

# Getting Practical with Interactive Tabletop Displays: Designing for Dense Data, “Fat Fingers,” Diverse Interactions, and Face-to-Face Collaboration

Stephen Volda, Matthew Tobiasz, Julie Stromer, Petra Isenberg, and Sheelagh Carpendale

Department of Computer Science, University of Calgary

2500 University Drive NW

Calgary, AB T2N 1N4 Canada

{svoida, tobiasz, petra.isenberg, sheelagh}@ucalgary.ca; julie.stromer@alumni.ucalgary.ca

## ABSTRACT

Tabletop displays with touch-based input provide many powerful affordances for directly manipulating and collaborating around information visualizations. However, these devices also introduce several challenges for interaction designers, including discrepancies among the resolutions of the visualization, the tabletop’s display, and its sensing technologies; a need to support diverse types of interactions required by different visualization techniques; and the ability to support face-to-face collaboration. As a result, most interactive tabletop applications for working with information currently demonstrate limited functionality and do not approach the power or versatility of their desktop counterparts.

We present a series of design considerations, informed by prior interaction design and focus+context visualization research, for ameliorating the challenges inherent in designing practical interaction techniques for tabletop information visualization applications. We then discuss two specific techniques, *i-Loupe* and *iPodLoupe*, which illustrate how different choices among these design considerations enable vastly different experiences in working with complex data on interactive surfaces.

## Author Keywords

Resolution discrepancy, interaction lenses, information visualization, *i-Loupe*, *iPodLoupe*

## ACM Classification Keywords

H5.2 [Information interfaces and presentation]: User Interfaces—Graphical user interfaces.

## INTRODUCTION

Information visualization research is often grounded in collaborations with scientists seeking to explore and analyze complex data sets. In many cases, these science teams utilize domain-specific tools to aid in their data exploration. It is also quite common for these groups to

collaborate and converse face-to-face over their data. All of these practices can potentially be a good match for the use of large, interactive tabletop displays, since these surfaces provide a large canvas on which data can be displayed, the capability to directly manipulate the data using touch input, and the opportunity to discuss and interpret the data with others without interrupting eye contact.


However, the strong match between these affordances and scientists’ collaborative information exploration practices also comes at a cost: it is difficult to design and implement practical tabletop applications that take into account the unique properties of interactive tabletop systems. These systems do not behave like traditional desktop applications and challenges arise when data sets become massive; when scientists’ fingers are much larger than the visualization details they want to manipulate; when interface controls for controlling the view into the data compete with interface controls for manipulating the data, themselves; and when collaborators sitting or standing around a tabletop each have a different perspective on the tabletop’s display contents. Although the research community has developed interface solutions to address each of these concerns independently, the design decisions that have to be made to address one challenge often directly impact the choices available for the others. Additionally, because the nuances of these decisions are often subtle, small differences in design choices can have a substantive impact in the overall user experience of a tabletop application and the usability of that application within a particular context.

In this paper, we present a series of design considerations, informed by an analysis of prior research and empirical studies of collaboration around tabletops, that address the three primary challenges that we encountered in designing practical applications for supporting collaborative, tabletop-based information exploration: ameliorating resolution discrepancies, enabling a diversity of interactions with both our tools and the underlying visualization, and fostering face-to-face collaboration. We then discuss two techniques that reflect different design choices, *i-Loupe* and *iPodLoupe*. These techniques illustrate how design considerations—and the different prioritization of design

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ITS ’09, November 23–25, 2009, Banff, Alberta, Canada.

Copyright 2009 ACM 978-1-60558-733-2/09/11...\$10.00



**Figure 1. The relationships among resolutions in interactive tabletop systems: high resolution data sets (bottom), medium resolution visual display (middle—note difficulty in resolving individual nodes), and effectively low resolution touch-sensitive input (top), due in large part to the “fat finger” problem.**

features—can lead to vastly different experiences in working with complex data on interactive surfaces.

### **DESIGN CONSIDERATIONS FOR SUPPORTING INFORMATION VISUALIZATIONS ON INTERACTIVE TABLETOP DISPLAYS**

There are many decisions that have to be made when designing interaction techniques to support collaborative information visualization on interactive tabletop displays. These decisions are also interrelated; the choices made for one design consideration affect the options available for making subsequent decisions.

Existing empirical and experimental research can inform these design choices, and this work benefits from a variety of perspectives since it exists at the intersection of several very active research domains. While working to develop a variety of applications to support teams of scientists, we identified three particularly significant challenges for supporting these kinds of collaborative visualizations on interactive tabletops. In this section, we enumerate these challenges and provide examples of the breadth of design solutions that have been proposed and, in some cases, demonstrated and evaluated in prior research from domains of study as diverse as information visualization, user interface software, and ubiquitous computing.

#### **Ameliorating Resolution Discrepancies among Data, Display Output, and Input**

Many dense visualizations contain more detail than can be clearly rendered in the available display space. With many real-world data sets, this can even be true when using large, high-resolution tabletop displays. As a result, information items that are adjacent (e.g., nodes in a graph layout) may overlap and become visually indistinguishable. The fact that most interactive tabletops rely heavily on touch-sensitive input complicates the situation: while direct interaction with fingers is appealing, fingers themselves are large enough compared to the size of the display pixels that even when the sensing substrate can detect touch locations

very precisely, the low *effective* input resolution makes interaction with visualization details difficult (Figure 1).

Although techniques like blurring [18] and semantic abstraction [2] can be used to address some of the discrepancies between visualization and display resolutions, these approaches do not extend easily to address the “fat fingers” problem with low-resolution input. An alternative approach is to provide flexible zooming or magnification capabilities that allow individuals to expand part (or all) of the visualization to show areas of interest at a resolution that is better supported by the display hardware *and* allow for more fine-grained interaction. A substantive body of research has explored this challenge; focus+context lenses are a particularly well-studied interaction technique aimed at ameliorating these kinds of resolution discrepancies.

Focus+context techniques have been studied extensively in the context of information visualization. While others have previously characterized the space of focus+context lenses (e.g., [7, 21]), these surveys generally focused on the presentation characteristics of these lenses. Notable frameworks for these techniques include Leung and Apperley’s early work on distortion-oriented presentation techniques [21], the Elastic Presentation Framework [5], and the Non-Linear Magnification framework [16]. Alternatively, in situations where distortion of the visualization may be undesirable, overlay regions [5, 35, 37] and varied levels of transparency [26] can be used to visually separate the detail region from its surrounding context.

#### *Design Considerations Related to Ameliorating Discrepancies among Visualization, Output, and Input Resolutions:*

**How to utilize zooming or magnification techniques in the application.** Possibilities include semantic zoom [2], non-distortion lenses or overlays (e.g., [35]), semi-transparent lenses (e.g., [26]), fisheye views (e.g., [10, 31]), or stretchable surfaces (e.g., [33]).

**How to control the level of magnification.** Possibilities include fixing the magnification level to a single value (e.g., [36]), providing “zoom in” and “zoom out”-style interaction widgets (e.g., [9]), allowing lenses and/or area-of-interest indicator objects to be directly resized (e.g., [17]), or basing magnification level on the position of the magnified view on the surface (e.g., [37]).

#### **Supporting Diverse View and Value Interactions Simultaneously**

When interactive tabletop displays utilize a flexible magnification approach, such as a focus+context lens, the level of magnification is just one of many *view* parameters [6] that may need to be controlled by the collaborators around the table. Each additional degree of freedom offered by an interaction technique, such as the position of the lens on the table, needs to map to a unique input mechanism so that it can be adjusted when necessary.

Additionally, since many of the information visualizations that we wish to support make extensive use of interaction, a major consideration is how best to facilitate interaction with the data *values* [6] without overloading previously used input mechanisms. The design of interaction techniques to support tabletop information visualization must strike a balance between supporting these *view* and *value* interactions.

Selection techniques are highly dependent on the type of input afforded by a display. Olwal et al. [25] developed two families of techniques for zooming and selection on *single-touch* displays: rubbing and tapping. Neither technique uses on-screen cues, but both combine previous techniques that include an offset cursor (a take-off) [28] and zoom selection [1].

For *multi-touch* and vision-tracked input devices, Benko et al. developed a single-finger (SimPress) and several dual-finger techniques including cursor offset and zoom techniques [3]. Empirical studies have confirmed the usefulness of these dual-finger approaches for zoom selection [1]. Pointing Lenses offer another option, using pen pressure as a cue to increase both visual and motor area under the cursor [29].

One well-known approach combines both focus+context lenses and interaction into one technique: Toolglasses and Magic Lenses [4]. These tools can not only visually enhance (e.g., magnify or highlight) underlying information but also allow interaction *through* the lens to create desired effects in the interface.

Gutwin and Skopik showed that navigation and efficiency of interaction were improved by the use of detail-in-context techniques in representations where magnification was required [12]. Nacenta et al. compared the ability of different tabletop interaction techniques to support collaboration and awareness [23].

*Design Considerations Related to Supporting Diverse Interaction Techniques:*

**How to enable interaction with the magnified values.** Possibilities range from allowing only passive viewing of the data (e.g., [26]) to tool-based interaction with pre-determined functionality (e.g., [4]) to direct interaction through the lens using the same techniques available elsewhere on the surface (e.g., [17]). Several of these options might be combined into a single interaction technique.

**How to control the parameters of the magnified view (e.g., its position and size)** while supporting interaction with the lens' contents. Possibilities include: fixing one or more of the parameters (e.g., tracking the interaction point with the focus visualization [29]), providing interface widgets to toggle among interaction modes (e.g., [9]), or leveraging multi-finger or multi-touch input (e.g., [1, 3]).

#### **Facilitating Face-to-Face Collaboration**

Supporting face-to-face collaboration is a key aspect of tabletop computing that distinguishes it from previous

interaction paradigms. Recent work has explored the role of information visualization on touch-sensitive devices with a focus on collaboration [14]. This research highlights the need for collaborators to easily share their discoveries or views of a data set.

Two of the most significant challenges in supporting face-to-face collaboration on interactive tabletops are that group members *share* a common visualization surface and that they can sit or stand at various locations *around* a table [19]. As a result, tabletop applications cannot be constructed using the single-user interaction paradigms established on the desktop, nor can they assume a direction that will consistently be perceived as “up” among all people using the system. While these design constraints do not adversely impact informal applications like photo sharing (e.g., [34]), they become problematic when trying to support structured collaborations, such as scientific data exploration.

*Design Considerations Related to Facilitating Face-to-Face Collaboration:*


**Whether to treat magnification and rotation as surface-wide or portal-based interactions.** Possibilities include treating the entire surface as a single, stretchable and/or rotatable sheet (e.g., [33]), providing tools to manipulate the scale and orientation of individual objects (e.g., [20]), or providing individual magnification lenses (when needed) for each collaborator (e.g., [35]).

**How to support multiple orientations of the magnified region** to support participants standing around different sides of the table. Possibilities include allowing free focus rotation (e.g., [27]), automatically rotating the magnified region based on its location on the table (e.g., [34]), or off-loading the magnified region onto a secondary or handheld device (e.g., [32]).

Based on these design considerations, we created two interaction techniques, each comprising a unique combination of features not seen in previous systems. Although we acknowledge that our techniques draw extensively from aspects of existing techniques and systems, we believe that the novel combination of design decisions better situate our techniques for supporting scientific information exploration, balancing among the many nuanced trade-offs inherent in designing for interactive tabletop displays.

#### **THE I-LOUPE INTERACTION TECHNIQUE**

*i-Loupe* is a lens-based interface that provides visual magnification of a region of interest, as well as an enlarged surface for affecting arbitrary, touch-based interactions with the underlying objects. The *i-Loupe* has two principal components: a *base* that indicates which region of the visualization is selected for interaction, and a *focus* where the chosen visualization region is made available for more



**Figure 2. The two primary components of the i-Loupe lens, the *base* and the *focus*, each defined by its center, size, and rotation.**

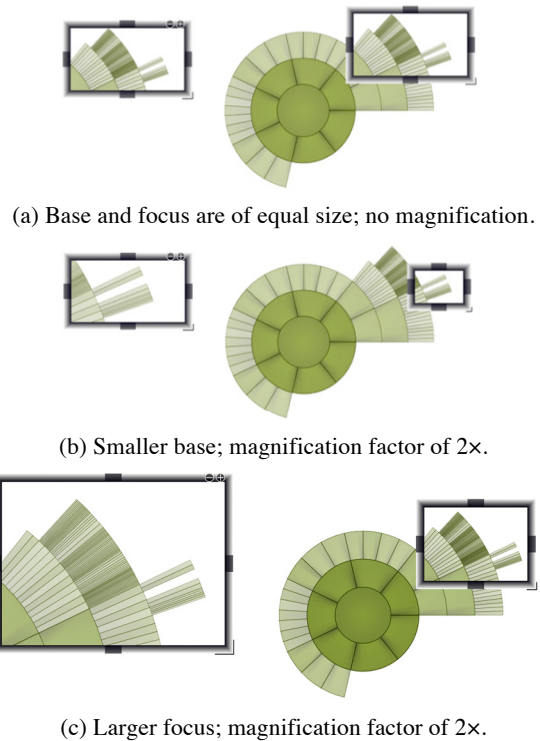
detailed interaction (Figure 2).<sup>1</sup> The i-Loupe serves as a portal into the visualization; it enables all of the touch-based interactions available on the underlying visualization, but at a higher resolution. In the language of the previous design considerations, the i-Loupe features *non-distorting magnification* with some properties of *semantic zoom*, primarily utilizes *direct-manipulation* and *relative component sizing* to control the level of magnification, enables *full touch-based interaction* with the underlying visualization, uses a *frame-and-content* metaphor to provide a distinction between *view* and *value* interactions, and facilitates collaboration by providing *portal-based*, *free-focus* rotation capabilities.

We incorporated the i-Loupe into an existing toolkit for creating information visualizations on large, interactive displays [15] and have included the technique in several scientific visualization applications.

#### i-Loupe and the Resolution Discrepancy Problem

In its simplest form—when the base and focus are the same size—the i-Loupe provides a straightforward mechanism for duplicating the visualization contents beneath the base region to a more conveniently located display position (or space), similar to the Shadow Boxes technique [35]. In this case, the magnification factor between focus and base is 1 $\times$ , as shown in Figure 3(a). Changing the focus or base size results in changes to the focus magnification, as illustrated in Figures 3(b) and 3(c). The degree of magnification provided in the focus is computed as the ratio between the widths of the two regions. (In our implementation, base and focus components’ aspect ratios are fixed.)

This interaction provides a direct magnification control where the interaction is coupled with the magnification,



**Figure 3. i-Loupe magnification is based on the relative size between the base and focus.**


thereby enabling an approximation of the degree of magnification at a glance (Figure 3). However, this magnification control is insufficient in itself. That is, for truly massive data sets or for a densely packed region of a visualization, there is not always enough screen space available to sufficiently enlarge the focus to support touch interaction within the focus at a pixel resolution. As a result, we augmented the focus size magnification interaction with “+” and “−” controls on the focus’ border. These controls allow the magnification ratio to be increased or decreased multiplicatively, making it possible to quickly reach extremely high magnification factors (e.g., 80 $\times$  normal resolution). To prevent the base from shrinking out of visibility at extreme magnifications, a pre-established minimum base size is enforced. Beyond this size, the areas in the base visualization that are no longer shown in the focus are de-saturated and reduced in brightness.

#### i-Loupe and Interactive Information Visualizations

With the magnification provided by the i-Loupe, touch-based interaction is less awkward. Figure 4 shows how the i-Loupe enables the selection of a graph node through the focus that would otherwise be too small to touch. Any interaction that is available within the context visualization is also available through the focus. Thus, fine-grained interaction with, organization of, and adjustments to the visualization can be performed in the focus, with the results instantly reflected in the surrounding context visualization.

Conversely, one might wish to make annotations that are not always fully visible, so as to avoid obstructing important areas of the visualization. In this case, choosing a

<sup>1</sup> The figures in this section are screen captures from an application that provides space-filling radial visualizations of clustered gene expression data that we developed together with biologists at our institution. To reduce visual clutter, we have omitted node labels.



**Figure 4. i-Loupe selection in the focus region with a 35x magnification factor.**

more extreme magnification level for the focus results in an annotation that is visible as only one or two pixels in the context visualization (Figure 5). Since we prevent our vector-based annotation markings from rendering smaller than a single pixel at normal magnification, they are in effect *elided*—filtered with residual evidence of their presence. Even with a small presence in the context visualization, their content is still readily retrievable with the i-Loupe.


#### **i-Loupe and Face-to-Face Collaboration around Tabletop Displays**

We designed the i-Loupe to be orientation-agnostic in order to foster co-located collaboration around the tabletop surface. The border regions of the i-Loupe base and lens components both feature a combination of *translate-only* and *rotate-n-translate* [20] interaction areas, making it easy to quickly re-position or re-orient the components to experience a new perspective on the underlying visualization.

If one member of a group is examining a visualization detail and wants to share a discovery with a colleague, he or she can create an i-Loupe instance, move the new base to the region of interest, and pass the focus over to the other person. Here, all the considerations of orientation as a communication facilitator come into play, and the focus can be passed using the lens' integrated rotation and translation capabilities [20]. This design allows the focus to be passed to the colleague in the orientation best suited for their viewing. Based on this interaction, both people can view the same visualization segment (using multiple i-Loupe instantiations, if necessary), even though they may be standing on different sides of a shared tabletop surface.

#### **IPODLOUPE, A HANDHELD, MULTI-TOUCH INTERACTION LENS**

While the i-Loupe provides a valuable set of affordances for interacting with high-resolution visualization data in a collaborative context, there are still some drawbacks to the approach. The most notable limitation is that the i-Loupe's lens component obscures at least some display space to provide an interaction portal. Consequently, we created an alternate version of the i-Loupe that illustrates different



**Figure 5. Creating an annotation with the i-Loupe that is unobtrusive—but still visible—at normal scale.**

choices from among our design considerations. In this technique, the lens component is off-loaded onto a multi-touch handheld computing device. We call this lens an *iPodLoupe* because it is based around the interaction and networking capabilities of an Apple iPod Touch device.


Unlike prior systems that coupled a handheld device with a large display (e.g., Chameleon [8], Total Recall [13], the a-Book Interaction Lens [22], or Ubiquitous Graphics [32]), iPodLoupe is not an augmented reality system that seeks to overlay virtual content on the physical world. In fact, the iPodLoupe application does not seek to track the iPod Touch's absolute position in 3D space at all, which reflects a design decision that we explicitly made in the interest of supporting face-to-face collaboration around our interactive tabletop surface. With iPodLoupe, the position, magnification, and orientation of the high-resolution viewport into the information visualization are controlled through explicit interactions on the iPod device or by manipulating the corresponding base component on the tabletop surface. In contrast to the a-Book's Interaction Lens [22], for example, this design allows a collaborator to pass an iPodLoupe across the table to her colleague without disrupting the magnified region of interest that is displayed on the device's screen.

#### **iPodLoupe and the Resolution Discrepancy Problem**

The iPodLoupe replicates the functionality of its tabletop-based counterpart, with two main differences: the focus is rendered only on the handheld device, and the handheld component of the system takes advantage of the multi-touch and accelerometer-based input capabilities of the iPod Touch. The rendering of the base is modified to more closely represent the form factor of an iPod and the entire "frame" around this virtual device's screen can be used to drag the base around the table, enabling quick changes to the location and orientation of the magnified view. The other interface components used in the i-Loupe lens (e.g., the "+" and "-" controls on the focus frame) are also included in the iPodLoupe's lens interface (Figure 6).

Like the i-Loupe, the iPodLoupe addresses the resolution discrepancy problem by providing a display and interaction surface at a higher resolution than the native tabletop





**Figure 6. The iPodLoupe lens, overlaid on a tabletop displaying several PhylloTree visualizations [24]. Although similar to the i-Loupe in function and design, iPodLoupe differs in two ways: (1) the focus component exists only on the handheld, and (2) the iPod enables gestural, multi-touch input.**

surface. When a small region of interest from the tabletop visualization is displayed on the device, it is possible to see more detail and manipulate items at a much finer scale than is possible on the tabletop surface. Since the display size of the handheld device is fixed, the relative magnification of the lens system is determined solely by the size of the base component on the interactive tabletop display.

#### **iPodLoupe and Interactive Information Visualizations**

Not only does iPodLoupe minimize the degree of visual distraction on the tabletop display by off-loading the lens' focus region to a secondary device, but the multi-touch and physical sensing capabilities of the iPod Touch also help to give primacy to the interactive capabilities of the underlying visualization. The iPodLoupe interface includes only three UI widgets: “+” and “-” buttons for rapidly zooming in and out and a toggle-style button to enable or disable annotation mode. Single-touch inputs anywhere else on the device's screen are captured and sent over the wireless network to the tabletop to be relayed on to the visualization application, providing a nearly full-screen surface for directly interacting with the visualization data.

The iPodLoupe's position, orientation, and magnification are controlled by multi-touch gestures and sensed physical inputs (i.e., accelerometer-measured device orientation). Panning the viewport is accomplished by dragging two fingers side-by-side across the screen in any direction or by tilting the entire device in the palm of the hand (e.g., [30]). The viewport orientation can be changed by rotating two fingers around an imaginary pivot point. Adjustments to magnification are made using the “pinch” and “stretch” gestures commonly found in other multi-touch systems.

Additionally, the iPodLoupe base component on the tabletop allows all of the same manipulations as its i-Loupe counterpart. This combination of design features take advantage of many of the unique affordances of the iPod Touch hardware and provide an enormous degree of

flexibility in specifying the area of interest and degree of magnification. However, the most direct form of interaction with the device—single-touch input—is dedicated to interacting with the data, not re-directed to control the configuration of the lens.

#### **iPodLoupe and Face-to-Face Collaboration around Tabletop Displays**

Because iPodLoupe off-loads the lens component onto the iPod device, the magnified area of the visualization can be easily manipulated as a tangible artifact above the tabletop display surface. This design enables side-by-side comparison of the visualization data and co-located collaboration without requiring any particular interface support, since the device itself can be passed among collaborators and re-oriented as needed.

To prevent confusion about who controls each iPodLoupe base component, we color-coded each of the iPodLoupe base-lens pairs: the color of each base component matches the border and UI widgets rendered on each iPod display. This design helps people around the table identify individual activity based on the color in which each base component is rendered.

#### **EMERGENT INTERACTIONS**

Although we explored particular decisions from among our design considerations as a consequence of the specific tasks that we wanted to support, we observed a number of unanticipated interactions that our techniques support as a result of these decisions. The presence of these emergent interactions suggests that the nuanced design decisions and trade-offs in this domain do, in fact, significantly affect the overall user experience.


#### **The i-Loupe as Both Precision Instrument and Collaboration Overview**

The relative magnification approach used by both the i-Loupe and iPodLoupe enables quick adjustment of the focus region's magnification factor and, as a result, the level of display detail and touch-based input precision that can be achieved. However, relative magnification also enables additional capabilities.

While testing our i-Loupe prototype, we discovered that the lens' base component can be quickly enlarged using the “-” button, allowing it to grow larger than the focus. If the context visualization extends beyond the available screen space, the i-Loupe can be used as an interactive overview window (Figure 7), providing similar capabilities as Radar Views, which have been shown to be powerful tools for collaboration [11].

#### **The iPodLoupe as a Zoomable User Interface Appliance**

When exploring an information space using iPodLoupe on an iPod Touch, the multi-touch pan-and-zoom navigation capabilities provide the same “infinite canvas” experience of zoomable user interface (ZUI) systems like Pad++ [2]. Additionally, because our rendering engine creates a visual display at the highest possible resolution on the handheld at any zoom level, the experience of zooming into an object to



**Figure 7. Magnification factors less than 1 $\times$  turn the i-Loupe into an interactive overview. In this image, a factor of 0.31 $\times$  is used. The focus is at the bottom and the base is on top.**

gain progressively more detail resembles semantic zooming.

While we do not support all of the advanced features of dedicated ZUI platforms (e.g., hyperlinks, bookmarks, and animated transitions between views), our multi-device implementation does offer one important advantage over interacting with a ZUI on a desktop. Because the iPodLoupe always displays a subset of the larger visualization space displayed on the tabletop, it is much more difficult to become lost in the infinitely zoomable interaction space displayed on the iPodLoupe. If one becomes disoriented, she can refer to the iPodLoupe's base component on the table to re-orient herself. The base component on the tabletop can also be used to resize and re-position the area of focus without having to first zoom all the way out on the handheld.

## CONCLUSIONS AND FUTURE WORK

We provide three main contributions in this paper.

First, we describe a set of practical design considerations for interactive lenses from which future visualization tools and interaction techniques might be designed. These design considerations articulate decisions that must be made in three primary areas: resolving resolution discrepancies between the visualized data, the display output, and the touch-sensitive input layers; supporting a variety of view and value interaction techniques; and facilitating face-to-face collaboration. Since different mappings of these parameters to interface elements can produce drastically different experiences, we anticipate that these design considerations might serve as a useful resource for visualization designers when considering the trade-offs inherent in different visualization applications or using different underlying display and input technologies.

Second, we demonstrate two different interaction techniques, i-Loupe and iPodLoupe, that both provide holistic solutions to the challenges of supporting co-

located, collaborative information exploration on large, interactive tabletop displays. Each of these systems illustrate different design choices from among our design considerations, resulting in different user experiences, but enable the same functional capabilities:

- exploration of detailed areas of interest within a large, high-resolution visualization;
- flexible levels of magnification for a variety of information exploration tasks;
- direct, general-purpose touch interaction and annotation of elements of the visualization at resolutions far finer than the resolution of a fingertip; and
- comparison of multiple visualization regions, even at different orientations or when displayed on a tabletop surface lacking a consistent notion of “up.”

Finally, we present systems that more closely match scientists' information exploration practices than existing approaches. Our systems achieve this goal due to the design decisions we made: allowing direct interaction with and annotation of fine-grained data, providing flexible rotation for co-located collaboration and side-by-side data comparisons, and a suite of interactions for controlling the lens parameters that do not consume much visual space or require use of a fixed number of input devices (in the case of i-Loupe).

This research represents an important step in overcoming some of the limitations of current tabletop systems in supporting real-world work. While the design of our techniques draw heavily from different aspects of existing systems, much of the research value is in the articulation of significant design considerations as a resource for creating more flexible, robust, and practical user experiences on interactive tabletop displays. We do not claim that these are the *only* design considerations that might play a role in the design of collaborative information visualizations on interactive tabletop displays, nor have we been able to present an exhaustive analysis of the relative strengths and weaknesses of the design solutions to these challenges in this short conference paper. We look forward to examining this design space in detail and unpacking the more nuanced implications of various design decisions on individual and collaborative use of tabletop-based information visualization tools as part of our future work.

## ACKNOWLEDGEMENTS

This research was supported in part by the NSERC/iCORE/Smart Technologies Chair in Interactive Technologies.

## REFERENCES

1. Albinsson, P.-A. and Zhai, S. High precision touch screen interaction. In *Proc. CHI 2003*, ACM Press (2003), 105–112.
2. Bederson, B.B., Hollan, J.D., Perlin, K., Meyer, J., Bacon, D. and Furnas, G. Pad++: A zoomable graphical sketchpad for exploring alternate interface physics. *Journal of Visual Languages and Computing*, 7, 1 (1996), 3–32.

3. Benko, H., Wilson, A.D., and Baudisch, P. Precise selection techniques for multi-touch screens. In *Proc. CHI 2006*, ACM Press (2006), 1263–1272.
4. Bier, E.A., Stone, M.C., Pier, K., Buxton, W., and DeRose, T.D. Toolglass and Magic Lenses: The see-through interface. In *Proc. SIGGRAPH 1993*, ACM Press (1993), 73–80.
5. Carpendale, M.S.T. and Montagnese, C. A framework for unifying presentation space. In *Proc. UIST 2001*, ACM Press (2001), 61–70.
6. Chi, E.H. and Riedl, J.T. An operator interaction framework for visualization systems. In *Proc. InfoVis 1998*, IEEE Computer Society (1998), 63–70.
7. Cockburn, A., Karlson, A., and Bederson, B.B. A review of overview+detail, zooming, and focus+context interfaces. *ACM Computing Surveys*, 41, 1 (2008), Article 2.
8. Fitzmaurice, G.W. Situated information spaces and spatially aware palmtop computers. *Communications of the ACM*, 36, 7 (1993), 39–49.
9. Forlines, C. and Shen, C. DTLens: Multi-user tabletop spatial data exploration. In *Proc. UIST 2005*, ACM Press (2005), 119–122.
10. Furnas, G.W. Generalized fisheye views. In *Proc. CHI 1986*, ACM Press (1986), 16–23.
11. Gutwin, C., Greenberg, S. and Roseman, M. Workspace awareness support with radar views. In *Ext. Abstracts CHI 1996*, ACM Press (1996), 210–211.
12. Gutwin, C. and Skopik, A. Fisheyes are good for large steering tasks. In *Proc. CHI 2003*, ACM Press (2003), 201–208.
13. Holmquist, L.E., Sanneblad, J., and Gaye, L. Total Recall: In-place viewing of captured whiteboard annotations. In *Ext. Abstracts CHI 2003*, ACM Press (2003), 980–981.
14. Isenberg, P. and Carpendale, S. Interactive tree comparison for co-located collaborative information visualization. *IEEE Trans. Visualization and Computer Graphics*, 12, 5 (2007), 1232–1238.
15. Isenberg, T., Miede, A., and Carpendale, S. A buffer framework for supporting responsive interaction in information visualization interfaces. In *Proc. C5 2006*, IEEE Computer Society (2006), 262–269.
16. Keahey, T.A. The generalized detail-in-context problem. In *Proc. InfoVis 1998*, IEEE Computer Society (1998), 44–51.
17. Khan, A., Fitzmaurice, G., Almeida, D., Burtnyk, N., and Kurtenbach, G. A remote control interface for large displays. In *Proc. UIST 2004*, ACM Press (2004), 127–136.
18. Kosara, R., Miksch, S. and Hauser, H. Semantic depth of field. In *Proc. InfoVis 2001*, IEEE Computer Society (2001), 97–104.
19. Kruger, R., Carpendale, S., Scott, S.D., and Greenberg, S. Roles of orientation in tabletop collaboration: Comprehension, coordination and communication. *Computer Supported Cooperative Work*, 13, 5–6 (2004), 501–537.
20. Kruger, R., Carpendale, S.T., Scott, S.D., and Tang, A. Fluid integration of rotation and translation. In *Proc. CHI 2005*, ACM Press (2005), 601–610.
21. Leung, Y.K. and Apperley, M.D. A Review and taxonomy of distortion-oriented presentation techniques. *ACM Trans. Computer-Human Interaction*, 1, 2 (1994), 126–160.
22. Mackay, W.E., Pothier, G., Letondal, C., Bøegh, K., and Sørensen, H.E. The missing link: Augmenting biology laboratory notebooks. In *Proc. UIST 2002*, ACM Press (2002), 41–50.
23. Nacenta, M.A., Pinelle, D., Stuckel, D., and Gutwin, C. The effects of interaction technique on coordination in tabletop groupware. In *Proc. GI 2007*, CHCCS (2007), 191–198.
24. Neumann, P., Carpendale, M.S.T., and Agarawala, A. PhylloTrees: Phyllotactic patterns for tree layout. In *Proc. EuroVis 2006*, Eurographics Association (2006), 59–66.
25. Olwal, A., Feiner, S., and Heyman, S. Rubbing and tapping for precise and rapid selection on touch-screen displays. In *Proc. CHI 2008*, ACM Press (2008), 295–304.
26. Pietriga, E. and Appert, C. Sigma Lenses: Focus-context transitions combining space, time and translucence. In *Proc. CHI 2008*, ACM Press (2008), 1343–1352.
27. Pinelle, D., Stach, T., and Gutwin, C. TableTrays: Temporary, reconfigurable work surfaces for tabletop groupware. In *Proc. Tabletop 2008*, IEEE Computer Society (2008), 45–52.
28. Potter, R.L., Weldon, L.J., and Schneiderman, B. Improving the accuracy of touch screens: An experimental evaluation of three strategies. In *Proc. CHI 1988*, ACM Press (1988), 27–32.
29. Ramos, G., Cockburn, A., Balakrishnan, R., and Beaudouin-Lafon, M. Pointing Lenses: Facilitating stylus input through visual- and motor-space magnification. In *Proc. CHI 2007*, ACM Press (2007), 757–766.
30. Rekimoto, J. Tilting operations for small screen interfaces. In *Proc. UIST 1996*, ACM Press (1996), 167–168.
31. Robertson, G.G. and Mackinlay, J.D. The Document Lens. In *Proc. UIST 1993*, ACM Press (1993), 101–108.
32. Sanneblad, J. and Holmquist, L.E. Ubiquitous Graphics: Combining hand-held and wall-size displays to interact with large images. In *Proc. AVI 2006*, ACM Press (2006), 373–377.
33. Sarkar, M., Snibbe, S.S., Tversky, O.J., and Reiss, S.P. Stretching the rubber sheet: A metaphor for viewing large layouts on small screens. In *Proc. UIST 1993*, ACM Press (1993), 81–91.
34. Shen, C., Vernier, F.D., Forlines, C., and Ringel, M. DiamondSpin: An extensible toolkit for around-the-table interaction. In *Proc. CHI 2004*, ACM Press (2004), 167–174.
35. Tang, A., Tory, M., Po, B., Neumann, P., and Carpendale, S. Collaborative coupling over tabletop displays. In *Proc. CHI 2006*, ACM Press (2006), 1181–1190.
36. Vogel, D. and Baudisch, P. Shift: A technique for operating pen-based interfaces using touch. In *Proc. CHI 2007*, ACM Press (2007), 657–666.
37. Ware, C. and Lewis, M. The DragMag image magnifier. In *Proc. CHI 1995*, ACM Press (1995), 407–408.