Petra Isenberg, Michael Sedlmair, Dominikus Baur,
Tobias Isenberg, Andreas Butz (Editors)

# Collaborative Visualization on Interactive Surfaces - CoVIS '09

Technical Report
LMU-MI-2010-2, April 2010
ISSN 1862-5207

University of Munich
Department of Computer Science
Media Informatics Group

# Preface

Message from the Organizers

## CONTENT

This report provides an overview of current research trends in the field of collaborative visualization on interactive surfaces. It is based on a workshop (CoVIS 2009) held at VisWeek 2009 and consists of three main parts:

(i) the preface of authors, motivating and clarifying the underlying challenges;

(ii) a research agenda which we derived from fruitful discussions during the workshop; and

(iii) the peer-reviewed papers which were presented at the workshop.

The report comprises interdisciplinary aspects form the fields of information visualization, scientific visualization, visual analytics as well as CSCW and HCI. It is meant to help other researchers better understand the role and the growing impact of interactive surfaces as an emerging technology for supporting collaborative visualization and visual analytics settings.

## MOTIVATION

It is common for small groups of people to gather around visual displays of information to discuss or interpret the information to form decisions. Groups can share the task load of exploring large and complex datasets and can share various interpretations of a dataset when working together. However, tools to support synchronous collaboration between several people in their data analysis are still relatively scarce. Traditionally, visualization and visual analytics tools have been designed from a single-user perspective and for desktop computers. While hardware such as multi-touch displays and network capabilities—that lend themselves especially well to collaboration—have emerged, software support for collaboration around visualizations is still relatively scarce. One of the reasons is that single-user systems do not necessarily translate well to collaborative scenarios or interactive surfaces and require specific re-design. The design of digital systems for collaboration around visualization and visual analytics systems, therefore, poses additional challenges: we need to better understand (a) how people collaboratively work with visual representations of data and which methods they use to solve information analysis tasks as a team, and (b) what the exact design requirements are for collaborative visual analysis scenarios.

## THE WORKSHOP

### Technical Scope

We discussed these topics during a workshop at VisWeek 2009 in Atlantic City, USA. The workshop was open for discussion on issues pertaining to the design of collaborative information visualization, scientific visualization, and visual analytics systems on interactive surfaces and specifically on topics related to:

- the design of information visualization, scientific visualization, or visual analytics interfaces and environments for co-located collaborative work,
- design of interactive visual representations for collaborative work on interactive surfaces,
- the use of interactive surfaces to visualize and interact with information and data
- social components in collaborative visual analysis environments
- aspects of cognition in multi-user visualization and visual analysis environments
- evaluation of collaborative information visualization, scientific visualization, and visual analytics systems,
- multiple and coordinated views for collaborative visualization and data analysis systems,
- design of multi-display environments for information analysis work,
- collaborative visualization and visual analytics applications,
- collaborative sensemaking,
- experiences with traditional collaboration in information and data intensive fields.

### Format

The workshop was held as a half day event and we received ten peer-reviewed position papers of up to four pages in length each presented by one of the authors. The workshop was organized as an open event previous to the VAST/InfoVis/Vis conferences. Approximately 90 participants from academia and industry joined us during the workshop sessions. The workshop included two sessions with different foci: Interaction and Design Considerations. The paper presentations were short (10 minutes) and after each session we included specific time for discussions about “future directions in collaborative visualization”. We discussed upcoming challenges, problems, and possibilities and how they could be classified as well as specific solutions and approaches that can help to address the main issues mentioned, including the feasibility of proposed solutions and their respective strengths and weaknesses. In doing so, we gathered a lot of interesting and generalizable aspects which we summarized into a research agenda (see next section).

# Research Agenda

## Collaborative Visualization on Interactive Surfaces (CoVIS)

Discussions at CoVIS'09 revealed several future directions in collaborative visualization on interactive surfaces. In the following, we describe the main findings derived from these discussions and show their specific meaning for the visualization research community. Based on this and on current literature we propose a set of valuable topics for future research:

### Role of Interaction / Touch

Interaction with Visualizations with multiple people on touch-sensitive displays poses several interesting new challenges. In collaboration it is common that people work together using different types of data, representations, and using unique views of the data. Streit et al. [9] discuss challenges related to this problem and propose a multi-user, multi-level interaction paradigm by specifically separating data and view domains and proposing specific transitions between both. Annotation of data and views is an important type of interaction for collaborative systems to support and was discussed by [7].
In addition, moving visualization interaction from mouse to touch-based interactions offers both advantages and challenges. Isenberg et al. [4] discuss advantages of intuitive interaction with a flow visualization, support for awareness of interaction in the group, and support for free-form data exploration. Touch interaction also offers the possibility and challenges of designing new data interaction widgets as shown by Flöring et al. with their TaP system [1]. Jeong et al. discuss the possibility to interact with existing representation types but allowing the modification of multiple data dimensions at once using multi-touch interaction [5] which may offer a new data analysis experience.
When appropriate interactions have been designed for collaborative visualization tools, it is important to consider the generation and visualization of interaction histories for group work. Collaboration often occurs asynchronously with people dropping in and out of an analysis session. Sarvghad et al. [8] discuss a set of design objectives and challenges for collaborative history tools.
Overall, these are just some of the multiple new challenges that arise as information visualizations are being moved from the single-user desktop to multi-user multi-touch environments.

### Role of Multiple Display Environments (MDEs)

Often, there will not be a single but multiple displays representing the data. These multiple display environments (MDEs), on the one hand, can be static installations of multiple, possibly heterogeneous displays such as interactive tables, desktop monitors and projectors. Waldner et al. [10] for instance use these three display types in a system for multiple view visualizations.
On the other hand, users will most likely bring their own devices such as laptops, smartphones or tablet PCs along rising up the question of how to harness these dynamically, formed device setups. Streit et al [9], for instance, outline a system where the user can bring their own laptops and incorporate them into a MDE. Fuchs et al. [2] discuss a strategy for automatic device-based adaptation of visualizations considering different affordances of displays in dynamically forming MDEs. This strategy can be used to distribute visualization onto heterogeneous displays.
In general, to support collaborative visualization it is important to better understand how designers can make optimal use of multiple display resources and what specific requirements have to be met in order to support visual analysis of massive datasets in MDEs.

### Role of Users

Another question about designing future visualization systems on interactive surfaces is who the users will be. It is important to consider different roles in collaborative visual analysis such as observers, interactors, latecomers, distractors or discussion leaders, to think about the relation between diverse devices and diverse users, and derive implications based on these aspects.
Additionally, previously found aspects of collaboration around shared surfaces have to be included, mainly the separation between public and private content and support for social protocols. Also, user identification allows following a single user's actions and providing undo- and other history-related mechanisms or merging these different interaction histories for summarizing etc.

Kim et al. [6] present a model for user identification and the subsequent idea of interaction workspaces that allow a clear separation and personalization for multiple users around an interactive surface.
Isenberg et al. [3] describe several factors of the data exploration context (public vs. work setting) which influence users' goals, tasks, and objectives when using a collaborative data exploration tool.
All in all, further progress in hardware and conceptual models is necessary to give the central aspect of user identification in virtual collaboration the same place as in real-world scenarios.

## ROLE OF DESIGN PROCESSES

The process of designing a visualization system has a strong influence on its utility, usability and usefulness. Iterative processes using paper, low fidelity and high fidelity prototyping for repeated evaluation have been proposed and evaluated as useful tools in other communities. Understanding whether current methods still work for visualization applications on interactive surfaces or if they have to be refined, rethought or renewed is an important future challenge. Paper prototyping, for instance, will definitely be influenced by the collaborative, highly interactive and surface-oriented nature of systems.
To learn more about the real-world utility of our tools, our community also has to think about user-centered, participatory and contextual design approaches. An important aspect will be to understand how collaborative visual analysis can contribute to solve real problems, from real users and in real environments. "Designing it with the people not for the people" should be a central design guideline for building valuable tools.

## ROLE OF EVALUATION

Evaluation is central to gaining an understanding of interactive systems. In the context of collaborative visualization this does not only include usability evaluation of final tools or interaction techniques but also studies that inform the direction of the field itself. It is important to explore task sets for which measurements can be performed that allow us to gain insight into people's use of and performance with the provided interfaces. Because collaborative visualization often targets information exploration with visual representation, however, evaluation cannot only design controlled studies but also needs to embrace methodologies such as qualitative and observational studies. Also important is the location, the results from field studies may differ considerably from laboratory studies. For example, Mayar et al. [7] conducted a qualitative study of off-screen note-taking behavior around a digital tabletop display during collaborative data analysis tasks. They identified several roles of notes and open questions in how to design for note taking in collaborative visualization.

## DIFFERENCES BETWEEN SCIVIS, INFOVIS, AND VAST

The three different domains in visualization, scientific visualization, information visualization, and visual analytics differ considerably in the tools and visualization means they typically employ and the interactivity they usually provide. This has effects on collaborative visualization as a field. For example, in scientific visualization it is often important to interact with the 3D visualization space itself and important visualization elements have a location in space. Interacting with these on an interactive 2D surface is a challenge. Also, multiple people interacting with a global 3D space will cause conflicts. In information visualization, in contrast, the data typically does not have an inherent mapping into spatial domains, so that the visualization creates this mapping. This means that visualizations can be designed to better work on interactive surfaces, and also to support multiple interacting people. Visual analytics as the third domain centers its research agenda on using computational support for discovering information in huge amounts of data. Here the question is how to control the computational support using interactive surfaces and provide the necessary interfaces to the people who are exploring the data. Finally, a challenge also consists in trying to combine elements from all three domains, to allow people to collaboratively and interactively gain insight into today's complex datasets.

## ACCEPTED WORKSHOP PAPERS

[1] Stefan Flöring and Tobias Hesselmann, *TaP: Towards Visual Analytics on Interactive Surfaces.*
⇨ **p. 9**

[2] Georg A. Fuchs, Conrad Thiede, Mike Sips, and Heidrun Schumann, *Device-based Adaptation of Visualizations in Smart Environments.*
⇨ **p. 32**

[3] Petra Isenberg, Uta Hinrichs, Mark Hancock, Matthew Tobiasz, and Sheelagh Carpendale. *Information Visualization on Interactive Tabletops in Work vs. Public Settings.*
⇨ **p. 28**

[4] Tobias Isenberg, Uta Hinrichs, and Sheelagh Carpendale. *Studying Direct-Touch Interaction for 2D Flow Visualization.* ⇨ **p. 17**

[5] Dong Hyun Jeong, William Ribarsky, and Remco Chang. *Designing a PCA-based Collaborative Visual Analytics System.* ⇨ **p. 24**

[6] Kyung Tae Kim, Tejas Kulkami, and Niklas Elmqvist. *Interactive Workspaces: Identity Tracking for Multi-user Collaboration on Camera-based Multi-touch Tabletops.* ⇨ **p. 1**

[7] Narges Mahyar, Ali Sarvghad, and Melanie Tory. *Roles of notes in co-located collaborative visualization.* ⇨ **p. 13**

[8] Ali Sarvghad, Narges Mahyar, and Melanie Tory. *History Tools for Collaborative Visualization.* ⇨ **p. 21**

[9] Marc Streit, Hans-Jörg Schulz, Dieter Schmalstieg, and Heidrun Schumann. *Towards Multi-User Multi-Level Interaction.* ⇨ **p. 5**

[10] Manuela Waldner, Alexander Lex, Marc Streit, and Dieter Schmalstieg. *Design Considerations for Collaborative Information Workspaces in Multi-Display Environments.* ⇨ **p. 36**

## PETRA ISENBERG
INRIA

is a PostDoc in the AVIZ research group at INRIA, France. She recently finished her Ph.D. on the topic of 'Information Visualization in Co-located Collaborative Environments' at the University of Calgary, Canada, under the supervision of Dr. Sheelagh Carpendale. Her research focuses on supporting collaborative work with information visualization on shared multi-touch displays. She has designed or helped to design several collaborative information analysis environments, studied information analysis behavior in traditional environments, and co-authored two book chapters on collaborative visualization.

## MICHAEL SEDLMAIR
BMW Research

studied Media Informatics at the University of Munich, where he received his Diplom degree in 2007. He is currently finishing his Ph.D. at BMW Group Research and Technology in cooperation with the University of Munich, where he is supervised by Dr. Andreas Butz. His Ph.D. deals with methods of information visualization for in-car communication processes and their practical application in the automotive sector. He has designed several visualization concepts targeted to this special domain and also works on collaborative solutions. He is particularly interested in combining applied research in the fields of information visualization, HCI, and CSCW.

## DOMINIKUS BAUR
University of Munich

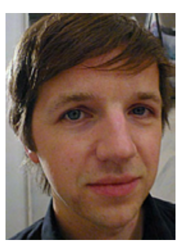

has a Diplom degree in Media Informatics from the University of Munich and is currently pursuing his Ph.D. under the supervision of Andreas Butz. In his research, he focuses on how people interact with media collections on tabletop displays and mobile devices. He also looks into ways to make personal media consumption understandable to casual users and interact with digital surfaces from a distance.

## TOBIAS ISENBERG
University of Groningen

received his doctorate in computer science in 2004 from the University of Magdeburg, Germany, working on shape analysis and non-photorealistic rendering. From 2004 to 2007 he was a postdoctoral fellow with Dr. Sheelagh Carpendale at the University of Calgary, Canada, where he combined computer graphics techniques and interaction research, resulting in a number of interfaces for direct-touch interaction for large-screen information displays. In 2007, he accepted a position as assistant professor for computer graphics and interactive systems at the University of Groningen, the Netherlands, and continues his work on interactive techniques, in particular, for non-photorealistic and illustrative rendering and visualization. He is interested in new scientific visualization applications for large interactive displays. He has previously been program chair, program committee member, and publicity chair for several events in computer graphics and visualization. He has been the program co-chair of NPAR 2009 and is currently organizing the upcoming CAe 2010 conference as general co-chair.

## ANDREAS BUTZ
University of Munich

received his Ph.D. in 1997 from Saarland University, Germany on the automatic generation of 3D animation clips. In 1998, he worked as a PostDoc at Columbia University, New York, USA on UIs for collaborative Augmented Reality. In 2000, he co-founded Eyeled GmbH, creating mobile information systems with a strong focus on UIs for small devices. In 2003, he returned to academia with a research group on UIs for ubiquitous computing, and soon assumed a tenured faculty position at Munich University (LMU) in 2004. In 2007, he won the Alcatel Lucent research award on "Technical Communication" for his work in ubiquitous user interfaces. Together with Antonio Krueger and Patrick Olivier, Andreas co-organizes the yearly Smart Graphics Symposium (http://www.smartgraphics.org/) since 2000. Andreas has published over 60 scientific papers and edited 12 conference and workshop proceedings, most of them in Springer's LNCS series. He was a Poster chair at IUI 2007 and ISMAR 2005, a video chair at Pervasive 2005, and a DC chair at UbiComp 2009.

## Program Committee

We thank our program committee for reviewing the papers and giving valuable feedback on our workshop design:

*Anastasia Bezerianos,* Ecole Centrale Paris

*Raimund Dachselt,* Otto-von-Guericke University Magdeburg

*David S. Ebert,* Purdue University

*Michael Haller,* Upper Austria University of Applied Sciences

*Nathalie Henry,* Microsoft Research

*Shahram Izadi,* Microsoft Research

*Daniel Keim,* University of Konstanz

*Johannes Schöning,* DFKI Saarbrücken

*Stacey Scott,* University of Waterloo

*Melanie Tory,* University of Victoria

# Table of Contents


*Session 1: Interaction*

*1.* Kyung Tae Kim, Tejas Kulkarni, Niklas Elmqvist
*Interaction Workspaces: Identity Tracking for Multi-User Collaboration on Camera-based Multi-touch Tabletops*
p. 1

*2.* Marc Streit, Hans-Jörg Schulz, Dieter Schmalstieg, Heidrun Schumann
*Towards Multi-User Multi-Level Interaction*
p. 5

*3.* Stefan Flöring, Tobias Hesselmann
*TaP: Towards Visual Analytics on Interactive Surfaces*
p. 9

*4.* Narges Mahyar, Ali Sarvghad, Melanie Tory
*Roles of notes in co-located collaborative visualization*
p. 13

*5.* Tobias Isenberg, Uta Hinrichs, Sheelagh Carpendale
*Studying Direct-Touch Interaction for 2D Flow Visualization*
p. 17

*Session 2: Design Considerations*

6. Ali Sarvghad, Narges Mahyar, Melanie Tory
*History Tool for Collaborative Visualization*
p. 21

7. Dong Hyun Jeong, William Ribarsky, Remco Chang
*Designing a PCA-based Collaborative Visual Analytics System*
p. 24

8. Petra Isenberg, Uta Hinrichs, Mark Hancock, Matthew Tobiasz, Sheelagh Carpendale
*Information Visualization on Interactive Tabletops in Work vs. Public Setting*
p. 28

9. Georg Fuchs, Conrad Thiede, Mike Sips, Heidrun Schumann
*Device-based Adaptation of Visualizations in Smart Environments*
p. 32

*10.* Manuela Waldner, Alexander Lex, Marc Streit, Dieter Schmalstieg
*Design Considerations for Collaborative Information Workspaces in Multi-Display Environments*
p. 36

# Interaction Workspaces: Identity Tracking for Multi-user Collaboration on Camera-based Multi-touch Tabletops

KyungTae Kim, Tejas Kulkarni, and Niklas Elmqvist, *Member, IEEE*


Fig. 1. Mockup screenshot of a bus route map application with two users (left), each one with a personal interaction workspace (right).

**Abstract**—The ability to distinguish between individual users in collaborative visualization is important for avoiding interference and facilitating cooperation. However, identification can be difficult on camera-based tabletops, where we typically cannot distinguish which tracked touch points belong to which user. Additional mechanisms are needed, such as video input, capacitance, or social and software protocols. We propose a model for the identification process on tabletop displays, and then use this model to suggest a software-based mechanism called interaction workspaces for supporting user identification. An interaction workspace is a movable and resizable region on the full visual space that belongs to a particular user, and inside which only that user is assumed to interact.

**Index Terms**—Digital tabletop displays, multi-touch, identity tracking, single-display groupware, horizontal interaction.

## 1 Introduction

Collaborative visualization may often benefit from the system being able to distinguish between the identities of the different users collaborating on the system [14], both in order to avoid interference between their respective actions as well as to faciliate role-specific operations nad cooperation between the users [3, 5, 15, 16]. Consider a scenario with two users working on developing a bus route map for an urban area on a single-display groupware [17] device (Figure 1). When user A tries to sketch a route from point P1 to P2, while user B sketches a route from point P3 to P4, a system without user identification might recognize these two gestures as coming from a single user and may perform an incorrect operation, such as zooming the whole map. (Refer to Figure 2)

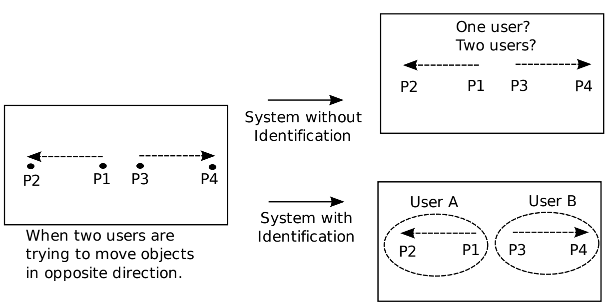


Fig. 2. Example showing that system could interpret the data in different ways.


- *KyungTae Kim is with Purdue University in West Lafayette, IN, USA, E-mail: kimk@purdue.edu.*
- *Tejas Kulkarni is with Purdue University in West Lafayette, IN, USA, E-mail: tkulkarn@purdue.edu.*
- *Niklas Elmqvist is with Purdue University in West Lafayette, IN, USA, E-mail: elm@purdue.edu.*


Another situation where user identity becomes critical is when multiple users are working on developing unique routes in the same visual space. Without user identification, if one user wants to undo modifications to the routes, the system will not be able to distinguish between modifications by different users. The system may inadvertently affect other users' work in that space, causing interference between the users.

Multi-touch [9] tabletop displays are becoming increasingly popular for this kind of co-located collaborative work [8] due not only to the natural user interface afforded by the user's own fingers [6], but also by virtue of camera-based tabletops being relatively inexpensive and easy to build [7]. However, while tabletop displays that rely on capacitance tracking, such as DiamondTouch [2], do support user identity, the more inexpensive camera-based tabletops [7] (of which Microsoft Surface is an example) do not. Camera-based tabletops merely track multi-touch data in the form of "blobs", which represent points of contact on the tabletop surface by objects such as fingers. In order to distinguish this multi-touch data as belonging to different users, these systems must be augmented with additional mechanisms. Examples of such mechanisms include simple distance-based heuristics, camera-based image processing [3], social protocols [10], or software mechanisms such as storage bins [12] or user territories [13]. However, all of these mechanisms are currently independent and cannot easily be combined, compared, or contrasted.

In this position paper, we propose to model the identification problem as a general transformation process which accepts multi-touch data and classifies this into multi-user multi-touch data. The identity classifier itself depends on the available hardware, software, and physical properties of the tabletop system. The model is general enough to describe all of the above known identification mechanisms for tabletop displays. Extending on previous work, we also propose a simple software-based classifier called *interaction workspaces* that permit fine-grained access control of the visual space by restricting users to personal work areas (Figure 1), but which relies on a social protocol for high-level access control (i.e., protecting your own workspace).

This position paper is structured as follows: First, we discuss the background of multi-user systems and the existing literature. We then describe the design of our process model that assigns unidentified blobs to their respective owners. We go on to present our extension to software-based identity classifiers: interaction workspaces. We close the paper with our conclusions and plans for future work.

## 2 BACKGROUND

As computers become more integrated in society, there has been increasing demands for co-located collaboration in multi-user environments [14]. However, traditional desktop computers typically use a standard keyboard and mouse setup that allow only a single active user at a time, rendering them unsuitable for multi-user collaboration [17].

Recent advances in both display and input technology has made a new generation of collaborative devices available for co-located collaboration [14]. Among these, digital tabletop displays are particularly suitable for co-located collaboration due to their similarity to traditional tabletops, their natural interaction using fingers or styli, as well as their obvious collaborative affordances [2, 18]. Furthermore, camera-based tabletop displays, such as FTIR [7], are inexpensive and relatively simple to build. However, a necessary—but, as we shall see below, not sufficient—requirement for these devices to support multiple concurrent users is that they first support multiple concurrent touch points on the display. This property is known as *multi-touch sensing* [7, 9].

### 2.1 Recognizing User Identity

Multi-touch systems are generally quoted as also being multi-user capable due to being able to distinguish between multiple points of contact, regardless of whether these points stem from a single user or multiple users. However, this is a qualified truth. While capacitance-based tracking, such as DiamondTouch [2], can trivially distinguish contact points for up to four users, camera-based tabletops generally cannot. They detect each point of contact (from a finger, hand, or other object) as a "blob" on the surface of the tabletop, but have no way of assigning a user identity to each individual blob [7].

However, the ability to recognize and track user identity in co-located collaboration may be critical to resolve interference and conflicts between users. For example, a system with no identity tracking will not be able to tell the difference between a single user interacting with both hands, or two users each interacting with a single hand. On a higher level, the lack of identity tracking may cause conflicts for shared objects on the display, and does not support assigning roles to users. Groupware has typically relied on social protocols—spoken or unspoken conventions between users about object ownership—for managing these problems, but this is not sufficient for many situations [5, 10].

### 2.2 Classification of Touch Technologies

As argued above, there is clearly a distinction between supporting multiple concurrent touch points and supporting multiple concurrent users. In Figure 3, we classify existing touch technologies based on the number of concurrent users (single or multiple), as well as the number of concurrent touches they support per user (single or multiple).

As we can see, camera-based multi-touch systems (as well as and mobile touch screen devices such as the iPhone), fall under the category of multi-touch single-user systems. These types of system face the problem of user identification as described above. A select subset of touch technologies, such as primarily DiamondTouch [2], are not only multi-touch but also multi-user systems.

### 2.3 Multi-user Multi-touch Tabletops

Most single-display groupware [17] projects in the literature which need multi-touch and multi-user support utilize the aforementioned

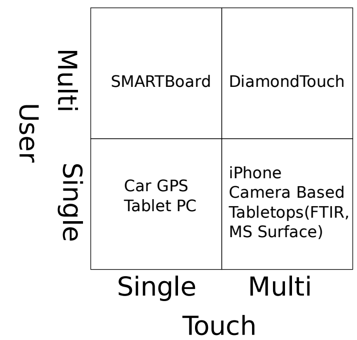


Fig. 3. Classification of existing touch technologies depending on their multi-touch and multi-user capabilities.

DiamondTouch [2] tabletop, developed by Mitsubishi Electric Research Laboratories (MERL). The device supports up to four unique users, each capacitively connected with a receiver in their chairs, and detects touch points as users close the circuit by touching the table surface. Because of this design, the display needs to be front-projected.

Despite the widespread success of the DiamondTouch device, both in research as well as industry applications, this tabletop is limited to a specific form factor, whereas camera-based tabletop technologies are more flexible in terms of size, and also more inexpensive to build [7]. Front-projection may also be a problem due to shadows from the users' hands obscuring the display, especially in crowded settings. Hence, for situations where a large interaction surface or support for many participants (or both) are required, DiamondTouch may not be the ideal choice, regardless of the excellent general features of the device.

### 2.4 Separating Multi-touch Data into Multiple Users

Clearly, the vast majority of multi-touch technologies are **not** true multi-user capable. At the same time, camera-based multi-touch tabletops of this kind have considerable potential for co-located collaborative work, given that the user identification problem can be solved. To achieve this, we propose a process model that transforms multi-touch data, in the form of touch points coordinates, into multi-user multi-touch data—separating blobs into their respective users—using an *identity classifier* (Figure 4). We define an identity classifier as a function to that maps unidentified finger blobs to its respective owners using different kinds of inputs.

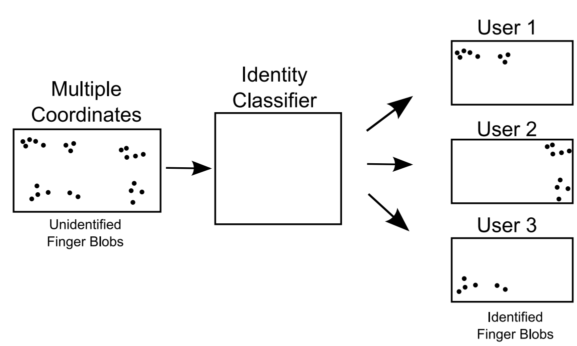


Fig. 4. Identity classification process for recognizing multiple users in a single-user multi-touch tabletop.

In the following section, we suggest several methods of implementing identity classifiers, most of which will be heuristic-based in nature.

### 2.5 Identity Classifiers

Distance-based: Perhaps the most straightforward identity classification heuristic is based on calculating distances between blobs and then connecting them according to the minimum distances. The intuition is that finger blobs belonging to a single hand are typically clustered close together, and two hands belonging to a single user are correspondingly clustered close together. According Epps et al. [4], users use their index finger or spread hand for about 80% of the time while interacting with a multi-touch tabletop. This behavior could be exploited to further enhance the performance of the classifier. Of course, this simple heuristic breaks down in the face of complex cooperative interactions between users, such as crossing arms or performing very large-scale gestures.

Camera-based: A powerful hardware-based classification heuristics rely on augmenting the camera-based tabletop with additional cameras—typically placed above or below (in this case, seeing through the surface to the participants hands above) the tabletop—and performing image processing on the live video feed. One possible implementation is to add an infrared camera above the surface and an infrared light source with different wavelength than that used in the acrylic of the FTIR tabletop, allowing the camera to track the movement of the users. Once rough coordinates of user hands are acquired, the classifier calculates the distance between blobs and these hands, assigning them to the right owner.

Dohse et al. [3] present a similar hardware setup based on a visible spectrum camera mounted above a FTIR tabletop that tracks hands using skin color segmentation techniques. This way, both the robustness of the multi-touch functionality can be increased, as well as user identities be tracked and used to differentiate touch points. The approach has the added capability of being able to detect hover gestures that are executed without actually touching the tabletop surface.

Territory-based: A software-based identification heuristic is to divide the surface of the tabletop into discrete *territories* [13], and to regard any interaction occurring in a territory as belonging to the owner of that territory. This relies on social protocols enforced not by the system but by the collaborating users themselves, something which may not be sufficient in many situations [10]. Furthermore, the strict and non-flexible division of the visual space detracts from the multi-user experience and may decrease the efficiency of the collaboration.

## 3 Interaction Workspaces

In this section, we propose a novel software-based identity classifier called *interaction workspaces* that supports fine-grained access control of the visual space by assigning movable and resizable regions to individual users. We first give a motivating scenario and then describe the details of the technique.

### 3.1 Scenario

Suppose Pete is working as part of a team redesigning the bus routes for Greater Lafayette, Indiana. John and Mary, Pete's co-workers, are familiar with the area and are already working collaboratively on a multi-touch tabletop to design the bus routes. Pete arrives, and proceeds to place his palm on an unoccupied part of the tabletop space for two seconds. In response, the system brings up a window to register his name. After registering himself, Pete defines his region of space—his *interaction workspace*—by using two fingers and dragging across the tabletop to specify the borders of the workspace. Once the workspace is instantiated, a simple tap on his name tag associates the interaction workspace with his identity.

At this point, the interaction workspace is now owned by Pete and any interaction inside its borders will be classified as belonging to him. However, Pete is also now responsible for physically protecting his workspace in order not to come into conflict with John and Mary.

In the middle of process, Pete makes a mistake and wants to undo his recent changes. By using the workspace toolbox for accessing the workspace history, the system will revert his work to any specified previous state. Then, after finishing work in his workspace, John needs Pete's help to improve his route, so John releases his workspace so that it is free. Pete can now destroy or release his current workspace and move to where John was working. By tapping on the released workspace, the system will ask Pete about his identity and then allow him to acquire the workspace to continue the collaborative editing.

### 3.2 Basic Concept

An *interaction workspace* is a well-defined sub-area of a visual display—its extents indicated using visual borders—that is considered to be owned by a particular user (Figure 1). Access to the workspace is assumed to be enforced by participants using a social protocol, so all interaction within the boundaries of the workspace is regarded as belonging to the workspace owner. Workspaces can be moved and resized. Because an interaction workspace can essentially be regarded as a mutual exclusion lock of the visual space, different workspaces are not allowed to overlap since this may cause ambiguities.

Beyond move and resize operations, interaction workspaces can be created and destroyed. They can also be released and acquired, allowing for sharing among users (see Section 3.3). Furthermore, each workspace may have an additional tool interface with operations local to the workspace (Section 3.4), as well as may maintain an interaction history to support undo and redo operations (Section 3.5).

Interaction workspaces are **not** independent desktops—they are intended to be used on full-display groupware applications (see scenario). Furthermore, by virtue of being resizable and movable, interaction workspaces are generally more flexible and provide more fine-grained control than many other territory-based mechanisms [13].

### 3.3 Workspace Management

**Creating** a new interaction workspace is the first step for a user to be able to collaborate in a multi-user environment (such as the bus route design scenario above). During this process, the user will be asked to authenticate himself and then indicate the position and extents of the new workspace. Analogously, **destroying** a workspace will remove it from the visual display (although it can be recovered again by a global undo operation, see Section 3.5). Beyond these basic management operations, users can also **release** a currently owned workspace to share it for others to use, and they can correspondingly **acquire** a workspace with no current owner.

Creation, authentication, acquisition, release, and destruction are all global operations that rely on a social protocol for enforcement. To augment this protocol and to decrease the burden on the participants, additional rules can be enforced on the visual display, such as limiting the number of workspaces that one user can own, as well as the maximum allowable size for a single workspace.

### 3.4 Workspace Tool Interface

Most groupware applications that support more than just basic navigation tasks require an application-specific tool interface. For example, in our bus route design scenario, there should be tools for drawing and deleting routes, as well as adding bus stops. Instead of having an interface toolbox associated with each user, every interaction workspace has its own toolbox. This will help both the users and the system. On the user side, this allows for customizing the toolbox of a particular workspace to fit the current task (Figure 5. It also provides a natural interface to accessing the interaction history of a workspace for undo or redo operations (see Section 3.5).

On the system side, minimizing the global interface controls and associating toolboxes to workspaces means that the system need not distinguish between multiple users accessing the same, global tool interface.

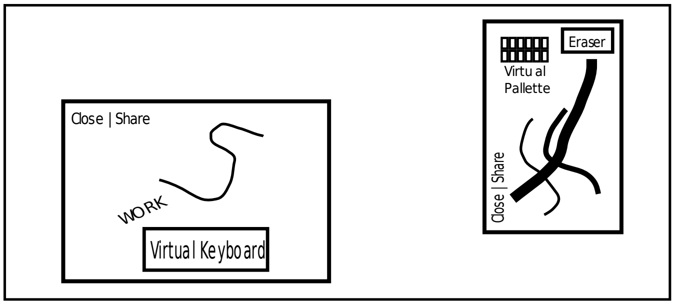


Fig. 5. Example of each interaction workspace containing a unique and independent tool interface.

### 3.5 Interaction History

The main purpose of interaction workspaces is to improve coordination on multi-touch tabletops. However, sometimes even our software and social protocols may not be sufficient to avoid conflicts. For such situations, we propose to add interaction histories that support redo and undo operations. Because of the structure of our workspace mechanism, we propose a two-tiered history mechanism:

1. **Global history:** Global event log, such as for workspace creation, deletion, release, and acquire operations (see Section 3.3).

2. **Workspace history:** Interaction log for operations conducted inside individual workspaces (see Section 3.4).

In other words, the system stores a history of workspace management operations since startup, and each interaction workspace in turn keeps track of changes made to the objects within that workspace. Accessing these histories can be done using redo/undo buttons placed on the global display, for the global history, and on each workspace tool interface, for the workspace history.

A problem might arise in situations when objects within different interaction workspaces are interdependent. If an outcome of an action within a workspace affects the outcome of an action within another workspace, user interference might restrict certain changes within those workspaces. We suggest developing a global dependency graph, compromising of objects irrespective of their parent interaction workspaces, to warn users in case of interference between changes being made. A dependency check flags a warning anytime changes to inter-related objects within different workspaces are performed. This mechanism provides a robust technique to resolve dependency issues between objects and prevents user from violating each other's work.

Furthermore, interaction histories in groupware systems are clearly suspectible to conflicts arising from interference between the actions of different users [1, 11]. Beyond the static dependency graph of domain objects described above, this will require maintaining a dynamic dependency graph for individual user operations in the interaction history. We anticipate exploring this issue further in our future work.

## 4 Conclusions and Future Work

We have proposed a general model to describe the user identification process for camera-based multi-touch tabletop displays. This model enables us to describe a wide variety of identification mechanisms in terms of an identity classifier that accepts input from sources such as overhead cameras, distances between touch points, and territories, and then utilizes this data to distinguish interactions made by different users collaborating on the tabletop. We have also proposed a straightforward classifier based on non-overlapping interaction workspaces that are owned by a particular user.

This is mostly a conceptual paper, and we anticipate implementing and evaluating many of the ideas put forth here in the future. For example, we are interested in replacing our "soft" social protocol for high-level access control of workspaces by a more strict hardware-based protocol based on image processing and an additional overhead camera, similar to the method employed by Dohse et al. [3]. We are also interested in studying conflict resolution and avoidance for our workspaces, particularly in the presence of conflicting histories.

## References


[1] G. D. Abowd and A. J. Dix. Giving undo attention. *Interacting with Computers*, 4(3):317–342, Dec. 1992.

[2] P. H. Dietz and D. Leigh. DiamondTouch: A multi-user touch technology. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 219–226, 2001.

[3] K. C. Dohse, T. Dohse, J. Still, and D. J. Parkhurst. Enhancing multi-user interaction with multi-touch tabletop displays using hand tracking. In *Proceedings of the International Conference on Advances in Computer-Human Interaction*, pages 297–302, 2008.

[4] J. Epps, S. Lichman, and M. Wu. A study of hand shape use in tabletop gesture interaction. In *Extended Abstracts of the ACM CHI 2006 Conference Human actors in Computing Systems*, pages 748–753, 2006.

[5] S. Greenberg and D. Marwood. Real time groupware as a distributed system: Concurrency control and its effect on the interface. In *Proceedings of ACM Conference on Computer-Supported Cooperative Work*, pages 207–217, 1994.

[6] V. Ha, K. M. Inkpen, T. Whalen, and R. L. Mandryk. Direct intentions: The effects of input devices on collaboration around a tabletop display. In *Proceedings of IEEE TableTop*, pages 177–184, 2006.

[7] J. Y. Han. Low-cost multi-touch sensing through frustrated total internal reflection. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 115–118, 2005.

[8] P. Isenberg and M. S. T. Carpendale. Interactive tree comparison for co-located collaborative information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1232–1239, 2007.

[9] S. Lee, W. Buxton, and K. Smith. A multi-touch three dimensional touch-sensitive tablet. In *Proceedings of the ACM CHI'85 Conference on Human Factors in Computing Systems*, pages 21–25, 1985.

[10] M. R. Morris, K. Ryall, C. Shen, C. Forlines, and F. Vernier. Beyond 'social protocols': multi-user coordination policies for co-located groupware. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, pages 262–265, 2004.

[11] A. Prakash and M. J. Knister. A framework for undoing actions in collaborative systems. *ACM Transactions on Computer-Human Interaction*, 1(4):295–330, Dec. 1994.

[12] S. D. Scott, M. S. T. Carpendale, and S. Habelski. Storage bins: Mobile storage for collaborative tabletop displays. *IEEE Computer Graphics and Applications*, 25(4):58–65, July/Aug. 2005.

[13] S. D. Scott, M. S. T. Carpendale, and K. M. Inkpen. Territoriality in collaborative tabletop workspaces. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, pages 294–303, 2004.

[14] S. D. Scott, K. D. Grant, and R. L. Mandryk. System guidelines for co-located, collaborative work on a tabletop display. In *Proceedings of the European Conference on Computer Supported Cooperative Work*, pages 159–178, 2003.

[15] C. Shen, N. Lesh, F. Vernier, C. Forlines, and J. Frost. Sharing and building digital group histories. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, pages 324–333, 2002.

[16] O. Ståhl, A. Wallberg, J. Söderberg, J. Humble, L. E. Fahlén, A. Bullock, and J. Lundberg. Information exploration using the pond. In *Proceedings of the International Conference on Collaborative Virtual Environments*, pages 72–79, 2002.

[17] J. Stewart, B. B. Bederson, and A. Druin. Single display groupware: A model for co-present collaboration. In *Proceedings of the ACM CHI 99 Conference on Human Factors in Computing Systems*, pages 286–293, 1999.

[18] P. Wellner. Interacting with paper on the DigitalDesk. *Communications of the ACM*, 36(7):86–96, July 1993.

# Towards Multi-User Multi-Level Interaction

Marc Streit, *Student Member, IEEE*, Hans-Jörg Schulz,
Dieter Schmalstieg, *Member, IEEE*, and Heidrun Schumann

**Abstract**— The necessity of incorporating experts from various domains in order to understand and draw meaningful conclusions from complex and massive amounts of data is an undisputed fact. In order to create and effectively use such a collaborative information workspace it is vital to understand the interaction processes involved. Established, high-level interaction patterns work well for single user, single data source scenarios. However, they cannot simply be applied to the collaborative analysis of heterogeneous data.
In this paper we propose a *Multi-User Multi-Level Interaction* concept which differentiates between operations in view and data domain while considering the relations and transitions between data on different levels of granularity. Hence, the users' interaction can be formalized as a seamless path of navigation. This in turn helps to gain a deeper understanding of the interaction process and allows to efficiently steer it to accelerate data analysis. We demonstrate the applicability and benefits of our concept by means of a clinical use case scenario which aims at finding the best treatment for cancer patients.

**Index Terms**—Interaction, Information Seeking Mantra, visual analysis, collaboration, multi-display environment.

## 1 Motivation and Background

Interdisciplinary applications require the integration of domain experts from various fields in the data analysis process. Each of these experts has a specific perspective on the data, pays attention to different details, and reasons along the lines of his/her own particular domain. Organizing this multifaceted interplay between large amounts of complex data, multiple domain experts from different areas, and the laborious back and forth between exploration and confirmation of the analysis process is a challenging task. This task is what collaborative environments have set out to support and to advance, as the results that can be gained from an interdisciplinary, collaborative data analysis outweigh technical problems. One essential problem is interaction with the complex, heterogeneous data spaces in these environments: due to the multidisciplinarity, data is available in various forms (full text documents, images, statistical tables, etc.), in various representations (tabular, tag clouds, visualizations, etc.), and on multiple levels of detail. Each of which is meaningful to at least one of the participating domain experts and all of them need to be integrated into one seamless, interactive analysis process to allow fruitful collaboration.

State-of-the-art applications and interaction paradigms mostly focus on single user interaction and are tailored to one specific application domain. Yet, for a scenario as described above the established tried and tested interaction patterns do not suffice. Hence, new multi-user interaction concepts for data from different domains and on multiple levels of detail must be established. We do so in this paper by introducing a novel concept that addresses the challenges posed by multi-user, multi-level interaction. The applicability of this concept is exemplified by the analysis of clinical data from cancer patients in a collaborative information workspace described in the companion paper by Waldner et al. [12] (see Figure 1). In this case, biomedical experts from different fields come together to collaboratively analyze their respective data to make a joint decision on a patient's diagnosis and further treatment plans. In detail, the experts and their data are:

- **the oncologist:** CT/MR-scan of the tumor, treatment history
- **the pathologist:** tissue samples of the tumor autopsy
- **the geneticist:** data on the genome-wide regulation of the genes
- **the biologist:** genes' regulation in the context of the cellular processes, i.e. pathway graphs

Although each expert has his/her core field of expertise (i.e. data), they often also have profound knowledge in related domains. The data forms a natural hierarchy shown in Figure 2. This illustrates the multi-level aspect of the data, emerging naturally from the interdisciplinary setup. This is not a special case, but occurs frequently, as other examples of such hierarchies show – e.g., of the assembly hierarchy of a whole network of electronic devices down to the individual logic gate in the field of electrical engineering [1] or the refinement process in software engineering from the specification documents down to the actual code.

The collaborative workspace being utilized by our use case is an adaptation of the Caleydo Biomedical Visualization Framework (http://www.caleydo.org) [9, 10] for the Deskotheque multi-display environment [5]. Similar setups have been described, e.g., for an office environment called *"The office of the future"* [6] or for an entertainment scenario called *"Smart Living Room"* [3]. Because of this diversity of possible applications, this paper discusses not only the implications of our concept for our special use case, but also its consequences for these setups in general.

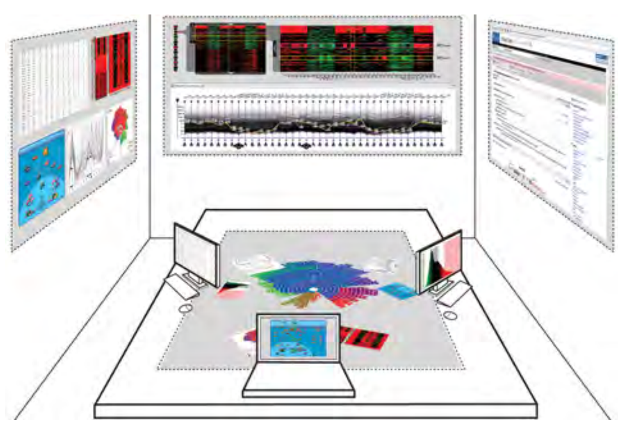

Fig. 1. Illustration of a collaborative information visualization scenario in a multi-display environment [12].

## 2 Multi-User Multi-Level Interaction

Most of the common data exploration patterns work well in a single user, single data source scenario. A particularly successful example is Shneiderman's *Information Seeking Mantra* (Overview first - zoom and filter - details on demand) [8]. While this pattern is of course also valid for more than one user and more than one data source, the characteristic properties of these scenarios are not captured by it. The two most crucial ones are:

- *Marc Streit and Dieter Schmalstieg are with Graz University of Technology, E-mail: {streit,schmalstieg}@icg.tugraz.at.*
- *Hans-Jörg Schulz and Heidrun Schumann are with University of Rostock, E-mail: {hjschulz,schumann}@informatik.uni-rostock.de.*

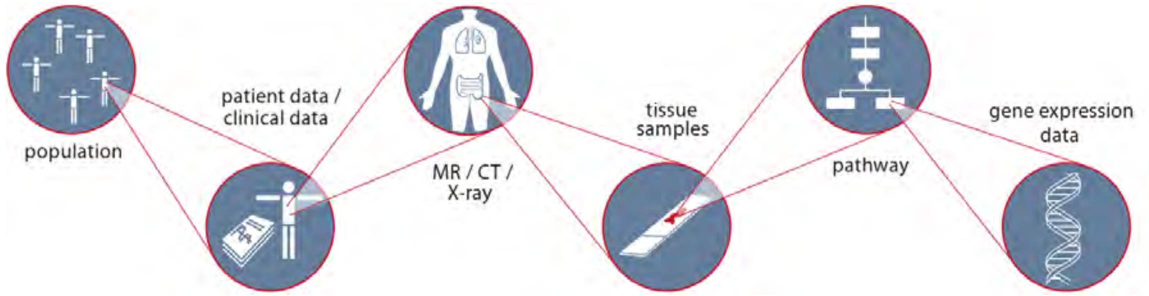


Fig. 2. Patient-centered data hierarchy starting with a population on the top and going down to the genes' regulation data of each patient. The detail of one level is the overview of the next level beneath.

- **Seamless navigation**, which describes the possibility to browse the data smoothly across the boundaries of the individual data sources. This property is important, as interaction patterns have to bridge the different data domains to ensure seamless navigation. For the *Information Seeking Mantra*, this means that for example, the detail of one data domain is the overview of another one and vice versa.
- **Integral data analysis**, meaning the common practice to intertwine visual and algorithmic analysis in the spirit of *Visual Analytics*. For Shneiderman's Mantra, e.g., the non-visual "zoom out" could be an aggregation operation, whereas the "zoom in" could be a query refinement. The users can then choose whether to do a visual analysis or to switch to the available algorithmic tools and use them in combination with the visual ones, according to which are the best fit to the analysis task at hand.

The latter of these two properties has been addressed by Keim in his *Visual Analytics Mantra* (Analyze first - show the important - zoom, filter and analyze further - details on demand) [4] and of course it also remains valid in the context of multiple users and data sources. Yet, one of the most important points for our scenario, namely the integration of the visualization from overview to detail across all data domains remains unspecified by Keim's mantra, making it not straight-forward to be applied here. Hence, this section embraces the *Information Seeking Mantra*, as well as the *Visual Analytics Mantra* and proposes a novel interaction concept that captures the above two points. The applicability of our concept and its practical benefits are discussed in a concrete biomedical application case.

### 2.1 Concept

The main idea, which is outlined in Figure 3, is to make a distinction not only between data and view domain [7], but also between the different application levels. This way, jumps and switches between data and view, and also between different application levels can be expressed by the concept. This is important, as different users in our collaborative scenario are responsible for different parts of the analysis and different levels of the data.

Basically, the concept applies the steps of the *Information Seeking Mantra* to each data level in the application hierarchy. And it does this not only for the view domain, but also for the data domain. This allows to differentiate in which domain operations are performed and yields the following categorization:

- **View operations** only affect the visual representation of the data. Examples are distortion based lens effects, geometric zoom, etc.
- **Data operations** affect the data by algorithmic means, from simple numerical operations to complex data mining methods.
- **Data+View operations** affect both domains. An example is any "Visual Query" mechanism that is triggered by the user in the visual representation, which carries out a query in the data domain, and reflects the result as a change in the view domain.

As a consequence, the overall multi-user interaction process of the different users' operations forms a path up and down the application levels and across the data and view domain. The seamless transition between multiple application levels is ensured by the assumed hierarchical nature of the data sources. The stippled lines in Figure 3 illustrate this natural shift from one level into the next.

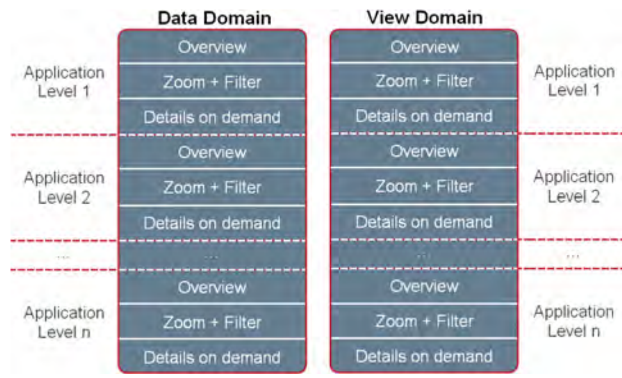


Fig. 3. The *Multi-User Multi-Level Interaction* concept separates the view and data domain while considering the relations and transitions between the application levels.

Conceptually, all analysis paths across multiple domain levels (in our case from the population of patients down to their individual gene expressions) and across multiple interaction levels therein (from overview to detail) are possible. However, in real world scenarios restrictions for the navigation are introduced. The constraints can either be implied by the nature of the data (e.g., missing data) or by the role of the analyst (e.g., security clearances). The knowledge about the constraints allows a guidance of the users through the domains and levels. In some situations multiple paths of interaction lead to the same result for the users. While one path could be potentially faster, another one might better support the users at keeping their mental map. In such cases the application designer can actively guide the user by promoting certain interaction paths, but without denying any of the other possibilities. While for example an expert user would take the faster way, the novice user should be guided along the path which supports the mental map best. Therefore, the awareness of the application designer of these constraints is vital in order to provide suitable visualization and interaction techniques. This knowledge also enables the application designer to preprocess most of the needed data along the most promising exploration path, in order to prevent time-consuming switches to the data domain and back. E.g., if a clustering is already precomputed, it is readily available to be included in the view domain and the interactive exploration process can continue instantly.

### 2.2 Use case

For demonstrating the proposed interaction concept we chose an exemplary use case from our clinical scenario: experts from four domains meet to discuss the treatment of a cancer patient. The use case bases on feedback from our medical partners on the Caleydo software as well as by studying their offline workflow in everyday collaborative situations. By accessing patient data from the whole spectrum of application levels (cf. Figure 2), the biomedical experts perform a collaborative analysis. Table 1 shows the interaction path through the data and view domain. In addition, the table states which domain experts actively interact on which level of the data hierarchy for each task.

When examining this and multiple other analysis paths from the biomedical application domain, we encountered certain reoccurring patterns within the extracted flow of visual data analysis. Three notable examples are:

- **The Information Seeking Mantra**, which often remains intact, if not interrupted by switches between different data sources. This can be observed, e.g., in steps 3-5 and 9-11.
- **Visual queries**, which are triggered by performing some action in the view representation. They carry out a query in the data domain and reflect the result as a change in the view domain. Again, this pattern occurs usually with the application level staying the same. This can be observed, e.g., in steps 11-13.
- **View to data switches** that occur by themselves and not as a part of a visual query pattern, are mostly switches between different application levels that are not seamlessly supported visually. If the switch would be seamless, the users could just use the detail view of the higher application level as the overview of the underlying one. If that is not possible, the users have to switch back to the data domain and generate a different representation before they can proceed. Examples for the seamless transition are, e.g., steps 5 and 7, examples for view to data switches can be seen in steps 13-14 and 16-17.

As a transition between application levels usually implies a switch of the analyst in charge, e.g., from the biologist to the geneticist and oncologist from step 16 to 17, once made explicit, these patterns help to effectively coordinate the analysis process throughout all different domain levels and between the different experts in a multi-user multi-level scenario like this.

## 3 Implications

As the concept is introduced and demonstrated by means of a real world analysis example, the next step is to discuss the hence resulting implications. While the first part addresses general considerations from the concept, the second section discusses the specific implications for smart environments.

### 3.1 General Implications

Although a seamless multi-level application hierarchy may exist, in some cases there can be data missing for one or more levels. The reasons can range from restrictions due to security concerns to the irrelevance of certain data sources for a specific use case scenario. For example, in the biomedical application scenario presented in Figure 2, an analysis task could aim at the discovery of new gene functions for which the magnetic resonance images and tissue samples are not needed. However, for providing a seamless (visual) transition from patients to pathways the missing levels are crucial to keep up the users' mental map. One possible approach to fill these gaps in the hierarchy is the integration of reference or sample data sources. This makeshift could be taken from external sources or alternatively also be extracted from available reference data sets. In our case this could be anatomical atlases or data from patients with similar medical records. These data sets bridging the gaps have to be explicitly marked as such, so that it becomes obvious that they are just means to facilitate a smoother exploration and analysis and are not part of a patient's data set.

The opposite to the absence of data in the hierarchy can also occur: the availability of multiple facets of the same data at the same level. Examples in terms of our use case are data sets on the organ level acquired by different imaging techniques – e.g., magnetic resonance, computer tomography, and X-ray images. These multiple facets of the same data introduce an ambiguous navigation path between the levels, where it is unclear which path to choose through the hierarchy. At this point the users' profiles and roles during the analysis can help

| # | Task description | User | App. | Data | View |
|---|---|---|---|---|---|
| 1 | Show overview of cancer patient | onc,pat | | | Ovv |
| 2 | Investigate patient's clinical data | onc,pat | | | Z+F |
| 3 | Show tumor CT scan | onc | | | Ovv |
| 4 | Investigate finer tumor structures | onc | | | Z+F |
| 5 | Show microscopy images of tissue | pat | | | DoD = Ovv |
| 6 | Search patients with similar tumor | onc,pat | | Z+F | |
| 7 | Show list of resulting patients | onc,pat | | | DoD = Ovv |
| 8 | Cluster gene expression of these patients | gen,pat | | Ovv | |
| 9 | Show hierarchical heat map | gen,pat | | | Ovv |
| 10 | Select cluster | gen,pat | | | Z+F |
| 11 | Select gene | gen,pat | | | DoD |
| 12 | Query gene in online database | gen,pat | | DoD | |
| 13 | Show info about gene in browser | gen,pat | | | DoD |
| 14 | Search for pathways in which gene occurs | bio,gen | | Z+F | |
| 15 | Show thumbnail list of pathways | bio,gen | | | Z+F |
| 16 | Investigate genes in specific pathways | bio | | | DoD |
| 17 | Filter patients where gene is deregulated | gen,onc | | Z+F | |
| 18 | Show list of filtered patients | onc,pat | | Z+F | |
| 19 | Investigate their treatment and outcome | onc,pat | | | DoD |

Table 1. *Interaction path in a sample use case where biomedical experts aim to select a cancer treatment for a specific patient. The decision is based on reference cancer patients data collected at the clinic. Based on the down-regulation of a gene known to be one of the causing factors of the tumor, similar patients are filtered and taken as a foundation for the treatment decision. For every task the involved expert is stated: oncologist (onc), pathologist (pat), geneticist (gen), and biologist (bio). The levels of interaction are referred as: overview (Ovv), zoom and filter (Z+F), and details on demand (DoD).*

to optimize the navigation path. Another optimization can be made by looking for recurring interaction patterns and adapting the application to make them readily available and easy to use. An example for the Caleydo Visualization Framework is the bucket representation with visual links. It was specifically introduced to make switching between different visual representations, a common pattern in our use case, easier and more intuitive.

### 3.2 Specific Implications for our Use Case

While the multi-level aspect is inherent in the application scenario, it is the multi-user aspect that distinguishes between the complexity of coordinating for a seamless interaction path through the multi-level data. In general, one can differentiate three cases:

- The **single-user** case, which is what the Caleydo framework is aimed at. It allows a seamless navigation through the multiple application levels by providing a linked multi-view visualization on a single output device.
- The **static multi-user** case, which is targeted by the adaptation of the Caleydo framework for the Deskotheque environment. This environment provides a fixed set of displays and projection areas to facilitate multi-user interaction.
- The **dynamic multi-user** case, where the set of the involved users is not static, but changes over time. In this so called smart environments, also the device ensemble of available displays is changing as users connect and disconnect their brought devices (netbooks, laptops, PDAs, etc.) with the environment during run-

time. A detailed discussion on this case' realization and its usage for a medical scenario is given in [11].

It can be observed that with each of these cases, the complexity of coordinating multiple data levels to be shown on multiple displays for multiple users is increasing. The challenges this poses are abundant and range from the distribution of the data to the available display devices (or views in the single-user case) to the assurance that privacy concerns are met. Our *Multi-User Multi-Level Interaction* concept provides a conceptual and concrete way to model all these complex dependencies and to derive solution approaches that finally achieve real seamless collaborative data analysis.

Collaborative information workspaces, such as described in [12], differentiate between private and public displays. In the simplest case, each domain expert displays his/her domain data on a private display – e.g., in Figure 1 three users from different domains are sitting around a table, each with a private view on a single monitor. Besides the plain distribution of views, the users' roles can further be facilitated to provide tailored visualizations, as a user's working domain influences the chosen visualization technique and terminology used for annotation purposes. Different domains can then be bridged either by a simple coordination of visualizations among the (private) displays or by the combination of data from different sources in public visualizations. Public displays, i.e., projection walls which are visible for multiple users, can host these integrative visualizations. This also allows multiple users to work on the same task.

The physical separation between public and private displays can also be used to circumvent privacy issues, by showing sensitive data only on private displays. In a clinical scenario, the biologist may not be allowed to see the clinical history of patients for privacy reasons. The control over the individual displays enables the collaborative environment to grant or deny access to experts depending on their role, either allowing them to roam freely within all available data sources or just within the absolutely necessary parts. Even annotations could differ, providing patient details in private views, but being anonymized in the public views. Thereby, the anonymization does not affect the linking of the individual views. Selections and other interactions are reflected throughout the whole ensemble of displays.

In dynamically changing environments, it is furthermore essential to have access to a wide range of information: the spatial model of the environment, the participating subjects and their roles, the underlying data, and the workflow of the analysis tasks. All these are essentially targeted by the proposed *Multi-Level Multi-User Interaction*, as it allows to specifically define in detail what (data set) is visible to whom (expert user) in which way (visualization technique) with which goal (aim of this analysis step) and in which order (workflow) – capturing the entire analysis session and going well beyond the pure definition of individual analysis tasks. Having this knowledge beforehand, enables the environment not only to provide a suitable data set from the specified application level to the experts who fit the role and have the necessary security clearance, as it is outlined in Table 1. But instead, the explicit knowledge of probable interaction paths and the resources needed for each step allow to adapt to a dynamically changing environment. E.g., if a certain analysis path requires an expert who is currently not present or a data source which is not available, a different path of analysis can be chosen, if one exists. To reach such a high level of coordination in a dynamic multi-user environment is a challenging task. Now, that the infrastructure as detailed in [11] is up and running, the first step for future research is to investigate how this solutions can be integrated with Caleydo to enhance the single-user scenario to a smart one that adapts to changing constraints.

## 4 Conclusions and future work

We presented the *Multi-User Multi-Level Interaction* concept as a way for formalizing the collaborative information seeking process of multiple domain experts working with heterogeneous data. The concept allows to model, analyze and consequently optimize and adapt the interactive workflow in complex environments. Although we introduced and demonstrated it by means of a static multiple user scenario, the presented concept can also be scaled down to a single-user, single-display setup and scaled up to a dynamic multi-user scenario, both being subject of future research.

In the single-user case, instead of deciding on which display to show which view of what kind of data, it can be used to decide which space of the screen (e.g., which wall of Caleydo's bucket representation) to use for which kind of data and how to link them appropriately. On an even smaller scale, extracted interaction paths and patterns do also help to automatically arrange and tailor the visualizations to data on different levels of granularity with the aim of providing a seamless exploration process. Once defined in terms of our interaction concept, this process can even be potentially accelerated by optimization (e.g., preprocessing) along predefined common interaction paths.

In the dynamically changing multi-user case, the extracted knowledge can potentially contribute to solutions for many of the challenges that dynamic smart environments face. With a holistic model of the entire workflow, it should be possible to overcome minor disturbances of the analysis process by an adaptation of the process according to the currently available resources and users.

So far, the *Multi-User Multi-Level Interaction* concept primarily focuses on the information seeking workflow, as defined by Shneiderman and Keim. However, the concept is most certainly applicable to a broader range of high-level interaction patterns, e.g., for data manipulation. Hereby, the *Information Seeking Mantra*, as it is embedded in our interaction concept, can be replaced with a different pattern, e.g., by Baudel's data manipulation process [2]: view adjustment, selection, and editing. Hence, it seems even possible to generalize our concept to any step-wise definable interaction pattern.


## Acknowledgements

The authors wish to thank Manuela Waldner and Alexander Lex for their work on collaborative information workspaces. This work was funded in part by the FIT-IT program (813 398), the Fonds zur Förderung der wissenschaftlichen Forschung (FWF) (L427-N15) and the DFG graduate school *dIEM oSiRiS*.


## References


[1] E. Andrianantoandro, S. Basu, D. K. Karig, and R. Weiss. Synthetic biology: new engineering rules for an emerging discipline. *Molecular System Biology*, 2, 2006.

[2] T. Baudel. From information visualization to direct manipulation: extending a generic visualization framework for the interactive editing of large datasets. In *Proc. of ACM UIST'06*, pages 67–76, 2006.

[3] J. L. Encarnação and T. Kirste. Ambient intelligence: Towards smart appliance ensembles. In *From Integrated Publication and Informations Systems to Virtual Information and Knowledge Env.*, pages 261–270, 2005.

[4] D. Keim, F. Mansmann, J. Schneidewind, and H. Ziegler. Challenges in visual data analysis. In *Proc. of IV'06*, pages 9–16, 2006.

[5] C. Pirchheim, M. Waldner, and D. Schmalstieg. Deskotheque: Improved spatial awareness in multi-display environments. In *Proc. of IEEE VR'09*, pages 23–126, March 2009.

[6] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: a unified approach to image-based modeling and spatially immersive displays. In *Proc. of SIGGRAPH'98*, pages 179–188, 1998.

[7] H.-J. Schulz and H. Schumann. Visualizing graphs – a generalized view. In *Proc. of IV'06*, pages 166–173, 2006.

[8] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. of IEEE VL'96*, pages 336–343, 1996.

[9] M. Streit, M. Kalkusch, K. Kashofer, and D. Schmalstieg. Navigation and exploration of interconnected pathways. *Computer Graphics Forum*, 27(3):951–958, 2008.

[10] M. Streit, A. Lex, M. Kalkusch, K. Zatloukal, and D. Schmalstieg. Caleydo: Connecting pathways with gene expression. *Bioinformatics, Oxford Journals*, 2009.

[11] C. Thiede, C. Tominski, and H. Schumann. Service-oriented information visualization for smart environments. In *Proc. of IV'09*, pages 227–234, 2009.

[12] M. Waldner, A. Lex, M. Streit, and D. Schmalstieg. Design considerations for collaborative information workspaces in multi-display environments. In *Proc. of CoVIS'09*, 2009.

# TaP: Towards Visual Analytics on Interactive Surfaces

Stefan Flöring and Tobias Hesselmann


Fig. 1. Hands on: Stacked halfpie menu for navigation in hierarchical data structures.

**Abstract**—Today the amount of collected data steadily grows. Today, larger amounts of data are collected than ever before. To cope with the problem of finding relevant information in collected data, we present TaP, a visual analytics system for visualization and gesture based exploration of multi-dimensional data on an interactive tabletop. Using TaP, users are able to control the entire analysis process by means of hand gestures on the tabletop's surface. TaP is focussed on multi-dimensional data and provides a novel menu design, stacked half-pie menus, to explore deeply nested hierarchical data structures, such as dimensions and measures of the OLAP data model. Considering that collaboration can play a key role in data analysis, we implemented several features into TaP to effectively support collaborative visual analysis. In this paper, we will present the different elements of the system and their role in collaborative analysis.

**Index Terms**—Visual Analytics, public health, OLAP, gesture based interfaces, collaborative analysis.

## 1 Introduction

In today's health care environments, especially in the epidemiological domain, the importance of data analysis steadily increases. Key topics of interest are the discovery of trends and correlations between diseases and potential influence factors, early recognition of infectious diseases and their diffusion and report generation based on geographical regions or influence factors [9].

Data-analysis systems have been an important issue in our research group for several years, initiated by the launch of the epidemiological cancer registry lower saxony (EKN), for which we developed CARESS, a tool for data-analysis. It supports the visualization of geographical, multi-dimensional measures on thematic maps. Based on the experience gathered in the project we created MUSTANG [7], a service based platform for report driven data visualization in health care applications named. MUSTANG facilitates numerous data-analysis applications and became the foundation of a variety of analytical systems. However the focus of current MUSTANG still is report driven analysis, i.e. predefined reports are filled with data from the epidemiological database. This rather static approach is not necessarily disadvantageous, as common epidemiological applications do not require for interactive analysis's. Nevertheless, as the amount of collected data increases the demand for more sophisticated analysis methods grows simply because there is potentially valuable information hidden in the database. Regarding to user feedback we received from continuous exchange with the experts working on our systems, there is a growing demand for ad-hoc analysis and explorative tasks in future. The idea is to visualize the collected data and then mingle data from other sources into the visualizations. For example the average amount of infections of a certain kind per region could be visualized on a thematic map, to possibly find regions with atypical high or low rates. Then data from other sources could be integrated interactively into the visualization to find possible environmental influence factors.

Tasks like this require visualizations which allow interaction to explore the database according to the visual analytics mantra (*Analyze First - Show the important - Zoom, Filter and Analyze Further - Details on Demand*) [6]. They also require the collaborative work of analysts with expertise in varying fields, for example people with medical or epidemiological background, people with demographic background and so forth.

In addition to that our user feedback shows that even in the reporting analysis style it is a common practice to discuss the results with other analysts before they are published. A typical approach is to either print out results and then seek for direct discussion or to generate electronic documents and transfer them via E-Mail to other experts, along with a certain set of questions. To us it is obvious that a more direct way of discussion and communication of results with immediate response and interaction on the visualized data itself might enhance the analysis process here. Therefore we determined two major challenges for the advance of epidemiological analysis: Increase the possibilities of interaction with visualizations to enable a more dynamic and spontaneous way of explorative data analysis and create the possibility of collaborative analysis with multiple experts working directly on the same scenarios in an interactive way.

This paper describes the TaP system, a visual analytics platform based upon a large scale interactive surface computer and explains how these two challenges are addressed by this system. Visual analytics on surface computers has only recently been recognized as a research topic and related work in this area is relatively sparse. In [2] Collins presents a system for file system interaction where a table top is used for a novel approach to associative-search in semi-structured data. In Isenberg *et al.* [5] a system for co-located collaborative work with information visualization based upon a large scale tabletop display is introduced. The paper provides guidelines for the design of col-

- *Stefan Flöring is with University of Oldenburg, Germany, E-mail: stefan.floering@informatik.uni-oldenburg.de.*
- *Tobias Hesselmann is with OFFIS Institute for Computer Science, E-mail: hesselmann@offis.de.*

laborative information visualization systems. Our system differs from existing approaches in two areas: We focus on multi-dimensional data in the OLAP data-model and it is our aim to create a gesture based interaction to directly manipulate the visualized data.

## 2 TAP: A NEW APPROACH TO EXPLORATIVE DATA ANALYSIS

OLAP databases are one of the most common ways of data management for multi-dimensional data. The MUSTANG platform provides a service based interface for accessing databases based upon this data-model. MUSTANG is the basis of TaP and is used as service layer for data access by means of OLAP operations.

In the multi-dimensional data model data is categorized into hierarchically structured dimensions (such as time, geographic location, products), each containing a number of distinct hierarchical levels, and measures (such as disease rates or sales). The data is represented in data-cubes which span across several dimensions and measures. In addition to data access, OLAP operations for generalization (roll-up) and specialization (drill-down) along the dimensional hierarchies are available.

This is a constraint to the TaP data-analysis platform. As all current applications are based on this multi-dimensional data management, it is necessary for the new system to allow the exploration of those structures. This is important for the first step in the analysis process, the selection of data to be analyzed, and for further steps such as zoom and filter, where the visualization has to support the corresponding OLAP operations.

Before we describe the user interface of the TaP system, we would like to expose our thought that lead to the use of an interactive table top computer for display and interaction with the data analysis platform.

## 3 TABLE TOP COMPUTERS FOR DATA-ANALYSIS

Typical computer workstations are optimized for single user interaction. This starts with the input devices (keyboard and mouse), which only fit well for single user interaction and can be extended to the display technology. Typical workstation displays offer a fairly good spatial resolution of up to 30-40 pixels / cm. However the display size is usually limited to something between 50 - 60cm and have a limited viewing angle and therefore are not well suited for more than one person at at time.

Another important aspect is territoriality in collaborative scenarios. Studies have shown that humans claim personal territories when they are performing tasks in collaborative workspaces. In [10] was analyzed how people work on collaborative tasks in desktop environments. The participants got different tasks which they had to fulfill. Their motions were recorded and afterwards their range of motion and the emphasis regions of their interaction was calculated. It was determined that people have personal territories, where they interact more frequently than in other regions. In collaborative tasks (such as sorting paper notes) there were areas where everyone would interact, the main regions around the object to deal with, and personal areas where some people would interact more frequently. The personal territories were automatically respected by other participants, they avoided to interact in the personal regions of the other participants. This shows that people have the exception that a personal area surrounding is their personal territory and will be respected by others.

Classical workstations do not provide a concept of personal workspaces or territoriality. To integrate this concept, artificial regions on the screen would have to be declared as workspace for specific persons. However humans define their personal regions on the base of absolute distances and it is questionable whether small 50-60cm screens provide enough space for a good territory metaphor. The concept of workspaces here would be absolutely artificial.

Based upon this assumption we created a large table top computer. Decisive for this was the idea, that a horizontal workspace provides better possibilities for collaborative work, as there is more space around the table than e.g. in front of a desktop or even of a wall of equal size and therefore natural personal territories based upon distance could automatically evolve. Another assumption that encouraged us to use a table top computer was that - if the data visualizations and their interaction methods are implemented well - a table top computer allows a more direct way of interaction with data than a classical computer workstation. The analyst could literally touch the visualized data and directly initiate operations on the data, following Shneidermans direct manipulation paradigm [12].

## 4 USER INTERFACE OF TAP

Unlike classical computer workstations, table top computers allow users to directly interact with the interface elements. Classical workstations are operated by input devices like mouse and keyboard, which only allow indirect manipulation of interface elements. Therefore it is necessary to research whether given interaction paradigms and interface elements are still suitable for the use on a table top computer.

Feedback by our users implied that the visual analytics system needed to allow the analysts to quickly select data and the display it in suitable visualizations. This resulted in the requirement for some sort of selection process for items of the multi-dimensional data model. As mentioned before the dimensions and measures in this data model are ordered in hierarchical structures. In our classical analysis applications the selection process is based on a dialog driven wizard, where the analyst makes selections for different aspects, such as the visualization to be used, the dimensions to be displayed, the hierarchical level of the dimension, the measures which should be displayed for this dimension and so forth. For selection of elements out of hierarchical structures classical tree views with textual labels are used. The selection itself is performed by a checkbox in front of a tree view label. Once all steps of the wizards are accomplished the visualization will be rendered. If the analyst wants to modify the visualized circumstances he has to re-enter the wizard and go through all of the selection steps.

This approach seemed to be unfortunate for usage in an interactive way on a table top computer because of two reasons. First of all the selection model in wizard style does not match the visual analytics mantra, where it is vital to quickly display something and then iteratively enhance the visualization with further aspects. And at least as important, tree views do not scale well on touch screens. The size of a tree view is determined by the length of the textual labels and the depth of the hierarchy. The longer the labels and the deeper the hierarchy, the more screen estate is used. This is not a big problem in user interfaces which are operated with a mouse, which very precise and where UI elements therefore may be quite small. Touch interaction based interfaces are far less precise than those with mouse interaction, because the size of the human finger is limiting the precision. Sears et al. studied the performance of participants typing on virtual keyboards on a touch based interface with varying key size [11]. It was shown that the larger the keys, the better the performance in terms of error rate and words per minute. Even from the second largest size to the largest size (2.27cm per side) they measured that the throughput was still increasing. Bender later confirmed these observations regarding to error rate [1]. Here was also shown that the target size influences the contact time, with larger targets (3cm per side) decreasing the contact time compared to smaller targets (1cm per side). This leads to the conclusion, that interface elements for touch screens should not be smaller than 2-3cm on the smallest side or in diameter. Transferred to the tree view this means that the height (the smallest size) of the menu should be at least 2cm. This means that a menu with 50 entries would already consume 100cm of vertical space. It also leads to very broad menus as the length of the textual labels increases likewise when the height of the textual label is adjusted to 2cm.

### 4.1 Pie Menu

In our current implementation the system is bound to a specific multi-dimensional database at compile time. To access the data the user needs to be able to explore the hierarchical dimensions and measures available in this database. With this information in mind we were seeking for a menu that would provide labels of a fixed size, large enough for touch screen usage and which scales with many hierarchies and entries in those hierarchies. We came up with the idea of the stacked half-pie menu [3], adapting the concept of round pie menus with fixed label sizes, as described by Hopkins [4] and modified this design for our purposes. Instead of a full circle we used a half circle design for the pie menu. This enables us to have a variable number of elements in each row of the pie menu. If there are more items in one hierarchy than would physically fit into one row, the items are hidden instead. Arrows on the edge of the row indicate, that more items are available in this hierarchy. An example Menu is shown in Fig. 1. The figure shows the menu in three different states. The leftmost state shows the menu fully collapsed. In the middle the menu is has two expanded levels, of which one is partly collapsed and only represented by it's textual label "dimensions". The image on the right shows the menu with four expanded levels, again with the inner one party collapsed. We have chosen to partly collapse the innermost level, to safe screen estate. In the multi-dimensional context the inner level will always provide the two items "dimensions" and "measures". We assume that a full visualization of those two items is not permanently necessary.

A touch on a label with children in the hierarchy (indicated by a little + in the label) will expand the next level. It is possible to drag items out of the pie menu, to the bubble chart, to visualize the underlying data.

### 4.2 Bubble chart

Figure 2 shows the chart component of our visual analytics application. We have chosen a bubble chart with additional coloration as first visualization because it is widely used and easy to implement. The chart component hast five axes of visualization:

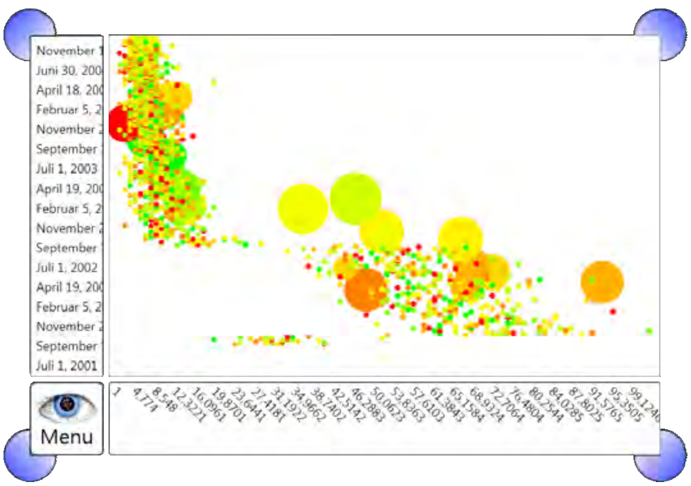


Fig. 2. Bubblechart with sample data.

**The y-axis and x-axis:** Those two axes are suitable for cardinal and ordinal values as well as for qualitative dimensions (such as products) with no natural order.

**The size-axis:** Values assigned to this axis will modify the bubble diameter relative to the other displayed values. This axis is suitable for cardinal values.

**The color-axis:** Values assigned to this axis will modify the bubbles coloration. This axis is suitable for qualitative values, where each value gets a certain color (e.g. male as blue and female as pink) as well as for cardinal values, where a color scale (e.g. red - yellow - green) is matched to the value scale.

**The animation-axis:** This axis is not shown in the figure, it is aimed for temporal data. The data-sets for distinct points of time will be rendered as individual charts and then displayed as animation.

To assign values to the chart we have chosen to use a touch and drag gesture. To be able to assign certain elements to an axis we created so called drop zones, which can be seen in Fig. 3. In the figure you can see how the element "products" is dragged to one of the drop-zones. There are six drop-zones, from left to right and top to bottom: y-axis, wild card, size, color, animation and x-axis. Each drop-zones is corresponding with the respective axis of the chart. The wild card drop-zone is special, values dropped onto this drop-zone will be assigned to one of the five available axis according to an integrated rule set which will take care that mandatory axis for the visualization are filled first and that ranges of values are matched between the axis and the assigned data-sets. The drop-zones of the active chart are shown on top of the chart as soon as an element from the pie menu is dragged into the direction of the chart.

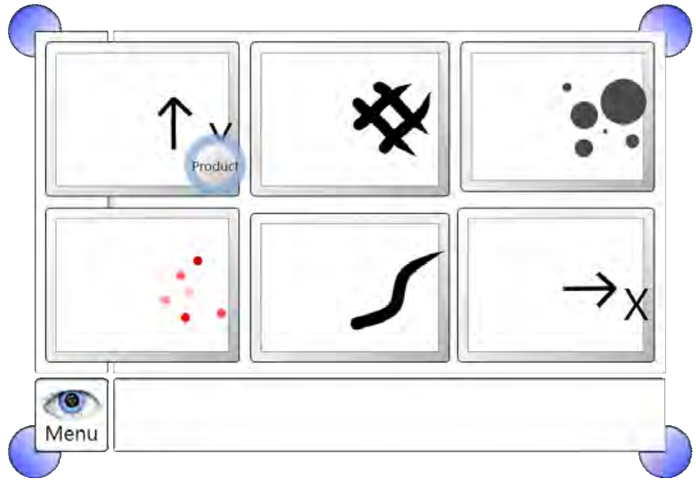


Fig. 3. Bubblechart with activated dropzones.

The chart supports gesture based interaction for OLAP operations on some of the axis. At the current stage of development we can perform the the dice operation by performing either a pinch or a spread gesture inside a chart. With a pinch from outside to inside the chart will zoom out, with a pinch from inside the outside chart will zoom into the data, accordingly performing the operations on the database. In addition to that we implemented the roll-up and drill-down operations on the x-axis and y-axis. For this the analyst has to perform the pinch gestures on the axis labels likewise and thus change the hierarchical level of a dimension or the granularity of a measure. If possible a matching OLAP operation is performed on the database and the chart is updated accordingly.

### 4.3 Interface Elements for Collaboration

Recognizing the importance of collaboration in visual analytics, we designed the system to support multi-user interaction in several ways. First of all we took care that all interface elements can be used anywhere on the surface. All elements contain round blue drag handles (see Figs. 2 and 3) which are accessible from all sides. Those handles allow to move and to resize the charts. In our first prototypes we experimented with fluid integration of rotation and translation, as described in [8]. However we had to deal with the problem that gestures performed to interact with the visualized data on the bubble chart would often cause the chart to move or rotate and therefore we decided to use handles for movement and rotation.

There was no simple way to attach such handles to the pie menu because the pie menu does not have any fixed corners which are accessible from all sides. Therefore we thought of a way to display the pie menu on any given place of the surface by a gesture. The touch recognition technique used in our table top computer does not only provide information about position and movement of the touches but also information about the size of the touches. The palm of a hand for example produces a significant larger touch-area than the touch of a finger. We listen to those events, which we call palm touches, and display the pie menu a few cm above the palm touches. In this way it appears underneath the fingers of the person who last performed a palm-touch. This enables everyone around the table to move the pie menu to a nearby position.

In standard applications a typical way to create another work item (another chart in our case) would be to select a menu entry from a drop down menu at the topmost position of the screen (e.g. File/New) or to click an item in a toolbar that has a fixed position on the screen. This approach is challenging on large touch screens for two reasons. First of all the user would possibly have to make a large movement to a distant area on the screen, for example to the top of the screen. This is no problem on small displays which are operated by a mouse, as most mouse drivers include some sort of acceleration mechanism where faster movement of the mouse will result in faster arrow movement. However on large touch screen, according to Fitt's Law, we can expect movements to distant areas to be either slow or imprecise. In addition to that continuous arm movement over large distances will most likely fatigue the user. The second problem with the drop down menu approach on a table top computer is, that it is possibly discriminating users who are at a position where it is not possible to access the menu quickly. Therefore we thought about a way to create a menu that is accessible from all positions of the table top and which would not cause large arm movements. For this we created a gesture layer, which can be displayed on demand by either touching into one of the four corners of the screen, where blue circles indicate a touchable region, or by knocking onto the border of the screen. In the moment either one of those actions is performed, a virtual layer will be displayed on top of all other interface elements.

Fig. 4. TaP working surface with activated gesture layer and example gesture path.

Fig. 4 shows the TaP system with activated gesture layer. The interface elements beneath the layer are slightly faded out, to visually indicate that the gesture layer is active. This layer allows path based gestures. To create a new chart for example a rectangular path has to be drawn. This gesture can be drawn by touching and dragging a finger across the surface. As soon as the gesture is finished (when the finger lifts up), a new chart will appear at the position where the gesture has been performed in the size of the drawn rectangle. Therefore the same interaction will create a new chart and assign size and position at once. This is an additional advantage over classical user interfaces.

## 5 Summary and Outlook

The TaP system is a first step towards an integrated approach for visual analytics on table top computers. It was designed under the assumption that visual analytics is a collaborative task and that table top computers are better suited to collaborative tasks than classical user interfaces. With the pie menu, a new menu type we developed, it provides mechanisms to explore multidimensional databases and to visualize elements of these databases. Interaction with the system is purely gesture based. We took special care that all interactions are possible from any place around the table top computer, to not discriminate people by their position. We have done first user evaluations of the pie menu and the results look promising [3].

At the current stage of development there are still a few shortcomings in the TaP system. First of all, the system does not distinguish between multiple users. Even though the underlying framework is capable to distinguish interactions by different users, the current hardware does not support user recognition. Future versions of the TaP system should provide the possibility to identify users or at least to allocate touch signals to specific users to deal with this issue. A challenging task here is to identify which user performed which gesture. Personal areas around a users position, according to [10] might be a way to assign certain interactions to users. Another shortcoming of the current implementation is, that all interface elements have a direction towards one side of the table top computer. However the framework used for graphical representation already supports rotation of interface elements and we are looking forward to quickly integrate these mechanisms into the system. For the pie menu, which does not incorporate handles for movement and rotation we want to extend the gesture recognition to recognize the orientation of the palm and then use this information to rotate the pie menu accordingly.

Besides the enhancement of the gesture based interface our main interest for further research is to better support the analysts in their work tasks. Therefore we are planning to integrate more visualizations, especially thematic maps, which are utterly important in the epidemiological domain. Additionally we want to integrate a history which allows to navigate trough the different states in the analytical process as suggested in [13] and [5] and find ways to extract knowledge from this history so it might be re-used in other tasks.

## References


[1] G. Bender. Touch screen performance as a function of the duration of auditory feedback and target size. 1999.

[2] A. E. Collins. Exploring tabletop file system interaction. In *CHI '07: CHI '07 extended abstracts on Human factors in computing systems*, pages 2171–2176, New York, NY, USA, 2007. ACM.

[3] T. Hesselmann, S. Flöring, and M. Schmitt. Stacked half-pie menus navigating nested menus on interactive tabletops. In *Proceedings of the ACM ITS 2009 (to appear)*, New York, NY, USA, 2009. ACM.

[4] D. Hopkins. The design and implementation of pie menus. *Dr. Dobbs J.*, 16(12):16–26, 1991.

[5] 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.

[6] D. A. Keim, F. Mansmann, J. Schneidewind, and H. Ziegler. Challenges in visual data analysis. In *IV 06: Proceedings of the conference on Information Visualization*, pages 9–16, Washington, DC, USA, 2006. IEEE Computer Society.

[7] S. Koch, J. Meister, and M. Rohde. Mustang - a framework for statistical analyses of multidimensional data in public health. In *The Information Society and Enlargement of the European Union*, pages 635–642, Metropolis Verlag, Marburg, 2003. Albrecht Gnauck, Ralph Heinrich.

[8] R. Kruger, S. Carpendale, S. D. Scott, and A. Tang. Fluid integration of rotation and translation. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 601–610, New York, NY, USA, 2005. ACM.

[9] J. Meister, M. Rohde, H.-J. Appelrath, and V. Kamp. Data-warehousing im gesundheitswesen. *it - Information Technology*, 45:179–185, 2003.

[10] S. D. Scott, M. Sheelagh, T. Carpendale, and K. M. Inkpen. Territoriality in collaborative tabletop workspaces. In *CSCW '04: Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pages 294–303, New York, NY, USA, 2004. ACM.

[11] A. Sears, D. Revis, J. Swatski, R. Crittenden, and B. Shneiderman. Investigating touchscreen typing: The effect of keyboard size on typing speed. *Behaviour and Information Technology*, 12:17–22, 1993.

[12] B. Shneiderman. Direct manipulation: A step beyond programming languages. *Computer*, 16(8):57–69, 1983.

[13] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations, 1996.

# Roles of notes in co-located collaborative visualization

Narges Mahyar, Ali Sarvghad, and Melanie Tory

Fig.1. Examples of note taking activities during our observational study. Note-takers are disconnected from the group activities.

**Abstract**—This paper focuses on the significant role that note taking plays in collaborative data analysis within the business domain. The discussion of note taking is based on preliminary observations from a user study in which co-located teams of business users worked on collaborative visualization tasks using large interactive surfaces. We propose an initial categorization of note taking activities and propose a list of research questions that need to be discussed and investigated in order to better understand note taking process in the context of collaborative visualization and analysis activities.

**Index Terms**—collaboration, computer supported cooperative work, Information visualization, note taking

## 1 Introduction

We discuss the importance of note taking activities during collaborative visualization on interactive surfaces. The need to support note taking arose from observations during a user study that we conducted to examine collaborative data analysis in the business domain.

Use of information visualization (InfoVis) tools to assist decision-making in the business domain is on the rise [8]. In order to better understand how software tools can support collaborative data analysis, we conducted an exploratory study to examine how people use visual representations of data collaboratively to solve a problem in the business domain and to observe behaviour and processes they use. We used an existing Business Intelligence (BI) application, "Polestar on Demand" proposed by SAP Business Objects. We believed that working with large displays and a specially made application for visualizing business data would help us to re-examine the process of collaborative visualization, as well as problems of current applications and their specific requirements to be customized for collaborative usage. One of the surprising results from this study was the observation that note taking is a critical process in collaborative data analysis and is not well-supported by current tools.

This paper is not intended to fully document our study and its results. Instead, we highlight some observations regarding note taking, and use them to raise questions about how to best support individual and group note taking activities for collaborative visualization on interactive surfaces.

In the following sections we present a concise review of related work, provide a brief description of our study, report some observations from the study, and finally raise a series of research questions that we believe will need to be addressed by future work on note taking for collaborative visualization.


- *Narges Mahyar is with the University of Victoria, E-Mail: nmahyar@cs.uvic.ca.*
- *Ali sarvghad is with the University of Victoria, E-Mail: asarv@cs.uvic.ca.*
- *Melanie Tory is with the University of Victoria, E-Mail: mtory@cs.uvic.ca.*


## 2 Related Work

While substantial research has been devoted to computer supported cooperative work (CSCW) in general, collaborative visualization is still under explored due to its unique challenges. It is still not fully clear how people collaborate to solve data analysis tasks, or how information visualization techniques and interaction methods need to change to support collaborative work. Recently, some research has begun to address this question. Several studies have identified processes or activities that contribute to the overall group analysis process [4] [7] [9] [10] [11], by using software supporting collaborative work [9] [10] or by using paper-based tasks [4] [11]. Findings of previous studies, regardless of whether the tasks were paper-based or software-based, suggested almost similar lists of processes involved in the collaborative data analysis. It also has been identified that very flexible tools to support co-located collaboration are needed [4] [11] [13]. This includes flexibility to change ordering of activities, work styles (from closely coupled to independent), role assignments, and workspace organization. It has also been pointed out that horizontal and vertical surfaces are suitable for different types of collaborative work [12].

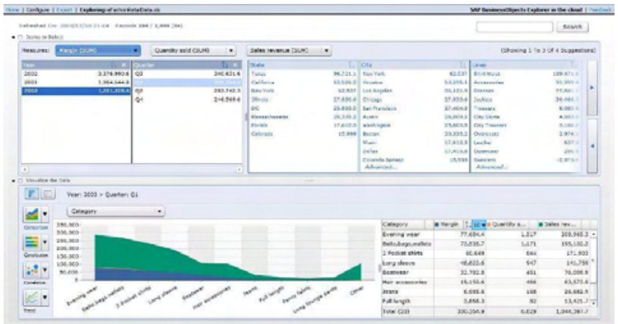

Fig. 2. Screen shot of Polestar, depicting a comparison chart that visualizes margin, quantity sold, and sales revenue over category, filtered based on a specific year and quarter.

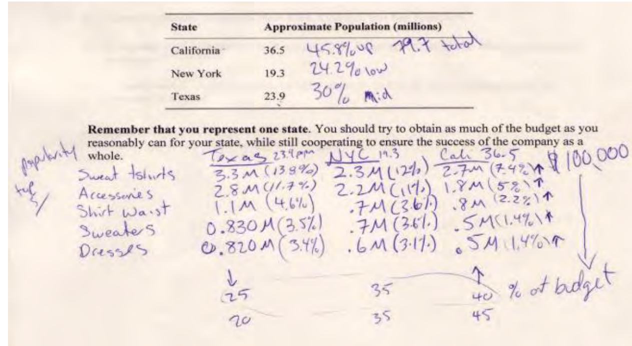


Fig. 3. Sample of notes taken for group use. Content is nicely structured and has a group scope.

To our knowledge, none of the work on collaborative visualization has explicitly focused on the need to support note taking activities. By contrast, this need became explicitly clear during our observational study.

## 3 Our Exploratory Study

Here we briefly describe our observational study, as background to help the reader interpret our observations and discussion.

### 3.1 Participants

Twenty-seven participants took part in our study, divided into nine groups of three. To increase collaboration effectiveness and to simulate common work situations, all the group members knew each other. Two of the groups were computer science graduate students and the other seven groups were 4th year BCom or MBA students.

### 3.2 Apparatus

Our apparatus were two identical Smart DViT (digital vision touch) screens, one in a wall configuration and the other in a tabletop. Both had four HD projectors with 3840 x 2160 resolution (8.3 Mpixels), and had a size of 61.2" x 34.4" (70" diagonal).

We used "Polestar on Demand" (Fig. 2) as our data visualization and analysis tool. Polestar allows users to upload any data set and then interactively browse through the information. Polestar has been developed as a single user application. It has a straightforward interface and is considered to be reasonably user-friendly. It can be accessed from https://create.ondemand.com/explorer.

Four groups used a tabletop display, four used a wall display and one used both displays. This gave us an opportunity to observe and obtain users' feedback on a variety of display configurations.

### 3.3 Task and Procedure

Each study comprised of two tasks, both using an e-fashion dataset. Task 1 included 6 warm up questions, which were focused questions designed in a way that users could learn important features of Polestar. These included selecting variables, filtering, creating different types of charts and saving. An example question from task 1 was, "How does the 2003 margin compare to previous years?"

Task 2 was a business case. Participants were asked to assume the roles of three top managers (representing different states) and together determine a marketing budget for the next year. Rationale for the budget was based on information within the data set. This task was competitive in nature: participants had to compete to obtain the maximum possible budget for their state.

Styli, paper, and pens were provided to help participants work with the system or to take notes. Initially, we provided a 10-15 minute introduction to Polestar, describing its features. Participants spent approximately 30 minutes on task 1 and 40 minutes on task 2. We offered an optional 5 minute break between two tasks. After task 2, they spent around 10 minutes to sum up and write down their results. We asked our participants to create a report of their results at the end of task 2. Our rationale was to have a record of how participants used charts to justify their decisions. Then we had an open-ended interview. All the sessions were audio and video recorded and all the screen logs were recorded for further analysis. During all sessions, one observer took notes about users' actions and problems they faced; she also helped them whenever they had a question regarding the tasks or the software.

## 4 Roles of Notes

Findings presented in this section are based on analysis of our recorded data, notes taken by the observer, as well as all the notes and reports made by users. The huge amount of note taking that we observed suggests that note taking is a significant activity in business data analysis.

### 4.1 Contents, purposes and usages of notes

Perhaps the most interesting and notable finding from our study was the importance and frequency of note taking. Participants in our study took notes at almost every single step of their data analysis. This might be related to the special requirements of business data analysis, which is usually dependent on numbers, percentages, calculations et cetera. Notes taken by participants often consisted of the following:

Numbers (e.g. data value)
Drawings (e.g. flag, chart)
Text (e.g. question, hypothesis, reminder)
Symbols (e.g. %, $)

Figures 3 and 4 are two samples of the notes taken by participants in our study. Figure 3 shows a note taken for group use. It has been nicely formatted and contains some calculated values. The person who took this note was assigned the role of note taking. He was sitting most of the time and observing others (who were exploring data and creating visualizations). He therefore was unable

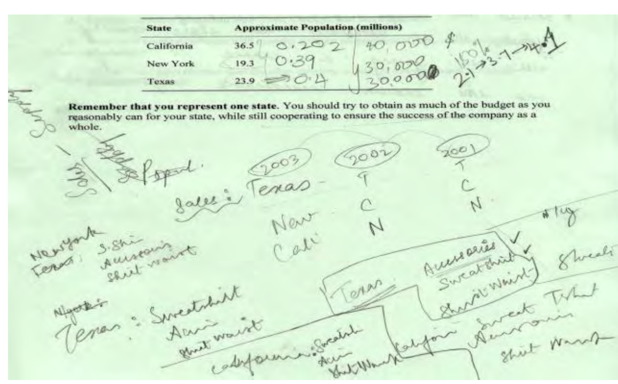


Fig. 4. Sample of a note taken for individual use. It is different in content and form compared to the group note in figure 3.

to work directly with the application a lot of the time. The group needed the content of the notes to help them further analyze data and solve the problem given in task 2. The small yet comprehensive tabular data that can be seen in figure 3 made the analysis task easier by saving important information; it was much more convenient and efficient to have this information recorded rather than revisiting previously created charts. The same person who was in charge of note taking also created the final report. Figure 4 depicts a sample of notes taken for individual use. It can be clearly seen that it has a less structured form compared to the sample shown in Figure 3.

In general, group notes were more carefully organized than individual notes, but this of course varied somewhat depending on the individual's note taking style. Individual notes were not always organized or written legibly or in a way that everybody at a glance could understand them. Again it depended on the individual who took these notes, and in rare cases, individual notes were nicely written, legible and structured. However, mostly individuals did not try to make it pretty or usable for the group. Sometimes they used some abbreviations or symbols that could be interpreted only by the note taker. Possibly they were witting as fast as possible to minimize distraction, since taking notes was not their primary focus.

In most sessions that we observed, one user assumed the role of note taker for the group. This role assignment was usually not discussed explicitly. It also did not necessarily remain fixed throughout the work session; sometimes the note taker changed part way through.

Generally, notes had different characteristics depending on their purpose and intended reader. Based on our preliminary analysis of the notes we suggest characterizing note taking and note use as shown in figure 5. Note creation shows that participants mainly took notes to save a value or artifact (e.g. a chart or the result of a calculation), to remind them to do something (e.g. review a chart) or to emphasize something important (e.g. what he/she or others find valuable). It also shows that both notes taken for group use and notes taken individually for private use can have the same purpose. The scope of notes is typically private when notes are taken for individual and public use when notes are taken for group use. However, in some cases, individual notes were shared with the group. Note use shows that notes' contents could be used for further analysis of data, creating a report, remembering an important artefact or value, or validating previous work. Validation here is mostly concerned with ensuring that a calculation result is acceptable.

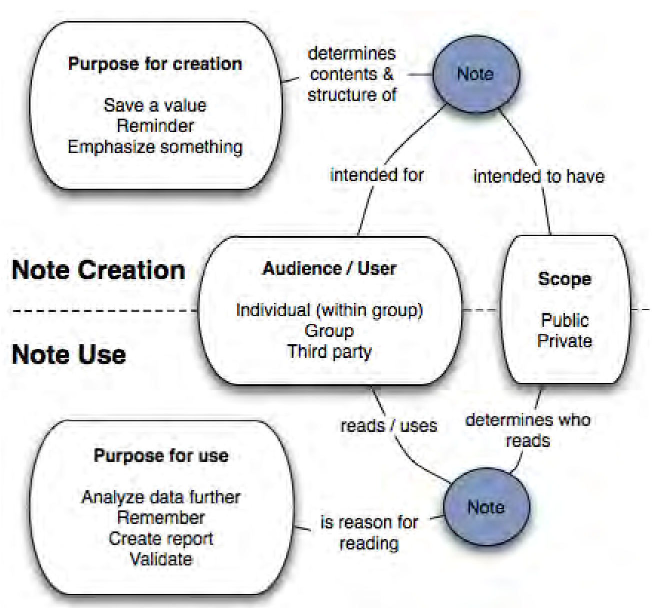


Fig. 5. Taxonomy of note creation and use.

We noticed that the manual note taking process impacted awareness. Participants lost a sense of what others were doing when taking notes, and consequently their awareness level was reduced. Each participant had to catch up with others after finishing taking notes. For example, Figure 1 illustrates how users who are taking notes on paper need to divert their attention from the group and the shared display. This drawback suggests that it may be good to integrate some types of notes with the visualization (as annotations). This feature could facilitate note taking in groups. It is obvious that not all the notes taken by users are appropriate to integrate with charts; we would still need to provide users with means of taking personal notes (e.g. a personal reminder) and notes that do not belong with any given chart (e.g. a "to do" item).

In some sessions where one person was in charge of note taking, others also took notes for themselves separately even though they had to stop working to take notes. This shows that participants needed to take notes individually and separately from the group. However, individual notes were not always solely used by the person who took them; sometimes they were shared by the group. This finding again emphasizes the necessity for software to support both individual as well as jointly coupled activities [5] [11] [14]. We also noticed that in task 1 (in which users were not saving charts), the amount of note taking was much higher than in task 2 (in which users were saving charts for comparison).

### 4.2 Note Taking in Competitive and Collaborative Situations

Our study suggests that nature of the task can affect both the process of collaboration and division of workspace. Task 1, which involved focused questions, required a highly-coupled collaborative style of work, while task 2, which required participants to compete for resources, led to a loosely-coupled collaborative work style. Here participants wanted to work individually to prepare the best possible arguments for increasing their state's resources. Hence, a competitive situation has a clear impact on user's collaboration style and process. Most of our participants said that they preferred to explore information for task 2 individually and later on share their results with other collaborators to have a discussion. Notes taken in task 1 had a public scope of use, while notes taken in task 2 had a combination of public and private scopes.

## 5 Discussion

Our findings suggest the importance of note taking for collaborative business data visualization and analysis. These findings raise further questions and issues such as:

**How can we best support note taking activities during collaborative work?** One probable answer to this question could be integrating note taking mechanisms into the software, which in turn raises issues such as how closely integrated note taking should be with the visualization tool, and whether it should be integrated with a history mechanism or should be a separate component. Some researchers [1] [2] [3] [6] have mentioned use of annotation (textual and graphical) to add information into visualization. But it is still not quite clear what the best strategy is to save information in a collocated collaborative visualization and analysis of business data where intensity of note taking is quite high.

**How can we support both individual and group notes?** Can this be accomplished by dividing work space into public and private areas?

**Does the process of note taking change by changing the underlying data?** For instance, working with business data might require larger amounts of note taking compared to working with scientific data. This is currently unclear.

**How complete is our list of note contents and purposes of use?** Will participants in a different domain or different situation need to save different information as notes, and will they have different purposes in creating and using notes?

## 6 CONCLUSION AND FUTURE WORK

In this paper, we identified note taking as a process that is intensively used by data analysts. More studies are required to answer questions about how exactly note taking support should be provided in collaborative visualization systems. In addition, we would like to conduct a field study to examine note taking activities in the context of real work. We would also like to explore the design of note taking support for collaborative work on interactive surfaces.


## ACKNOWLEDGEMENTS

This research was supported by grants from SAP Business Objects and the Natural Sciences and Engineering Council of Canada (NSERC). We wish to thank Kellogg Booth and our colleagues at SAP Business Objects for their many thoughts and suggestions. Many thanks to Dr. Rebecca Grant for helping us to find business students for this study and for her valuable comments to improve our business case. We also would like to thank all the VISID lab members for their feedback and participation in pilot studies.


## REFERENCES


[1] S. E. Ellis, and D. P. Groth, "A collaborative annotation system for data visualization", *Proc. Working Conference on Advanced Visual interfaces,* 2004

[2] J. Heer, F.B. Viegas, and M. Wattenberg, "Voyagers and voyeurs: Supporting asynchronous collaborative visualization," Commun. ACM, vol. 52, pp. 87-97, 2009.

[3] P. Isenberg, and S. Carpendale, "Interactive Tree Comparison for Co-located Collaborative Information Visualization", *IEEE TVCG*, Vol. 13, no. 6, pp. 1232–1239, 2007.

[4] P. Isenberg, A. Tang, A., and S. Carpendale, "An Exploratory Study of Visual Information Analysis", *Proc. CHI 2008*, pp. 1217-1226, 2008.

[5] P. Isenberg and D. Fisher, "Collaborative Brushing and Linking for Co-located Visual Analytics of Document Collections", *Proc. EuroVis, Computer Graphics Forum*, Vol. 28, no. 3, pp. 1031–1038, June 2009.

[6] T. Isenberg, P. Neumann, S. Carpendale, S. Nix, and S. Greenberg, "Interactive Annotations on Large, High-Resolution Information Displays", *Proc. 2006 Conference Compendium of IEEE VIS, IEEE InfoVis, and IEEE VAST*, pp. 124–125, 2006.

[7] T. Jankun-Kelly, K.-L. Ma, and M. Gertz, "A Model and Framework for Visualization Exploration", IEEE Transactions on Visualization and Computer Graphics, Vol. 13, no. 2, pp. 357–369, March/April 2007.

[8] G. Mark, K. Carpenter, and A. Kobsa, "A Model of Synchronous Collaborative Information Visualization", *Proc. Seventh international Conference on information Visualization* , 2003.

[9] G. Mark, and A. Kobsa, "The Effects of Collaboration and System Transparency on CIVE Usage: An Empirical Study and Model", *Presence*, Vol. 14, no. 1, pp. 60–80, 2005.

[10] K.S. Park, A. Kapoor, and J. Leigh, "Lessons Learned from Employing Multiple Perspectives In a Collaborative Virtual Environment for Visualizing Scientific Data", *Proc. CVE 2000*, pp.73–82. 2000.

[11] A.C. Robinson, "Collaborative Synthesis of Visual Analytic Results", *Proc. VAST 2008*, pp. 67-74, 2008.

[12] Y. Rogers and S. Lindley, "Collaborating around large interactive displays: Which way is best to meet?", *Interacting with Computers*, Vol. 16, no. 6, pp. 1133–1152, 2004.

[13] S. D. Scott, K. D. Grant, and R. L. Mandryk, "System guidelines for co-located collaborative work on a tabletop display", *Proc. ECSCW'03*, pp. 159–178, 2003.

[14] S.D. Scott, M. S. T. Carpendale, and K. M. Inkpen, "Territoriality in collaborative tabletop workspaces", *Proc. CSCW 2004*, pp. 294-303, 2004.

# Studying Direct-Touch Interaction for 2D Flow Visualization

Tobias Isenberg, Uta Hinrichs, and Sheelagh Carpendale

**Abstract**—Traditionally, scientific visualization research concentrates on the development and improvement of interactive techniques to support expert data analysis. While many scientific visualization tools have been developed for desktop environments and individual use, scenarios that go beyond mouse and keyboard interaction have received considerably less attention. We present a study that investigates how large-display direct-touch interaction affects data exploration and insight generation among groups of non-experts exploring 2D vector data. In this study, pairs of participants used interaction techniques to customize and explore 2D vector visualizations and collaboratively discussed the process to develop their own understanding of the data sets.

**Index Terms**—Interactive scientific visualization, direct-touch interaction, wall displays, qualitative evaluations, 2D vector visualization.

## 1 Introduction

Research in scientific visualization has made enormous progress in recent years, allowing us to get a better understanding of complex datasets. With the exception of co-located VR applications and a number of distributed visualization environments (including VR ones), however, most research in scientific visualization has focused on developing and improving techniques that are aimed toward individual data exploration and analysis by expert users. Moreover, such analysis and exploration typically occurs in desktop environments using keyboard-and-mouse interaction. However, the advance of large, touch-sensitive display hardware has enabled us to explore other forms of interaction techniques within scientific visualization environments.

Most notably is direct-touch interaction on large displays which potentially has a number of advantages for scientific visualization. Large displays offer more space for high-resolution data and support co-located collaborative data analysis, adding the possibility to actively present, discuss, and explore hypotheses or findings about data in-place. Direct touch also enables direct interaction with visual elements as well as gestures or hand postures as alternatives to mouse-based interfaces. Such alternative interaction techniques may be more suitable in collaborative analysis tasks since direct-touch provides rich awareness cues. These cues are important during collaborative analysis and seamlessly integrate with group discussion. It has been found, e. g., that 'hands-on' interaction can enhance engagement and understanding, especially within learning environments [1], showing promise for scientific visualizations targeted toward non-expert audiences as well.

The idea of analyzing scientific visualizations on large displays using direct-touch interaction raises several research questions. For instance, how do scientific visualization techniques need to be designed to support large-display direct-touch interaction? And how does the use of large display hardware and direct-touch interaction influence the way how people approach the analysis of scientific datasets?

We present a qualitative study that explores the potential of an interactive 2D vector analysis visualization tool [8] as used by groups of non-experts. The study was conducted on a direct-touch enabled large-display. We asked non-expert pairs to analyze two different 2D vector field datasets. The results of this study further our understanding of how people made use of this interactive visualization tool and of the methods they used to explore the data individually and collaboratively.

We found that the ability to directly experiment with the data visualizations of previously unknown datasets helped participants to develop an understanding of the data's behavior and to detect local phenomena and causalities. All participant groups engaged in lively discussions and made varied use of the customization and interactivity offered in the tool. We observed frequent turn-taking, communicated through movement in front of the display and gesturing. Customizable glyphs, animation, and direct manipulation of visualization elements were used extensively to create personalized visualizations and supported various data analysis strategies including local and global data exploration. Participants did not follow a linear sequence of analysis strategies but fluidly went back and forth between different exploration activities. This lack of temporal sequencing in tasks parallels observations in a previous study on co-located collaborative work [7].

## 2 Related Work

Most systems for scientific visualization are developed for single-user desktop systems, thus interaction with the data typically occurs using keyboard or mouse. With the advent of large, high-resolution displays some time ago, more specific virtual analysis environments such as the CAVE or the Responsive Workbench were developed that required new interaction metaphors for work with visualizations. The interaction design of these environments aimed at creating a natural mapping of physical input, for example from tracked gloves or wands, to convey a feeling of embodiment in the virtual world.

The progression of large touch-interactive screen technology offers new interactive environments for scientific data analysis that do not require people to wear special equipment such as glasses or gloves to perform interactions. Such direct interaction techniques can be more accessible because of their resemblance with real-world interaction, and lend themselves more easily to collaboration [15]. However, with large display direct-touch interaction, we face a number of challenges: we need to design adequate mappings between input and interaction to support scientific data exploration and learn how people can adapt to these mappings on large display surfaces in general. Some research has been done in this direction. For example, Forlines and Shen [4] visualize geospatial data and explore multi-user zoom-in-context interaction. They map the user input to data manipulations by providing dedicated elements, DTLenses, that represent the data manipulations. Our study is based on [8] where similar exploration objects and hand postures are used to visualize 2D vector field data. We build upon this work, looking specifically how groups of people make use of these features during their data exploration process.

The type of data used in our study, vector or flow data, is relevant to many application domains such as physics or meteorology. Several visualization techniques have been developed to help analyze 2D and 3D vector data. These techniques include direct visualization using, for instance, glyphs, texture-based approaches, geometric techniques, and feature extraction [10]. Even though many recent approaches address more complex issues in three-dimensional vector visualization, two-dimensional techniques still play an important role. This is particularly true for datasets that are presented to a non-expert audience, for example, in weather reports and forecasts or in educational environments such as geography classes. Closely related to the visualization techniques employed in our study are methods that allow interactive exploration of vector data such as selecting a specific view of the data or changing global attributes. Some techniques go beyond such straightforward interaction and explore interactive probing and annotation of

• *Tobias Isenberg is with the University of Groningen, E-mail: isenberg@cs.rug.nl. Uta Hinrichs and Sheelagh Carpendale are with the University of Calgary, E-mail: {uhinrich|sheelagh}@ucalgary.ca.*

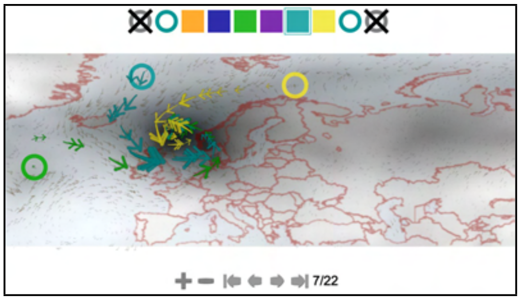


Fig. 1. Screenshot of the adapted interface for collaborative interactive 2D vector data exploration and visualization, with example visualization.

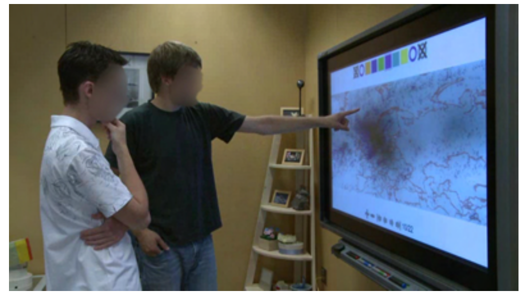

Fig. 2. Participants in front of the touch-sensitive wall display, collaboratively working with our data exploration and visualization tool.

flow fields with glyphs to show local properties [3] or use customized glyphs [9, 13, 17]. Also, techniques have been developed for interactive particle sources placement [11, 18]. The techniques we use in our study are related to these approaches in that all are based on visualization of 2D vector fields using particle sources while the tool used in our study [8] provides ways for direct-touch interaction with static and animated, customized glyphs as outlined in the next section.

## 3 Hand-Posture-based 2D Vector Visualization

Combining traditional 2D vector visualization techniques with direct touch display technology, Isenberg et al. [8] have previously developed a vector visualization technique that allows people to explore 2D vector data using several hand postures on a large interactive wall display. This approach is unique in three ways. First, it allows people to draw their own glyphs to represent the vector data. Second, the tool enables global exploration of data behavior by filling larger areas of the data display with glyphs. Alternatively, data can be explored locally by using ‘glyph sources’ that constantly emit animated glyphs. Third, the interface of the visualization tool is also based on hand postures. These postures rather than typical buttons and sliders are mapped to adding or removing glyphs to and from the data display.

These interactive visualization techniques can be used as follows. After a 2D vector dataset (e. g., wind data) is loaded, people can bring up a drawing canvas where they create (draw) the desired glyph to represent the data (e. g., an arrow or a straight line) using different hand postures to determine the stroke width. When the drawing canvas is closed, glyph instances can be added to the data display with the loaded 2D vector data using different hand postures (fist for adding several glyphs or finger for adding individual glyphs). The orientation and size of the glyphs reveal the character of the loaded data, namely the local direction and strength of the vector field. To explore the flow behavior of the data, glyph sources can be added to the data display. These sources can be moved around within the data display to explore different areas. Different shapes, line widths, and colors of glyphs can be used simultaneously to highlight certain aspects of the dataset.

For the study we added a number of features to the interactive visualization to account for the study tasks as well as to address some usability issues uncovered during pilot studies. Those features included controls to add or removes globally from the data display, and cycling through the different time steps of the data set (Fig. 1).

## 4 Exploratory Study

The goal of our study was to better understand how pairs of people work with customizable vector visualization on a large touch-interactive display. Specifically, we investigated the potential of hand-posture-driven interaction with vector data to ease the process of exploring and visualizing such datasets and how people approach 2D vector data analysis tasks using our visualization tool on a large wall display. Quantitative study methods that rely on controlled study scenarios and the collection and analysis of numerical data are less adequate for answering open-ended questions like these. We therefore conducted an exploratory laboratory study where pairs of participants were asked to complete a series of experimental tasks. The study was based on a mixed methods approach which allowed us to combine qualitative with quantitative data collection and analysis to shed light into people’s analysis processes while using our visualization tool.

**Participants and Setting.** Sixteen university students (seven females, nine males) participated in our study. They were asked to work together in groups of two as a strategy to increase verbal explanations and ‘thinking aloud,’ resulting from discussions with the partner. This helped us to gain insights into participants’ thought processes during the data analysis. Groups performed the study tasks on a 5’ × 3’ plasma wall display (1360 × 768 pixels) with direct-touch interaction enabled using a SMART DViT overlay. The experimental software ran on an Intel 2.4 GHz Windows XP PC. The display was large enough to comfortably accommodate both participants standing next to each other (Fig. 2). While participants had to interact with the display standing, they were able to sit down on a couch located in front of the display. Participants were free to move around the display while solving the experimental tasks. Due to technical reasons our visualization tool is a single-touch interface which forced participants to take turns with display interaction. While this condition led to some interferences between group members during the visualization analysis, it did not hamper collaboration as we will describe in the findings.

**Experimental Data, Tasks, and Procedure.** Participants were asked to use our visualization tool to analyze two different real-world 2D vector datasets: wind data from a storm that hit Europe in the spring of 2008 (22 time steps, enhanced with a map of Europe showing high and low air pressure zones in form of a gray-scale color scheme; Fig. 1), and a moving fluid simulation where an obstacle causes turbulences (22 time steps, enhanced with a line integral convolution visualization [2] of the vector data, showing flow direction; Fig. 3, left).

Participant groups were asked to work on five tasks in total—four tasks based on the wind dataset and one task involving the fluid dataset. After a short introduction to the tool and a practice task, each study task involved answering one or two open ended questions concerning relations within the data (e. g., relation of low/high air pressure and wind speed) as well as the assignment to illustrate certain particularities within the data using the visualization tool. The assignment for the second dataset (moving fluid) was a free-form exploration of data followed by a presentation of the general behavior of the data. Groups were given 10–15 minutes to solve each of these experimental tasks. A semi-structured interview concluded the study where we elicited subjective perceptions from participants regarding their general experiences with the visualization tool. We also followed up on certain analysis strategies we had observed which participants used to gain insights into how they approached the experimental tasks. Each study session took approximately 1.5 hours in total.

The purpose of this setup was to observe how the participant groups would approach the tasks in general, using the visualization tool. We were, in particular, interested in the different strategies that groups would apply to solve the tasks. We used the answers that groups provided for each question/task as an indicator how well they understood the dataset after this short time of exploration. Furthermore, we closely observed groups conducting the experimental tasks to be able to evaluate the interaction techniques the visualization tool provides.

**Data Collection and Analysis.** Two examiners oversaw each study session to minimize bias, with one of them being external to the project team. Each study session was videotaped and screen captures were taken every 10 seconds. Both experimenters took notes of their observations, highlighting particular events and participants' comments. These field notes informed the semi-structured interview to ask participants about their strategies or events that happened while participants were solving the tasks. To answer our research questions, the video data was analyzed in-depth regarding exploration strategies that participants applied for the experimental tasks and the use of particular features of the visualization tool, e. g., animations and customized glyphs. We also analyzed how participants interacted with and in front of the large touch-interactive display, individually and with the partner.

## 5 Findings

Our findings focus on the strategies that groups applied for exploring the provided datasets using large-display direct-touch interaction. We first provide a brief overview of how groups generally approached their analysis tasks before addressing the influence of the physical display characteristics on exploratory data analysis and collaboration.

### 5.1 Exploration Approaches

In general, all participant groups took a similar approach for exploring and visualizing the data which can be categorized into four activities: drawing, overview, local exploration, and temporal exploration:

**Drawing.** Our tasks and visual tool required groups to first draw a glyph in order to explore the data, hence this was the activity participants started with. Participants frequently experimented with different glyph shapes and hand postures when drawing glyphs, creating on average 2.75 ($\sigma = 1.581$) different glyphs for the illustration task and 1.594 ($\sigma = 0.946$) for all other tasks.

**Overview.** After a representative glyph was drawn, participants of all groups first tried to gain an overview of the dataset. To achieve this they equally distributed the drawn glyph on the display, generally using the '+' button (in 97.5% of all tasks) and, in addition to this, occasionally the fist posture (in 35.0% of all tasks) to adjust glyph density in some areas of the data display.

**Local Exploration.** All groups explored local features of the dataset, in particular by placing sources in various locations (in 82.5% of all tasks). Sources were dynamically moved to probe and explore local aspects of data. Occasionally, groups used the one finger posture for fine-grained continuous local exploration (in 22.5% of all tasks).

**Temporal Exploration.** Since our study tasks required exploration of temporal changes, all groups made heavy use of the time controls, frequently stepping back and forth in time. This temporal exploration sometimes required rearrangement of sources or adding glyphs.

These activities did not follow a linear sequence, paralleling previous work [7]. While participants would typically start by drawing a glyph and distributing it evenly across the data display for overview, local and temporal exploration happened without a visible sequence, with participants often switching back and forth between both activities. Also, participants sometimes decided in the middle of a task to go back to the drawing stage and to bring in a new glyph shape.

We also noticed that no group went through a pre-discussion on how to solve the task in general, what steps to take, or what an answer should look like, before actually approaching the task. Instead, discussion occurred in parallel to the data exploration and visualization activities and evolved naturally from the task. Participants would either just start an activity and discuss observations in the data display during or after the action took place. They sometimes had just a brief exchange about next exploration steps, for instance, when they wanted to draw a new glyph or move on to the next time step.

### 5.2 The Role of Large Display Direct-Touch Interaction

We observed close collaboration among participants of all groups and noticed a high engagement of participants (for all groups and tasks), visible in active discussions of ideas or hypotheses and lively interaction with the data display. This is apparent in the extent to which they discussed ideas concerning the glyph drawing, exploration strategies, and hypotheses about the data. In two groups, one participant was more dominant and took the lead in activities such as glyph drawing. In the remaining six groups, participants actively took turns with glyph drawing and data exploration. Typically, during non-active moments, a partner would participate in activities by providing verbal feedback to actions carried out by the other participant. Constraining participants to single-person interaction sometimes led to interferences with both participants trying to interact with the display at the same time and the system ignoring the second input. However, conflicts like this were usually resolved quickly and did not limit collaboration among group participants. In fact, we observed frequent turn taking among all groups and study tasks. A participant of a group would step back from the display, handing over the exploration to the partner. Meanwhile, he or she would actively follow the interactions of the partner and observe changes on the display, always prepared to jump in and take over if some new idea occurred (e. g., Fig. 3, left).

This contrasts previous studies involving collaborative tasks on large vertical displays where groups were found to usually elect a person 'in charge' of the interaction while other group members would stay rather passive [14]. We attribute this high engagement of both group participants to their similar level of expertise with the datasets, the exploration techniques we provided, and the physical study setting. The datasets were relatively unknown to participants, thus close collaboration and discussion was important for coming up with hypotheses and causalities. This was facilitated by the large display and rather broad interaction techniques, easily visible to both participants. The visualization tool did not force participants to linearly follow a predetermined sequence of exploration activities, but allowed participants to explore based on their interest. Consequently, participants repeatedly and without any effort switched between different exploration strategies discussed above, including the observation of the data visualization. Thoughts, ideas, and hypotheses were collaboratively discussed and explored, through interaction and discussion, ultimately leading to a basic understanding of the data, evident in the answers that participants provided for our experimental questions.

Fluid collaboration and maintaining awareness of the partner's exploration activities require the support of deictic communication means, such as gestures and body movement (e. g., [6, 16]). We found that the large display and the direct-touch interaction supported these communication mechanisms well, allowing participants to frequently switch back and forth between data exploration and gestures without having to worry about external input devices. All groups frequently used hand gestures, such as pointing with a finger, a hand, or both hands, to communicate and exemplify certain insights (e. g., Fig. 3). Participants expressed ideas and thought processes to their partner via gestures, directly speaking to each other, and accompanying an action with speech. Turn-taking often happened non-verbally, communicated by stepping toward the display or reaching out for it.

## 6 Discussion

Reflecting critically on our findings and interpreting the results, we consider our study to be a step toward understanding the use of direct-touch interfaces for data exploration and visualization and discuss in the following new research questions that arise from our findings.

We found that the **possibility to manipulate elements through direct touch** enabled participants to quickly test different hypotheses. Most interaction techniques had a direct local impact on the visualization elements. Moving the finger while touching a source, for example, would move the source to a different location or running the fist across the display would locally add glyphs. Participants had no problems understanding this direct mapping, evident in their fluid interaction with the visualization tool. Another advantage of direct touch in our large-display setting is that it enabled the temporarily non-active participant to visually track the activities carried out by their partner, contributing to the awareness during collaboration. When supporting direct-touch manipulations of elements, however, it is important to map interaction techniques in a consistent way. While our tool allowed to directly move sources within the data display, this was not possible for already created glyphs. This caused some confusion among participants who tried to move glyphs via direct-touch interaction.

With the exception of drawing a glyph as a first step which was required by the program, participants did not explore data following a

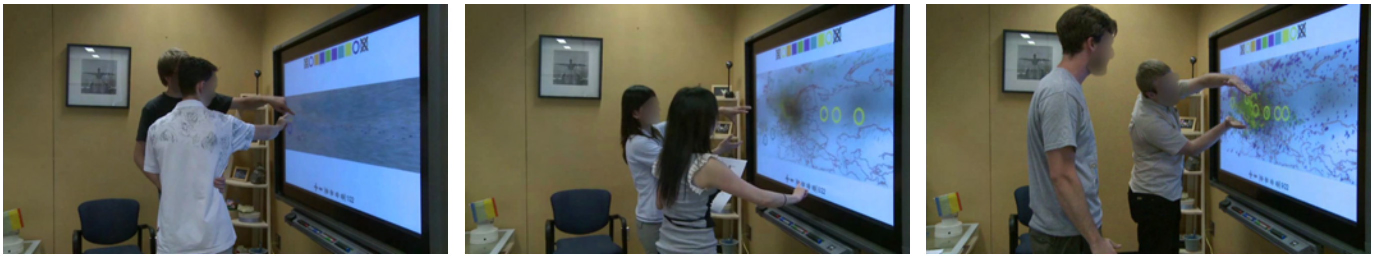

Fig. 3. Participants using various gestures to explain ideas or thoughts to each other, and/or to reference locations on the data display.

particular sequence of activities. Exploration strategies such as gaining an overview of the data, examine local regions or probing, creating additional glyph shapes, or exploring data along temporal dimensions were applied in various sequences. The tool did not enforce a certain sequential order in which exploration activities had to be applied (except for the initial drawing), but participants decided when to apply them. We believe that the support of this kind of **free-form data exploration** is important to enable interest-based exploratory data analysis. That is, visualization tools should provide certain basic functionalities that can be used by people as needed and which even can be 'appropriated.' Appropriation happened with our visualization tool, for instance, when some groups 'invented' continuous probing. While both the sources and the single-finger-posture were initially intended for other activities, participants appropriated them for their own purposes. While future work needs to explore how scientific visualization tools can support free-form data exploration in general, we hypothesize that it is, in particular, facilitated through large-display direct-touch interaction because this form of interaction is evocative of how people 'handle' and appropriate basic tools in real life.

It has been shown that collaborative data exploration can lead to better results and insights than analysis by individuals [12]. In our study, participants collaboratively discussed ideas and hypotheses while exploring the data at the same time and adjusting their analysis based on this discussion. We believe that this combination of discussion and exploration activities positively affected participant's understanding of the data and was enabled by the interactive large display technology we employed. However, further studies are needed to confirm this. The use of **large displays to support co-located collaboration for scientific visualization**, thus, needs to be explored further—in scientific, educational, and other domains. While our study setup did not support simultaneous multi-touch interaction, it would also be interesting to investigate exploration techniques that make use of this technology and its impact on collaboration strategies during scientific data analysis.

## 7 Conclusion and Future Directions

We discussed aspects of a study that investigated how pairs of people explore customizable vector visualizations on a large, touch-sensitive screen. Several factors in this visualization lead our participants to in-depth and insightful explorations of the data. This is a promising result since our participants were not experts in scientific visualization. Overall, we observed a high degree of engagement that was evident in the frequent turn-taking and exchange of ideas by participants. We attribute this high engagement of our non-expert participants and the quick understanding they were able to gain about the unknown datasets in part the direct-touch interface, as well as the possibilities to customize and personalize the data display. We believe that tools such as this one could be particularly useful in classroom settings, enhancing traditional teaching methods with 'hands-on' learning. Students exploring data using the tool would, similar to our participants, directly experience certain correlations such as the specific rotation directions of low and high pressure zones, which could foster learning.

The findings from our study point toward numerous interesting future research directions in the application of interactive and animated visualizations. In particular, the potential of large-display direct-touch interaction should be explored further with regard to different scientific datasets including three-dimensional ones. 3D direct-touch techniques using multi-handed interaction have been developed for large horizontal direct-touch displays [5, 19] and could be applied to various scientific visualization techniques, such as volume renderings or three-dimensional vector datasets. In addition, our study suggested the benefit of co-located collaboration for scientific data analysis. While most visualization and interaction techniques for scientific visualization concentrate on individual analysis scenarios or distributed collaboration, future research could consider co-located collaborative settings with several people discussing and interacting with scientific data together.

## References


[1] S. Allen. Designs for Learning: Studying Science Museum Exhibits that do more than Entertain. *Science Education*, 88(S1):S17–S33, July 2004.

[2] B. Cabral and L. C. Leedom. Imaging Vector Fields using Line Integral Convolution. In *Proc. SIGGRAPH*, pp. 263–270, New York, 1993. ACM.

[3] W. de Leeuw and J. van Wijk. A Probe for Local Flow Field Visualization. In *Proc. VIS*, pp. 39–45, Los Alamitos, 1993. IEEE Computer Society.

[4] C. Forlines and C. Shen. DTLens: Multi-User Tabletop Spatial Data Exploration. In *Proc. UIST*, pp. 119–122, New York, 2005. ACM.

[5] 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.

[6] W. C. Hill and J. D. Hollan. Deixis and the Future of Visualization Excellence. In *Proc. VIS*, pp. 314–320, Los Alamitos, 1991. IEEE Comp. Soc.

[7] P. Isenberg, A. Tang, and S. Carpendale. An Exploratory Study of Visual Information Analysis. In *Proc. CHI*, pp. 1217–1226. ACM, 2008.

[8] 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.

[9] M. Kraus and T. Ertl. Interactive Data Exploration with Customized Glyphs. In *Proc. WSCG*, pp. 20–23, 2001.

[10] R. Laramee, H. Hauser, H. Doleisch, B. Vrolijk, F. Post, and D. Weiskopf. The State of the Art in Flow Visualization: Dense and Texture-Based Techniques. *Computer Graphics Forum*, 23(2):203–221, June 2004.

[11] K.-L. Ma and P. J. Smith. Virtual Smoke: An Interactive 3D Flow Visualization Technique. In *Proc. VIS*, pp. 46–53, Los Alamitos, 1992. IEEE Computer Society.

[12] G. Mark, A. Kobsa, and V. Gonzalez. Do Four Eyes See Better than Two? Collaborative versus Individual Discovery in Data Visualization Systems. In *Proc. IV*, pp. 249–255, Los Alamitos, 2002. IEEE Computer Society.

[13] W. Ribarsky, E. Ayers, J. Eble, and S. Mukherjea. Glyphmaker: Creating Customized Visualizations of Complex Data. *Computer*, 27(7):57–64, July 1994.

[14] Y. Rogers and S. Lindley. Collaborating around Vertical and Horizontal Large Interactive Displays: Which Way is Best? *Interacting with Computers*, 16:1133–1152, Dec. 2004.

[15] C. Shen, K. Ryall, C. Forlines, A. Esenther, F. D. Vernier, K. Everitt, M. Wu, D. Wigdor, M. Ringel Morris, M. Hancock, and E. Tse. Informing the Design of Direct-Touch Tabletops. *IEEE Computer Graphics and Applications*, 26:36–46, Sept./Oct. 2006.

[16] J. C. Tang. Findings from Observational Studies of Collaborative Work. *Intern. Journal of Man-Machine Studies*, 34(2):143–160, Feb. 1991.

[17] R. van Teylingen, W. Ribarsky, and C. van der Mast. Virtual Data Visualizer. *IEEE Transactions on Visualization and Computer Graphics*, 3(1):65–74, Jan.–Mar. 1997.

[18] D. Weiskopf, R. P. Botchen, and T. Ertl. Interactive Visualization of Divergence in Unsteady Flow by Level-Set Dye Advection. In *Proc. SimVis*, pp. 221–232, Erlangen, 2005. SCS European Publishing House.

[19] A. D. Wilson, S. Izadi, O. Hilliges, A. Garcia-Mendoza, and D. Kirk. Bringing Physics to the Surface. In *Proc. UIST*, pp. 67–76. ACM, 2008.

# History Tools for Collaborative Visualization

Ali Sarvghad, Narges Mahyar, and Melanie Tory

**Abstract**—In the context of collaborative data visualization and analysis, history tools can play an important role. We present a compilation that characterizes users' probable objectives when using history tools for collaborative work, as well as operations commonly performed on histories. We further characterize user objectives according to the likely time/space setting in which they would be used, and whether they are likely to be used by individuals, groups, or both. We conclude by compiling a list of design and implementation challenges and research questions that need to be discussed and investigated in order to make history tools adequately support collaborative visualization activities.

**Index Terms**—History tool, collaboraton,visualization, analysis.

## INTRODUCTION

In this paper, we present a preliminary list of history operations and users' most common objectives in the context of collaborative data visualization and analysis. We then identify a set of key research challenges that will need to be addressed in order to make history tools effective for collaborative visualization tasks. Though the lists of history operations and objectives presented in this paper are detailed, they are neither final nor complete. They are preliminary lists assembled to trigger discussions and raise questions regarding collaborative use of histories in data visualization and analysis.

Many researchers have mentioned advantages of history tools and their importance for collaborative data visualization and analysis [5][6][10][12][14]. However, to date, visualization histories have been designed only for individual use, not communal use. Histories for group use will demand a new set of functionalities and design considerations.

Several applications provide a general-purpose undo/redo tool but this simplest form of history reuse is inadequate for participants of a collaborative data visualization and analysis task. They need to use history items, individually or collectively, to coordinate their work, try a different course of visualization and analysis, recover from a system crash, train naive users, and so on. Scientific workflow management tools such as Vistrails [1] and Taverna [16] capture very detailed information about scientific workflows. This information consists of data, created visualizations, and their manipulation. Though these systems maintain rich historical (provenance) information, they are designed primarily for expert users who are able to understand and manipulate complex workflows for creating visualizations from scientific data. More importantly, these systems have not been designed with collaborative work in mind.

History items can be browsed [5][12], searched [5][12], edited [5][12], filtered [5] and exported [5] for different purposes such as analysis, decision-making, validation and correction. As the name "history tool" suggests, users can revisit and reuse historical items. This reuse involves enacting specific operations to achieve specific objectives. In the following sections we will point out what we expect to be the most common operations and objectives performed and intended by history tool users.


- *Ali Sarvghad is with the University of Victoria, E-Mail: asarv@cs.uvic.ca.*
- *Narges Mahyar is with the University of Victoria., E-Mail: nmahyar@cs.uvic.ca.*
- *Melanie Tory is with the University of Victoria, E-Mail: mtory@cs.uvic.ca.*


## 1 MOST COMMON OPERATIONS ON HISTORY REPOSITORIES

Heer et al. [5] list a number of operations that a history tool should support. We built our list of operations largely based on their work, but we make some alterations. We expect the most common operations on history repositories to be:

- Browse
- Search
- Filter
- Edit
- Delete
- Export

We consider an *editing* operation to be changing the content of a history item, such as adding metadata, and we consider *deleting* history items to be independent from editing. We also consider *searching* and *filtering* as two different operations. Other researchers also point out the importance of *browsing* [3][12], *searching* [12][13] and *editing* [4] operations for history tools and some other researchers [4][8] mention the necessity of a tool to *export* and communicate history.

There is no one to one dependency between operations and user objectives. In other words, an operation, solely or in conjunction with other operations, can be performed to achieve a number of objectives. For example, searching and filtering both are required to accomplish analysis and validation objectives.

## 2 MOST COMMON OBJECTIVES

Based on a literature survey and our own experience, we expect history operations would be mainly used to achieve the following objectives:

- **Analysis** [3][8]: Users can traverse a history item repository and revisit different data visualizations to investigate data. Products of this analysis can vary from making a decision to verifying a hypothesis. We define analysis as investigating data with a specific goal in mind, in contrast to exploration.

- **Validation** [5][8][10][14]: Correctness and admissibility of decisions/findings or appropriateness of a single visualization can be examined by using history items. For instance, analysts may review visualizations created in the course of an analysis process to double-check that their findings are correct, or they may revisit a particular visualization to ensure that it is the result of correct mapping and filtering of data. This might be more helpful when users' collaboration style changes over time such as autonomous collaboration. Participants may need to corroborate the outcomes on individual works that will be concatenated later.

- **Memory aid**: The limitation of humans' short-term memory is a known fact, and a history tool can act as external memory aid [12]. Data analysts can add important notes, observations, calculations et cetera to history items for future referral.
- **Correction/Recovery**: If data analysts find their current visualization undesirable for any reason, they can perform a selective undo/redo [3][5][8][11][13][5]. It is also possible to continue a visualization and analysis process from the last point in the history repository after a system failure.
- **Exploration**: Exploration involves investigating data without a specific goal in mind. Having a repository of history items enables data analysts to try different courses of visual analysis by revisiting a history item and trying a different possible path. "Insight often comes from comparing the results of multiple visualizations that are created during the data exploration process" [2].
- **Reporting** [5][12]: A history repository, wholly or partially, can be sent to peers or upper management as a progress report, indication of the amount of work done, or formal report of findings.
- **Presentation** [5]: History items can be summarized and presented in a meeting situation. Presentation is similar to reporting, but typically occurs synchronously, as shown in Table 1.
- **Coordination** [4][8][11][12][14]: History items can help collaborators coordinate their effort by increasing awareness in situations such as autonomous collaborative work or remote synchronous/asynchronous situations. Also, viewing another users' history can bring a person up-to-speed on the work done so far.
- **Training** [12]: Novice data analysts can learn from experts by reviewing the history of visualizations created and decisions made.

It is quite possible that users have a combination of objectives when working with history items. For instance, users might review visualizations created in the course of an analysis process to both ensure their validity (i.e. correctness/admissibility) as well as make a decision.

**Research** is an interesting additional objective offered by rich history tools. Researchers can survey users' behaviour or assess a system's usability by observing the history of analysts' actions [5]. We do not include it in Table 1 because it is not performed directly by visualization users; nonetheless, it is worth mentioning.

## 3 EFFECTS OF TIME/PLACE SETTING

Table 1 predicts the most likely time/place settings in which each objective might occur. As shown in the table, most of the objectives are likely to occur in all of the different time/place settings. However, we suspect that history records may need to be more explicitly displayed for synchronous distributed work in order to help users maintain awareness of others' activities. Additionally, using histories in asynchronous work may require different functionality than synchronous work. For instance, when sharing a history with another user who will take over the work later, a person may want to highlight particularly important findings to ensure they are noticed, or remove an unsuccessful path of analysis and replace it with a simple note to say that investigating that direction was not fruitful.

Table 1: Objectives' most likely time/place setting. ST = same time, DT = different time, SP = same place, DP = different place

| | ST, SP | ST, DP | DT, SP | DT, DP |
|---|---|---|---|---|
| **Analysis** | √ | √ | √ | √ |
| **Validation** | √ | √ | √ | √ |
| **Memory aid** | √ | √ | √ | √ |
| **Correction/Recovery** | √ | √ | --- | --- |
| **Exploration** | √ | √ | √ | √ |
| **Reporting** | --- | --- | √ | √ |
| **Presentation** | √ | √ | --- | --- |
| **Coordination** | √ | √ | √ | √ |
| **Training** | √ | √ | √ | √ |

## 4 INDIVIDUAL VS. COLLABORATIVE USE OF HISTORIES

Reporting, collaborating, coordinating and training are inherently collaborative objectives and require engagement of more than one person; the rest of the objectives could apply to both individuals and collaborating users. Though individuals and groups share most of the objectives, design of a history tool might need to be quite different to support the activities of a group as compared to one person. To adequately support group activities, we anticipate that history tools may need to provide the following functionality:

- Representation of *who* was responsible for each action recorded in the history.
- Both individual and shared histories. This will hopefully prevent users from being overwhelmed with history items from all members of the group. In addition, privacy control may be needed so that some items can be kept private.
- Additional awareness mechanisms, such as an indication that another user has worked on a similar chart or has looked at the same data. This might be similar to awareness mechanisms previously used in collaborative document search [9].
- Extensive editing, highlighting, and annotation capabilities. These will help users to communicate what they have done, or convert a history into a series of visual items and descriptions suitable for a report, presentation, or tutorial.
- Ability to export elements of a history to a document or presentation format for further manipulation.

## 5 DESIGN CHALLENGES/QUESTIONS

There are some important issues to be considered in designing and developing history tools. These issues need to be resolved before history tools can effectively support collaboration:

**What content should a history item contain?** Researchers have suggested and examined a variety of probable contents such as user interactions (or commands) [3][15][17], software states [5], a combination of commands/states [13] and states plus users' augmented information [5]. User information (which user was responsible for each action) may also be needed for collaborative objectives such as coordination. However, it is still unclear exactly which content is needed to support different collaborative tasks (e.g. training vs. shared analysis) and collaboration styles (e.g. loosely coupled to closely coupled work).

**What data structures should be used?** Histories can rapidly grow in size and need appropriate data structures and scaling tools [5].

**How should a history be represented?** Selecting the form that best suits users depends partly on form of the content [5]. For instance, a repository of executed commands can be represented as list of textual commands, a history consisting of a number of graphs can be represented as a comic strip [7], and for hybrid content of commands/states, text and graphics can be used jointly [13]. The ideal representation will also depend on the task, display and input hardware, and setting. For instance, a history that can support distributed awareness during joint analysis may look very different from a history that can support co-located training.

**How can we support fluid interaction with histories?** Especially for co-located collaboration, where interactive touch surfaces may be used, new mechanisms may need to be developed for interaction with histories.

**What are underlying data challenges?** It is important to pay attention to the underlying data. Volatile or streaming data add additional challenges for history tools [5]. Moreover, we might need to closely survey different data (e.g. business data and scientific data) to understand their effect on content and representation of history repositories and functionalities they should provide to facilitate collaborative work.

**What features of a history tool are needed to support different collaborative activities?** Can a single architecture support all of the different time/place settings and user objectives?

## 6 Conclusion

In this paper we compiled a list of operations and objectives related to history tools, and described the importance of such tools for the process of collaborative data visualization and analysis. History tools to support collaborative work are not merely instruments for correcting errors but also provide users with some vital functionality necessary for coordination, training, sharing information, and many other objectives. Designers of software for collaborative work need to take into consideration operations that a history tool must support and objectives that users are most likely to desire. Open research questions include what content to include in histories, how to store histories efficiently, and how histories should be best represented to support different collaborative tasks and situations.

## Acknowledgements

This research was supported by grants from SAP Business Objects and the Natural Sciences and Engineering Council of Canada (NSERC). We also thank Kellogg Booth and our colleagues at SAP Business Objects for their many thoughts and suggestions.

## References

[1] S.P. Callahan, J. Freire, E. Santos, C.E. Scheidegger, C.T. Silva, and H.T. Vo, “VisTrails: visualization meets data management,” *Proc. ACM SIGMOD international conference on Management of data*, Chicago, IL, USA: ACM, pp. 745-747, 2006.

[2] S.P. Callahan, J. Freire, E. Santos, C.E. Scheidegger, C.T. Silva, and H.T. Vo, “Managing the Evolution of Dataflows with VisTrails,” *Proc. 22nd International Conference on Data Engineering Workshops, IEEE Computer Society*, p. 71, 2006.

[3] W.K. Edwards, T. Igarashi, A. LaMarca, and E.D. Mynatt, “A temporal model for multi-level undo and redo,” *Proc. 13th annual ACM symposium on User interface software and technology*, San Diego, California, United States: ACM, pp. 31-40, 2000.

[4] W.K. Edwards and E.D. Mynatt, “Timewarp: techniques for autonomous collaboration,” Proc. SIGCHI conference on Human factors in computing systems, Atlanta, Georgia, United States: ACM, pp. 218-225, 1997.

[5] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala, “Graphical Histories for Visualization: Supporting Analysis, Communication, and Evaluation,” IEEE Transactions on Visualization and Computer Graphics, vol. 14, pp. 1189-1196, 2008.

[6] J. Heer, F.B. Viegas, and M. Wattenberg, “Voyagers and voyeurs: Supporting asynchronous collaborative visualization,” Commun. ACM, vol. 52, pp. 87-97, 2009.

[7] R.R. Hightower, L.T. Ring, J.I. Helfman, B.B. Bederson, and J.D. Hollan, “Graphical multiscale Web histories: a study of padprints,” Proc. 9th ACM conference on Hypertext and hypermedia: links, objects, time and space---structure in hypermedia systems, Pittsburgh, Pennsylvania, United States: ACM, pp. 58-65, 1998.

[8] J. Heer, F. Ham, S. Carpendale, C. Weaver, and P. Isenberg, “Creation and Collaboration: Engaging New Audiences for Information Visualization,” Information Visualization: Human-Centered Issues and Perspectives, Springer-Verlag, pp. 92-133, 2008.

[9] P. Isenberg and Danyel Fisher, “Collaborative Brushing and Linking for Co-located Visual Analytics of Document Collections,” Proc. Eurographics / IEEE-VGTC Symposium on Visualization, pp. 1031-1038, 2009.

[10] P. Isenberg, A. Tang, and S. Carpendale, “An exploratory study of visual information analysis,” Proc. 26th annual SIGCHI conference on Human factors in computing systems, Florence, Italy: ACM, pp. 1217-1226, 2008.

[11] G. Johnson, “Collaborative Visualization 101,” ACM SIGGRAPH - Computer Graphics, vol. 32, May, pp. 8-11. 1998.

[12] A. Komlodi and W.G. Lutters, “Collaborative use of individual search histories,” Interact. Comput., vol. 20, pp. 184-198, 2008.

[13] C. Meng, M. Yasue, A. Imamiya, and X. Mao, “Visualizing Histories for Selective Undo and Redo,” Proc. 3rd Asian Pacific Computer and Human Interaction, IEEE Computer Society, pp. 459, 1998.

[14] G. Mark and A. Kobsa, “The effects of collaboration and system transparency on CIVE usage: an empirical study and model,” Teleoper. Virtual Environ, vol. 14, pp.60-80, 2005.

[15] B.A. Myers and D.S. Kosbie, “Reusable hierarchical command objects,” Proc. SIGCHI conference on Human factors in computing systems: common ground, Vancouver, British Columbia, Canada: ACM, pp. 260-267, 1996.

[16] T. Oinn, M. Addis, J. Ferris, D. Marvin, M. Senger, M. Greenwood, T. Carver, K. Glover, M.R. Pocock, A. Wipat, and P. Li, “Taverna: a tool for the composition and enactment of bioinformatics workflows,” Bioinformatics, vol. 20, pp. 3045-3054, 2004.

[17] 12-J.S. Vitter, “US&R: A New Framework for Redoing,” IEEE Software., vol. 1, pp. 39-52, 1984.

# Designing a PCA-based Collaborative Visual Analytics System

Dong Hyun Jeong, William Ribarsky, and Remco Chang

**Abstract**—In visual analytics, collaboration is viewed as a knowledge sharing process that helps people perform analytical reasoning tasks effectively. In this paper, we present a collaborative visual analytics tool, iPCA-CE, that supports interactive data analysis using principal component analysis (PCA) on a tabletop display. We define three data analysis scenarios that are addressed when designing the collaborative data analysis system. With the system, users are able to collaboratively analyze data, share ideas or knowledge, and divide their work-load.

**Index Terms**—Collaborative data analysis, Touch-table, Multi-touch interaction.

## 1 Introduction

In knowledge management literature, socialization is defined as a process in which people communicate with each other in order to share their ideas or personalized (tacit) knowledge [11, 18]. In visualization, this is achieved through the collaboration process of sharing knowledge, learning, and building consensus through the use of computers [22]. Several researchers have studied users' behavior through collaborative environments in order to better understand this knowledge-sharing process. Mark and Kobsa [17] performed an empirical study to understand the differences between group and individual behavior within collaborative information visualization environments. They found that a group solves the given questions more accurately and spends less time doing so. However, it is still unknown what features should be supported within a collaborative data analysis system on a touch-table in order to reliably gain these benefits.

Analyzing data is a complicated task. If people can combine their efforts in an analytical task, they might have a better chance of solving complex problems or finding obscured information. In this paper, we focus on designing a collaborative visual analytics environment to support interactive data analysis on a touch-table. Since previous research shows that with a more user-friendly collaborative visualization system, people find results more easily and accurately [17], we choose our existing visual analytics system (called iPCA - interactive principal component analysis) and extend it to work on a multi-touch tabletop display. We named the extended version of iPCA as iPCA-CE (interactive PCA within collaborative environments). When designing the collaborative visual analytics system, we carefully consider addressing three different types of collaborative data analysis scenarios (see Section 3 for detail).

The rest of this paper consists of four sections. First we discuss related research in collaborative visualization environments. Then we explain our system's interface design and multi-touch interactions. In section 3, we introduce three collaborative data analysis scenarios supported by our system, and conclude with discussion and future work.

## 2 Previous Work

In the past, many notable studies have been done in collaborative visualization. There are roughly three main research trends: building collaborative visualization environments, sharing knowledge through web-based collaborative workspaces, and interactively sharing tacit knowledge with people on a touch surface. In this section, we introduce some of the existing literature.

Collaboration has been described as the process of sharing tacit knowledge between people [18]. Mark and Kobsa [17] defined collaborative information visualization behavior as a social process as well as a cognitive process because it involves both interpreting visualization and coordinating complex social activities. Although the knowledge sharing process and the cognitive process are both broadly regarded as important research topics [22], limited study has been done in visualization. However, building collaborative visualization environments has a long history [5, 13]. Coleman et al. [5] provided four general reasons why collaborative visualization is compelling. (1) Experts' knowledge can be available any time and at any place. (2) The expertise can be transferred to others, improving the local level of knowledge. (3) Based on the supported accessibility, visualization products can be reviewed and modified as they are produced, reducing turn-around time. (4) Remote accessibility also reduces the need to relocate the expertise physically. Johnson [13] defined collaborative visualization as a subset of computer-supported cooperative work (CSCW) in which control over parameters or products of the scientific visualization process is shared.

More recently, Grimstead et al. [8] reviewed 42 collaborative visualization systems in terms of five attributes: number of simultaneous users, user access control, communication architecture, type of transmitted data, and user synchronization. They found that the synchronous system has the benefits of bringing groups of individuals together over a distance, bridging the knowledge gaps among them, and building their knowledge structure concurrently. But, they noticed that the synchronous system is still limited in that people have to be in front of computer machines at the same time. However, in an asynchronous collaborative visualization system, collaboration occurs at different times. If people are in different time zones and different places, an asynchronous collaborative system might be beneficial [16]. Once important knowledge is found, it can be shared with others asynchronously at their own convenience. However, it is still unclear how collaborative visualization should be designed. Because of this, Heer and Agrawala [9] provide design considerations for asynchronous collaboration in visual analytics environments. Ma [15] noted that sharing visualization resources will provide the eventual support for a collaborative workspace. He discussed existing web-based collaborative workspaces in terms of sharing high-performance visualization facilities, visualizations, and findings. Burkhard proposed a collaboration process of transferring knowledge between at least two persons or group of persons [2].

Although much research has been done in collaborative visualization, there has been less work in collaboration on touch surfaces. Isenberg and Fisher [10] designed a system (called Cambiera) to support collaborative search through large text document collections on a touch surface. They considered collaborative activities to involve not just searching through documents, but also building individual's findings and maintaining awareness of another person's work. North et al. [19] studied how users approach a multi-touch interface and what

- *Dong Hyun Jeong is with Charlotte Visualization Center at UNC Charlotte, E-mail: dhjeong@uncc.edu.*
- *William Ribarsky is with Charlotte Visualization Center at UNC Charlotte, E-mail: ribarsky@uncc.edu.*
- *Remco Chang is with Charlotte Visualization Center at UNC Charlotte, E-mail: rchang@uncc.edu*

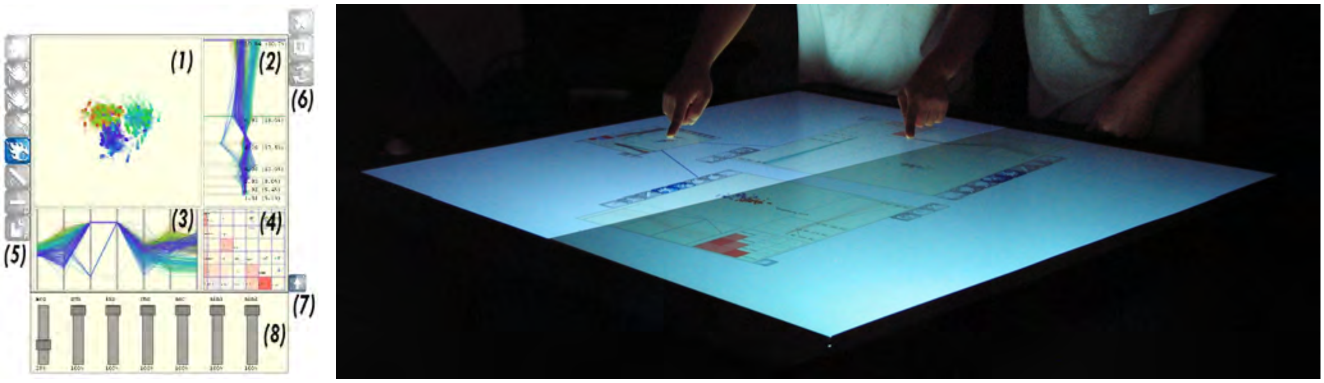


(a) (b)

Fig. 1. The system overview (a) showing four views ($1 \sim 4$), several selectable buttons ($5 \sim 7$), and a set of sliderbars (8) with E.Coli dataset ($336 \times 7$ matrix). The system supports changing the scale and the location as well as manipulating the projected data item(s). With using the system, multiple people can collaborate each other on a multi-touch tabletop display interactively (b).

types of gestures they are willing to use. In the study, they performed object manipulation tasks on a physical table, a multi-touch table, and a desktop computer. From the study, they found that people completed the tasks significantly faster with multi-touch interactions on a multi-touch table than with mouse interactions on a desktop computer. Furthermore, they found that subjects are significantly faster on a physical surface than on a touch surface.

However, to the best of our knowledge on collaborative visualization, what we should support when designing a collaborative data analysis system has not been broadly studied. Developing this design philosophy requires understanding how people act on a touch table, especially when analyzing data. In the following sections, we will provide a detailed explanation about how our system is designed and what features are supported.

## 3 Collaborative Visual Analytics Application

Understanding how people act when analyzing data is an important research topic in visual analytics, but it is also extremely challenging [22]. In this paper, we focus our research to the understanding of analytical behavior to be strictly within the context of a collaborative environment. Our work begins with designing a useful collaborative visual analytics application with which people can easily share their ideas and knowledge. Our design philosophy is not to develop a new visual analytics application, but instead we extend an already known and useful visual analytics application to work on a touch table. Specifically, we choose our existing visual analytics application (iPCA) because studies have shown that user-friendly visualizations in a collaborative environment enable users to find results more accurately [17]. In our previous study [12], we found that iPCA is both easy to use and effective in helping users learn about PCA and the datasets they are using.

Data analysis is often considered as a stand-alone analytical task. However, as previous research has shown, analysis of (empirical) data in collaborative environments is important and should be considered while developing visualization applications [4, 7]. While collaborative analytics can occur in different interaction modalities, we focus specifically on collaboration on a multi-touch table based on existing work that demonstrated potential increase in analysis performance [10].

### 3.1 System Design

iPCA is designed to help the user understand the complex black box operation of Principal Component Analysis [14] and interactively analyze data [12]. We extend this application to support collaborative data analysis on a touch table.

Figure 1(a) shows the system overview, which includes four views, touchable buttons, and a set of dimension sliders. The overall interface is developed with OpenGL. It supports multiple-touch interactions on a horizontal display. The multi-touch display system was designed at the Renaissance Computing Institute (RENCI) [1]. It provides a 62" diagonal work surface (42 x 46), in which two HD resolution projection displays create images on the surface to support multiple people working together. Figure 1(b) represents the overall workspace, in which two people are collaborating on a touch table.

Like the original iPCA, our extended application (iPCA-CE) consists of four views: Projection view (Figure 1(a-1), Eigenvector view (Figure 1(a-2), Data view (Figure 1(a-3), and Correlation view (Figure 1(a-4)). In the Projection view, all data items are projected based on the first and second principal components by default. The Eigenvector view displays the calculated eigenvectors and eigenvalues in a vertically projected parallel coordinate. The distances between the eigenvectors in the parallel coordinate view vary based on their eigenvalues, separating the eigenvectors based on their mathematical weights. The Data view shows the original data points in a parallel coordinate. The Correlation view represents Pearson-correlation coefficients and relationships between variables as a matrix of scatter plots and values. All views are closely connected, so that an action in one view can affect the other views. If the user interactively changes the elements in one view, its corresponding results are updated in other views (*brushing & linking*). This interactivity thus allows the user to infer relationships between the coordinated spaces (see [12] for detail).

There are a total of 12 touchable buttons designed: 8 buttons are for interacting with represented data items (Figure 1(a-5)), 3 buttons are for controlling the application (Figure 1(a-6)), and the last button (Figure 1(a-7)) is for making the sliderbars appear and disappear. Table 1 represents the touchable buttons and their meanings.

Table 1. Touchable buttons and their meanings

| Button | Meaning | Button | Meaning |
|---|---|---|---|
| | Go back to the initial state | | Delete the selected item(s) |
| | Individual item selection | | Partition the selected item(s) into a new workspace |
| | Range item(s) selection | | Close the application |
| | Manipulation | | Create a new application |
| | Trail enable – on/ off | | Rotate the application |
| | Cancel the selected item(s) | | Make the sliderbar panel appear / disappear |

The system supports basic multi-touch operations such as zooming, panning, and rotation. The zooming operation is activated by making two finger touches closer (zoom-in) and farther apart (zoom-out). The panning operation is initiated by dragging a finger on the surface. However, the rotation only works when the rotation option (a touchable button) is enabled. We adopt this passive operation because if the user accidently changes the angle between two touches during analysis, the rotation operation is activated unintentionally, and sometimes distracts people from concentrating on analyzing the data.

Additionally, the system provides several data operations such as individual item selection, range item(s) selection, deletion, and manipulation. Both the individual item selection and the range item(s) selection operations are allowed in all four views. In Data View and Eigenvector View, where the visualizations are parallel coordinates, selection means clicking on a single line or brushing a range of items. In Projection View and Correlation View, the user can either click on a single dot or draw an enclosed space upon which all data items within the space will be selected. In analysis using PCA, a common task is for the user to remove outliers. The deletion operation is to remove the selected data item(s) from the PCA calculation. The manipulation is the operation, which allows the user to see the relationship between principal component(s) and data dimensions.

### 3.2 Multi-touch Interactions

As shown in Figure 1(a), the E.coli dataset has 7 dimensional attributes. But it is not linearly separable by a PCA calculation since PCA assumes that the input data are always linear. Because of this, weighted principal component analysis (WPCA) is often considered, which allows different weights on different variables as $s_1, s_2, ..., s_n$ [14]. This approach assumes that data are not always linearly increasing or decreasing, and there may be reason to allow different observations to have different weights. To provide the ability to analyze the data non-linearly, iPCA has a set of dimension sliderbars, which allow the user to change the dimension contributions of each dimension. However, with a mouse-based interface, the user has to try all possible combinations of dimension contribution changes with a series of single mouse inputs to fully understand and analyze the data. iPCA-CE gives the user the ability to change several dimensions at once on a multi-touch table, thus permitting much more effective exploration of the high dimensional space and how the dimensions correlate. Figure 2(a) shows an example in which the user changes dimension contributions by moving the sliderbars with two finger touches.

iPCA-CE also allows the user to alter the values of data items. For instance, if the user drags a data item in the Projection View towards the positive direction along the $x$-axis (increasing the data point's value in the first principle component), the user should be able to immediately observe in the Data View how that change affects the values of that data item in the original data space, thus shedding light on the relationship between the first principle component and all dimensions in the original data space. Figure 2(b) shows the user manipulating the selected data item in the Data view with two finger touches.

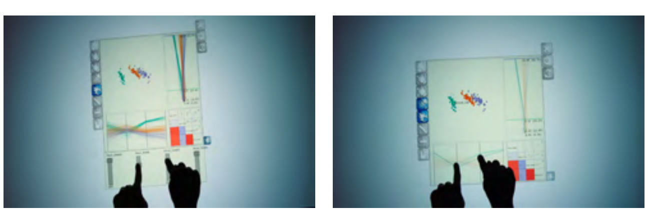

(a) (b)

Fig. 2. Multi-touch interactions. (a) The user changes the dimension contributions using sliderbars and (b) the user directly modifies the values of a data item in the Data view.

## 4 Collaborative Data Analysis

A collaboration process can occur through the use of collaborative visual environments. However, the most natural method for sharing tacit knowledge is still direct communication between users. In either case, the users are actively sharing their discoveries and tacit knowledge and incorporating each other's domain expertise into their own. However, understanding and addressing analytical procedures are important when designing a useful collaborative visual analytics application. In general, collaborative environments on a touch table support either tightly coupled collaboration (having a shared workspace and working together) or loosely coupled collaboration (having independent workspaces and working alone for long periods of time) [21]. In our collaborative visual analytics application, we considered addressing three types of analytical scenarios: 1) people are collaborating with others by looking at the same results (tightly coupled collaboration); 2) people are analyzing the same dataset with their own individual workspaces (loosely coupled collaboration); and 3) people are working with a partitioned dataset within their own workspaces (tightly and loosely coupled collaboration). The third scenario, however, is especially important because it supports both tightly and loosely coupled collaboration (see Section 3.3 for detail).

### 4.1 Looking at the Same Results

In visual analytics, people are often working together by looking at the same results displayed on a screen, which is a common analytical procedure when collaborating with others. Most visual analytics applications support this analytical procedure, as it works in any types of display system. However, on a touch table, existing visual analytics applications allow multiple people to work at the same time (tightly coupled collaboration). Butkiewicz et al. [3] designed a geospatial analysis tool running on a touch table, with which people can interactively create multiple probes based on their regions of interest. In such an environment, people can easily share ideas, findings, and their expertise with others by looking at the same results. This is also somewhat related to a learning system, in which an expert explains interesting results or his personalized knowledge to novice users so they can come up with solutions and analyze the data effectively on their own.

Figure 3(a) shows two users working together by looking at and interacting with the same representation displayed on a touch surface. In this example, the user (left) is trying to show the effectiveness of data value changes to the other user (right). Within this environment, users can directly communicate with each other focusing on the same visual representation and results.

### 4.2 Working with the Same Dataset

In collaborative visualization applications, a common analytical procedure is to work with the same dataset synchronously and asynchronously. Because of this, most existing collaborative visualization applications support both synchronous and asynchronous knowledge sharing. However, in our collaborative visual analytics application, we only consider synchronous collaboration.

On a multi-touch table, people can analyze the dataset by looking at different representations. Once a person finds an interesting result, he can directly communicate it by passing or showing the result to a colleague. This is somewhat related to the analytical procedure described in Section 3.1. However, having individual workspaces may increase the overall performance of finding hidden information and analyzing the data (loosely coupled collaboration). Figure 3(b) shows an example in which people collaboratively analyze the public Iris dataset with their own workspaces. In Figure 3(b), the user (left) analyzes the data by changing the dimension contribution of the first (Sepal length) and second (Sepal width) variables, and the other user (right) manipulates the values of the selected data item in the parallel coordinates within the Data view to understand how the selected data item(s) are placed in a certain cluster.

### 4.3 Working with the Partitioned Datasets

In data analysis, data partitioning is an important pre-processing operation. For instance, a Bayesian phylogenetic analysis tool (MrBayes 3 [20]) partitions data according to the data type by default, and then analyzes the partitioned datasets separately. This is because most real-world datasets do not exist in the form of a combined dataset. Also,

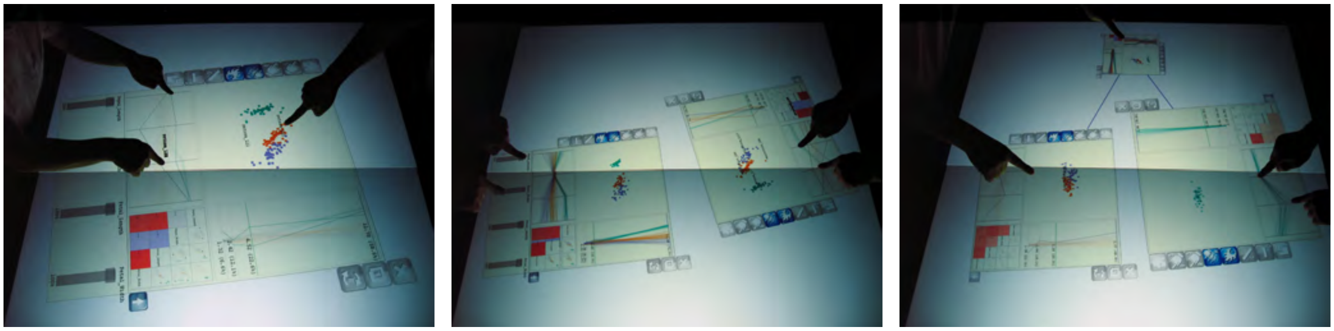

(a) (b) (c)

Fig. 3. The pictures show people performing multiple collaborative data analysis scenarios in iPCA-CE system with the Iris dataset ($140 \times 4$ matrix). (a) People are working together by looking at the same tool and results, (b) working with the same dataset, but in different workspaces, and (c) working with partitioned datasets in their own workspaces. The lines between workspaces in (c) indicate the independent workspaces (the partitioned datasets (left and right)) from the shared workspace (the original dataset (top middle)).

people often tend to focus on analyzing a specific dataset based on their interests or personalized (tacit) knowledge. In financial fraud analysis, analysts tend to investigate specific financial datasets (e.g. the transactions between two specific countries) based on their experience [6].

In iPCA-CE, users are able to interactively partition the dataset in order to collaborate with others. Once the dataset is partitioned, the partitioned dataset creates a (blue) connected line to its original dataset. In this analytical scenario, the system supports both loosely coupled collaboration and tightly coupled collaboration. The system is designed to support creating multiple independent workspaces from a shared workspace. Figure 3(c) shows a shared work space and two independent workspaces. The unpartitioned dataset is projected in the shared workspace and the partitioned datasets are displayed in the independent workspaces.

## 5 Conclusion and Future work

Since data analysis is a complex analytical task, many useful visual analytics applications are designed to assist users analyzing data effectively. However, limited research has been done on understanding how to support data analysis on a touch table. In this paper, we described three important analytical scenarios that should be supported when designing a collaborative data analysis application on a touch table. We also designed a collaborative data analysis application (iPCA-CE) based on these analytical scenarios.

Since how people share ideas or personalized (tacit) knowledge on a touch-table when solving complex analytical tasks is still not known, our future work includes understanding the human knowledge sharing process on a touch table.


## Acknowledgements

The authors wish to thank Caroline Ziemkiewicz for her help on revising the drafts.


## References


[1] Renaissance computing institute, 2009. http://www.renci.org.

[2] R. A. Burkhard. Learning from architects: The difference between knowledge visualization and information visualization. In *IV '04: Proceedings of the Information Visualisation, Eighth International Conference*, pages 519–524, Los Alamitos, CA, USA, 2004. IEEE Computer Society.

[3] T. Butkiewicz, D. H. Jeong, W. Ribarsky, and R. Chang. Hierarchical multi-touch selection techniques for collaborative geospatial analysis. In *SPIE Defense, Security, and Sensing 2009*, 2009.

[4] I. Carlbom, W. M. Hsu, G. Klinker, R. Szeliski, K. Waters, M. Doyle, J. Gettys, K. M. Harris, T. M. Levergood, R. Palmer, L. Palmer, M. Picart, D. Terzopoulos, D. Tonnesen, M. Vannier, and G. Wallace. Modeling and analysis of empirical data in collaborative environments. *Commun. ACM*, 35(6):74–84, 1992.

[5] J. Coleman, A. Goettsch, A. Savchenko, H. Kollmann, W. Kui, E. Klement, and P. Bono. Teleinvivotm: towards collaborative volume visualization environments. *IEEE Comput. Graph. Appl.*, 20(6):801–811, 1996.

[6] W. Dou, D. H. Jeong, F. Stukes, W. Ribarsky, H. R. Lipford, and R. Chang. Recovering reasoning processes from user interactions. *IEEE Comput. Graph. Appl.*, 29(3):52–61, 2009.

[7] J. Ekanayake, S. Pallickara, and G. Fox. A collaborative framework for scientific data analysis and visualization. In *Collaborative Technologies and Systems, 2008. CTS 2008. International Symposium on*, pages 339–346, May 2008.

[8] I. J. Grimstead, D. W. Walker, and N. J. Avis. Collaborative visualization: A review and taxonomy. *Distributed Simulation and Real-Time Applications, IEEE International Symposium on*, 0:61–69, 2005.

[9] J. Heer and M. Agrawala. Design considerations for collaborative visual analytics. *Information Visualization*, 7(1):49–62, 2008.

[10] P. Isenberg and D. Fisher. Collaborative brushing and linking for co-located visual analytics of document collections. *Computer Graphics Forum (EuroVis 2009)*, 28(3):1031–1038, 2009.

[11] D. H. Jeong, R. Chang, and W. Ribarsky. An alternative definition and model for knowledge visualization. IEEE Visualization 2008 Workshop on Knowledge Assisted Visualization, 2008.

[12] D. H. Jeong, C. Ziemkiewicz, B. Fisher, W. Ribarsky, and R. Chang. ipca: An interactive system for pca-based visual analytics. *Computer Graphics Forum (EuroVis 2009)*, 28(3):767–774, 2009.

[13] G. Johnson. Collaborative visualization 101. *ACM SIGGRAPH - Computer Graphics*, 32(2):8–11, 1998.

[14] I. Jolliffe. *Principal Component Analysis*. Springer, second edition, 2002.

[15] K.-L. Ma. Creating a collaborative space to share data, visualization, and knowledge. *SIGGRAPH Comput. Graph.*, 41(4):1–4, 2007.

[16] F. T. Marchese and N. Brajkovska. Fostering asynchronous collaborative visualization. In *IV '07: Proceedings of the 11th International Conference Information Visualization*, pages 185–190, Washington, DC, USA, 2007. IEEE Computer Society.

[17] G. Mark and A. Kobsa. The effects of collaboration and system transparency on cive usage: An empirical study and model. *Presence*, 14:2005, 2005.

[18] I. Nonaka and H. Takeuchi. *The Knowledge Creating Company*. Oxford University Press, 1995.

[19] C. North, T. Dwyer, B. Lee, D. Fisher, P. Isenberg, G. Robertson, and K. Quinn. Understanding multi-touch manipulation for surface computing. In *In: Proceedings of 12th IFIP TC13 Conference in Human-Computer Interaction (Interact 2009)*, pages 236–249. Springer, 2009.

[20] F. Ronquist and J. P. Huelsenbeck. Mrbayes 3: Bayesian phylogenetic inference under mixed models. *Bioinformatics*, 19(12):1572–1574, 2003.

[21] A. Tang, M. Tory, B. A. Po, P. Neumann, and M. S. T. Carpendale. Collaborative coupling over tabletop displays. In *CHI*, pages 1181–1190, 2006.

[22] J. J. Thomas and K. A. Cook. A visual analytics agenda. *IEEE Comput. Graph. Appl.*, 26(1):10–13, 2006.

# Information Visualization on Interactive Tabletops in Work vs. Public Settings

Petra Isenberg Uta Hinrichs Mark Hancock Matthew Tobiasz Sheelagh Carpendale


Fig. 1. Groups of people engaging in information exploration and analysis in work (left) and public settings (right).

**Abstract**— Digital tabletop displays and other large interactive displays have recently become more affordable and commonplace. Due to their benefits for supporting collaborative work—when compared to current desktop-based setups—they will likely be integrated in tomorrow's work and learning environments. In these environments the exploration of information is a common task. We describe design considerations that focus on digital tabletop collaborative visualization environments. We focus on two types of interfaces: those for information exploration and data analysis in the context of workplaces, and those for more casual information exploration in public settings such as museums. We contrast design considerations for both environments and outline differences and commonalities between them.

**Index Terms**—Digital tabletop displays, information visualization, information exploration, collaboration

## 1 Introduction

Groups of people often form decisions or gain knowledge about a topic by coming together in physical environments to discuss, learn, interpret, or understand information. These groups often make use of physical tables to view, share, and store visual information. These types of group tasks or goals commonly occur in meeting rooms, research labs, classrooms, museums, and other public settings. Digital tabletop displays can augment information exploration and analysis in these physical spaces; they can support the collaborative and interactive exploration of digital information beyond the possibilities that printed paper, projected slide shows, or non-interactive media such as posters, black-boards, or bulletin boards can offer.

In the remainder of the paper, we discuss the role of tabletop displays for collaborative information exploration or analysis in two specific contexts: work environments and public spaces. In work environments, such as meeting rooms or research labs, teams of analysts can be characterized by a vast amount of domain-specific knowledge, while in public spaces, such as museums or art galleries, people's level of knowledge on a certain topic varies and is difficult to predict or expect. Nonetheless, both contexts invite the possibility of gaining insight through the process of exploring and analyzing information. By looking at existing examples of information visualization in both contexts, we discuss their commonalities and differences in order to arrive at practical considerations for designing tabletop interfaces to support information exploration in each context.

*Petra Isenberg, Uta Hinrichs, Mark Hancock, Matthew Tobiasz, and Sheelagh Carpendale are with the University of Calgary. E-mail: {pneumann|uhinrich|mshancoc|tobiasz|sheelagh}@ucalgary.ca.*

## 2 Tabletop Displays in the Workplace

In many areas, domain experts perform data analysis on a daily basis. For example, molecular biologists frequently analyse huge datasets from lab experiments, business analysts look at trends in financial data, or historians explore large document databases to bring historical events into context. With the rapid growth of the complexity and size of datasets in many work scenarios the need to support multiple people simultaneously viewing and manipulating data is increasing. This growth means that domain experts from different disciplines and with different skill sets are often required to collaborate, to make informed decisions about a dataset, and to improve the quality of an analysis result. Datasets on which decisions and discoveries are based may not only be too large to handle by a single analyst but may also be susceptible to a variety of interpretations, in which case experts may need to discuss and negotiate their interpretations of the data.

Digital tables offer great potential to support this type of work. In the near future digital tabletops may be installed in offices, meeting rooms, or research labs where today's domain experts already meet to discuss, interpret, and analyse data. One of the great advantages of tabletop displays in the workplace is their ability to support such collaborative work. Analysis systems that use digital tables can enable in-situ discussion, exploration, and interpretation—in close contact with the data and its visualization. Team members can work independently and together while being able to spontaneously react to findings in the data and to resolve data conflicts as a group. The design of interfaces, visualizations, and interaction techniques for visual analysis by teams of domain experts around tabletops is an active research area. At the time of this writing, examples of systems for exploring information at a tabletop display in the workplace have been limited mostly to research prototypes. As the cost of such systems goes down, we expect to see more commercial examples arise. Nonetheless, the research prototypes demonstrate the viability of tabletop systems for improving people's ability to collaboratively explore information.

The authors have previous experience building system prototypes for information analysis in the workplace. Lark [12] was designed to help multiple analysts coordinate their individual and joint analysis activities on a tabletop (Figure 2-left). Cambiera [6] supported several collaborators foraging for information in large text document collections, highlighting overlap in found and accessed documents (Figure 2-middle). In another project [5] guidelines for collaborative information visualization were tested as a case study in a tabletop tree comparison system (Figure 2-right).

Fig. 2. Collaborative infovis systems built by the authors.

## 3 Tabletops in Public Spaces

Tabletop displays have started to become more common outside of research labs and work environments. For instance, we can find them in museums and art galleries where they are used to convey information to people in an interactive and potentially engaging way. The use of horizontal digital surfaces to present interactive data visualizations has several advantages, especially for more casual public settings where people gather in their spare time. Information visualizations presented on digital tabletops can turn abstract data into interactive exhibits that evoke attention and curiosity and entice visitors to take a closer look. The physical setup of tabletop displays enables visitors to approach the presented information from all sides; several groups or individuals can collaboratively explore, share, and discuss the data visualization. The ultimate goals of large horizontal information displays in public spaces are to attract people's attention, draw them closer to the installation, and promote lightweight information exploration that leads to serendipitous discoveries, some reflection on the presented data and/or active discussion of this data with peers.

We have previously presented and exhibited information visualization systems in public spaces. memory [en]code (Figure 3-left) is a tabletop system that visualizes the dynamics of human memories in an interactive way [9]. Visitors are invited to type their own thoughts or memories into the system. The participatory aspect of memory [en]code positively influenced people's engagement with the installation. The fact that all information was created by other visitors and the ability to leave personal traces within the system added a personal touch to the installation. EMDialog [3] is an interactive information installation that that was developed to enhance an art exhibition showing paintings from the artist Emily Carr. The installation presents two interlinked information visualizations that invite museum visitors to explore the extensive discourse about Emily Carr along temporal and contextual dimensions (Figure 3-right).

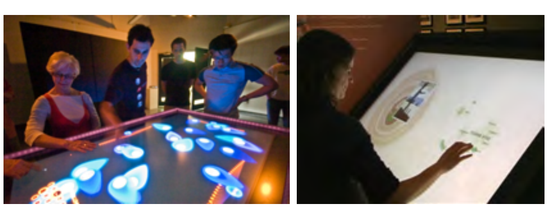

Fig. 3. Infovis systems by the authors exhibited in public spaces.

## 4 Designing for Work vs. Public Spaces

When designing visualization systems for collaborative information exploration, we are faced with a number of challenges in common with other tabletop work: the need to support awareness and common ground formation, perceptual problems, as well as collaborative interaction issues. However, several challenges also arise due to the nature of interaction with information visualizations. In this section, we discuss these challenges and point out the differences that need to be considered when designing for workplace and public settings.

### 4.1 Contextual Challenges

One of the main differences to consider when designing tabletop applications for workplace or public settings is the context in which the information is being accessed. While the context for workplace systems often goes hand-in-hand with well-defined tasks and goal-oriented analysis, the context for public settings can vary dramatically. We discuss design challenges for both situations next.

**Work Environments:** Domain Experts typically perform information exploration and analysis in small groups whose members are already acquainted. There are also typically well defined analysis goals. These goals must be supported by the tabletop software and, hence, the development of specific software may be necessary when datasets and tasks change. In contrast to tabletop systems designed for public spaces, the expectations about interaction techniques and data representations differ in the workplace. The questions in work scenarios are typically quite complex and difficult. Also, the data analysis results might be vital to make important (sometimes time-critical) decisions with many variables to consider. Information visualization interfaces, therefore, typically have a large number of parameters to manipulate. Work teams are often prepared to invest time in learning, and tabletop interfaces designed for these settings can, therefore, often include new interactions and visual designs if they might improve the efficiency and quality of collaborative information exploration. Work teams also often may spend considerable time using an interface, making the effort to learn new techniques worthwhile.

Several information exploration sessions are often necessary to come to a common understanding of a particular dataset in the workplace. Tabletop software for collaborative information exploration should, therefore, support capturing of interaction histories with the information in order to allow groups to interrupt their analysis and continue at a later stage. At the same time, it is often the case that individual group members may drop in and out of a running collaborative information exploration session. For these group members it may also be useful to implement history and summarization mechanisms to show what has been missed. First approaches are incorporated in Lark and Cambiera (see above) [6, 12].

**Public Spaces:** The audience gathering around a tabletop in a public space can be highly diverse. Visitors of museums and art galleries, for instance, not only differ in age but also in social and cultural background, knowledge, and interests [10]. Furthermore, people often visit exhibitions without clearly defined questions or goals in mind but explore them serendipitously based on spontaneous interest [10]. Interaction with exhibits tends to be brief and usually only occurs once per visitor. This means that tabletop interfaces for information exploration in public settings need to be designed differently from workplace systems. Interaction techniques need to be designed with a walk-up-and-use scenario in mind. Visitors of public spaces are not likely to read elaborate instructions on how to interact with the system but will try to figure out exploration techniques and capabilities of the visualization on the fly. Interaction with the tabletop system therefore should be accompanied by direct feedback mechanisms that encourage further interaction or lead visitors to try different interactive mechanisms. The diversity of people visiting public spaces is often reflected in a variety of interaction times and exploration styles. Some people will only interact with the tabletop installation for a few moments, while others will explore information in detail for a longer amount of time. Therefore, the design of information visualizations on public tabletop systems should reward both short- and long-term information exploration. Furthermore, some people prefer guided exploration, while others like to follow their personal interests using more open exploration techniques.

### 4.2 Technological Challenges

In both workplaces and public spaces, hardware challenges exist for the setup of information exploration environments. These challenges relate to size and resolution of the table but also its spatial placement, robustness, and form factor.

**Workplace Environments:** Domain experts often have to do fine-grained analysis of large and detailed datasets. For the visualization of this data, the size and resolution of a tabletop is critical. As datasets increase in size, it becomes more and more difficult to display them in their entirety. Large and high-resolution tables allow more data to be displayed and support several people working together—either with multiple copies of a data representation or with different parts of a shared visualization. However, detailed and large datasets may require the rendering and reading of small textual labels and other data items. With growing resolution, the displayed information items can become physically smaller resulting in selection difficulties. Using fingers or pens may no longer be sufficient to select small data items and alternative selection techniques may have to be used or designed. Also, when large datasets have to be rendered on high-resolution tabletop screens, combined with several simultaneous inputs, response time may become more important. It is necessary to develop algorithms that can support multi-person interaction on very high resolution tables. Groups of domain experts may also often meet around a digital table to perform long analysis sessions. Therefore, the form factor of the table should be such that it supports comfortable seating positions similar to current meeting spaces in conference rooms or offices.

**Public Spaces:** Similar to the workspace, public settings can benefit from the availability of large and high-resolution tabletop displays. In public settings, the size of a group wanting to access a table may be much larger than in a workplace. For example, it is not unusual for school classes to gather around a tabletop to interact with and explore information in a museum. In such situations, it is critical that the whole system remains responsive and that the software does not crash, even if 40 hands are touching the table at the same time or even issue conflicting information exploration commands. Tables for public settings also need to be robust in their physical design, be spill-proof and resistant to scratching or pushing. In contrast to domain expert information exploration sessions, one cannot expect children or large groups of adults to treat a public tabletop display with care. It is important to consider that the physical setup of the display (size, orientation, and location) can influence the group size and number of different groups of people interacting with it. Physical form factors also need to be considered with regard to physical accessibility. For instance, all visitors need to be able to see and access the display surface, including children and people in wheelchairs.

### 4.3 Perceptual Challenges

The environment suggested by a tabletop display is particularly unique to computing systems. In particular, the display has a horizontal orientation and affords multiple people standing at different sides of the table. These properties are compelling for a variety of reasons, but also introduce some unique perceptual challenges. Specifically, the assumption common to desktop computing that there will be one viewer directly in front of the display is no longer valid. For example, Wigdor *et al.* [14] performed a study that suggests that visual variables (e. g., angle, length, shape) are perceived differently on a horizontal surface than on a vertical one. In 3D, the problem is exacerbated, as the projection from 3D onto the 2D surface requires an assumption about the point of view of the (one and only) observer. Thus, a projected image may appear drastically different to observers standing at opposite sides of the table. Several systems have explored solutions to the problem of multiple points of view [1, 7] but the degree of this problem on digital tables has still been largely unexplored.

Some visual elements in both 2D and 3D are particularly sensitive to changes in orientation (e.g., text). Some studies have shown that people are still capable of reading short bits of text at non-zero orientations [13], but they are still slower, and so larger bits of text are best to read in the correct orientation. Other research suggests that the act of orienting visual elements is often used to communicate with others [8] and a variety of methods to perform this act have been introduced to tabletop display environments (see [2] for an overview). Thus, perception of visual elements that have an intrinsic orientation may play an important role in the collaboration that occurs in a tabletop display environment. These perceptual challenges exist in both workplace as well as public settings, but the types of problems that may arise vary somewhat.

**Work Environments:** Here, the perception of the visual information may be relevant for a variety of reasons. The visual variables used to represent the information may need to precisely depict a value to be judged by the observer, or it may be important to compare two (or more) visual elements. A person on one side of the table may also need to be able to trust that someone across the table can perceive a visual variable in a predictable way (i. e., that their view is not warped in some way). At present, there is little work to suggest how to design systems that address these issues. However, the current work points to the fact that the simple solution of using the same design criteria for vertical displays may not suffice for horizontal ones [14].

**Public Spaces:** In more artistic or learning environments found in public spaces, the precise value of a particular visual element may not be as important as in systems designed for domain expert analysis in the workplace. Instead, it may be more important for the designer to consider the fact that the perceptual experience of two observers standing at opposite sides of the table will differ. This difference in experience can be thought of as an additional challenge for the designer; the system can be made to either mitigate these perceptual differences, or to take advantage of them in order to create a unique experience for the observers. Nonetheless, the consideration of the orientation of the visual elements can be particularly important in a public space. Grabbing the attention of someone passing by will involve the consideration of how the display looks from both far away and from close proximity. Orientation-sensitive elements, such as text, may play an important role in drawing attention, indicating a suitable viewpoint, or to help encourage communication between multiple simultaneous observers.

### 4.4 Collaborative Challenges

Several previous studies of collaborative information exploration, both for work environments [11] as well as public spaces [3], suggest a need to support a wide range of collaboration styles. People may be interested in exploring parts of the information by themselves without interfering with other people but may, at any given time, switch from this parallel work to a phase in which they work more closely together, sharing information items, and discussing them closely. Despite these initial similarities, the information exploration goals and contextual exploration scenarios for information visualization in work environments and public spaces form different design challenges.

**Work Environments:** If one wants to support collaborative information exploration, one has to either design visual representations that support synchronous interaction or that allow for the ability to create several interactive views of the same dataset. Global changes to views and encodings of data are fairly common in single-user visualization systems and if one is interested in re-designing such an application for tabletop use, the re-design of these features for synchronous group work is critical [6].

Since, the datasets used in expert systems are often large, complex, uncertain, and subject to different interpretations, people have to pay close attention to the data they may be working with in order to keep their exploration context and intermediate findings in memory. Thus, for information exploration tasks, the physical cues naturally available in a co-located environment only provide limited support for awareness and common ground formation. Team members may still be able to see each others' hand and arm movements, gestures, and hear their incidental comments about data, but when the complexity of the information visualization requires increased concentration, these awareness cues may be missed. For example, a person may be pointing to a specific data item in a visualization and make a comment about it but another person may be too focused to pay attention to which item it is, what its context is within the dataset, or even which dataset it is from. When designing interfaces and visual representations for collaborative information exploration, we thus need to ensure that people can simultaneously concentrate on the complex data and maintain an awareness of each others' work and activities. Mechanisms may have to be put in place to support better contextual understanding for the reference of data items.

Large and complex datasets place a high cognitive load on the viewers. It is, therefore, important that collaborators can externalize some of their findings easily and, for example, annotate the data to mark a finding or to rate the reliability, trustworthiness, or certainty of a data item. This externalization is particularly important for collaborative data analysis because individuals may, on a momentary notice, switch context, work with another person, and then have to return to their previous work. Keeping an integrated exploration history together with data annotations could greatly support this type of expert information exploration.

**Public Spaces:** Museum studies have found that people often visit public exhibitions in groups. The studies conducted by Hinrichs et al. [3] and Hornecker et al. [4] confirm this finding for tabletop installations within museum settings. The physical setting of a tabletop display allows different visitor groups to approach the installation from all sides. When several people interact with a tabletop display at the same time, however, it is hard to maintain awareness of who is exploring what part of the visualization. In a public setting, this awareness is even more compromised since it is less likely for visitors who do not know each other to communicate or pay attention to each other and, hence, the possibility of interaction conflict is high. Different public tabletop systems deal with this problem in different ways. floating.numbers (http://www.artcom.de) and memory [en]code [9] both involve visualizations that consist of independent information objects; people can interact with different objects without interfering with each other. The visualization in EMDialog [3] was not designed to support several people exploring it in parallel, hence, the physical setup of the installation did not to invite parallel information exploration among unacquainted people. As a third example, information presented on the Tree of Life table is divided in four quadrants [4] to allow four different groups of people to explore it without interfering with each other. These examples show that there is a variety of ways to enable parallel independent information exploration.

Group interaction in public settings also is less focused around maximizing insights from the visualization and more about experiencing information collaboratively in a social way. When collaboratively exploring a museum exhibit, social interaction and sharing information can play an important role. Parents, for instance, often use information exhibits to explain causalities within the information to their children [4]. While in this situation often only one person is interacting at a time, the process of information exploration is still highly collaborative. Similar forms of collaboration can be observed among adults when they are still unclear of what an installation has to offer and how to interact with it. Groups also explore visualizations in parallel and, from time to time, share their insights through discussion, whereas others go through all information together.

## 5 Summary

It is likely that future technology will become even more ubiquitous in our environments and that it will come in many different form factors. Humans have considerable experience and expertise working together on physical tables, making this form factor a particularly promising one to promote. At the same time, we are collecting more diverse sets of information than ever before. Much of this information is being collected for the purpose of being explored interactively. Tabletop displays combine the benefits of a large display area for information, enough space for several people to share, and a seating or standing arrangement that allows for easy discussion and interaction among group members. Supporting collaborative information exploration will become an extremely important task for future systems in a large number of different settings.

We have discussed contextual, technological, perceptual and collaborative challenges arising when designing tabletop systems for information exploration in two different contexts: workplace settings where domain experts gather to explore and analyse often large and complex datasets, and public spaces where the design has to support a much more diverse set of people, tasks, and goals. While several issues are common in both settings, other challenges are unique to workplace environments or public spaces and need to be addressed accordingly.

## References


[1] Hancock, M. and Carpendale, S. Supporting multiple off-axis viewpoints at a tabletop display. In *Proc. Tabletop*, IEEE Computer Society, 2007, pp. 171–178.

[2] Hancock, M. S., Vernier, F. D., Wigdor, D., Carpendale, S., and Shen, C. Rotation and translation mechanisms for tabletop interaction. In *Proc. Tabletop*, IEEE Computer Society, 2006, pp. 79–86.

[3] Hinrichs, U., Schmidt, H., and Carpendale, S. EMDialog: Bringing information visualization into the museum. *T-VCG 14, 6* (2008), 1181–1188.

[4] Hornecker, E. "I don't understand it but it is cool": Visitor interactions with a multi-touch table in a museum. In *Proc. Tabletop*, IEEE Computer Society, 2008, pp. 121–128.

[5] Isenberg, P. and Carpendale, S. Interactive tree comparison for co-located collaborative information visualization. *T-VCG 13, 6* (Nov./Dec. 2007), 1232–1239.

[6] Isenberg, P. and Fisher, D. Collaborative brushing and linking for co-located collaborative visual analytics of document collections. *Comput Graph Forum 28, 3* (Nov.Dec. 2009), 1031–1038.

[7] Kitamura, Y., Nakayama, T., Nakashima, T., and Yamamoto, S. The illusionhole with polarization filters. In *Proc. VRST*, ACM Press, 2006, pp. 244–251.

[8] Kruger, R., Carpendale, S., Scott, S. D., and Tang, A. Fluid integration of rotation and translation. In *Proc. CHI*, ACM Press, 2005, pp. 601–610.

[9] Schmidt, H., Hinrichs, U., Dunning, A., and Carpendale, S. memory [en]code—Building a collective memory within a tabletop installation. In *Proc. CAe*, Eurographics Association, 2007, pp. 135–142.

[10] Screven, C. G. *Information Design*, chapter Information Design in Informal Settings: Museums and other Public Spaces, pp. 131–193. MIT Press, 2000.

[11] Tang, A., Tory, M., Po, B., Neumann, P., and Carpendale, S. Collaborative coupling over tabletop displays. In *Proc. CHI*, ACM Press, 2006, pp. 1181–1190.

[12] Tobiasz, M., Isenberg, P., and Carpendale, S. Lark: Coordinating co-located collaboration with information visualization. *T-VCG 15, 6* (2009). To appear.

[13] Wigdor, D. and Balakrishnan, R. Empirical investigation into the effect of orientation on text readability in tabletop displays. In *Proc. ECSCW*, Kluwer Academic Publishers, 2005, pp. 205–224.

[14] Wigdor, D., Shen, C., Forlines, C., and Balakrishnan, R. Perception of elementary graphical elements in tabletop and multi-surface environments. In *Proc. CHI*, ACM Press, 2007, pp. 473–482.

# Device-based Adaptation of Visualizations in Smart Environments

Georg A. Fuchs, Conrad Thiede, Mike Sips, and Heidrun Schumann

**Abstract**—Smart environments are beginning to have a large impact to collaborative group work in business and science. The multi-user and multi-display character of these group work environments presents a novel challenge for information visualization, namely, the adaptation of graphical representations of data to specific target devices in the environment. In this paper, we discuss a general strategy for an automated device-based adaptation of visualizations. We report interesting preliminary results of our adaptation strategy for conventional scatterplots used within a service-oriented visualization framework.

**Index Terms**—Information Visualization, Display Adaptation, Smart Environment, Collaboration, Scatterplot.

## 1 Introduction

Smart environments facilitate collaborative work of a group of users, e.g., in the analysis of massive customer databases to achieve better business decisions. A typical device ensemble in a smart room environment consists of stationary devices such as desktop computers, projectors, light, or motion trackers, but also strives to integrate mobile devices such as laptops, PDA or smart phones which are often carried by the users. In contrast to classical meeting room environments, smart environments augment sensor devices to monitor the environment and its users to enable a "smart" interaction between the users and the environment.

These novel environments present a number of challenges for information visualization, namely, (a) to support different user goals and data sources, (b) to utilize multiple displays, and (c) to facilitate interaction among a group of users. In this paper, we consider the adaptation of graphical representations of data to specific target devices in smart environments. The adaptation of graphical representations gives rise to the following two visualization challenges:

- In collaborative work sessions, users usually share a visualization on a wall-sized display to analyze/discuss potentially interesting features of the data, but also use the same graphical representation on their personal output devices to look at the data. A smart room environment should allow a dynamic adjustment of the requirements such as the task at hand and the visualization needed to foster insight into the data, but also should support a interaction of the users with the environment such as moving around to join different subgroups of the users. Thus, a graphical representation of the data often needs to be distributed to different output devices.

- The diversity of output devices/display sizes is often quite high in smart environments. To maintain visual effectiveness of a graphical representation under different display sizes, i.e. important features of the data are faithfully communicated to the user, a visual interface should apply a *device-based adaptation* to the visual output.

To facilitate a smart interaction between the users and the environment, the adaptation of a visualization to a specific target device should be performed automatically. This requires suitable metrics to measure and assess the effectiveness of a visual representation for the current output device and task. In this paper, we focus on automatic device-based adaptation of visual representations to support the dynamic assignment of visualizations to varying display sizes.



- *G. Fuchs, C. Thiede and H. Schumann are with University of Rostock, E-mail: {georg.fuchs|conrad.thiede|heidrun.schumann}@uni-rostock.de.*
- *Mike Sips is with MPII Saarbrücken, E-mail: msips@mpi-inf.mpg.de.*



We review related work and discuss principal distribution scenarios in Section 2. As the main contribution, we propose a general adaptation strategy for automatic device-based adaptation of visualizations, and discuss the key challenges in Section 3. Section 4 presents preliminary results on our ongoing work on integrating this strategy for conventional scatterplots with a service-oriented framework for distributed visualizations. Section 5 concludes with a discussion of open research challenges and gives an outlook on future work.

## 2 Background & Related Work

In smart environments the devices form a loosely coupled network, which allows for the necessary communication to accomplish tasks in a coordinated fashion [1]. Specifically, this enables distribution of visual output among several output devices depending on the current situation, based on an automatic situation assessment [7]. However, multi-user collaboration in these environments is explicitly *not* limited to showing a single user's content on several displays. Instead, available output devices need to be assigned to visual representations according to the users' current situation and task requirements. Two principal distribution scenarios can be distinguished: multiple users work jointly on a single representation, or one or more users require individual visualization specific to their current task.

In the first scenario, all users view the same information representation that is either *distributed* to show simultaneously on different displays, or *split* to show sub-regions among several (neighboring) displays. The latter distribution scheme is useful if an array of small displays (such as PDA and TabletPC) is available to the user, or if the data set is extremely large. In these situations, using well-established concepts such as Overview & Detail and Focus & Context can help in exploring the data. Here, one "public" display (e.g. a whiteboard display) shows an overview representation of the data, while other "private" displays (e.g. Laptops) show additional data for regions of personal interest. In the second scenario, users require different information representations for their individual tasks. If there are fewer displays available than visual representations required, the available display space must be shared by *combining* representations. Note, both distribution scenarios require that visualizations are scaled to fit a particular display area on the target device.

Many approaches found in literature deal with scalable representations of graphical content such as video streams and vector graphics [8], 2D maps [4] or 3D virtual models [5], but do not address information visualization. Most adaptive visualization approaches, on the other hand, consider the properties of the data and the visualization goal (e.g., [6, 17]), but only few approaches adapt to the available resources on different target devices (e.g. [12]). Others address issues related to distributed visualization, e.g. multiple client platforms [12] or the use of web services as the output distribution mechanism [16]. In [13], the authors propose a more general approach for distributed visualization. It uses a service-oriented architecture (SOA) to generate visual representations in a distributed fashion, including mobile devices that can enter or leave the smart room's ensemble.

Fig. 1. Example for lost effectiveness of a conventional scatterplot due to transfer to a smaller display: the perceived number of classes may change from three to two.

These architectures support, in principle, dynamic multi-user and multi-display settings as outlined above. However, as the generated visual representations get dynamically reassigned to different displays, adaptation may become necessary to maintain their effectiveness (see Fig. 1). Thus, a distributed visualization architecture should incorporate suitable adaptation mechanisms to adapt graphical representations of data to specific output devices.

## 3 General Device-based Adaptation Strategy

The basic idea of our general device-based adaptation in this setup is illustrated in Figure 2. The first step is to estimate the visual effectiveness of the representation. This might be done either as an initialization step prior to the use of the output device or after a visual representation that is already displayed has been re-assigned to a different display. If the effectiveness score indicates the visual encoding becomes poor for the assigned display area, in a second step an appropriate adaption is selected to increase the visual saliency of potentially interesting features of the data.

For this general adaptation scheme to work, a number of questions need to be addressed. Namely, (a) how to evaluate visual effectiveness, (b) what methods and constraints exist for adaptation, and (c) what are the requirements for a suitable infrastructure to support both output distribution and device-based adaptation in smart environments.

### 3.1 Visual Effectiveness

The effectiveness of a graphical display to faithfully present potentially interesting features of data to the user can be degraded in either of two cases: (a) when the visual representation has been transferred to a display smaller than originally intended, or (b) to a much larger screen, e.g. a whiteboard display. In the first case, a problem occurs when too much data is displayed on too small displays. The resulting ”visual clutter” is well researched (e.g. [2, 11]) and can have a significant impact on the effectiveness of a visual representation. However, there is also evidence that in some situations, upscaling a visual representation to a larger display area may also lead to a degradation in effectiveness as e.g. perception of overall data distribution and local densities change [2].

The definition of measures of goodness to score the visual effectiveness of a particular visualization is still a open research problem [9]. Tufte [15] proposes some measures to estimate the quality of 2D representations of static data, like the *data-to-ink* ratio or the *data density*, which take into account the size of the visual representation in relation to the amount of data displayed. Other approaches measure the overall ”visual clutter” [2] or consistency [11] to evaluate how faithfully a visual representation communicates data characteristics. Moreover, to enable automatic adaptation reliable thresholds specifying the perceptual boundaries on what constitutes effective visualizations are required. This will necessitate user experiments to determine

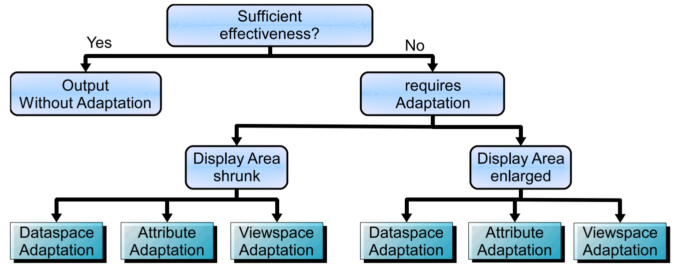


Fig. 2. Scheme of the adaptation process after a visual representation has been reassigned to a new display of different size.

those thresholds, and to get a good estimate on the perceived quality of adapted visualizations.

### 3.2 Adaptation Options & Constraints

The visualization process can be understood as a pipeline of four data stages plus intra-stage and transform (inter-state) operators [3]: raw data ($1^{st}$ stage) is transformed into analytical abstractions ($2^{nd}$ stage), e.g. by calculating statistical moments, which are then further mapped to visual abstractions ($3^{rd}$ stage), e.g. 2D points with position and color. Finally, the rendering process generates the image data ($4^{th}$ stage). The first two stages constitute the *data space* that is transformed by the visualization's *mapping* into the *view space* comprised of the last two pipeline stages. This visualization model yields starting points for adaptation both in data space and in view space, depending on the modified pipeline stage. In addition, the mapping parameters of data values to visual attributes, i.e. the transformation from data to view space, can also be modified (attribute adaptation).

If a visual representation is transfered to a smaller display, the level of detail may need to be reduced. Here, adaptation of the representation in data space includes filtering or using a higher abstraction level (e.g. clusters or statistical aggregates) to reduce the amount of data items displayed. View space adaptation aims to reduce visual clutter (e.g., by employing density binning [10]). Contrary, on larger displays, more details can be shown, by selecting a lesser degree of abstraction (e.g. another level from hierarchical clustering) or adjusting filter settings accordingly.

Another question is which visual attributes and aggregations are eligible for adaptation. Visualization techniques encode data values to different visual elements and their associated attributes, thus requiring specific adaptation mechanisms based on the set of used attributes. Adaptation is also inherently task-dependent, i.e., what view space aggregations and abstractions of the raw data are admissible for a given visualization goal? Identifying outliers has different requirements for a visual representation than analyzing complex relations, for instance.

Providing adequate solutions to these questions is not trivial. As a proof of concept, we chose 2D scatterplots to derive concrete procedures from the general adaptation scheme (see Section 4, using the infrastructure described next.

### 3.3 Infrastructure

A suitable infrastructure generates device-driven visual representations of the data for the available output devices, utilizing computing devices in the smart room's device ensemble, and distributes these according to the current requirements of the users. We chose a service-oriented framework called SSC from [13] as the infrastructure for our experiments. It uses a visualization pipeline composed of distributed services implementing pipeline operators. Adaptation mechanisms can be integrated into this general framework through service parameterization, or through extensions of the basic pipeline with additional services, such as a filtering service to sample data prior to the mapping stage of the pipeline. Furthermore, in line with the principal distribution scenarios identified in Section 2, we implemented the following distribution mechanisms for visual representations into SSC to facilitate testing:

To enable *distribution* of a single visual representation to multiple devices, the final rendering stage of the pipeline is forked to multiple rendering services [13], one for each device. This allows device-based adaptation in view space and attribute adaptation on each device individually, while the earlier data stages are processed only once for all devices[1].

*Splitting* a visual representation across multiple displays is achieved by forking a second pipeline to render a detail view of the user-selected area on another device. Currently, the framework supports detail view of regions of interest that are interactively selected by the user.

The capability to *combine* multiple visualizations on a single display is provided by an *aggregator service* that partitions the physical display into a corresponding number of viewports. The visual representation for each viewport is generated by the respective visualization pipeline that feeds its output to the aggregator service.

The general framework augmented with these distribution mechanisms provides the basis for our experiments.

## 4 Demonstrating Example

As a starting point for our test implementation, we created a visualization pipeline with the SSC framework that defines the necessary operators to create a 2D scatterplot from multivariate data. We assume the data points have class labels assigned, and that the visualization goal is to communicate the class structure of the data. We further assume a suitable 2D projection is used that captures the high-dimensional class structure in the 2D scatterplot [11].

The scatterplot representations generated by this pipeline are then distributed to displays in the environment using SSC's distribution mechanisms as described in Section 2. We discuss implementation details for the different steps of the adaptation process as summarized by Figure 2 next.

### 4.1 Efficiency Evaluation

The first step is to estimate the effectiveness of the (newly created or re-assigned) scatterplot. For this purpose, we use two measures. First, the class consistency score [11] is calculated. Second, the *visual density* of the scatterplot is determined to measure the impact of the new display area's size. Here we define the visual density of a scatterplot as the average ratio of cluster members to the screen space occupied by the cluster. The area is conservatively estimated by calculating the size in pixels of the convex hull of all points belonging to that cluster.

We established approximate thresholds for the two measures in an informal pilot study for the two principal cases of assigning a scatterplot to a smaller and a larger display area, respectively. In the first case, a drop of the consistency score and an increase in visual density can be observed. Low consistency scores ($60-80\%$) in conjunction with density values between 0.1 and 0.7 suggest the scatterplot display is saturated with points. At this point, clusters begin to mix visually (see Figure 1). This can be countered by suitable attribute adaptation, e.g. by using different shapes or by increasing the color contrast between points belonging to different clusters. When the consistency is even below 60% and visual density is above 0.7, mixing of clusters has become so severe that overplotting has likely occurred, and attribute adaptation does not help much to improve effectiveness. We propose to switch to a density plot in this situation since individual data points are no longer discernible anyway. The density representation at least allows the user to faithfully extract the cluster structure. We chose to integrate a binning approach (view space adaptation) further supplemented with an alternate color coding (attribute adaptation).

The second case – assignment to a larger display area – exhibits no drop in consistency. However, plots with a good consistency score ($> 80\%$) but visual density below 0.005 (i.e., only about one out of 200 pixels within the convex hull is set) describes a situation in which data points spread too much. An important finding of our preliminary user study is that the extraction of potentially interesting features by the human is biased toward low visual densities. In our scenario, the participants in the user started to identify sub-regions as individual clusters due to diminishing visual densities. This effect was also observed by Bertini and Santucci [2]. This can be addressed by encoding cluster membership into unused visual attributes (e.g. shapes, color) or by deliberately downscaling the representation to use only a fraction of the available area. To find appropriate adaptation strategies for this situation, however, is still an open research problem we did not yet pursue further. Therefore, the following subsections discuss examples for the three adaptation types (cf. Figure 2) specifically for the case of shrunk display sizes.

[1]Note that for adaptation in data space, the pipeline would have to be split into parallel services even sooner, at the corresponding data stage.

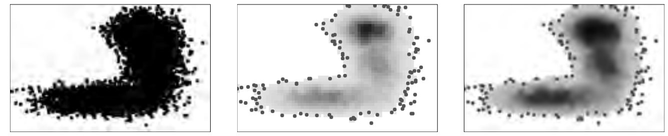

(a) (b) (c)

Fig. 3. **Results of density binning for a synthetic data set** – (a) standard scatterplot, (b) density binning with outlier preservation and linear color scale, (c) bin frequencies are mapped to logarithmic color scale.

### 4.2 View Space Adaptation

Our density binning approach borrows from [10], which has been proposed as a Focus & Context technique for crowded parallel coordinates. The basic idea can be summarized as follows. Both scatterplot axes are divided into $b$ regular intervals. The resulting set of $b \cdot b$ bins represents a so called bin map and can be thought of as a 2D histogram of the data point distribution in view space. Every non-empty bin is represented as a rectangle in the adapted scatterplot, with the bin frequency color-coded into its fill color.

However, to faithfully extract the cluster structure, the user should be able to discern the cluster centers from the frequency representation. Ideally, each cluster should register as a high-frequency region in the plot that is visually distinguishable from peaks of neighboring clusters. To facilitate these properties, we introduce an extension of the approach based on the following ideas.

**Automatic binning resolution adjustment:** First, we adjust the bin resolution along the scatterplot axes with respect to cluster center locations to determine a good binning. Initially, the scatterplot area is partitioned into $b_x \times b_y$ bins according to a given starting bin size. Next we check if a bin contains more than one cluster center. The binning subdivision is then refined by increasing $b_x$ ($b_y$) by 1, and the check is repeated. This subdivision continues by alternately increasing $b_y/b_x$ until (a) all cluster centers are located in individual bins, or (b) predetermined bin size is reached. We found that a bin sizes between $5 \times 5$ pixels (starting value) to $2 \times 2$ pixels (minimum threshold) yield a good compromise between clutter reduction and faithful reproduction of clusters on small displays. (see Figure 3(a, b)).

**Local magnification with sub-binning:** After the bin size and the resulting bin frequencies have been determined, we optionally apply a rectangular fish-eye distortion aligned with the bin grid centered about those bins containing the cluster centers (see Figure 4(b)). This increase in screen space available for the clusters center regions allows sub-binning these regions. The sub-binning factor is thereby proportional to the magnification factor, e.g. a magnification of focused bins by factor two results in a two-fold subdivision of these bins. The locally increased bin resolution reproduces frequency variations around the cluster centers with higher fidelity and thus can improve visual separation of clusters with low separation (Figure 4(c)).

### 4.3 Attribute Adaptation

Moreover, instead of using the originally proposed, linear color map from [10], we use a logarithmic scale for bin colors. The skewed distribution of bin frequencies between dense centers of compact clusters and sparse regions where clusters mix suggests a non-linear color scale is better suited for this kind of data [14]. Figure 3(c) illustrates the difference between a linear and a logarithmic color scale. The latter scale

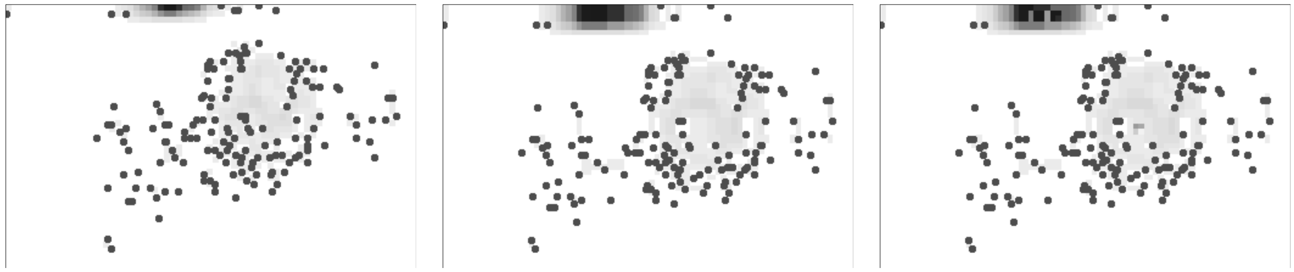


Fig. 4. **Use of rectangular fish-eye distortion** – (a) undistorted density plot, (b) bins containing cluster centers magnified, (c) magnified regions sub-binned. The increased density sampling rate in (c) reveals that the dense region at the top is actually comprised of three peaks (i.e., clusters). The location of the center peak within the low-density cluster at the middle-right is also emphasized.

assigns more gray levels to the low end of the frequency value range, thus further enhancing visibility of the peaks around cluster centers.

### 4.4 Data Space Adaptation

A third option for smaller displays is to reduce the number of data items prior to mapping them to the visual representation. This can be achieved by employing data sampling. To preserve the class structure, however, the sampling process should maintain local densities [2]. Although random sampling schemas are, in principal, of value in many adaptation situations, it is rather useless in our scenario. For this reason, implementation of a suitable density-preserving sampling service has not been pursued yet.

## 5 Summary & Discussion

Multi-user, multi-display settings in smart environments present novel challenges for information visualization. In particular, the varying sizes and capabilities of different output devices require device-based adaptation of generated visual representations. In this paper, we proposed a general strategy to guide this adaptation based on the notion of visual representation effectiveness and a visualization pipeline model. As a proof-of-concept, we implemented corresponding adaptation mechanisms for 2D scatterplots in the SSC framework [13].

We believe the general strategies proposed in this paper are valid and can be employed to many visualization techniques, however, *smart* device-based adaptation (i.e., minimal user intervention) requires more work. First, thresholds for consistency and visual density need further evaluation in controlled user experiments, which is the subject of our current efforts. This specifically includes cases with low visual densities (display size is too large for a given visualization), as this branch of the adaptation scheme (Fig. 2) was not pursued in detail so far. Also, the problem of meaningful effectiveness measures requires more research. Consistency is applicable only for scatterplots of clustered data. Metrics striving to capture the amount of visual clutter in visual representations [2, 11] seem promising candidates for a more generic effectiveness evaluation.

The distribution of visual representations is still work in progress as well. Smart distribution requires detection of the current user situation followed by inferring the individual goals as well as the group's intention, which is not the focus of our work. However, we plan to integrate these schemes with an existing inference module (cf. [7]) in the future.

Additionally, we plan to further investigate task-driven aspects of the adaptation process. A typical smart room scenario is a decision making process where several domain experts look at the same data, albeit with different goals and requirements to the visual representation. So far, we only considered a single visualization goal, namely communicating the class structure of multi-dimensional data.Note that the visualization goal respectively the user's current task have a direct impact on device adaptation. The task determines what data abstractions are permissible or how visual attributes should be modified e.g. through color coding. Using a suitable task description in the adaptation process would therefore allow to integrate different task-specific aspects in a single collaborative visualization on the same screen, rather than just juxtaposing several independent representations. Our initial studies in this direction included task-based adaptation of graphical content using enriched task models [5] and a task taxonomy for color coding [14].

## References


[1] E. Aarts and J. Encarnação. *True Visions: The Emergence of Ambient Intelligence*. Springer, 2006.

[2] E. Bertini and G. Santucci. Quality metrics for 2d scatterplot graphics: Automatically reducing visual clutter. In *Proc. Smart Graphics (SG'04)*, volume LNCS 3031, pages 77–89. Springer-Verlag, 2004.

[3] E. H. Chi. A taxonomy of visualization techniques using the data state reference model. In *INFOVIS*, pages 69–76, 2000.

[4] J. Follin, A. Bouju, F. Bertrand, and P. Boursier. Management of multi-resolution data in a mobile spatial information visualization system. In *Proc. 4th Intl Conf. on Web Information Systems*, pages 92–99, 2003.

[5] G. A. Fuchs, M. Holst, and H. Schumann. 3D Mesh Exploration for Smart Visual Interfaces. In *Proc. 10th Intl. Conference on Visual Information Systems*, volume LNCS 5188, pages 80–91, 2008.

[6] I. Fujishiro, Y. Ichikawa, R. Furuhata, and Y. Takeshima. GADGET/IV: a taxonomic approach to semi-automatic design of information visualization applications using modular visualization environment. In *Proc. IEEE InfoVis*, 77-83, 2000.

[7] T. Heider and T. Kirste. Minimizing Cognitive Load by Automatic Display Management. In *Proc. Ubicomp Workshop'07*, 2007.

[8] M. Klima, P. Halabala, and P. Slavik. Semantic information visualization. In *Proc. 19th Intl. CODATA Conference*, 2004.

[9] C. North. Toward Measuring Visualization Insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, 2006.

[10] M. Novotny and H. Hauser. Outlier-Preserving Focus+Context Visualization in Parallel Coordinates. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):893–900, 2006.

[11] M. Sips, B. Neubert, J. P. Lewis, and P. Hanrahan. Selecting good views of high-dimensional data using class consistency. In *Proc. EuroVis 2009*, volume 28 of *Computer Graphics Forum*, 2009.

[12] L. Skorin-Kapov, H. Komericki, M. Matijasevic, I. Pandzic, and M. Mosmondor. MUVA: a Flexible Visualization Architecture for Multiple Client Platforms. *Journal of Mobile Multimedia*, 1(1):3–17, 2005.

[13] C. Thiede, C. Tominski, and H. Schumann. Service-oriented information visualization for smart environments. In *International Conference Information Visualisation (IV)*, Barcelona, Spain, 2009.

[14] C. Tominski, G. Fuchs, and H. Schumann. Task-Driven Color Coding. In *Proc. 12th International Conference Information Visualisation (IV'08)*, London, 8 - 11 July 2008. IEEE Computer Society.

[15] E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 2nd edition edition, May 2001.

[16] Y. Zhao, C. Hu, Y. Huang, and D. Ma. Collaborative Visualization of Large Scale Datasets Using Web Services. In *Proc. of Intl. Conf. on Internet and Web Applications and Services*, 2007.

[17] M. X. Zhou, M. Chen, and Y. Feng. Building a Visual Database for Example-based Graphics Generation. In *IEEE Information Visualization*, pages 23–30, 2002.

# Design Considerations for Collaborative Information Workspaces in Multi-Display Environments

Manuela Waldner, Alexander Lex, *Student Member, IEEE*,
Marc Streit, *Student Member, IEEE*, and Dieter Schmalstieg, *Member, IEEE*

**Abstract**—The incorporation of massive amounts of data from different sources is a challenging task for the conception of any information visualization system. Especially the data heterogeneity often makes it necessary to include people from multiple domains with various fields of expertise. Hence, the inclusion of multiple users in a collaborative data analysis process introduces a whole new challenge for the design and conception of visualization applications. Using a multi-display environment to support co-located collaborative work seems to be a natural next step. However, adapting common visualization systems to multi-display environments poses several challenges.

We have come up with a number of design considerations for employing multiple-view visualizations in collaborative multi-display environments: *adaptations of the visualization* depending on display factors and user preferences, *interaction techniques* to facilitate information sharing and to guide the users' attention to relevant items in the environment, and *the design of a flexible working environment*, adjustable to varying group sizes and specific tasks.

Motivated by these considerations we propose a system relying on a spatial model of the environment as its main information source. We argue that the system design should be separated into basic multi-display environment functionality, such as multiple input handling and the management of the physical displays, and higher level functionality provided by the visualization system. An API offered by the multi-display framework thereby provides the necessary information about the environment and users to the visualization system.

## 1 Introduction

Modern information workers need to explore large information spaces to reach crucial decisions, such as those with strong influences on people's well-beings. Those decisions are rarely made by a single person but are rather discussed and evaluated by a team of experts. Examples are doctors deciding for treatment courses after exploring and discussing the diagnostic data of patients, architects and other stakeholders discussing on urban planning issues [15], emergency services having to react to ongoing crises, scientist discussing patterns and findings in data, or engineers collaborating with their peers when designing the car of the next generation. All these scenarios are accomplished by a small group of experts and involve massive amounts of data.

Information visualization software helps to cope with large amounts of data by letting the user interactively explore the information space. Especially multiple-view visualization can prove useful in collaborative information analysis situations, where users might prefer different visualization styles based on their personal preferences and knowledge backgrounds. However, most software solutions have two major shortcomings impeding effective collaborative work: First, they are designed as single-user applications not able to distinguish input from multiple users, even if the underlying operating system is capable of handling multiple input devices. Second, single-machine software has to cope with limited screen space, typically a single or dual monitor setup.

It seems natural to match the multiplicity of displays in a multi-display environment (MDE) to the multiplicity of users and visualization views in a multiple-view visualization system. MDEs combine displays of various form factors to a unified interaction space. Traditional collaboration in small groups, where participants discuss print-outs on a table, take notes in private notebooks, and sketch ideas on a white board, can be emulated by turning unused wall and table spaces into interactive workspaces and integrating brought-in personal devices into the interactive environment.

Building a visualization system that makes optimal use of an MDE is not simply a question of providing a very large number of pixels. Collaboration requires that users can manipulate application content simultaneously and tailored to their personal preferences, and that tools for guiding the users' attention in the large workspace are provided. Visualization styles, placements, and detail-levels should differ depending on the used display and the users interacting with the visualization. Tasks such as choosing the appropriate display for a visualization or the appropriate level of detail of a visualization for a particular display can be solved manually. However, we believe that an automated approach can facilitate the usefulness of such systems. To automate these operations the system requires knowledge of the geometric and topological properties of the display setup, the locations of the users within the environment, and their backgrounds and preferences.

(a) (b)

Fig. 1. Examples for collaborative information analysis in multi-display environments: (a) analysis of biomolecular data and (b) urban planning.

In this paper we present a set of design considerations for visualization and interaction techniques tailored to collaborative multi-display situations, as illustrated in Figure 1. Subsequently, we will propose a system design for a co-located collaborative information workspace incorporating multiple displays of varying form factors. We will show that the system's detailed knowledge of spatial display arrangements and user locations is crucial when building collaborative information workspaces in MDEs.

• *Manuela Waldner, Alexander Lex, Marc Streit and Dieter Schmalstieg are with the Institute for Computer Graphics and Vision at Graz University of Technology, E-mail: {waldner|lex|streit|schmalstieg}@icg.tugraz.at.*

## 2 Related work

In their "rule of diversity", Baldonado et al. [1] suggest that multiple views should be employed if users' preferences and knowledge backgrounds differ. Convertino et al. [4] proposed a single team view and role-specific private views for each team member to ease the group analysis task of a map-based visualization. Tang et al. [20], as well as Forlines and Shen [6] demonstrated systems providing each user with tools for filtering a single, shared view. Isenberg et al. created a collaborative visualization system for a multi-touch table [9]. They also presented a set of design guidelines for collaborative information visualization systems, which was extended by Heer et al. [8]. Our design considerations differ, as we are more focused on MDEs with special emphasis on the influence of display geometries, display topologies, user locations and user preferences on the visualizations.

In an MDE, Forlines and Lilien [5] distributed multiple coordinated 3D views of a protein to an interactive touch display, two wall displays, and a tablet PC for fine-grained interaction. Although their system supports multiple users by facilitating a multi-touch table, they do not provide special collaborative interaction features. Shen et al. [16] developed a taxonomy of multiple-view visualization styles in multi-display environments . They proposed three visualization styles differing in their synchronization method. In contrast to their taxonomy, we propose the separation of the system into a multi-display framework and a visualization framework to not limit the MDE's functionality to the visualization system and likewise, not to limit the visualization system's applicability to the specific MDE.

## 3 Design Considerations

Special requirements for co-located collaborative information visualization systems arise from the implicit components: *multiple users* are operating on *multiple data sets* by using *multiple visualizations* on *multiple displays*. A discussion of the requirements in terms of visualization and interaction is followed by considerations about the needed properties of a multi-display environment supporting co-located information seeking.

### 3.1 Visualization Techniques

In a collaborative MDE, not only inherent display factors (i.e., the display geometries) have an influence on the subjective quality and perceptibility of the visualization. The users' preferences and knowledge backgrounds play an equally important role.

**Visualization LOD: The visualization's level of detail should vary with the display.** Display parameters, such as size, resolution, distance to the users, and viewing angles of the users, influence the required or possible level of detail a visualization can show. The level of detail adjustment depends on the view, but generally affects the level of abstraction, the number of text labels, the size of the remaining text labels, and how much data entities are shown in the view. If the number of available pixels does not allow a visualization to show all data entities simultaneously, abstractions such as clustering or focus+context methods can be applied to convey an acceptable compromise between overview and detail. As an example, consider a low-resolution wall projection serving as contextual information space, while users conduct individual work on their private workstations. The local visualizations on the private displays show a large number of elements in a plot, while context views on the wall show only contextually relevant elements.

**Personalized views: The visualization's level of detail should vary with the user.** Experts from different domains might not only prefer different data representations, but also specific terminology. When reviewing data collaboratively, a shared information space easily gets cluttered with extensive text labels and alternative representations. In an MDE, private monitors provide a convenient space for visualizations adjusted to the users' background and preferences without affecting shared or other users' private views. We hypothesize that users prefer more sophisticated and interactive views on private displays, while views, which require less precise interaction and convey information in a more obvious way, are preferred on shared display spaces. In addition, visualizations on public displays can combine information from different data domains and therefore bridge the knowledge gaps between experts from different fields, as explained in a companion paper by Streit et al. [19]. However, we believe that it is crucial that the user retains control over which information should be visible on which display.

### 3.2 Interaction Techniques

Typical activities when using information visualization systems include interactive filtering [17] and brushing [10] to understand the data and its relations. These actions are equally important in a collaborative multi-display setting, but there are some issues which have to be considered: First, multiple discontinuous display spaces make relating linked elements and arranging the multiple views more difficult. Second, having multiple users frequently shifting between a loosely coupled and a tightly coupled work style [7, 20] poses challenges to make these shifts fluid, while preserving sufficient privacy for undisturbed individual work.

**Visual Linking: Guide the users' attention.** When information is scattered in an MDE, relevant items, for instance data elements related to the user's current selection, might not be in the user's direct field of view. Subtle highlighting of related elements might not be sufficient to guide the user's attention to secluded display spaces. An approach to show relationships between items more explicitly is visual linking [3], which draws line connections between related elements in two views. However, in an MDE the visual links between views need to bridge display space potentially covered by other applications, as well as display-less space between two adjacent displays. The path across discontinuous displays becomes ambiguous as soon as the displays are not located on the same depth level. Figure 2 shows two possibilities how to design cross-display visual links: A static determination of entry- and exit points for lines connecting two screens makes the visual links predictable and works equally well from every perspective (Figure 2a). When drawing visual links from a single user's perspective, the complexity for the user is reduced (Figure 2b), while other users perceive discontinuities in the links.

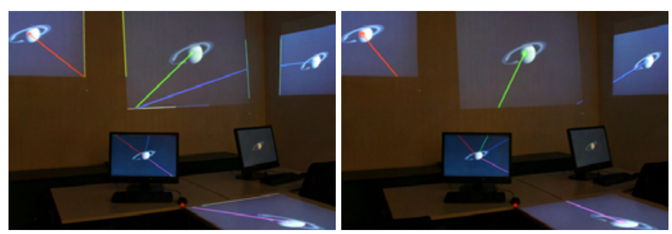

(a) (b)

Fig. 2. Visual links demonstrating the shortest navigation paths from the current mouse pointer location to potential target displays to ease navigation in an MDE: (a) through static entry- and exit points, and (b) by incorporating user perspective.

**View relocation: Provide semi-automatic mechanisms to share information.** Especially when working in an MDE with private monitors hidden from other participants, there are several situations where users would like to move visualizations from one display to another: They may want to retrieve a copy of a public visualization view for detailed investigation on their private workstations, they may want to send a copy or reference of their visualization directly to a single collaborator or to a group of collaborators for discussion, or they could also place visualization views on a public display space without the intention to immediately discuss or present findings. Moving objects (e.g. visualization views or application windows) across display boundaries using interaction techniques like drag and drop [13, 11] is a non-trivial task, especially if the display topology is complex and user interaction can potentially interfere with other participants. However, with the system's knowledge of display arrangements and user locations, we can provide high-level interaction techniques like *pub-*

*lish*, *deposit*, or *obtain* which intelligently move views to displays. For example, users wanting to discuss findings identified on their private workstations, press a *publish* button in the graphical user interface. The system then identifies the most suitable target display by considering properties of the visualization (*Can the level of detail be adjusted? Does it contain rotation-sensitive elements?*) and the potential target displays (*Is it visible for all affected users? Is there someone interacting with the display?*). The view is then relocated to the best suited display. To allow for manual adjustment, conventional drag and drop relocation techniques should additionally be supported.

**Privacy: Ensure uninterrupted individual work.** Linking & brushing assures that selections made in one visualization are reflected by all other views. For instance, if a user selects an element in one view, this event can modify all other views in the environment – shared views, as well as views on private display spaces. If other participants' private views are modified in a loosely coupled collaboration situation, their ongoing activities might get interrupted. The system thus needs to treat views placed on private display spaces with special care. Linking & brushing events for private views need to be restricted to avoid sudden, unexpected changes interrupting individual work. Likewise, visual links should not connect elements between shared views and private views, unless invoked by the owner of the private view.

**Personalized Interaction: Make individuals' actions distinquishable.** Personalized interaction techniques, for example user-based color-coding visual links and highlights, helps users to distinguish actions from different collaborators. System responses tailored to the users' preferences (e.g. by providing customized mouse-over information [14]) are especially important when experts from different domains collaborate. Such features require a system capable of distinguishing input from multiple users.

### 3.3 Environment

In an MDE, the working environment can be tailored to the information analysis task to be accomplished and the group being involved. This affects the physical display arrangements as well as the display form factors. Additionally, the users should be able to customize their system by freely choosing their supportive software tools.

**Display setup: Make the display environment (re-)configurable.** By mixing displays of varying affordances, collaborative tasks can be simplified. For instance, people can gather around a tabletop display to discuss information, while a wall display is used for presentation purposes. Private displays introduce an implicit task separation and foster a loosely coupled work-style. An MDE has to be carefully designed to find the perfect balance between providing sufficient display space, arranged in a fashion to best support the group activities, while not overwhelming the users with a seemingly endless amount of visible information. A collaborative information workspace has to accommodate for these situations by being adjustable to task requirements and group size. It should be easily reconfigurable to support a changing group size and to incorporate brought-in mobile devices, such as personal laptops.

**Display geometry: Make the displays configurable.** For detail-rich visualization representations, it is not only important to provide high-resolution displays. In certain cases, visualizations can benefit from unconventional displays in terms of aspect ratio or display geometry. Consider, for instance, the parallel coordinates visualization shown in Figure 3: With limited screen space, horizontal scrolling or panning is required to explore all dimensions. By combining multiple projectors to a very wide, high-resolution projected display, even degenerated visualizations with aspect ratios not conform with conventional monitor dimensions can be explored without scrolling, panning, or zooming.

**Application Transparency: Provide supportive applications.** A collaborative information analysis session clearly benefits from a rich visualization system support such as cross-machine linking & brushing. However, conventional software tools, such as web browsers, e-mail clients, or presentations tools can further enrich the collaborative session. It is therefore important to also allow legacy applications to function as usual in such a setup.

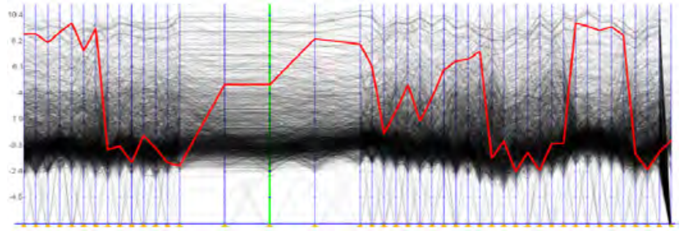

Fig. 3. On a conventional monitor showing all dimensions of this parallel coordinates view would result in visual clutter, due to the limited space between the axis. On a very wide horizontal display more dimensions can be visualized simultaneously with sufficient spacing.

## 4 System Design

Based on the design considerations discussed above we propose a system design that includes a detailed spatial model of the environment as its main information source. As shown in Table 1, many of the visualization and interaction techniques proposed cannot be provided without knowledge of geometric display properties, the display topology (i.e., the spatial relationship of displays to each other), and the location of the user within the environment with respect to the display locations. However, the spatial model is not only necessary for visualization and interaction techniques, but is also required to build a flexible, configurable display environment. The aforementioned displays with unconventional aspect ratios and geometries can be built from multiple casually aligned projectors or projections onto non-planar surfaces, using geometric compensation and edge blending (see [2] for an overview on seamless multi-projector displays).

| | Display geometry | Display topology | User locations | User preferences |
|---|---|---|---|---|
| **Visualization LOD** | x | | x | (x) |
| **Personalized views** | | | | x |
| **Visual linking** | | x | (x) | |
| **View relocation** | x | (x) | x | |
| **Privacy** | | x | x | x |
| **Personalized interaction** | | | x | x |

Table 1. Information sources required to provide display- and user-adaptive visualization and interaction techniques.

At our institute, we have developed *Deskotheque* [12], a distributed multi-display framework which acquires a three-dimensional model of the environment in a camera-assisted offline calibration step. Figure 4 shows a multi-display setup coordinated by the Deskotheque framework and the corresponding spatial model. Based on this model, we derive geometric compensation and edge blending for projected displays to support the construction of large high-resolution displays from multiple projectors and projections on multi-planar surfaces. Geometric compensation is applied in a 3D compositing window manager, thus transparent to any applications run in the environment [21]. From the spatial model we can also roughly estimate user locations, by assuming the users to be located at a static distance in front of a personal workstation monitor. This information is employed for providing spatially consistent cross-display mouse pointer navigation, which is crucial to access all display spaces in an intuitive fashion.

As Deskotheque is designed in an application-transparent manner, any information visualization application can be operated on the MDE framework without further adaptations (c.f. Figure 1a). However, to implement all the design considerations discussed in the previous sections, knowledge about the environment is required by the visualization application. We are therefore currently working on extending *Caleydo* [18], a multiple-view visualization system from the biomedical domain developed at our institute, to a distributed system which will make use of information provided by *Deskotheque*.

We anticipate a clear separation of MDE- and visualization framework. The multi-display framework has to provide the basic technology to create a shared workspace – irrespective of the anticipated con-

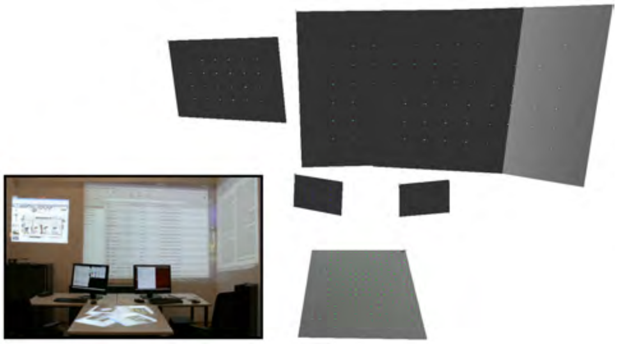

Fig. 4. The spatial model of a multi-display environment. Mind that the right multi-planar wall display is composed of two overlapping projections and all projected displays are geometrically compensated for projective distortion.

tent. This includes – but is not limited to – geometric compensation of projected displays, cross-display mouse pointer navigation, multiple input support, and object relocation facilities on window level, as well as the creation and maintenance of the spatial model of the environment. The visualization framework keeps records of user profiles and is responsible to provide appropriate visualizations adapted to display factors and user preferences. It also has to take care that multiple views distributed on multiple displays, and machines respectively, are synchronized and events are forwarded to all instances.

To adapt the visualization style, to calculate automatic placement positions, and to distinguish multiple collaborators, it can rely on an API exposed by the MDE framework which provides access to the display geometries, arrangements, user locations, and users associated with input events received by the visualization framework. The MDE framework furthermore has to take care to provide interfaces for cross-display painting of visual links (c.f. Figure 2), which is accomplished by a window manager plugin.

## 5 Conclusion

With increasing power and popularity of projectors and large-scale monitors, as well as the availability of massive amounts of data, extending visualization systems to MDEs seems to be a logical step. In this position paper we have presented a set of design considerations for adopting visualization and interaction techniques to this new situation and for what the environment for a collaborative information workspace should look like. Based on these, we have proposed a system design for such an information workspace with a clear separation between multi-display- and visualization framework. As a major requirement for a collaborative information workspace we have hypothesized the availability of a spatial model of the environment, describing the individual displays' geometries, the display topology, and the location of the users within this environment. Only with this knowledge we believe that the system can sufficiently support the users in their collaborative analysis task by adapting visualizations to the display form factors, providing highly sophisticated interaction techniques, and guiding their attention to relevant information.

## 6 Acknowledgments

The authors would like to acknowledge Christian Pirchheim for his contributions to the Deskotheque framework. This work was funded by the European Union (IST-4-27571), the *Fonds zur Förderung der wissenschaftlichen Forschung (FWF)* (Y193 and L427-N15), the *FIT-IT* program (813 398), and the *Zukunftsfonds Steiermark* (3007).

## References

[1] M. Q. W. Baldonado, A. Woodruff, and A. Kuchinsky. Guidelines for using multiple views in information visualization. pages 110–119, 2000.

[2] M. Brown, A. Majumder, and R. Yang. Camera-Based Calibration Techniques for Seamless Multiprojector Displays. *IEEE Transactions on Visualization and Computer Graphics*, 11(2):193–206, 2005.

[3] C. Collins and S. Carpendale. Vislink: Revealing relationships amongst visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1192–1199, 2007.

[4] G. Convertino, C. H. Ganoe, W. A. Schafer, B. Yost, and J. M. Carroll. A multiple view approach to support common ground in distributed and synchronous geo-collaboration. In *CMV '05: Proceedings of the Coordinated and Multiple Views in Exploratory Visualization*, pages 121–132, Washington, DC, USA, 2005. IEEE Computer Society.

[5] C. Forlines and R. Lilien. Adapting a single-user, single-display molecular visualization application for use in a multi-user, multi-display environment. In *AVI '08: Proceedings of the working conference on Advanced visual interfaces*, pages 367–371, New York, NY, USA, 2008. ACM.

[6] C. Forlines and C. Shen. Dtlens: multi-user tabletop spatial data exploration. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 119–122, New York, NY, USA, 2005. ACM.

[7] C. Gutwin and S. Greenberg. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In *CSCW '98: Proceedings of the 1998 ACM conference on Computer supported cooperative work*, pages 207–216, New York, NY, USA, 1998. ACM.

[8] J. Heer, F. van Ham, S. Carpendale, C. Weaver, and P. Isenberg. *Information Visualization*, volume 4950/2008, chapter Creation and Collaboration: Engaging New Audiences for Information Visualization, pages 92–133. Springer Berlin / Heidelberg, 2008.

[9] 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.

[10] D. A. Keim. Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):1–8, 2002.

[11] M. A. Nacenta, S. Sakurai, T. Yamaguchi, Y. Miki, Y. Itoh, Y. Kitamura, S. Subramanian, and C. Gutwin. E-conic: a Perspective-Aware Interface for Multi-Display Environments. In *Proceedings of the 20th annual ACM symposium on User interface software and technology*, pages 279–288, 2007.

[12] C. Pirchheim, M. Waldner, and D. Schmalstieg. Improving spatial awareness in multi-display environments. In *Proceedings of IEEE Virtual Reality*, pages 123–126, 2009.

[13] J. Rekimoto and M. Saitoh. Augmented Surfaces: A Spatially Continuous Workspace for Hybrid Computing Environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 378–385, 1999.

[14] K. Ryall, A. Esenther, C. Forlines, C. Shen, S. Shipman, M. R. Morris, K. Everitt, and F. D. Vernier. Identity-differentiating widgets for multiuser interactive surfaces. *IEEE Comput. Graph. Appl.*, 26(5):56–64, 2006.

[15] M. Sareika and D. Schmalstieg. Urban sketcher: Mixing realities in the urban planning and design process. In *ACM Computer Human Interaction 2008 (CHI 2008) Workshop on Urban Mixed Realities - Technologies, Theories and Frontiers. Florence, Italy*, 2008.

[16] C. Shen, A. Esenther, C. Forlines, and K. Ryall. Three modes of multi-surface interaction and visualization. In *Information Visualization and Interaction Techniques for Collaboration across Multiple Displays Workshop associated with CHI'06 International Conference*, 2006.

[17] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *VL '96: Proceedings of the 1996 IEEE Symposium on Visual Languages*, page 336, Washington, DC, USA, 1996. IEEE Computer Society.

[18] M. Streit, A. Lex, M. Kalkusch, K. Zatloukal, and D. Schmalstieg. Caleydo: Connecting pathways with gene expression. *Bioinformatics*, 2009.

[19] M. Streit, H.-J. Schulz, D. Schmalstieg, and H. Schumann. Towards multi-user multi-level interaction. In *Workshop on Collaborative Visualization on Interactive Surfaces (part of VisWeek'09)*, 2009.

[20] A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative coupling over tabletop displays. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 1181–1190, New York, NY, USA, 2006. ACM.

[21] M. Waldner, C. Pirchheim, and D. Schmalstieg. Multi projector displays using a 3d compositing window manager. In *IPT/EDT '08: Proceedings of the 2008 workshop on Immersive projection technologies/Emerging display technologiges*, pages 1–4. ACM, 2008.