users that we informally tested with reported problems with unexpected location at activation or with target being unreachable. Early users didn’t report problems with bi-manual use of the technique and several users expressed the need to have large fisheye lenses. We finally chose a trade-off between the direct and the lazy activation. When the first two fingers hit the surface, the fisheye appears in the middle between the thumb and the middle finger. If any later third contact point is detected in the focus of the fisheye it is considered as a precise “click”. If any later third contact point is detected outside of the fisheye lens, along the normal vector illustrated in *Figure 1*, the fisheye is re-centered to this point. The radius of the fisheye is then modified to have the first two contact points close to the edge of the lens. This new radius is computed proportionally to the distance between the first two fingers. As the focus is already under the index finger of the first hand, such action can only be triggered by a bimanual gesture.




Figure 3: 2D Fisheye on a subway map with train stations selectable (“Vert-Galant” station selected).

This transition to a bimanual mode of operation re-centers the fisheye to a very predictable location. The new radius of the fisheye allows the user to setup the size of fisheye as big as necessary and the proportional aspect means the first two fingers can still further modulate the size of the lens. Users who know they won’t use the first fisheye can use their index finger to replace the middle finger as anchor on the edge of the lens.

CONCLUSION

We presented the Multitouch Magic Fisheye which leverages multitouch capabilities of direct touch surfaces to enhance fisheye input precision and flexibility. We applied this generic technique to an 1D fisheye menu and described how the Magic Fisheye Combobox can smoothly combine 1, 2, and 3-finger operation modes. We also applied this technique to a 2-D fisheye lens and found bimanual operations more appropriate. We explained why MMF dramatically improves precision despite it is implemented with a combination of an imprecise fisheye lens and an imprecise direct touch surface. However this paper also contributes by showing fine setup necessary to integrate the selection in a fluid gesture. A good setup depends on the task and the shape of fisheye so future works should be conducted to produce generic guidelines for designers.

ACKNOWLEDGMENTS

We want to thank the AVIZ and In-Situ teams of INRIA for their precious comments on this work in particular Jean-Daniel Fekete and Emmanuel Pietriga.

REFERENCES

1. Appert C., Chapuis O. and Pietriga E., High-Precision Magnification Lenses. *To appear in proceedings of CHI 2010 (Atlanta, GA, USA, April 10-15, 2010)*, ACM, NY
2. Bederson B. Fisheye Menus, In *Proceedings of UIST'00 (San Diego, CA, USA, November 5-8, 2000)*, ACM, NY, pp. 217–226.
3. Benko, H., Wilson A. and Baudisch P., Precise Selection Techniques for Multi-Touch Screens. In *Proceedings of CHI'06 Human Factors in Computing (Montréal, Québec, Canada, April 22 - 27, 2006)*, ACM, NY, pp. 1263–1272.
4. Esenther A. and Ryall K., Fluid DTMouse: better mouse support for touch-based interactions. In *Proceedings of AVI'06 (Venezia, Italy, May 23-26, 2006)*, ACM, NY, pp. 112–115.
5. Forlines C. and Shen C. DTLens: multi-user tabletop spatial data exploration. In *Proceedings of UIST'05 (Seattle, WA, USA, October 23-26, 2005)*, ACM, NY, pp. 119–122.
6. Gutwin C. Improving Focus Targeting in Interactive Fisheye Views, In *Proceedings of CHI'02 Human Factors in Computing (Minneapolis, USA, April 20 - 25, 2002)*, ACM, NY, pp. 267–274.
7. Magic Mouse. Apple Inc. October 2009 <http://www.apple.com/magicmouse/>
8. Olwal, A. and Feiner S. Rubbing the fisheye: Precise touch-screen interaction with gestures and fisheye views. In *Proceedings of UIST'03 (Vancouver, Canada, November 2-5 2003)*, pp. 83–84
9. Volda S., Tobiasz M., Stromer J., Isenberg P. and Carpendale S. Getting Practical with Interactive Tabletop Displays: Designing for Dense Data, “Fat Fingers,” Diverse Interactions, and Face-to-Face Collaboration. In *Proceedings of ITS 2009*.

Towards Utilizing Novel Interactive Displays for Information Visualization

Christian Tominski Heidrun Schumann
Institute for Computer Science
University of Rostock
{ct|schumann}@informatik.uni-rostock.de

Martin Spindler Raimund Dachsel
User Interface & Software Engineering Group
Otto-von-Guericke-University Magdeburg
{spindler|dachsel}@ovgu.de

ABSTRACT

Visualization research has yielded a number of useful techniques to generate visual representations and to allow users to explore data interactively on computer screens. Yet, the constantly growing complexity of today's data calls for new ways of visual and interactive data analysis. This is where new display and interaction technologies offer promising possibilities yet to be explored. In this work, we identify three gaps to be addressed by future research. (1) The *technology gap* is about the lack of a systematic mapping of common interaction tasks onto interaction techniques. (2) The *integration gap* concerns the integration of novel interaction techniques and technologies with existing information visualization approaches to create new powerful visualization solutions. (3) The *guidelines gap* criticizes the shortage of support for users to choose suitable solutions for the task at hand. We outline related challenges and ideas for future work by the example of our work on tangible views for information visualization.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Design, Human Factors

Keywords: Tangible views, information visualization, magic lenses, gestural input, multitouch, pen, gestures, diagrams.

INTRODUCTION

The goal of visualization is to support people in forming mental models of otherwise difficult-to-grasp subjects, such as massive data, complex models, or dynamic systems [12]. The term *forming* implies that visual output is not the end product of visualization. It is rather the process of adapting the visual output and interacting with the data in order to gain insight.

For several decades, visualization researchers have developed a wealth of visualization and interaction concepts for many different types of data and tasks. What most of the existing techniques have in common is that they are targeted for regular desktop workplaces with a computer display, a keyboard, and a mouse. With the advent of new display technologies, such as large high-resolution displays or small hand-held displays, it became necessary to adapt existing visualization approaches or to devise new ones. Recently, modern interaction technologies, such as multi-touch interaction or tangible interaction have considerably broadened the spectrum of what's possible and created a need for rethinking ex-

isting visualization solutions with regard to interaction. The seamless integration of both display and interaction in a single touch-enabled device, such as interactive tabletops and tablets, makes direct manipulation [11] truly direct. By exploring information directly under your fingertips, the *forming* of mental models for interactive visualizations seems to be particularly well supported and promising.

In this paper, we aim to describe several issues concerning the future development of information visualization in the context of new interactive surfaces and interaction technologies. We identify gaps in the state of the art, illustrate them with our own previous work [3, 13] and motivate possible next steps for future research. Here, our main concern is related to the systematic investigation of the possibilities of the classic as well as the promising new technologies, on the one hand, and the well-justified application of these possibilities to solve visualization and interaction tasks, on the other hand.

IDENTIFYING THE GAPS

In the following, we identify three research gaps worth being addressed by future research.

Technology Gap

Visualization research builds upon commonly accepted strategies for visualizing data. In his classic book, Bertin [1] introduces visual variables and defines how data is to be mapped onto them. Cleveland and McGill [2] investigate the effectiveness of visual encodings for data exploration. Thus, on a conceptual and on an experimental level, we have backed knowledge how to transform data D to visual representations V using the mapping $vis : D \rightarrow V$.

However, there is no such commonly accepted mapping in terms of interaction. So far, mouse and keyboard have been the basic and dominant devices for user interaction. Advances in technology have recently added new modalities to the interaction toolbox. Multi-touch interaction, gestural interfaces, pen-based interaction and tangible interaction are only a few examples. What still has to be developed is the mapping $interact : T \rightarrow I$ that defines how interaction tasks T are effectively and efficiently carried out with the different interaction techniques I available. Obviously, specifying appropriate sets T and I turns out to be a necessary and challenging condition for successfully developing novel interactive visualizations. Therefore, a repertoire of suitable interaction techniques has to be defined and described in a consistent way, which eventually allows for an easy task mapping.

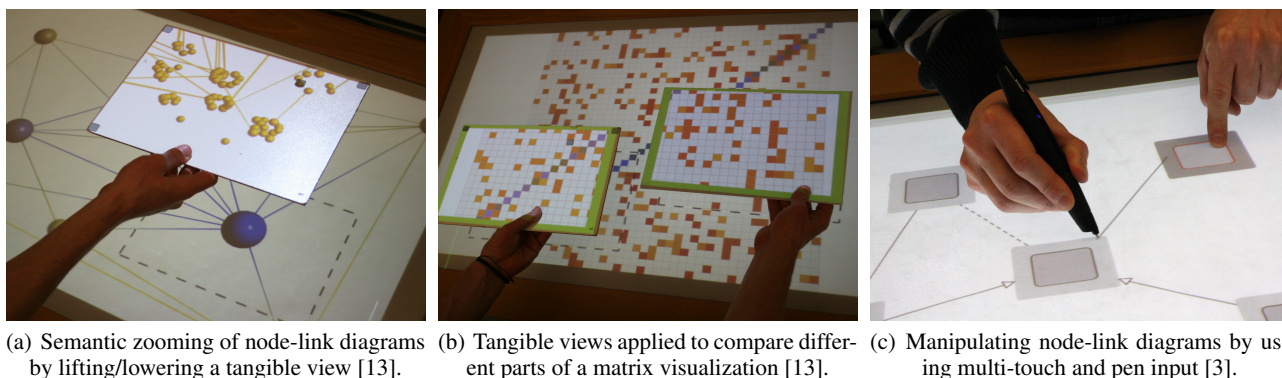


Figure 1: Utilizing novel interactive displays for information visualization.

Implementing the new techniques will also form the basis for future visualization applications.

Integration Gap

Closing the technology gap will result in a new repertoire of interaction techniques. However, by now users of interactive visualizations mostly apply techniques that are designed for classic desktop computers. Utilizing novel interactive displays for visualization purposes has not received much attention so far. So there is a gap in terms of promising new possibilities on the one hand, but only little integration of these possibilities into visualization research and applications on the other hand.

Yet, there are first approaches that specifically address the integration of visualization and interactive display technology. For instance, Isenberg et al. [5, 6] utilize interactive tabletop displays for collaborative visualization, Volda et al. [14] discuss design considerations for interacting with data on interactive surfaces, Spindler et al. [13] contribute the concept of tangible views for information visualization, Heilig et al. explore multitouch input for interacting with scatterplots [4], and Kosara [7] investigates multi-touch brushing for parallel coordinates. However, these first visualization approaches using novel interactive displays primarily address very specific problems. The broad range of possibilities of the new technology have by far not been explored sufficiently nor analyzed appropriately.

Closing this gap by systematically exploring the design space for combining modern visualization approaches with recent interaction technologies will lead to novel solutions for today's data exploration challenges.

Guidelines Gap

With the combination of different visualization techniques and interaction technologies, a vast body of possible solutions becomes available. This immense variety of existing and possible new approaches makes it difficult for users to decide which techniques to use. What is needed in the future are guidelines or rules for choosing effective approaches for the data, tasks, and device context at hand.

An excellent example of systematically choosing “good” visualizations is Mackinlay's [10] pioneering work on automat-

ing the design of visual representations. The beauty of this approach is that it enables automatic suggestion of visual variables based on a given data type (quantitative, ordinal, nominal). This is possible thanks to the well-defined sets of data types and visual variables, which abstract from the subtle details of real world problems. It is part of ongoing research how the details of today's often complex visualization application scenarios can be integrated.

Wouldn't it be great if we had a similar system to which we input our data D and our tasks T , and the system would tell us which visualization techniques V and interaction techniques I to use given a particular input and output setup? Obviously, the required mapping $guide : D \times T \rightarrow V \times I$ will be difficult to define. We consider solving this research question a formidable and rather long term task.

DISCUSSING THE GAPS

Narrowing and eventually closing the aforementioned gaps will require much research. It is beyond the scope of this article to comprehensively suggest directions for future work. We would rather like to use an example to illustrate possible avenues of investigation.

We chose the example of *tangible views for information visualization* [13] for the following reasons. Tangible views illustrate quite well the new possibilities of advanced technology with a set of different interactive visualizations (see Figures 1(a) and 1(b)). They also serve as a good illustration of what is still missing. Finally, since tangible views are our own previous work, it is easier to criticize and to envision research goals.

Conceptually, tangible views are spatially aware lightweight displays that serve two purposes simultaneously: visual output and direct input. Multiple of such tangible views are used together with an interactive tabletop display to build a multi-display multimodal visualization ensemble that supports both interacting *on* the views (by using touch and pen input) and interacting *with* the views (by moving them through the physical space above a tabletop and performing gestures). An interaction vocabulary (see Figure 2) has been compiled as the basis upon which manifold applications can be developed. Several example visualization cases illustrate how tangible views can be utilized to display and to interact with




Figure 2: Extract of the interaction vocabulary provided by tangible views. The figures show *with-the-view* interactions only, *on-the-view* interactions, such as pen and touch input, can be found in [13].

data. These use cases include scatter plots and parallel coordinates plots of multivariate data, node-link representations and matrix representations of graph data, as well as map-based visualization of spatiotemporal data.

Addressing the Technology Gap

In order to arrive at a mapping $interact : T \rightarrow I$, we first need a specification of the set of interaction tasks T . There are approaches that provide first categorizations of interaction. Yi et al. [15] describe a list of general interaction intents in visualization. Besides these general descriptions, more specific categorizations exist. For instance, dedicated interaction tasks for exploring graphs are described by Lee et al. [9]. These are valuable starting points for defining a comprehensive set of interaction tasks. Most likely, this set will contain tasks of different complexity ranging from very basic selection to common brushing and linking to the more complex applications of logically combinable dynamic filtering.

Secondly, defining an interaction vocabulary as in [13] is a valid first step for closing the technology gap. Such a vocabulary serves as a container that holds technically possible solutions to be utilized for interaction tasks. The tangible views vocabulary focuses on interaction with spatially aware lightweight displays. However, it is not comprehensively addressing the different classes of interactive displays in general. So, future work has to systematically extend the interaction vocabulary with further interaction techniques I .

An example for a successful *interact* mapping can be given for the task of exploring spatio-temporal data with tangible views. Such data can be mapped to the virtual volume above a tabletop, where the XY-dimensions encode spatial location and the Z-dimension represents time (i.e., space-time-cube [8]). In order to control which part of the geo-space and which time step are visible ($\in T$), the user can translate the tangible view horizontally and vertically ($\in I$), as shown in Figure 2(a).

Another example is the adjustment of a visualization parameter, e.g., the distortion factor of a fisheye lens ($\in T$), which can be mapped onto rotating the view horizontally ($\in I$) as shown in Figure 2(b).

Addressing the Integration Gap

Closing the integration gap, that is bringing together visualization research and new interactive displays, involves many different aspects. To name only a few, integration is nec-

essary on a conceptual level (e.g., utilizing tangible views for focus+context visualization), on a software level (e.g., combining different visualization and interaction toolkits), as well as on a hardware level (e.g., integration of lightweight displays with touch and pen input and tabletop displays). Because we cannot detail all aspects here, we will resort to illustrating the integration of *exploration* and *manipulation* of node-link diagrams of graphs as an example.

Usually, exploration tasks and manipulation tasks are considered separately from each other. While exploration is largely addressed in the visualization community, manipulation tasks are more relevant in the realm of human-computer interaction. For instance, with tangible views we mainly support the exploration of node-link diagrams by utilizing the *with-the-view* interaction (see Figure 1(a)). Other works address the authoring and manipulation of the underlying graph data, e.g., by using multi-touch and pen input for diagram editing as shown in Figure 1(c) [3]. Taking advantage of both worlds by integrating them into a single unified system would clearly be useful, not only because users could accomplish multiple tasks within the same familiar environment, but also because data exploration often involves data manipulation (at least temporary) for testing different “what if” scenarios.

However, such integration also implicates several challenges. On a conceptual level, distinctive features of different interaction modalities and visualization techniques need to be combined appropriately for different tasks. This could be achieved, for example, by utilizing *with-the-view* interaction for exploration aspects, while the more precise *on-the-view* interaction could be used for manipulation tasks. Seamless switching between these tasks could be accomplished by choosing different tools or even different interaction modalities, e.g., touch input for zooming/panning and pen input for graph editing. On a software level, different software worlds need to be consolidated into a single framework that addresses issues such as distributed rendering required for a multi-display setup, state synchronization between different devices, and most importantly the incorporation and adaptation of visualization techniques that meet the requirements of such a setup.

Addressing the Guidelines Gap

The developed example cases of tangible views indicate that there is much potential in utilizing new technologies for information visualization. Although being interesting exam-

ples, it remains unclear why and how tangible views are used under which circumstances and when alternative solutions might be better suited (as one reviewer of [13] once pointed out). So, there are often questions like *Would you really carry out this task with tangible views?* or *Wouldn't this be easier to accomplish with basic mouse interaction?*

Even though the introduction of an interaction vocabulary is an important step, there are still no definite rules for its application. In order to make information visualization on modern interactive displays a viable approach, we should strive to provide concrete answers and guidelines much like in the spirit of Bertin [1], Cleveland and McGill [2], and Mackinlay [10].

However, developing approaches for guiding the user in choosing the “right tool” is ongoing research, which is challenging for the following reasons. First, it is usually more difficult to categorize the data, because today’s data sets are increasingly complex and heterogeneous. Furthermore, one has to take the users’ tasks and goals into account with regard to both: what the users want to see and how they would like to interact. In terms of the output, a step has to be made from simple visual variables to more complex visualization techniques, and possibly to combinations thereof. The aspect of interaction is entirely missing in classic works. Given some data and a suitable visualization, how can the user effectively interact to accomplish the tasks and to achieve the goals? And finally, it is no longer just a question of which visualization technique to use for which data and task, but rather one of which display and interaction technologies to use for which visualization techniques, data, and tasks.

CONCLUSION

For taking full advantage of novel display and interaction technologies for information visualization, several gaps have to be addressed as identified and illustrated in this paper. First, a categorization of interaction tasks and a repertoire of novel interaction techniques have to be described, which then allows for mapping specific tasks to particular techniques. Second, the design space of combining novel interaction concepts and existing visualization approaches has to be explored appropriately. Third, guidelines have to be developed for choosing appropriate and effective approaches within a vast body of possible solutions. Filling these gaps step by step is a formidable task that can only be accomplished by a vivid research community bringing together visualization and interaction experts.

REFERENCES

1. Jacques Bertin. *Semiology of Graphics: Diagrams, Networks, Maps*. University of Wisconsin Press, 1983.
2. William S. Cleveland and Robert McGill. Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods. *Journal of the American Statistical Association*, 79(387):531–554, 1984.
3. Mathias Frisch, Jens Heydekorn, and Raimund Dachelt. Diagram Editing on Interactive Displays Using Multi-Touch and Pen Gestures. In *Proc. of Diagrams*, pages 182–196. Springer, 2010.
4. Mathias Heilig, Stephan Huber, Mischa Demarmels, and Harald Reiterer. Scattertouch: a multi touch rubber sheet scatter plot visualization for co-located data exploration. In *Proc. of the ACM International Conference on Interactive Tabletops and Surfaces (ITS)*, ITS ’10, pages 263–264. ACM, 2010.
5. Petra Isenberg and M. Sheelagh T. Carpendale. Interactive Tree Comparison for Co-located Collaborative Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6), 2007.
6. Petra Isenberg and Danyel Fisher. Collaborative Brushing and Linking for Co-located Visual Analytics of Document Collections. *Computer Graphics Forum*, 28(3), 2009.
7. Robert Kosara. Indirect Multi-Touch Interaction for Brushing in Parallel Coordinates. In *Proc. of the Conference on Visualization and Data Analysis (VDA)*, pages 786809–1–786809–7. SPIE/IS&T, 2011.
8. Menno-Jan Kraak. The Space-Time Cube Revisited from a Geovisualization Perspective. In *Proceedings of the 21st International Cartographic Conference (ICC)*, pages 1988–1995, Newcastle, UK, 2003. The International Cartographic Association (ICA).
9. Bongshin Lee, Catherine Plaisant, Cynthia Sims Parr, Jean-Daniel Fekete, and Nathalie Henry. Task Taxonomy for Graph Visualization. In *Proc. of the AVI workshop BELIV*, pages 1–5. ACM, 2006.
10. Jock Mackinlay. Automating the Design of Graphical Presentations of Relational Information. *ACM Transactions on Graphics*, 5(2):110–141, 1986.
11. B. Shneiderman. Human-computer interaction. chapter Direct Manipulation: A Step Beyond Programming Languages, pages 461–467. Morgan Kaufmann, 1987.
12. R. Spence. *Information Visualization: Design for Interaction*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2nd edition, 2007.
13. Martin Spindler, Christian Tominski, Heidrun Schumann, and Raimund Dachelt. Tangible Views for Information Visualization. In *Proc. of the ACM International Conference on Interactive Tabletops and Surfaces (ITS)*, pages 157–166. ACM, 2010.
14. Stephen Volda, Matthew Tobiasz, Julie Stromer, Petra Isenberg, and Sheelagh Carpendale. Getting practical with interactive tabletop displays: designing for dense data, “fat fingers,” diverse interactions, and face-to-face collaboration. In *Proc. of the ACM International Conference on Interactive Tabletops and Surfaces (ITS)*, pages 109–116. ACM, 2009.
15. Ji Soo Yi, Youn ah Kang, J.T. Stasko, and J.A. Jacko. Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, 2007.



**RESEARCH CENTRE
SACLAY – ÎLE-DE-FRANCE**

Parc Orsay Université
4 rue Jacques Monod
91893 Orsay Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399