Information Visualization and Proxemics: Design Opportunities and Empirical Findings

Mikkel R. Jakobsen, Yonas Sahlemariam Haile, Søren Knudsen, and Kasper Hornbæk

Abstract—People typically interact with information visualizations using a mouse. Their physical movement, orientation, and distance to visualizations are rarely used as input. We explore how to use such spatial relations among people and visualizations (i.e., proxemics) to drive interaction with visualizations, focusing here on the spatial relations between a single user and visualizations on a large display. We implement interaction techniques that zoom and pan, query and relate, and adapt visualizations based on tracking of users' position in relation to a large high-resolution display. Alternative prototypes are tested in three user studies and compared with baseline conditions that use a mouse. Our aim is to gain empirical data on the usefulness of a range of design possibilities and to generate more ideas. Among other things, the results show promise for changing zoom level or visual representation with the user's physical distance to a large display. We discuss possible benefits and potential issues to avoid when designing information visualizations that use proxemics.

Index Terms—Proxemics, information visualization, user study, large displays, user tracking, movement, orientation, distance.

Introduction

Information visualization uses interactive graphics to amplify cognition [5]. It can improve many aspects of dealing with large sets data: Visualizations help explore and navigate large information spaces [40], analyze and make discoveries in high-dimensional data [44], and discuss data within on-line communities [52].

Most information visualizations—commercial products and research prototypes alike—are designed for a setting where the user interacts using a mouse on a desktop-sized display. Recent research has explored how visualizations should be designed for non-desktop settings [27], in particular for large high-resolution displays. Examples of visualizations designed for this setting include using tangible input controllers [21], sensing body movements as implicit navigation input [9], and adapting interaction techniques for large displays [19].

We extend this work by using the notion of proxemics to identify design opportunities. Proxemics studies the relation between people as it is expressed in the use of space [14,15]. Compared to early work on proxemics, recent work [13] as well as this paper extend the notion of proxemics to describe also the relation between people and objects (often user interfaces). In research on human-computer interaction (HCI), proxemics has for instance been used to design interaction techniques that change user interface layout based on the user's position [3], and to study orientation and distance among devices and doctors in neurosurgery [31]. Previous research has also demonstrated how body orientation and position can be used with visualizations: for implicit interaction with ambient displays [53] and for coarse 3D navigation in microseismic visualizations [32]. We build on previous work to explore how the notion of proxemics can be applied to interaction with information visualization.

The opportunities for proxemics in information visualization are manifold. First, it may be used to adapt visualizations based on the users' position and orientation relative to the display. Second, it could use movements in front of a display to have visualizations follow users' movements or blend as two users get close. Third, we could augment users' backing away from a large display by even further zooming out or abstracting the visualizations. Many other uses of proxemics in information visualizations may be imagined.

This paper explores in particular design opportunities for information visualization based on movement and distance to large high-resolution displays. We focus on using movement and distance because earlier work has emphasized physical navigation as important when using large displays [2] and in group work [20]. We explore spatial relations only between a single user and visualizations; exploring relations between people would provide more opportunities, but is beyond the scope of this paper. The opportunities are illustrated with a design space and with sketches; the opportunities focus both on supplementing other input techniques and on replacing them. We also show how earlier work that has not explicitly used the notion of proxemics (e.g., [53]) can be understood through proxemics and potentially benefit from its analytic framework. We select a subset of design opportunities to implement and test in three user studies: (1) navigation by physical movement, (2) querying coordinated views by movement, and (3) adapting visual representations to distance. We do so to generate design ideas, but also to provide initial data on the usefulness of combining information visualization and proxemics. Our approach is to ground some opportunities in empirical data rather than to give an exhaustive systematic review of the opportunities or to present in-depth data on a single case.

We contribute (a) an initial analysis of using proxemics for information visualization, (b) prototypes of information visualizations that adapt based on tracking of their users, and (c) an evaluation of a set of proxemic visualizations. The argument is that proxemics may offer promising design opportunities for non-desktop visualizations; we think such opportunities are valuable to both researchers in visualization and to designers for large displays.

1 Related Work

The term proxemics is due to Edward T. Hall [14,15], who used it to describe the study of "how man unconsciously structures microspace—the distance between men in the conduct of daily transactions, the organization of space in his houses and buildings, and ultimately the layout of his towns". Among other contributions, he related physical and social distance in a set of four zones, from intimate space (less than 46cm between people) over personal and social space to public space (more than 3.7m). Hall discussed how social, gender, and cultural factors may mediate this relation. Much research has built on and extended Hall's work, applying it for instance to design [48], human-robot interaction [33], and HCI.

- *Mikkel R. Jakobsen is with University of Copenhagen. E-mail: mikkelrj@diku.dk.*
- *Kasper Hornbæk is with University of Copenhagen. E-mail: kash@diku.dk.*
- *Yonas Sahlemariam Haile is with University of Copenhagen. E-mail: yoniclassic@yahoo.com.*
- *Søren Knudsen is with University of Copenhagen, E-mail: sknudsen@diku.dk.*

Manuscript received 31 March 2013; accepted 1 August 2013; posted online 13 October 2013; mailed on 5 October 2013.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

Proxemics is increasingly used in HCI, both as (1) a notion to understand and analyze collaboration and interaction, and (2) a notion to drive the interaction among users and devices. The first point has been studied in computer supported collaborative work (CSCW), where the relation between physical distance and perception of social distance has been a key issue [38]. Applications in CSCW include a study by Hawkey et al. [17] investigating the relation between proxemics and collaboration success with a large wall display. Stretching proxemics to include the relation among users and devices has led to several descriptive accounts. Mentis et al. [31] studied orientation and distance among devices and doctors in neurosurgery using notions of proxemics. Jakobsen and Hornbæk [20] used proxemics to describe interaction in front of a large display.

The second point above has in particular been inspired by Marquardt and Greenberg's notion of proxemic interactions [13]. Their work extends the notion of proxemics so that it pertains not only to relations among persons, but also to relations among people, digital devices, physical objects, and the environment. They consider five categories of proxemic dimensions particularly for ubiquitous interaction (which is relevant more broadly for HCI):

- *Distance*, the physical distance between entities, either given as a continuous measure or relative to discrete zones. In Lean and Zoom, for instance, semantic zooming is based on the user's distance to a laptop screen [16].
- *Orientation* concerns which direction a person (or other entity) is facing. This has been used, for instance, to adapt presentation software to different views depending on which way the presenter is facing [13].
- *Movement* concerns the changes in distance and/or orientation over time. For instance, personal territories on tabletops can be adapted when one user approaches another user's space [26].
- *Identity* concerns distinguishing between entities. For instance, a display may respond differently to the movement of a mobile phone than to the movement of a person [13].
- *Location* describes the place of interaction. A simple instance is the presence of a person in a room.

A recent toolkit helps detect and react to these dimensions [29].

Some earlier work has used related types of movement to control interaction, without explicitly using the notion of proxemics. Vogel and Balakrishnan [53] presented a display system that supported a smooth transition from public use of the display, through implicit interaction at a distance, to up close, personal interaction. Ju et al. [23] presented an interactive whiteboard that sensed users' distance to the board for switching between modes of using a whiteboard, in particular between authoring and ambient use. Marquardt describes gradual engagement in providing connectivity, information exchange and transfer as a function of proximity [28]. Marquardt and colleagues give many other examples of using movement to control interaction [13,30]. Work on navigating virtual environments has also used movement and orientation extensively. For instance, Souman et al. [49] described how an omnidirectional treadmill allowed participants to walk in any direction they wanted in a virtual environment, with information in a head-mounted display being updated based on their walking. Such work differs from the focus of the present paper in that movement and orientation are used to generate a view (say, in a head-mounted display) of a virtual environment corresponding to a particular position of the user's head; instead, we consider uses of proxemics data for changing visualizations of abstract data.

The present paper uses the notion of proxemics to drive innovation in interaction with information visualizations. One reason to do so is that the notion of proxemics might help generate interesting designs, beyond those described in the literature. Another reason is that to our knowledge, no paper has attempted to relate proxemics and visualization, despite the interest in using visualization on large displays and despite the frequent observation that movement [2] and orientation [4] play key roles in interaction with large displays. A third reason is that even though earlier papers have used movement to control interaction (e.g., Vogel and Balakrishnan [53]) they rarely relate to the information visualization literature and do not evaluate visualization tasks. Thus we proceed to discuss the relation between proxemics and visualization.

2 Design Opportunities

As argued earlier, a variety of design opportunities may be generated from the proxemics literature. Because these have not been explored in relation to visualization activity, we next discuss some design opportunities, in part summarized as the design space in Table 1. Some of the opportunities are implemented as prototypes and evaluated in user studies in the second half of the paper (marked #1, #2, or #3). Some entries in the table are blank, either because they are uninteresting or because we have yet to come up with, or find in the literature, a compelling example.

The design space is organized from established views of key characteristics of proxemics and information visualization. To this end we choose categories from earlier work on proxemics [13] and information visualization tasks [18].

Many alternatives to these two choices exist. With respect to

Table 1. Combinations of information visualization tasks (excerpt from [18]) and proxemics categories (excerpt from [13]). The symbols #1, #2, and #3 refers to design opportunities that are tested in the second part of the paper.

Proxemics category \ Information visualization task	Visualize	Filter	Sort	Select	Navigate	Coordinate	Organize
Distance	Show details when close/ aggregates when far (#2)	Filter items depending on the physical distance to user (body fisheye)	-	Distance increases selection scope (#2)	Focus and demagnified context at distance	Brush-and-link close data	Distance-dependent workspaces
Orientation	Visualize for different viewing angles	-	Sort by variable selected by orientation	Coarse selection by orientation	Head orientation controls zoom center (#1)	Indicate related areas through orientation	-
Movement	Switch between encodings by moving (#2)	Dynamic querying when moving (#3)	Sort by variable selected by movement (#3)	Coarse selection by movement	Zoom and pan by moving relative to display plane (#1)	Selected views move along with user (#3)	Reorganize windows in workspace
Location	Contextual visualizations	Switch between subsets	-	-	Overview and detail in left to right	-	Location-dependent perspectives or activities

proxemics, earlier definitions emphasize different types of proxemics. We chose the much cited taxonomy of proxemic interaction [13], because it captures the relations between people and devices like large displays, which is our focus. With respect to information visualization, a host of alternative models exist. We decided against relatively low-level models (e.g., [1]) of information visualization because we think the initial promise of proxemics is to enhance higher-level tasks. We also decided against taxonomies focused on data (e.g., [24,46]), because they were not easy to combine with the proxemics taxonomy. Finally, the visualization taxonomy that we have chosen to use integrates many aspects of earlier work; for example, it includes most of the tasks in Shneiderman's task by data type taxonomy [46].

The resulting design space does not include all categories: Some categories of proxemics are less applicable to single-person interaction with visualizations on a large display (e.g., Identity). Similarly, some visualization tasks do not map well to proxemics (e.g., Derive). The opportunities presented here are intended to generate design ideas. Other possibilities exist that could be more useful than the examples given here.

2.1 Distance

Viewing distance is important in using information visualization on large high-resolution displays: Users can step back to get an overview and to navigate [2] or to see patterns in data [8]. However, earlier work has mainly studied visualizations that do not change with user's distance. Vogel et al. [53] is a notable exception as they adapt visual representations and interaction modes to discrete distances. Below we describe how visualizations can adapt and react to distance for particular visualization tasks.

Visualize. Visual encodings may dynamically change with the user's distance. Different tasks can thus be supported at varying distances, for instance by showing aggregate representations at a distance and details up-close. This is illustrated in Fig. 5, where each level of aggregation is associated with a discrete distance zone. The alternative, combining the data in the same static visualization, can overload the display and potentially overwhelm the user. While we focus on the spatial distance between user and display, the distance of a hand-held display relative to a large display could similarly be used for semantic zooming in for instance graph visualizations [50].

Filter. Distance can be mapped to a variable so as to allow filtering out data. For instance, adapting the generalized fisheye view [11] to a large display, could help users focus on the most relevant items; items are filtered out if they have a degree-of-interest below a threshold that grows proportional to the user's distance to the display. More interesting items can be made prominent or shown in detail at a distance while other items are aggregated.

Select. Distance can influence the scope or granularity of user's selections. For instance, Peck et al. [39] describe a multi-scale interaction technique that "chang[es] the user's scale of interaction depending on their distance from the current object(s) of interaction."

Navigate. One possible visualization that adapts to large displays for supporting multi-scale navigation is focus+context: As the user steps back from the display, selected elements in focus can be magnified to remain a constant size in the user's field-of-view; in effect those elements are brought closer together, for instance to support comparison, while the context is demagnified (rather than being filtered out as done in the generalized fisheye view discussed above). This is illustrated in Fig. 1. Another idea is to relate distance to zoom-level, so that when a user moves away from the display, the zoom level changes.

Coordinate. Whereas most coordination of views relies on explicit actions [37], the user's distance to particular views in the display may provide for implicit coordination. For instance, depending on which graphs that are close to the user, they could automatically become linked, so that data points selected in the one are highlighted in the others.

Organize. Views can be reorganized for interaction when the user stands within touching distance of the display (e.g., showing data views and widgets for dynamic querying), while larger overview-providing views are shown when the user is standing at a distance.

2.2 Orientation

Although orientation is used extensively in virtual reality, it is rarely seen in research on information visualization. Research that comes close are the ChairMouse [9], which used the users' rotation on a chair to control cursor movement, and the study by Bezerianos and Isenberg [4], who looked at the role of angle and movement in perception on large displays. Neither study used orientation to adapt visualizations.

Visualize. Visual encodings that become distorted at extreme viewing angles cause problems [4]. A visualization can dynamically change to a visual encoding that is more robust to extreme viewing angles, based on its orientation toward the user. Related techniques are E-conic, which dynamically corrects the perspective of windows [34], and Screenfinity, which rotates, translates, and zooms content to ease reading while users pass by large displays [43].

Sort. Ordering data helps reveal trends or clusters of values. The most common method of ordering, sorting records by one or more variables [18], could be supported by detecting the user's orientation toward a particular variable (e.g., a column in TableLens).

Select. Orientation may supplement other pointing input for selecting data points in visualizations. For instance, a user's motor space with a pointing device can map to a particular view, which is selected by changing orientation (see Fig. 2).

Navigate. Orientation may control navigation by giving additional information about the user's current focus. For instance, orientation may be used to enrich the parts of the display that the user focuses on or (as will be experimented with in study #1) to control the point around which zooming is performed.

Coordinate. Orientation can support exploration across views. For instance, body and head orientation can be used together for indicating distinct areas of interest, so that relations between data in those areas can be visualized.

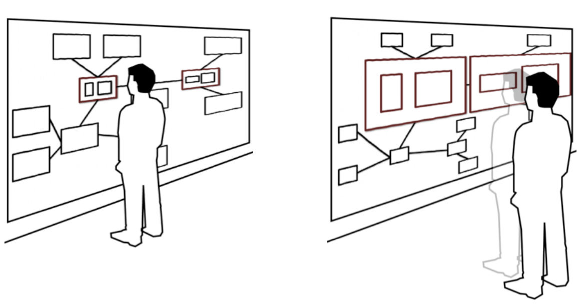

Fig. 1: Distance-based focus+context: Focus elements are selected (outlined in red) while up close (left). As the user steps back, the focus elements are magnified (right).

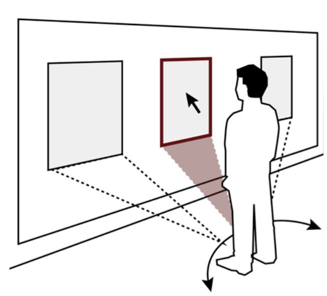

Fig. 2: Selecting view by changing orientation relative to the display.

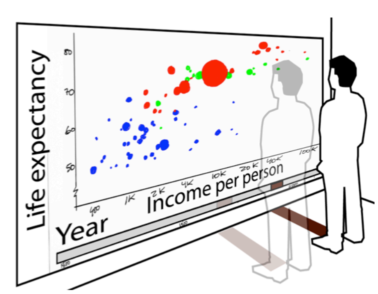

Fig. 3: Changing a dynamic query slider by moving.

2.3 Movement

Visualize. Study #2 will present an example where movement is used to change the encoding of visual representations.

Filter. The spatial relation between the user and a dynamic query slider can be used for filtering. By mapping the user's position to the slider in the display, the user can move relative to the slider in order to change the value. For instance, in Fig. 3 the user's lateral position maps to a timeline: in that way, the user can move right towards the most recent data.

Sort. In study #3 we explore the use of movement to select a variable for sorting a table of data items.

Select. Movement could be used for coarse selection of a view in order to help users select data points in a visualization.

Navigate. The user's physical navigation around a large display can be further supported through view manipulations. For instance, physical navigation can be extended through movement-based zooming and panning: moving forward to zoom in and back to zoom out; moving sideways to pan. This is explored in study #1. This is related to work in virtual reality that have used omnidirectional treadmills to allow movement (e.g., [42]); such studies have typically strived to make rendering of the virtual reality smooth and realistic, not to use movement to adapt interactive visualizations.

Coordinate. Selected views could move with the user's position, for instance to allow comparison across views that are otherwise too far apart to be viewed simultaneously.

Organize. Manually reorganizing visualization views, legends, and controls can be tedious, particularly on a wall-display. However, related views and legends could be automatically reorganized depending on the user's movement relative to the workspace in order to fit the user's focus in a task.

2.4 Location

Visualize. Facilities for creating new visualizations could leverage contextual information from the location so that a new visualization is tailored to that particular context.

Filter. Visualization views could be filtered to show different subsets of the data as the user switches between different locations.

Navigate. To aid navigation, different visualizations that are aimed at taking a broad view of the data (overview) and at specific, detailed investigations of parts of the data (details) may be anchored to different physical locations. For instance, having an overview perspective on the left part of a large display would provide the user with custom visualizations tailored for coordinating several detailed investigations going on in the right part of the display.

Organize. Different configurations of views may be shown at different locations in order to give different perspectives of the data (e.g., when the user stands near the left side of the display, the rest of the display changes to show information related to the views at that location) or to provide stations for different activities (e.g., monitoring while seated in a certain part of the room).

2.5 Prototyping and testing opportunities and options

The techniques that we prototype and test in the next section present a sample of the design space (see Table 1) selected to probe interesting options. First, we wanted to study one of the simplest cases of linking proxemics and visualization: linking movement of the body to zooming and panning. It is unclear whether continuous or discrete measures are the most appropriate in that case, or whether to base interaction on absolute or relative movement. Second, we wanted to compare continuous measures of proximity (e.g., controlling filtering through movement) to discrete measures (e.g., levels of aggregation for discrete distances). Third, proxemics may be used to control fluid visual transitions (e.g., zooming, panning) and discontinuous changes (e.g., change encoding, linking movement to selection of variables). We wanted to see if either is more useful or more sensible when linked to proxemics data. Fourth, a potential use of proxemics data is to make things appear to be constant size (adapting for instance a graph based on distance) or in the same relative location (e.g., always near the users right arm). We wanted to explore such effects. In sections 4.1, 5.1, and 6.1, we explain the designs we have studied in detail.

3 Overview of User Studies

Whereas the exploration of design opportunities has identified novel and interesting designs, it has not provided any data about the usefulness of such designs. Next, we therefore present three user studies aimed at obtaining such data. The studies aim to provide initial, qualitative data about usefulness by having participants use and compare designs. The studies are lightweight (i.e., each participant interacts for about 40 min) and formative (i.e., qualifying and developing design opportunities rather than finding a "best" option).

This choice of method requires justification. The overall aim of the present paper is to explore design opportunities. We therefore decided against running a controlled experiment, as done in many evaluations of information visualizations and of proxemics [19,22,56]. Instead we wanted to gain empirical insight on a range of design possibilities. We also wanted to avoid rushing to experimentation (as warned about by Shadish et al. [45] and Greenberg and Buxton [12]). We decided against some of the other methodologies for evaluating information visualizations [6] because they mostly assume a hi-fidelity and well-defined design or require a specific application domain, task set, or user base. The former is not the case for the combination of information visualization and proxemics, and the latter seemed to constrain finding and developing design opportunities.

3.1 Commonalities of the studies

The three user studies presented next have a common structure (see Table 2). First, they all have six participants. This number is often recommended for formative user studies [36] and while it gives low power (in the sense of being able to detect quantitative differences, see [7]), it does allow us to gain qualitative insights about usefulness.

Second, all studies use one or two combinations of proxemics/visualization and a reference interaction style. It has been shown that users generate more comments when exposed to several alternatives than to just one [51].

Third, we collect qualitative data from the studies. In addition to

Table 2. Overview of user studies. Categories refer to the information visualization tasks and proxemics categories in Table 1.

Study	Categories	Users	Interfaces	Tasks	Data
#1	[Navigation] + [Move, Location]	6	(a) Absolute: Navigation by absolute movement (b) Relative: Navigation by location (c) Baseline: Virtual navigation with gyro mouse	Three tasks involving maps, adapted from [19,47]	Map from OpenStreetMap
#2	[Visualize] + [Dist, Move]	6	(a) Distance-controlled detail/aggregation (b) Baseline: Interaction with gyro mouse	Five tasks, some adapted from [54]	Data sets of 1000-3000 homes (5 attributes)
#3	[Filter, Sort] + [Dist, Move]	6	(a) Position-controlled variable selection and brushing (b) Baseline: Interaction with gyro mouse	Five multi-variate analysis tasks [55]	406 cars (8 attributes) [41]

capturing preference data, we have at least two persons observing users while interacting: the observers take time-stamped notes that can be referenced and coupled to video recordings during analysis.

Fourth, while the studies are formative, we prescribe tasks for users to solve. The idea is to ensure that they engage in demanding tasks so as to experience and be able to discuss the usefulness of the interaction styles. All tasks were adapted from previous studies of information visualizations.

3.2 Participants

In all, 18 participants (4 female), ages between 23 and 37 years (M = 29.8), were recruited by word of mouth; six participants for each study.

3.3 Procedure

The procedure was similar across studies. Participants were welcomed to the study, and informed of its purpose. They were introduced to the wall-display and the interfaces, and the tasks were explained to them. Participants then completed a set of tasks with each interface. For each interface, the experimenter first explained its use and participants were given time to try using it. Participants were then given the tasks, one at a time. They were encouraged to ask questions during the experiment. After completing the last task with an interface, we asked participants about their experience with the interface they had just used, including its benefits and drawbacks. Finally, after having completed all the tasks, participants were interviewed about each of the forms of proxemic interaction provided by the interfaces.

3.4 Data analysis

Sessions were video recorded and the experimenter and one or two additional data loggers took notes. Each study was analyzed immediately following its last session using the Instant Data Analysis technique [25]. For the analysis, the experimenter and the data loggers gathered in front of a whiteboard. Observations from the notes and comments from interviews were discussed. When an important issue was identified, it was written on a post-it note and put on the whiteboard. The notes were categorized into themes. Based on the clusters of post-its on the whiteboard, the most important findings were written down with clear references to the observations and any supporting video recordings. On average, the analysis session lasted around two hours.

3.5 Technical setup

Participants used a 24 megapixel display that measures 3m×1.3m. The display consists of 4×3 tiles projected from the back by 1920×1080 projectors. Projectors are manually aligned so as to minimize seams between tiles. The display was run by a single computer running Microsoft Windows 7. The room in which the display was set is 3.5m wide and the distance from the display to the back wall is 2.95m.

For input we used a NaturalPoint OptiTrack motion capture system (www.naturalpoint.com/optitrack/) that tracks, via reflective infrared markers attached to a baseball cap, the location and orientation of the participant's head. Participants also used a wireless gyroscopic mouse. The mouse cursor was enlarged to maximum size.

4 Study #1: Navigation by physical movement

The first study investigates the potential of using physical movement in the zoom+pan visualization technique.

4.1 Conditions

Three variations of a zoom+pan interface were used for navigating geographical maps. In all conditions, a Gyro mouse was used for interacting with targets in the tasks.

4.1.1 Absolute: Navigation by absolute movement

This interface uses a direct mapping between participants' movement and movement of the map. The user moves toward the display in order to zoom in (i.e., to see details) and away from the display to zoom out. This is illustrated in Fig. 4 (a-c). Movement is combined with head orientation for zooming. A crosshair indicates the point where the ray cast from the cap worn by the user intersects the display, and zooming is centered on that point. Lateral movement controls horizontal panning: Moving left causes the map to move right; moving right causes the map to move left. Our initial intent was to map floor position directly to map position. However, to keep panning speed at a reasonable pace when the user is close to the display (i.e., at high zoom factors), we reduced the floor-to-map movement ratio. This restricts the panning range when close to the display. Head orientation is used for panning up and down. Pitching the cap so that the ray intersects the display plane above or below the display causes the map to pan vertically at a fixed rate.

4.1.2 Relative: Navigation by location

In this interface, participants control zooming and panning by moving relative to a 75x75 cm rectangular region in the center of the floor, illustrated by Fig. 4 (d-f). The map moves right when the user's body is left of the region; moves left when the user's body is right of the region. Similarly, the map zooms in when the user has stepped toward the display from the center region; zooms out when the user has stepped backward from the center. The zoom rate is inversely proportional to the zoom level so that when zoomed in to a detailed level, the zoom rate is lower. The use of head orientation for zooming and for vertical panning is similar to Absolute.

4.1.3 Baseline: Virtual navigation using mouse

In this condition, the user operates the interface using only the gyro mouse. The interface resembles widespread mouse-operated map interfaces (e.g., Google Maps): To pan the map, the user clicks and drags the mouse opposite the panning direction (i.e., so that the map follows the mouse cursor); to zoom the map, the user rolls the scroll wheel on the mouse forward (for zooming in) or backward (for zooming out). Mid-air input techniques for zoom+pan interfaces [35] allow more efficient navigation than the baseline interface we used here. However, we did not aim for performance, but rather a simple-to-use mouse-based interaction style that we expected users to be familiar with.

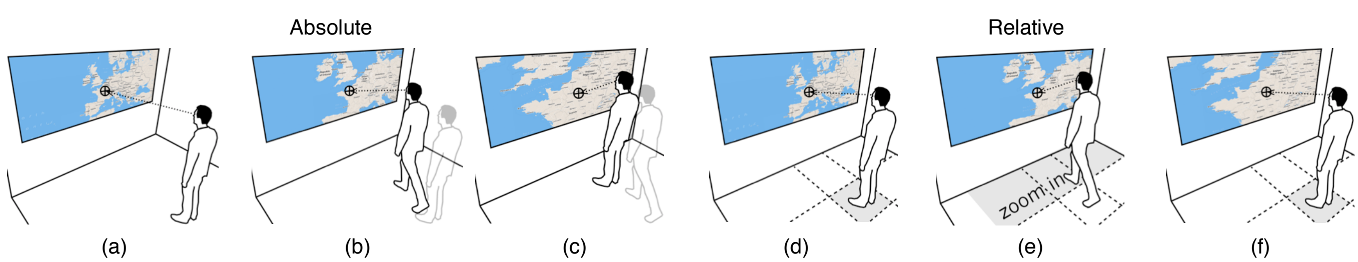

Fig. 4: Zooming in the two conditions that use proxemics in Study #1. In Absolute (a-c), the zoom level increases as long as the user keeps moving toward the display, and stops zooming when the user stands still. In Relative (d-f), the zoom level increases as long as the user is within the zoom zone (e). Zooming is centered on a crosshair, which indicates the point where the ray cast from the user's head intersects the display.

4.2 Tasks

Participants performed a series of tasks using a map obtained from OpenStreetMap (www.openstreetmap.org) at different scale levels. The following types of task adapted from [19,47] were used:

- *Navigate*: Participants had to navigate to a clearly marked target and click on it with the mouse. Then a new target was shown, until participants had navigated to ten targets.
- *Trace*: Participants had to trace a railway, where targets were placed close to ten selected stations. Participants had to move each of the pins onto the station using the mouse.
- *Search*: Participants were handed a description on paper of a location (e.g., "Near 'city' find 'lake'") and they had to point out the location. Participants were given three locations to search for.

4.3 Results

We present only results that relate to the use of movement and location to control navigation. In the instant data analysis, four themes emerged.

4.3.1 Using your body for navigation was liked

Several participants said they liked controlling navigation with their body: it is a "nice concept to use your body to move" and "it is nice that you move a lot, particularly in a work environment". Reasons for this view varied. Two participants mentioned that movement was intuitive, three that movement required less effort than the mouse, and two perceived movement to be faster than using the mouse.

4.3.2 Observed benefits and drawbacks of using body

We saw much movement in the Absolute and Relative conditions. Body movement was expected as it controlled navigation. Some observations were nevertheless surprising. One participant transformed the navigation task of finding and clicking an object at high magnification to a smooth movement from the back of the room (zoomed out) to the display (zoomed in). Several participants moved to the back of the room in preparation for receiving the next task.

We noticed a lot of awkward movement. Some participants moved very slowly, some expressed uncertainty about the size of the steps to take. Also, movement of your body is difficult to use for fine-grained navigation and it is hard to stop panning as quickly as with a mouse. Some participants adopted particular movement types to deal with these limitations. Three participants leaned rather than moved to control location; in the Relative condition, two participants kept a foot in the center region while lunging forward or to the sides (one participant mentioned the similarity to dance-mat games).

4.3.3 Movement versus location

A key difference among conditions was the use of movement for navigation versus using location for navigation. Participants were split in their preference for either technique (movement: 3; location: 2; one undecided).

Navigation by movement was well received. Two participants commented that this technique was intuitive, in particular because there was a direct relation between your movement and what happened on the screen. Another difference was the freedom to move around. With Absolute, one participant found "a lot of freedom to move all over the place"; two participants contrasted this with feeling "restricted" and unable to "move freely" with Relative.

Navigation by location was liked for several reasons. One reason was that "zooming was nice here" because one could zoom without getting too close to the screen; when using movement to zoom, participants by definition were close to the screen when they had zoomed a lot. One participant mentioned the benefit of a stable center, in contrast to navigation by movement where the display was changing much of the time. However, participants had to keep track of their position relative to the center. They described how you were "fixed to the center" and that it "requires concentration to keep track of zones".

4.3.4 Design ideas and variations

Several design ideas came up. Rate control was mentioned as an improvement for Relative, so that the speed at which panning and zooming was done depended on your distance to the center point. This would increase the issue of small movements causing large steps in navigation, which is why we did not implement it in the first place.

Movement did not control all aspects of navigation in Absolute or Relative. Head pitch was used to control panning up/down, which caused unintended panning when participants looked down. Participants suggested the use of alternative means for controlling panning, for instance by using gestures.

5 Study #2: Adapting Representations to Distance

The second study investigates the adaptation of visualizations based on the user's distance and movement.

5.1 Conditions

Two variations of a map-based visualization of real-estate data were used. The visualization allows the user to vary the visual representation of the data (individual homes or geographic areas) and to select an area for which to call up details. A diverging color scale is used to indicate how the value of an attribute, which the user can select from a menu (e.g., price per m^2), is above or below the mean value of that attribute.

5.1.1 Distance-based aggregation and details on demand

This condition uses distance and movement. First, *distance-based aggregation* changes the visual representation based on the user's distance to the display (see Fig. 5). At less than .75m, individual homes are shown as points. As the distance increases, the representation changes to show data aggregated on geographic areas

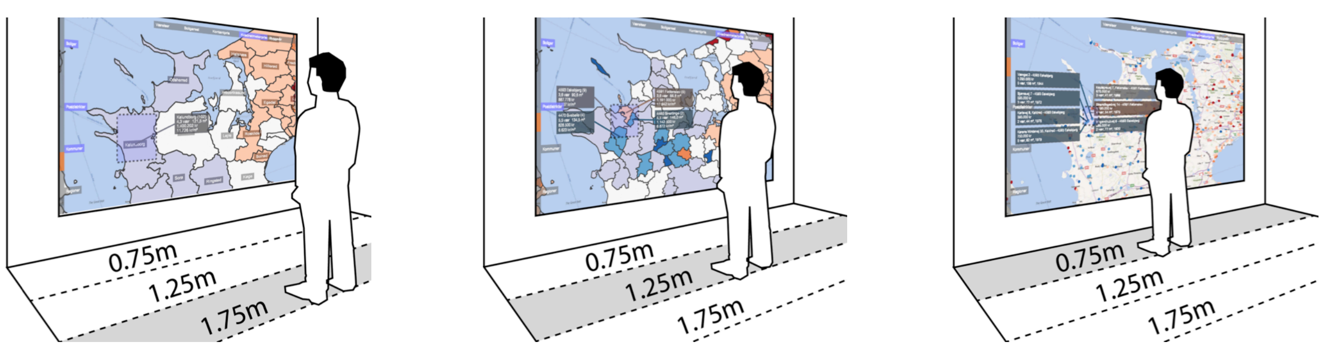

(a) Aggregate, municipalities (b) Aggregate, postal districts (c) Individual homes

Fig. 5: Techniques used in Study #2: Distance-dependent aggregation of real-estate data by geographic area in (a) and (b); details on demand for geographic areas in (a) and (b), and for individual homes in (c); multi-scale selection of map area.

(.75m: postal districts; 1.25m: municipalities; 1.75m: regions), and using larger font sizes. Transitions between representations use alpha blending over a 20cm distance range. Second, *movement-based Excentric Labeling* [10] gives details about homes within a selection box that follows the user's position horizontally and moves vertically with the pitch of the user's head. Third, for *multi-scale interaction* [39], the selection box grows in size with increasing distance and details are shown for data at higher scales: homes, districts, or municipalities. Fourth, *movement-based change of color encoding*. When the user is more than 2.5m away from the display, the attribute menu (shown in the top-center area of the display) responds to the user's lateral movement: Moving left or right causes an indicator to move to another attribute that will be used for color encoding.

5.1.2 Baseline: Gyro mouse

In this condition, the user operates the interface using only the gyro mouse, that is, for changing the visual representation of home data and for selecting the area of the map for which details are shown. The mouse scroll wheel maps to distance as it is used in the other condition: scrolling the wheel forward corresponds to walking forward (and vice versa). This changes the representation, the size of the selection box, and the level of details that are shown. As feedback to the user, the four representations of home data are placed on a vertical slider in the left side of the visualization, with an indication of the representation that is currently shown. The selection box is moved with the mouse cursor (that is, while the mouse trigger button is pressed); and details about homes within the selection box remain fixed when the user stops moving the mouse cursor.

5.2 Tasks

Participants performed the following tasks, some adapted from [54], with subsets of a real-estate database:

- Find the region that has the lowest average price per m^2 (or lowest average number of rooms).
- Find the municipality in a given region that has the highest average asking price (or largest average area).
- Find the home in a particular postal district that has the largest area (or smallest area).
- Find the postal district in a particular municipality that has the highest average price per m^2.
- Find the most expensive house in two (geographically remote) municipalities.

5.3 Results

Three themes that are related to distance and movement emerged in the analysis.

5.3.1 Use of distance makes sense and "works well"

Four participants described the Distance condition as natural, intuitive, and making good sense. For instance, one said it was "natural to use the body", another that it was "intuitive to get more information in less space when up close. It works very well."

In relation to aggregation of data with increasing distance, one participant said that it was nice that there was not much data when standing at the back.

Several participants seemed to change between representations with ease by moving. In particular, we observed three participants that moved back and forth repeatedly to switch between representations for solving tasks that involved relating homes or districts to municipalities. Changing representations using the mouse seemed less fluid, and participants glanced more often at the slider at the left side of the interface.

5.3.2 Discrete distance zones versus free movement

However, using distance did not work equally well for all participants. For instance, one participant said that although it was natural to move, he had to think more while moving than while using the mouse. Another said that she had to remember to stand still at a distance.

One drawback, which was clear from our observations and from participants' comments, relates to the discrete distance zones: To see certain information, the user is bound to a certain distance. From our observations this was a problem for one participant in particular, who said that it is "natural to step back for overview, but then the data I want to overview disappears." In the mouse condition, this participant solved the tasks while standing noticeably farther away from the display than the other participants: He read details about individual homes from around 1.5m distance. Other participants made related comments. One said you have to get close to see details on individual homes, but then "up close, I had trouble keeping an overview of it all."

5.3.3 Details-on-demand too sensitive to movement

All the participants said they liked the mouse better for selecting the area to show details. One reason is that the mix of using body position and head orientation for selection was confusing.

Participants suggested different ways of improving details-on-demand based on movement. Three participants said that they wanted to use their hands to "lock" the view of details or for selecting houses, when they were within reaching distance. Also, two participants suggested leaning toward the display as a way to lock the view of details. *Details on proximity*, or using head position relative to body position, could be a promising design variation.

6 Study #3: Dynamic Query by Movement

The last study investigates the use of movement for attribute selection, brushing and linking, and filtering of multivariate data.

6.1 Conditions

Participants used two variations of an interface containing multiple coordinated views of data about cars. The interface comprises a

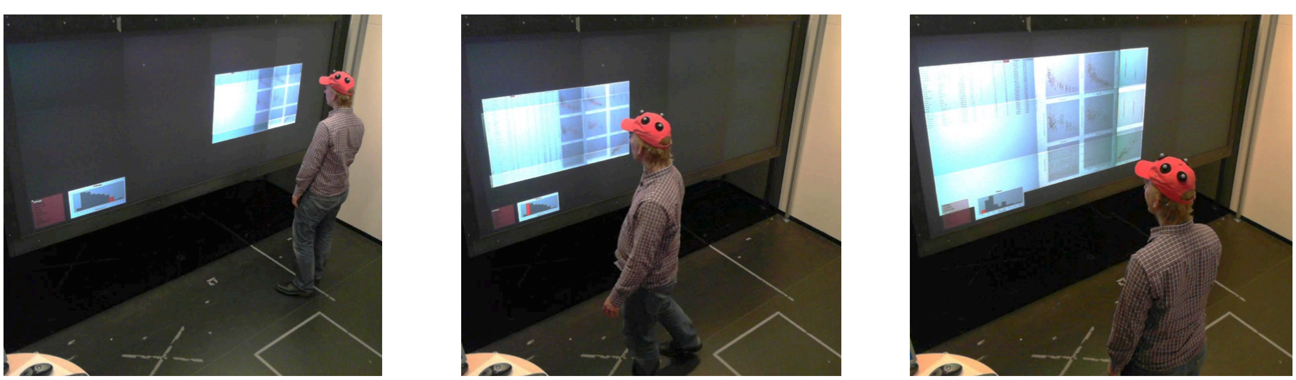

(a) (b) (c)

Fig. 6: Techniques used in Study #3: The user brushes the bars in a histogram by walking sideways, (a) to (b); the views move to stay in front of the user. The user moves backwards in order to select another attribute (c); the views scale to remain at a readable size.

window containing nine scatterplots and a data table, a view showing a histogram for an attribute, and a view listing the available attributes. If the user selects an attribute from the list, the histogram for that attribute is shown and the data table is sorted by that attribute. For visualizing the histogram, the values of most attributes were binned to produce 10 bars. For attributes with less than 30 values, each value had its own bar (e.g., model year of cars spans 12 years). Histogram bars can be selected: the table is filtered to show only the corresponding data points, which are also marked red in the scatterplots. The two variants of the interface differ in the way that the user can select attributes from the list or select bars in the histogram.

6.1.1 Position-controlled variable selection and brushing

This condition uses distance and movement. First, the attributes in the list are mapped to discrete distance zones, 1m (the first attribute) to 2.5m (the last attribute) from the display. The user selects an attribute by moving closer or farther from the display, shown in Fig. 6 (b-c). In the attribute list, a circle indicates the user's position relative to the attribute zones. Hysteresis tolerance is used for transitions between the zones of two variables: The user enters and exits a zone at separate distances. This helps avoid unintentional switching back and forth between two attributes. Users' sideways movement is used for brushing over bars in the histogram: The user's position along an axis parallel to the display maps to the x-axis of the histogram, see Fig. 6 (a-b). One bar is selected at a time. The physical space for brushing (from the leftmost to the rightmost bar) spans 1.65m in the center of the display. To enable users to read the data while they move, the views are scaled depending on the user's distance, see Fig. 6 (b-c). Also, the window containing the table and the scatterplots is positioned according to the user's position. The other views remained fixed.

6.1.2 Baseline: Gyro mouse

In this condition, the interface is operated using only the gyro mouse. Attributes can be selected from the list by pointing and clicking with the mouse cursor. Histogram bars can be brushed by clicking on the bars. Views are fixed in a size corresponding to standing 1.5m from the center of the display in the Position-controlled condition.

6.2 Tasks

Participants performed five types of task adapted from [55], using a dataset with eight attributes for 406 cars [41]:

- Find the car that has the most power among Ford cars.
- Is there a correlation between engine power and weight?
- Does Dodge make more car models than other American manufacturers?
- Please categorize car models into two types: one consisting of cars with poor mileage and one consisting of cars with good mileage. Try to take model year into account. Which has most models?
- State the conditions for your ideal car and identify it using the interface.

6.3 Results

Three themes that are related to distance and movement emerged from our analysis.

6.3.1 Physical mapping of data

Participants liked the idea of mapping physical space to data space. After having used both conditions, one participant said: "Distance for selection of variables seems very natural"; another described it as fun, but said he felt more efficient when using the mouse.

Participants were split on preference for using movement and using the mouse; all suggested combining the two forms of interaction, one reason being that they could change variables using the mouse. They also suggested adding a lock to position tracking so as to be able to approach the display or step back from it. One said "[I would like to] be able to lock such that I can walk closer to something and then unlock it again"; another that "[I would like to] be able to lock variable choice such that you don't change in error, when you are busy." One participant demonstrated this by taking off the tracking cap so that he could move without changing a variable.

One reason why participants wanted such a lock was because they found it difficult to keep the attribute selected while moving sideways to brush bars. Participants were observed to "drift" in distance to the display while brushing; this could result in abrupt changes of selection. It seems this issue caused some participants to move more cautiously and to look at the histogram while moving.

6.3.2 Scaling

Four participants disliked the way the views were scaled and positioned depending on location. They suggested instead a fixed size (and using a locking mechanism as suggested above to be able to look closer at an item). Three participants suggested that the location-dependent scaling and positioning could be improved by moving and scaling in discrete steps, instead of continuously.

One participant got confused when pointing at the scatterplots, because it scaled when he walked closer to the display while doing so. This participant proposed zooming in when approaching the display (similar to the absolute condition in Study #1). In the baseline condition, several participants moved close to the data to point at it.

6.3.3 Thinking physically about the data space

Two participants used physical descriptions of the data space. For example, one participant said: "Let me see what is out here", another: "I was in kind of a lane where I could filter instead of clicking with a mouse." That participant added: "It feels navigable," and considered that the way he had the attributes mapped to the floor space, he would be able to "Go to cars with large engines".

7 Discussion

We have explored opportunities for using body movement to interact with visualizations on large high-resolution displays and we have tested several of them. In particular, we have relied on the notion of proxemics [15] and a particular set of visualization tasks [18]. Overall, the three user studies provide initial data in support of the idea of using movement and distance to change visualizations. Participants in all studies said that using body movement was intuitive or natural.

Specifically, changing the visualization in response to changes in the user's distance to the display seemed useful. In Study #1, participants moved closer to the display for zooming in; in Study #2, participants moved closer to see data represented in higher detail ("more data in less space"). Changes to zoom level and representation made sense to several participants, maybe because it relates to the experience of physically zooming out and seeing less detail (due to limits of visual acuity). In contrast, scaling views with user's distance worked contrary to the expectations of some participants (Study #3).

Based on observations and feedback from participants, potential benefits of proxemics-based zooming and aggregation are reduced effort and more smooth interaction compared to mouse control. Proxemics-based control also seems to allow navigation in or manipulation of many variables at a time in a natural way.

Another opportunity is the use of body movement for dynamic querying: In Study #3, we mapped the user's movement to selection of attribute values. One benefit observed for several participants, was that they could fix their focus on the data views while changing the selection by moving their body.

The studies also showed how using proxemics and visualizations together may give a distinct physical sense to abstract data. Study #3 differed from the other two in that movement was mapped to abstract data rather than spatial data. We note that the proxemics mappings used here did not directly reflect spatial relation between the user and the on-screen data range (as does Fig. 3), rather the data range was mapped onto the floor. The study revealed some interesting interactions nonetheless: You can step back to get an overview or walk to the left-hand side of the display to re-find previously seen details. The purpose of our empirical studies was not to provide detailed experimental data on the cognitive benefits of proxemics in visualization, but we think exploring this is important future work.

Our studies also suggested a need to get the fine details of interaction right. Participants needed a way of locking, both when using orientation and when using their body to change views: For instance, leaning forward, close to the screen, could lock the screen. Such interactions could derive from more sophisticated proxemics data for distinguishing between relative poses of different parts of the body (e.g., shoulder relative to torso or hip) in addition to distance. Alternatively, users could have discrete zones for interacting through touch (close to the display) and for navigating through movement of the body (farther from the display). Also, proxemics-enhanced visualizations in our studies occasionally had unintended consequences: When participants in Study #3 moved to brush coordinated views, they sometimes changed the attribute unintentionally. Giving users more feedback on the sensed proxemics data might alleviate some of these problems. Vogel and Balakrishnan [53] also found that users were sometimes unsure about the exit threshold of a distance zone.

The idea of using proxemics for interacting with information on large displays is not new. Recent work has for instance demonstrated use of discrete distance zones for changing layout and representation of information [3,53]. The present work differs from previous work by explicitly relating proxemics to information visualization tasks; the studies demonstrate mapping of movement, orientation, and distance (continuous measures as well as discrete zones) to visualize, filter, sort, select, navigate, and coordinate tasks [18].

Also, whereas previous work has investigated mainly *static* visualizations on large high-resolution displays [2,8,56], the present research has investigated physical navigation for *interactive* visualizations, which presents new opportunities. For instance, Endert et al. showed that different encodings offer varying support for visual aggregation and thus impact the effectiveness of large-display visualizations [8]: "To support physical navigation, encodings need to have a balance between the expressiveness of glyphs and good visual aggregation properties." However, the findings from the present studies suggest that alternative designs are possible that allow users to benefit from different encodings at different distances and from more generally changing visualizations through movement.

Our studies suggest several avenues of future work; in particular we want to highlight four of these: (a) We have prototyped and evaluated uses of movement and distance for information visualization, but uses of other proxemics categories need to be explored in more depth, as well as combination of proxemics-driven interactions with other input (as already discussed above); (b) our aim was not to empirically understand the cognitive benefits of proxemics in visualization, this is important future work; (c) we have focused on single-user interaction, but proxemics may help us design visualizations for multiple users—to help doing so, future work should relate proxemics to research on collaborative visualization; (d) we have not looked at combining proxemics with other emerging interaction styles, such as mid-air pointing (e.g., [35]) or free hand gestures, which is another promising avenue of future research.

8 Conclusion

The present paper has presented findings from initial probing into proxemics-based interactions with visualizations. We intend to experiment further with combining proxemics-driven interactions and other input for information visualization; the studies presented here are intended to lend credibility to the hypothesis that it is useful (and even pleasant) to control and interact with visualizations using ones body movements.

Acknowledgments

This work has been supported in part by the Danish Council for Strategic Research grant 10-092316.

References

[1] R. Amar, J. Eagan, and J. Stasko, "Low-Level Components of Analytic Activity in Information Visualization," in *Proceedings of the 2005 IEEE Symposium on Information Visualization*, IEEE Computer Society, Washington, DC, USA, pp. 15–21, 2005.

[2] R. Ball, C. North, and D. A. Bowman, "Move to improve: promoting physical navigation to increase user performance with large displays," in *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM, New York, NY, USA, pp. 191–200, 2007.

[3] T. Ballendat, N. Marquardt, and S. Greenberg, "Proxemic interaction: designing for a proximity and orientation-aware environment," in *ACM International Conference on Interactive Tabletops and Surfaces*, ACM, New York, NY, USA, pp. 121–130, 2010.

[4] A. Bezerianos and P. Isenberg, "Perception of Visual Variables on Tiled Wall-Sized Displays for Information Visualization Applications," *IEEE Transactions on Visualization and Computer Graphics (Proc. InfoVis)*, vol. 18, no. 12, Dec. 2012.

[5] S. K. Card, J. D. Mackinlay, and B. Shneiderman, *Readings In Information Visualization: Using Vision To Think*. Academic Press, 1999.

[6] S. Carpendale, "Evaluating Information Visualizations," *Information Visualization: Human-Centered Issues and Perspectives*, pp. 19–45, 2008.

[7] J. Cohen, "A power primer," *Psychol Bull*, vol. 112, no. 1, pp. 155–159, Jul. 1992.

[8] A. Endert, C. Andrews, Y. H. Lee, and C. North, "Visual encodings that support physical navigation on large displays," in *Proceedings of Graphics Interface*, Canadian Human-Computer Communications Society, pp. 103–110, 2011.

[9] A. Endert, P. Fiaux, H. Chung, M. Stewart, C. Andrews, and C. North, "ChairMouse: leveraging natural chair rotation for cursor navigation on large, high-resolution displays," in *CHI '11 Extended Abstracts on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 571–580, 2011.

[10] J.-D. Fekete and C. Plaisant, "Excentric labeling: dynamic neighborhood labeling for data visualization," in *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM Press, New York, NY, USA, pp. 512–519, 1999.

[11] G. W. Furnas, "Generalized fisheye views," in *CHI '86: Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM Press, New York, NY, USA, pp. 16–23, 1986.

[12] S. Greenberg and B. Buxton, "Usability evaluation considered harmful (some of the time)," pp. 111–120.

[13] S. Greenberg, N. Marquardt, T. Ballendat, R. Diaz-Marino, and M. Wang, "Proxemic interactions: the new ubicomp?," *interactions*, vol. 18, no. 1, pp. 42–50, Jan. 2011.

[14] E. T. Hall, "A System for the Notation of Proxemic Behaviour," *American Anthropologist*, vol. 65, no. 5, pp. 1003–1026, 1963.

[15] E. T. Hall, *The Hidden Dimension*. Garden City, N. Y.: Doubleday, 1966.

[16] C. Harrison and A. K. Dey, "Lean and zoom: proximity-aware user interface and content magnification," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 507–510, 2008.

[17] K. Hawkey, M. Kellar, D. Reilly, T. Whalen, and K. M. Inkpen, "The proximity factor: impact of distance on co-located collaboration," in *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*, ACM, New York, NY, USA, pp. 31–40, 2005.

[18] J. Heer and B. Shneiderman, "Interactive dynamics for visual analysis," *Commun. ACM*, vol. 55, no. 4, pp. 45–54, Apr. 2012.

[19] M. R. Jakobsen and K. Hornbæk, "Sizing up visualizations: effects of display size in focus+context, overview+detail, and zooming interfaces," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 1451–1460, 2011.

[20] M. R. Jakobsen and K. Hornbæk, “Proximity and physical navigation in collaborative work with a multi-touch wall-display,” in *CHI EA ’12: Extended abstracts on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 2519–2524, 2012.

[21] Y. Jansen, P. Dragicevic, and J.-D. Fekete, “Tangible remote controllers for wall-size displays,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 2865–2874, 2012.

[22] W. Javed, S. Ghani, and N. Elmqvist, “Polyzoom: multiscale and multifocus exploration in 2d visual spaces,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 287–296, 2012.

[23] W. Ju, B. A. Lee, and S. R. Klemmer, “Range: exploring implicit interaction through electronic whiteboard design,” in *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, ACM, New York, NY, USA, pp. 17–26, 2008.

[24] D. A. Keim, “Information Visualization and Visual Data Mining,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 8, no. 1, pp. 1–8, Jan. 2002.

[25] J. Kjeldskov, M. B. Skov, and J. Stage, “Instant data analysis: conducting usability evaluations in a day,” in *Proceedings of the third Nordic conference on Human-computer interaction*, ACM, New York, NY, USA, pp. 233–240, 2004.

[26] D. Klinkhammer, M. Nitsche, M. Specht, and H. Reiterer, “Adaptive personal territories for co-located tabletop interaction in a museum setting,” in *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, ACM, New York, NY, USA, pp. 107–110, 2011.

[27] B. Lee, P. Isenberg, N. H. Riche, and S. Carpendale, “Beyond Mouse and Keyboard: Expanding Design Considerations for Information Visualization Interactions,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 12, pp. 2689–2698, Dec. 2012.

[28] N. Marquardt, T. Ballendat, S. Boring, S. Greenberg, and K. Hinckley, “Gradual engagement: facilitating information exchange between digital devices as a function of proximity,” in *Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces*, ACM, New York, NY, USA, pp. 31–40, 2012.

[29] N. Marquardt, R. Diaz-Marino, S. Boring, and S. Greenberg, “The proximity toolkit: prototyping proxemic interactions in ubiquitous computing ecologies,” in *Proceedings of the 24th annual ACM symposium on User interface software and technology*, ACM, New York, NY, USA, pp. 315–326, 2011.

[30] N. Marquardt and S. Greenberg, “Informing the Design of Proxemic Interactions,” *IEEE Pervasive Computing*, vol. 11, no. 2, pp. 14–23, 2012.

[31] H. M. Mentis, K. O’Hara, A. Sellen, and R. Trivedi, “Interaction proxemics and image use in neurosurgery,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 927–936, 2012.

[32] A. E. Mostafa, S. Greenberg, E. Vital Brazil, E. Sharlin, and M. C. Sousa, “Interacting with microseismic visualizations,” in *CHI ’13 Extended Abstracts on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 1749–1754, 2013.

[33] J. Mumm and B. Mutlu, “Human-robot proxemics: physical and psychological distancing in human-robot interaction,” in *Proceedings of the 6th international conference on Human-robot interaction*, ACM, New York, NY, USA, pp. 331–338, 2011.

[34] 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*, ACM, New York, NY, USA, pp. 279–288, 2007.

[35] M. Nancel, J. Wagner, E. Pietriga, O. Chapuis, and W. Mackay, “Mid-air pan-and-zoom on wall-sized displays,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 177–186, 2011.

[36] J. Nielsen and T. K. Landauer, “A mathematical model of the finding of usability problems,” in *Proceedings of the INTERACT ’93 and CHI ’93 Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 206–213, 1993.

[37] C. North and B. Shneiderman, “Snap-together visualization: a user interface for coordinating visualizations via relational schemata,” in *Proceedings of the working conference on Advanced visual interfaces*, ACM, New York, NY, USA, pp. 128–135, 2000.

[38] G. M. Olson and J. S. Olson, “Distance matters,” *Hum.-Comput. Interact.*, vol. 15, no. 2, pp. 139–178, Sep. 2000.

[39] S. M. Peck, C. North, and D. Bowman, “A multiscale interaction technique for large, high-resolution displays,” in *Proceedings of the 2009 IEEE Symposium on 3D User Interfaces*, IEEE Computer Society, Washington, DC, USA, pp. 31–38, 2009.

[40] E. Pietriga and C. Appert, “Sigma lenses: focus-context transitions combining space, time and translucence,” in *CHI ’08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, ACM, New York, NY, USA, pp. 1343–1352, 2008.

[41] E. Ramos and D. Donoho, “The 1983 ASA data exposition dataset: Cars.” 1983.

[42] R. A. Ruddle, E. Volkova, and H. H. Bülthoff, “Walking improves your cognitive map in environments that are large-scale and large in extent,” *ACM Trans. Comput.-Hum. Interact.*, vol. 18, no. 2, pp. 10:1–10:20, Jul. 2011.

[43] C. Schmidt, J. Müller, and G. Bailly, “Screenfinity: extending the perception area of content on very large public displays,” in *Proceedings of the 2013 ACM annual conference on Human factors in computing systems*, ACM, New York, NY, USA, pp. 1719–1728, 2013.

[44] J. Seo and B. Shneiderman, “Knowledge Discovery in High-Dimensional Data: Case Studies and a User Survey for the Rank-by-Feature Framework,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 3, pp. 311–322, 2006.

[45] W. R. Shadish, T. D. Cook, and D. T. Campbell, *Experimental and Quasi-experimental Designs for Generalized Causal Inference*. Boston, MA, USA: Houghton Mifflin, 2002.

[46] B. Shneiderman, “The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations,” in *Proceedings of the 1996 IEEE Symposium on Visual Languages*, IEEE Computer Society, Washington, DC, USA, p. 336–343, 1996.

[47] L. Shupp, C. Andrews, M. Dickey-Kurdziolek, B. Yost, and C. North, “Shaping the Display of the Future: The Effects of Display Size and Curvature on User Performance and Insights,” *Human-Computer Interaction*, vol. 24, no. 1, pp. 230–272, 2009.

[48] R. Sommer, *Personal Space: The Behavioral Basis of Design*. Prentice Hall, 1969.

[49] J. L. Souman, P. R. Giordano, M. Schwaiger, I. Frissen, T. Thümmel, H. Ulbrich, A. D. Luca, H. H. Bülthoff, and M. O. Ernst, “CyberWalk: Enabling unconstrained omnidirectional walking through virtual environments,” *ACM Trans. Appl. Percept.*, vol. 8, no. 4, pp. 25:1–25:22, Dec. 2008.

[50] M. Spindler, C. Tominski, H. Schumann, and R. Dachselt, “Tangible views for information visualization,” in *ACM International Conference on Interactive Tabletops and Surfaces*, ACM, New York, NY, USA, pp. 157–166, 2010.

[51] M. Tohidi, W. Buxton, R. Baecker, and A. Sellen, “Getting the right design and the design right,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY, USA, pp. 1243–1252, 2006.

[52] F. B. Viegas, M. Wattenberg, F. van Ham, J. Kriss, and M. McKeon, “ManyEyes: a Site for Visualization at Internet Scale,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1121–1128, Nov. 2007.

[53] D. Vogel and R. Balakrishnan, “Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users,” in *UIST ’04: Proceedings of the 17th annual ACM symposium on User interface software and technology*, ACM, New York, NY, USA, pp. 137–146, 2004.

[54] C. Williamson and B. Shneiderman, “The dynamic HomeFinder: evaluating dynamic queries in a real-estate information exploration system,” in *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval*, ACM, New York, NY, USA, pp. 338–346, 1992.

[55] J. S. Yi, R. Melton, J. Stasko, and J. A. Jacko, “Dust & magnet: multivariate information visualization using a magnet metaphor,” *Information Visualization*, vol. 4, no. 4, pp. 239–256, 2005.

[56] B. Yost and C. North, “The Perceptual Scalability of Visualization,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 837–844, Sep. 2006.