Using Spatially Aware Tangible Displays for Exploring Virtual Spaces

Martin Spindler, Jana Sieber, Raimund Dachselt

Faculty of Computer Science, Otto-von-Guericke University, Magdeburg, Germany

Abstract

To solve the challenge of exploring large information spaces on interactive surfaces such as tabletops, we introduce *SpaceLens*, a spatially aware handheld display that provides elegant three-dimensional exploration of rich datasets. This can either be volumetric data, a layered information space or a zoomable information space, which are mapped to the virtual exploration volume above a stationary surface. By moving the lens through the volume, corresponding data is displayed, thus serving as a window into virtuality. Various interaction techniques are introduced, which especially utilize the Z-axis (lens height) in a novel way, e.g. for zooming or displaying various information layers. The SpaceLens implementation uses an optically tracked, passive, top-projected paper lens, which is cheap, lightweight and flexible. A formative user study gave valuable insight and confirmed an intuitive interaction. In addition, a marker-tracked, active UMPC lens allows data exploration without a contextual display.

1 Introduction

The past years have seen rapid improvements in research and technology for large displays, interactive surfaces, such as walls or tabletops, and associated natural interaction techniques using hand gestures, multitouch, interactive pens, or tangibles. Simultaneously, mobile and miniaturized devices allow people to move away from stationary desktop user interfaces to a ubiquitous everyday use. However, even though research approaches exist that combine small personal devices with large displays, this area is far from being fully explored. In almost any case, the interaction is constraint to a single – small or large – 2D display.

At the same time, information spaces are ever growing and become more and more complex. This results in challenging data filtering and exploration tasks, such as planning a brain surgery using medical imagery or working with geographic information systems. Often, various information layers need to be handled, and the early idea of *Magic Lenses* (Bier et al. 1993) is one promising solution for it, which has been extended to more advanced detail & context techniques. Two approaches using handheld displays to explore complex data shall be mentioned as examples: *Tangible Magic Lenses* (Ullmer & Ishii 1997) and *Peephole Displays* (Yee 2003). Both serve as windows into Virtuality, one with a contextual display and the other without. However, these and comparable systems only focus on certain aspects of spatial data exploration, do rarely use the 3D space above a display and often employ complicated or heavy hardware (e.g. tablet PCs).

To improve on that, we developed *SpaceLens*, a lightweight handheld display solution that enables the user to explore virtual information spaces in a more natural manner. These can consist of volumetric data or 2D image data as part of large continuous information spaces. To achieve this, the real 3D space above a stationary surface (e.g. a tabletop) is used as a virtual 'container' for the data to be explored. This will be done by moving the tracked SpaceLens through the interaction volume, which results in a change of the displayed content. We developed two SpaceLens prototypes: first, a tech-free, passive, top-projected paper lens as a cheap, lightweight and flexible solution affording various shapes (see Figure 1) and second an active UMPC lens, which does not require a display context at all. Considering the envisioned application domains and feedback from a first user study, we assume that tangible SpaceLenses will allow an elegant and smooth spatial exploration of virtual spaces.

The remaining paper is structured as follows. After a discussion of related work, the SpaceLens concept is introduced in Section 3 along with interaction techniques for exploring volumetric, abstract and large continuous data. In Section 4, two technically different implementations of SpaceLenses are presented. This is followed by a section that introduces various application scenarios which benefit from SpaceLenses. Section 6 reports on the conducted user study. The paper is concluded with a discussion and an outlook on future plans.

Figure 1: Exploration of a layered information space with the paper lens. Depending on the height of the paper lens, different layers of information are displayed, i.e. skeleton (left), muscles (middle) and nerves (right).

2 Related Work

The use of spatially aware displays for the exploration of virtual spaces has been a field of study for more than a decade. As one of the first approaches, *Tangible Magic Lenses* were introduced with the *metaDESK* system by (Ullmer & Ishii 1997). Besides a passive lens that is bound to the tabletop surface, an active lens consisting of a flat-panel LCD was proposed for 3D interaction. Thereby, the Z-axis is simply utilized for zooming in a 3D geospace. In addition, the movement of the active lens is restricted by its attachment to a mechanical arm.

(Lee et al. 2005) came up with a more flexible, handheld lens solution in form of either a passive, projective display or a tablet PC. Similar to our work, the developers suggest using the lens to explore 2D datasets with multiple layers, e.g. geographic information data. Yet, a 3D exploration of this data is not proposed. A tablet PC is also used in the work of (Sanneblad & Holmquist 2006). We picked up their idea of using a lens to access details that are unrecognizable in a large image, even on large screens. However, they do not exploit the space in front of the wall display to allow for controlling the level of detail. (Brown & Hua 2006) developed an embedded, surface-bound lens as well as a freely movable lens for augmented virtual environments. Their solution was an important step because the user is given more choices on the lens effect. Again, changes in the distance of the lens to the tabletop are not taken into consideration. Instead, a less natural interaction via buttons and menus on the lens is provided to allow for zooming, clipping or layering the context.

All mentioned solutions require the users to hold rather heavy devices. Two recent technical approaches allow a more lightweight solution. *UlteriorScape* (Kakehi & Naemura 2008) uses a simple paper lens in various shapes and sizes. It is optically tracked and back-projected through the tabletop display underneath by using a Lumisty film. Unlike our lens, only the user's position at the tabletop but not the height of the lens above the surface influences what is displayed. The *SecondLight* solution (Izadi et al. 2008) uses a back-projection of both the surface and the lens, too, but by applying a shutter concept and multiple projectors instead. Similar to our approach, the developers suggest the utilization of height changes for zooming or altering rotation to reveal new layers. While both approaches were inspiring for our solution, we aimed at a considerably less complex technical solution that is more affordable.

The work that resembles our intention most closely is the *peephole display* (Yee 2003), where the virtual information space around a user can be explored with a tracked PDA. With it, the metaphor of the Z-axis is taken into account, e.g. for zooming and layering, but the focus is only on common desktop applications, such as a calendar. In summary, none of the available solutions allows a spatially aware display to truly utilize the 3D space above a reference surface in order to provide a rich set of natural 3D exploration techniques for virtual spaces, such as 3D volumetric data or large zoomable 2D information spaces.

3 The SpaceLens Concept

In this section, we introduce *SpaceLenses* as spatially aware displays that are suitable for exploring virtual spaces in a more natural manner. Similar to a suggestion of (Fitzmaurice 1993), our principle idea is to find a set of meaningful mappings that translates the movement of a SpaceLens in real world space (the *interaction space*) into an intended action in a virtual space (the *exploration space*). For SpaceLenses, navigation and zooming techniques will be proposed in this section, and navigation aids will be briefly discussed.

3.1 Spaces and Mappings

We specify the *interaction space* as the continuous three-dimensional real-world space where the user holds and moves the SpaceLens in order to explore parts of the exploration space. The interaction space is defined by a two-dimensional *reference surface* that is aligned to the X_i- and Y_i-axis of the *interaction space coordinate system* (see green-colored coordinate axes in Figure 2). A third dimension is described by the Z_i-axis pointing into the direction perpendicular to the reference surface. Typically, the interaction space could be the space above a horizontal surface such as a tabletop or desk, but vertical or slanted reference surfaces are also possible (e.g. a wall or a speaker's desk). Please note that the reference surface can also provide a visual context that helps guiding the user during the interaction process.

With the *exploration space* we describe a virtual three-dimensional volumetric space that is virtually “filled” with the data the user intends to explore. According to the type and extent of the data, we distinguish between three different classes of exploration spaces: *virtual three-dimensional volumes*, *layered information spaces* and *zoomable information spaces*. In order to let the user spatially explore these classes in a more natural way, appropriate mappings from the interaction space coordinate system to the exploration *space coordinate system* (see red coordinate axes in Figure 2) need to be found. In general, all three types can be similar in the X_eY_e-space, but may differ conceptually in the third dimension (Z_e-axis). In the following we will discuss each class and its corresponding mappings in more detail.

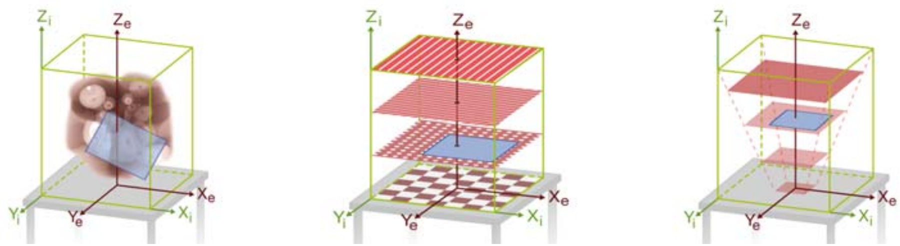

(a) Virtual 3D Volumes *(b) Layered Information Spaces* *(c) Zoomable Information Spaces*

Figure 2: Structural sketch of the different spaces and coordinate systems. The interaction space is depicted in green (XiYiZi), the associated exploration space in red (XeYeZe). Blue rectangles show a possible lens position.

VIRTUAL 3D VOLUMES. A *virtual 3D volume* is a set of 3D samples (voxels) with a volumetric nature. Typical examples are datasets acquired via CT or MRI techniques used in the realm of medical or scientific visualization. Volumetric data exhibit a continuous form in all three dimensions and thus allow for a direct linear mapping from the interaction space ($X_iY_iZ_i$) onto the exploration space coordinate system ($X_eY_eZ_e$), see Figure 2a.

LAYERED INFORMATION SPACES. We define a *layered information space* as a set of two-dimensional information layers (raster data and/or vector data) with each layer representing a unique feature of a model. Layered information spaces are commonly employed in various application fields, such as geographic information systems or interactive maps. In general, these spaces are continuous within a 2D plane and thus allow for a direct linear mapping from the X_iY_i-space onto the X_eY_e-space (see Figure 2b). The mapping from Z_i to Z_e is not as straightforward. We can obtain a non-linear mapping by dividing the Z_e-axis into discreet intervals and by associating these intervals with distinct layers. This way, we construct a “volumetric” layered information space - with each layer covering a distinct “height” in the exploration space - that can be mapped directly onto the interaction space.

ZOOMABLE INFORMATION SPACES. Based on the concept of space-scale diagrams (Furnas & Bederson 1995) we define the *zoomable information space* as a paradigm for the exploration of large 2D information worlds where zooming and panning plays an essential role (e.g. gigapixel images or maps). We can easily map the interaction space onto the exploration space in such a way that the pyramid’s Z_e-axis representing the level of detail/zoom is aligned with the Z_i-axis (see Figure 2c). The Z-axis is thereby interpreted as a scaling function and the XY-plane describes the position in 2D space as required for panning operations.

3.2 Navigation and Zooming Techniques

With the two spaces and their mappings defined we can now propose a *SpaceLens* held by the user in the interaction space serving as a window into the exploration space. A lens as a spatially aware display typically provides six degrees of freedom (DOF) that describe location and orientation in the interaction space. For the purpose of navigation and zooming, we can make use of this information by directly mapping the six values onto the exploration space. In order to do so, we identified three techniques that restrict the degree of freedom depending on the intended task: *depth translation*, *XY-translation*, and *tilting*.

DEPTH TRANSLATION. The *depth translation* describes a change of Z_i that is interpreted differently depending on each of the three types of exploration spaces. For the first type (a, virtual three-dimensional volumes) this simply leads to a *shift of height*, such as a selection of a different slice in a medical volume dataset. For the second type (b, layered information spaces), a modified height translates to a *change of the selected layer*, e.g. from “satellite image” to “street map”, when above the assigned space range. For the third type (c, zoomable information spaces), the depth translation describes a *zooming operation.* When the user intends to examine a detail of the dataset (e.g. a high-resolution photograph), he or she simply needs to lift up the SpaceLens closer to her face.

XY-TRANSLATION. We define the *XY-translation* as being restricted to the X_iY_i-space, i.e. all planes parallel to the reference surface with the simplest case of the surface itself. By directly mapping the X_iY_i- onto the X_eY_e-space, this technique is well-suited for all kinds of panning operations with only minor differences for the exploration space types. For types (a) and (b), a *shift in the X_eY_e-plane* is equivalent to navigating through a single slice of a volume (e.g. the Visual Human) or an information layer (e.g. a street map). Although this is also true for the third type (c), here the level of zoom (depending on the depth-translation) determines the scaling factor of the mapping between the X_iY_i and X_eY_e-space.

TILTING. With *tilting* we define a technique that allows the user to horizontally or vertically rotate or “tilt” a SpaceLens in the interaction space in order to study the exploration space in more detail. In particular, this makes sense for the first type (a), virtual three-dimensional volumes, where tilting provides a natural way of how to arbitrarily "slice" through a dataset. How tilting might be utilized for the other two types (b) and (c), depends very much on the specific application type and requires further investigation.

When exploring virtual spaces with a spatially aware display, orientation support is necessary to help the user to not getting lost. Naturally, the extent of the *reference surface* provides a good estimate for the size of the interaction space. It can also serve as a visual context and supplementary tabletop display. With it, the interaction metaphor changes from *interacting in a 3D space* to *interacting "on top" of a contextual display*. For the Z_i-axis, the "lower" boundary is well-defined by the reference surface. For the "upper" boundary human capabilities and eye sight suggest an approximate height of 35 to 45 cm above the reference surface depending on the user's size. The display of the SpaceLens can also be augmented with a *height indicator* in order to provide a better visual feedback for the minimum, maximum and current Z_e-value, i.e. "vertical" position of the lens, within an information space.

4 Implementation

PAPER LENS. With the *paper lens* we introduce a simple and affordable *passive display* solution. This is the first of two solutions serving as a proof of the SpaceLens concept. The general setup consists of a *table*, a *projector* above, an *infrared camera with infrared emitter,* and the *paper lens* (see Figure 3). In this deliberately simple setup, the table serves both as horizontal reference surface and visual context. The paper lens itself is made of a rectangular-shaped piece of pressboard (30x20cm) with three IR-reflective markers (5x5mm) on three corners (see Figure 5) that are tracked with an infrared camera at 100Hz and a 640x480 pixel resolution. By taking the rectangular shape of the paper lens into consideration, the position of the fourth corner can easily be extrapolated. The determination of the corner's Z-coordinates is more difficult, since 2D tracking is not sufficient here. Instead, we determine the triangular-shaped area spanned by the markers and derive the lens' height from the increasing/decreasing size of the area calibrated earlier. This method works very well for horizontally held paper lenses, but needs to be extended to also detect the height of individual corners for tilting. For the work presented here, we use tilting with the paper lens to a limited extent only. With the known position and orientation of the paper lens, the top-projector can be used to not only project images onto the table surface but also onto the paper lens as long as it resides within the projection volume. For that, a dynamically calculated area of the overall projected screen will be replaced with the content to be displayed on the lens.

Although our paper lens implementation appears to be convincing, the passive display approach and marker-based tracking yield some inherent limitations. First of all, the projection and tracking volumes become smaller with an increased height above the table and therefore limit the size of the interaction space. Second, the tracking is only stable as long as the user does not occlude one of the markers with his hands or fingers. Third, the projector's shallow depth of field causes blurring artifacts on the paper lens the higher the lens is held. Another problem is unwanted shadows on the tabletop surface caused by the paper lens. This could be fixed by a second projector underneath the tabletop (leading to a more complex setup).

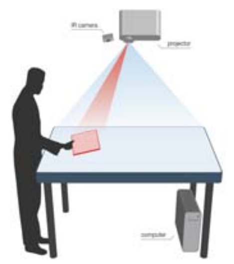

Figure 3: Paper Lens Setup. A top-projector projects the image content onto the lens as well as the reference surface. The paper lens is tracked via an infrared camera.

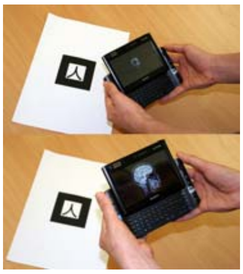

Figure 4: UMPC Lens Setup. The build-in UMPC camera tracks a printed marker (ARTag) and displays slices of a 3D volume depending on the distance to the marker.

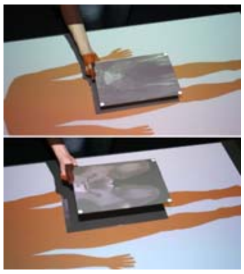

Figure 5: Exploration of virtual 3D volumes with the paper lens. By moving the paper lens in all three dimensions the user can explore the volumetric CT scan of a human.

UMPC LENS. With the *UMPC lens* we implemented an *active display* variant that utilizes a common Ultra Mobile PC with a touch sensitive TFT (1024x600 pixel) and a build-in camera of 640x480 resolution (see Figure 4). We follow the approach of (Fiala & Roth 2007) by using an optical marker-based tracking that is based on the (ARToolKit 2008). For this, we attach printed markers on an arbitrary flat reference surface (e.g. a desk or a wall) or display them dynamically on a tabletop surface. The distance between lens and surface (Z-value) can easily be derived from the reference coordinate system that is provided by the underlying ARToolKit. Compared to the paper lens, the UMPC lens does not restrict the working space to a fixed location, but lacks a visual context, though this restriction does not apply by using dynamic markers on the tabletop display. Furthermore, the UMPC's dimension and weight diminish the user's experience, which is far more "magical and airy" with the paper lens.

5 Application Scenarios

VIRTUAL 3D VOLUMES. For the field of medical visualization, we have tested the exploration of 3D volumetric MRI data with both the paper and the UMPC lens. Here, we follow an idea suggested by (Brown & Hua 2006). As shown in Figure 4, we use the UMPC lens to scroll through a stack of brain scans by lifting and lowering the lens above a visual marker. In Figure 5, we use the paper lens to explore a volumetric CT scan of a human. Here, the tabletop display provides a contextual reference image (the man's silhouette) that helps the user in orientation.

LAYERED INFORMATION SPACES. The SpaceLens can also support medical students in studying the anatomy of the human body. For example, in Figure 1, students examine different layers of information showing the skeletal, muscular or nervous system of a human body. Another promising application is the exploration of layered maps or GIS where users can smoothly navigate and select desired layers by lifting or lowering the lens.

ZOOMABLE INFORMATION SPACES. The *SpaceLens* is also a useful tool to simplify navigation in large 2D information spaces, such as g*igapixel images*. In Figure 6, the paper lens is utilized to show details of a large image. Here, users can pan and zoom in a natural way: by horizontally moving the lens (pan) and lifting and lowering it (zoom). Since the complete image is always visible underneath, the user does not get lost easily. Another promising application is semantic zooming with visual changes depending on zoom level.

Figure 6 Exploration of a zoomable information space with the paper lens. Depending on the height of the paper lens, the level of detail on the lens is changed from small (left) to large (right).

6 Evaluation

We conducted a formative user study in order to gather initial feedback from users working with SpaceLens. Twelve participants (3 female) between the age of 25 and 35 took part. Only three of them had prior experiences with interactive tabletops. The study was conducted with the paper lens prototype described in Section 4 and was videotaped. Two of the authors observed each session and took notes, particularly concerning the think-aloud data.

PROCEDURE AND TASKS. At the beginning, the paper lens was briefly introduced verbally without explaining the interaction modes, which participants had to figure out for themselves. Afterwards, they had to complete five tasks addressing all three classes of exploration spaces. For the layered information spaces participants were asked to count the number of colored dots in two different layers. For virtual 3D volumes, they were requested to find an unknown number of colored stains (tumors) in an MRT scan of a human torso. For zoomable information layers, participants were asked to figure out two details in a gigapixel image that were too small to be recognizable on the tabletop display. The order of the tasks was counterbalanced. Each task was succeeded by a small questionnaire regarding the subjective impressions of how easy the task was to be fulfilled (11 questions in total). At the end, the participants completed a post-test questionnaire on their subjective overall impressions (6 questions) as well as on demographic and computer usage information. All qualitative questions used a 7-point Likert scale from strongly disagree (-3) to strongly agree (3).

RESULTS. All of the 12 participants had been able to instantly accomplish the given tasks almost without errors. Considering that the participants had not been taught how to use the system, we conclude that our technique can be labeled as “intuitive”. This is also supported by the questionnaire, where participants agreed to the statements that “it was easy to learn the system” (2.83, SD=0.39) and “it was easy to use the system” (2.33, SD=0.49). However, some negative feedback was given regarding the layered information space, where the participants found it easier to keep the focus in a layer close to the tabletop (2.25, SD=0.62) than in a higher layer (1.58, SD=1.08). Some participants (n=3) suggested that the selected layer should be “thicker than other layers to make it easier to stay inside”. Our observations also revealed some evidence that a non-linear stacking of layers might be more appropriate than a linear one. Our idea of including navigational indicators was supported by (n=3) participants, who explicitly asked for them. Many participants (n=7) complained about the abrupt changes of layers and some (n=4) suggested to add a blending between layers. An estimated 60% of the participants had problems holding the lens stable when zooming into the gigapixel image. One participant commented that a slight tilting posture seemed to be more natural, especially in greater heights. Some (n=2) participants mainly used one hand for holding the paper lens, which underlines its superiority to the heavier UMPC lens.

7 Conclusion

With *SpaceLens* we have introduced a natural approach for the spatial exploration of volumetric and large 2D information spaces. This was accomplished by mapping data to a virtual space above a stationary surface which serves as a spatial reference and optionally as a contextual display. Thus, the interaction and exploration space are “*physically*” brought together, with the SpaceLens serving as a window into the virtual space. With the introduction of the three basic types of data spaces *virtual 3D volumes*, *layered information spaces* and *zoomable information spaces*, we identified novel ways of accomplishing exploration tasks that are more difficult to solve with traditional 2D interfaces. We developed a set of natural lens navigation techniques, with the special focus on exploiting the Z-dimension of the interaction space. The concept has been implemented in various application scenarios using a UMPC and a paper lens. We favor the latter as it is a cheap, lightweight, flexible, and shape-variable solution allowing an almost “magical” user experience. This was also reported by a number of participants in the formative user study, who enjoyed it and gave valuable feedback.

Acknowledgements

This work was partly funded by the “Stifterverband für die Deutsche Wissenschaft”. We thank Sandro Bosio, Lars Uebernickel, and Ricardo Langner for their help and support.

Contact

Martin Spindler, Otto-von-Guericke University, Institute of Simulation and Graphics, POB 4120, 39016 Magdeburg, Germany.{spindler, sieber, dachselt}@isg.cs.uni-magdeburg.de

References

ARToolKit (2008). *http://www.hitl.washington.edu/artoolkit/*.

Bier, E. A., Stone, M. C., Pier, K., Buxton, W., and DeRose, T. D. (1993). Toolglass and magic lenses: the see-through interface. In *Proceedings of the 20th Annual Conference on Computer Graphics and interactive Techniques*. SIGGRAPH '93. New York, NY: ACM, pp. 73-80.

Brown, L. D. and Hua, H. (2006). Magic Lenses for Augmented Virtual Environments. In *IEEE Comput. Graph. Appl.* Volume 26, Issue 4. Los Alamitos, CA: IEEE Computer Society Press, pp. 64-73.

Fiala, M. (2005). ARTag, a Fiducial Marker System Using Digital Techniques. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cvpr'05) - Volume 2 - Volume 02*. Washington, DC: IEEE Computer Society, pp. 590-596.

Fiala, M. and Roth, G. (2007). Magic lens augmented reality: table-top and augmentorium. In *ACM SIGGRAPH 2007 Posters*. SIGGRAPH '07. New York, NY: ACM, p. 152.

Fitzmaurice, G. W. (1993). Situated information spaces and spatially aware palmtop computers. *Communications of the ACM*. New York, NY: ACM 36, 7, pp. 39-49.

Furnas, G. W. and Bederson, B. B. (1995). Space-scale diagrams: understanding multiscale interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*). I. R. Katz, R. Mack, L. Marks, M. B. Rosson, and J. Nielsen (Hrsg). Conference on Human Factors in Computing Systems. New York, NY: ACM Press/Addison-Wesley Publishing Co., pp. 234-241.

Izadi, S., Hodges, S., Taylor, S., Rosenfeld, D., Villar, N., Butler, A., and Westhues, J. (2008). Going beyond the display: a surface technology with an electronically switchable diffuser. In *Proc. of the ACM Symposium on User interface Software and Technology*. UIST '08. ACM, pp. 269-278.

Kakehi, Y. and Naemura, T. 2008. UlteriorScape: Interactive optical superimposition on a view-dependent tabletop display. In *Horizontal Interactive Human Computer Systems*, 2008. TABLETOP 2008, pp. 189-192

Lee, J. C., Hudson, S. E., Summet, J. W., and Dietz, P. H. (2005). Moveable interactive projected displays using projector based tracking. In *Proceedings of the 18th Annual ACM Symposium on User interface Software and Technology*. UIST '05. New York, NY: ACM, pp. 63-72.

Mackay, W. E., Pothier, G., Letondal, C., Bøegh, K., and Sørensen, H. E. (2002). The missing link: augmenting biology laboratory notebooks. In *Proceedings of the 15th Annual ACM Symposium on User interface Software and Technology*. UIST '02. New York, NY: ACM, pp. 41-50.

Sanneblad, J. and Holmquist, L. E. (2006). Ubiquitous graphics: combining hand-held and wall-size displays to interact with large images. In *Proceedings of the Working Conference on Advanced Visual interfaces*. AVI '06. New York, NY: ACM, pp. 373-377.

Ullmer, B. and Ishii, H. (1997). The metaDESK: models and prototypes for tangible user interfaces. In *Proceedings of the 10th Annual ACM Symposium on User interface Software and Technology*. UIST '97. New York, NY: ACM, pp. 223-232.

Yee, K. 2003. Peephole displays: pen interaction on spatially aware handheld computers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '03. New York, NY: ACM, pp. 1-8.