

A Review of Focus and Context Interfaces

ANDY COCKBURN^{*}, AMY KARLSON[†], BENJAMIN B. BEDERSON[†]

There are many interface schemes that allow users to work at, and move between, focused and contextual views. We review and categorise these schemes according to the interface mechanisms used to separate and blend views. The four approaches are spatial separation, typified by overview+detail interfaces; temporal separation, typified by zoomable interfaces; seamless focus+context, typified by fisheye views; and cue-based techniques which selectively highlight or suppress items within the information space. Critical features of these categories, and evidence of their success, are discussed. The aim is to provide a succinct summary of the state-of-the-art, to illuminate successful and unsuccessful interface strategies, and to identify potentially fruitful areas for further work.

Categories and Subject Descriptors: D.2.2 Design Tools and Techniques—*User Interfaces*; H.5.2 User Interfaces—*Graphical User Interfaces (GUI)*

General Terms: Human Factors

Additional Key Words and Phrases: Information display, information visualization, focus and context, overview and detail, zoomable user interfaces, fisheye views, review paper.

1. INTRODUCTION

In almost all computer applications, users need to interact with more information and with more interface components than can be conveniently displayed at one time on a single screen. This need is dictated by pragmatic, technological, and human factors. The pragmatic issues concern form-factors such as the size, weight, and fashion of displays that are used for varied tasks in diverse locations, as well as the cost of construction. Technological limitations constrain the ability of displays to match the breadth and acuity of human vision. Humans, for instance, can see through a visual angle of approximately $200^{\circ} \times 120^{\circ}$ (without turning the eyes or head), while typical computer displays extend approximately 50° horizontally and vertically when viewed from normal distances (Woodson and Conover 1964). Similarly, the ‘resolution’ of the human eye, with approximately 180 cones-receptors per degree in the fovea (Ware 2000), can perceive

^{*} Department of Computer Science and Software Engineering, University of Canterbury, Christchurch, New Zealand, andy@cosc.canterbury.ac.nz

[†] Human-Computer Interaction Laboratory, Computer Science Department, Institute for Advanced Computer Studies, Univ. of Maryland, College Park, MD 20742, {akk; bederson}@cs.umd.edu

Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2006 ACM 1073-0516/01/0300-0034 \$5.00

around 200 dots per cm at a viewing distance of 50cm—far more than the typical display resolution of 40 dots per cm. Finally, even if a ‘perfect’ display could be constructed, displaying all information simultaneously is likely to hinder the user’s ability to distinguish between salient and irrelevant information.

The traditional interface mechanisms for dealing with these display trade-offs involve allowing information to be moved (typically through paging, scrolling and panning) or spatially partitioned (through windowing, menus, and so on). Although scrolling and windowing are standard in almost all user interfaces, they introduce a discontinuity between the information displayed at different times and places. This discontinuity can cause cognitive and mechanical burdens for the user who must mentally assimilate the overall structure of the information space and their location within it, and manipulate controls in order to navigate through it. For example, O’Hara and Sellen (1997) note that scrolling is “irritatingly slow and distracting”, and that it harms user’s incidental memory for the location of document features (O’Hara, Sellen and Bentley 1999).

An alternative to moving or spatially partitioning the data space is to vary the scale at which the data is displayed. There are many interface variations on the scaling theme, including overview+detail techniques, zooming, and fisheye views. In this paper, we use the term “focus and context interfaces” to describe this wide range of interaction techniques. These schemes promise several potential advantages derived from their ability to allow users to rapidly and fluidly move between focused views and contextual overviews. Note that our use of the term “focus and context” encapsulates a superset of the research normally described under the banner term “focus+context”, which traditionally describes just distortion-oriented visualisation techniques. Additionally, we include in our review cue-based techniques, which modify and supplement the information space with cues regarding focal items and their context.

1.1 Road Map

The objective of this paper is to summarise the state of research on focus and context interfaces, to identify effective and ineffective uses of them, and to identify promising areas for further work. This review is motivated by three factors. First, advanced focus and context interaction techniques are increasingly being deployed in desktop interfaces, often ahead of research into the implications of doing so: for example, Apple deployed the Mac OS X Dock[†] providing a fisheye-distortion of items in its icon panel prior to

[†] www.apple.com/macosx/theater/dock.html

research showing that its target-moving behaviour is detrimental to efficiency (Gutwin 2002). Second, there has been a dramatic increase in the range of tasks being supported through an increasingly wide range of focus and context interaction techniques, and the relationship between many of these tasks and techniques have not been explored. Finally, although there have been previous reviews of related work (Section 2), they are now dated and cover a subset of focus and context research. In particular, they were published when the research focus was on developing new interaction techniques rather than on empirically derived understanding of the techniques' effectiveness.

Our review is organized into four categories that distinguish systems according to the nature of the user's interaction. First, we describe work on interfaces that use a spatial separation to display focal and contextual information—these are best exemplified by prior work on “overview+detail” interfaces. The second category of interfaces uses a temporal separation between focus and context—best exemplified by zooming. The third ‘seamless’ category simultaneously reveals focus and context within a continuous display—best exemplified by distortion-oriented visualisations such as ‘fisheye views’. Unlike the other categories, the fourth ‘cue’ category does not necessarily alter the scale at which the information is displayed; rather, it modifies the way in which items are depicted in order to highlight, suppress, or contextualise them. For example, search results can be highlighted, or non-results could be visually deemphasized, to draw the user's attention to elements of interest. The cue category depends on the availability of semantic information about elements in the data-set, and can be used to enhance any of the other three techniques.

Within each category we present an overview of the discriminating features, a review of their history, foundations, and objectives, and we identify commercial exemplars where they exist. We also review research systems that demonstrate novel variations on the theme. After the description of the categories, we summarise the empirical work that identifies the strengths and weaknesses of the techniques. The review of empirical work is presented in two broad categories of user task: low-level ‘motor’ tasks such as target acquisition; and high level ‘cognitive’ tasks such as the user's ability to comprehend the information space. We finish the paper by presenting summary guidelines and agendas for further research.

2. PREVIOUS REVIEWS OF FOCUS+CONTEXT RESEARCH

There have been several previous reviews of focus+context research. Leung and Apperley (1994) provided the first comprehensive review of distortion-oriented visualisations, using a metaphor of a rubber sheet to unify the theoretical properties of

different focus+context schemes. The rubber-sheet analogy had previously been used to describe visualisations by Tobler (1973) Mackinlay, Robertson and Card (1991), and Sarkar, Snibbe and Reiss (1993). More recently, Carpendale and Montagnese (2001) presented a mathematical framework, called the ‘Elastic Presentation Framework’ that formalises the relationship between diverse focus+context visualisation schemes. Herman, Melancon and Marshall (2000) review graph visualization schemes, including focus+context views, with an emphasis on the algorithms used to generate the graphs. Following the research emphasis at the time, the focus of all of these reviews is on describing the properties of the visualisations generated, rather than on empirically comparing their effectiveness at supporting the users’ tasks.

Plaisant, Carr and Shneiderman’s (1995) review of focus+context systems included a taxonomy of users’ tasks with image-browsers, but the lack of empirical evidence at the time inhibited their ability to provide strong guidance on the task-applicability of the schemes reviewed.

In a side-bar to their paper, Kosara, Miksch and Hauser (2002) provide a taxonomy of focus+context research that is similar to the one presented in this paper. They present three groups of focus+context methods: spatial methods, dimensional methods, and cue methods. Unfortunately, their side-bar presents only a few sentences on each technique.

Focus+context research is also reviewed in several books on information visualisation, notably the following: Ware (2000), which excels in the psychological foundations of visualisation; Spence (2001), which provides a strong system-oriented review; and Card, Mackinlay and Shneiderman (1999), which provides a collation of the seminal papers in information visualization.

3. OVERVIEW+DETAIL — SPATIAL SEPARATION

An overview+detail interface design is characterized by the simultaneous display of both an overview and detailed view of an information space, each in a distinct presentation space. Due to the physical separation of the two views, users interact with the views separately, although actions in one are often immediately reflected in the other.

Many forms of overview+detail interfaces exist, both in the standard desktop environment and in research systems. Their key discriminating features are as follows:

- Minimum and maximum scale ratios. The ratio between the proportion of the workspace shown in the overview-area and in the detail-region is called the scale-ratio. Most overview+detail systems have minimum and maximum scale ratios beyond which they become ineffective.

- Geometry. The relative size and position of the overview and detail regions is often related to the range of possible scale ratios. As well as using x, y-coordinate displacement for the overview and detail region, some systems also make use of the z-coordinate by overlaying either the overview or detail on top of the other.
- Configurable scale ratios. Many overview+detail interfaces allow the user to tailor the scale-ratio. A wide variety of interface metaphors have been used to configure the scale ratio.
- View abstraction. Although most overview+detail interfaces provide a miniaturised representation of the workspace inside the overview region, some use abstract or ‘semantic’ representations.
- Navigation controls. The control mechanisms for navigating through the information space differ widely. In some interfaces, the overview region is the primary control for navigation, while in others the overview region provides a passive view.
- Synchronisation of views. Although it is potentially powerful to allow users to see different document regions in the overview and detail areas, it is also potentially disorientating. Some systems allow separation of views, while others do not.

In the remainder of this section we describe exemplar overview+detail interfaces and discuss the design issues involved in mapping their user-support to these discriminating features.

3.1 Scrollbars, Embellished Scrollbars and Thumbnail Overviews

Scrollbars are a familiar component of graphical user interfaces, and they can be considered to provide overview+detail functionality. The position of the scroll-thumb within the scroll-trough represents the current location of the detail-view within the overall document. The length of the scroll-thumb also depicts the proportion of the document currently displayed.

The minimum scale-ratio for scroll-bars is 1:1, and they are often removed from the display at this ratio because they serve no clear purpose with the entire document shown in the detail view. Their maximum ‘truthful’ scale-ratio is on the order of 100:1, but they are robust in that they remain fairly effective at very high scale-ratios: for example, scrollbars can be used to navigate through documents that are several hundreds of pages long. At high scale-ratios the length of the scroll-thumb does not accurately reflect the

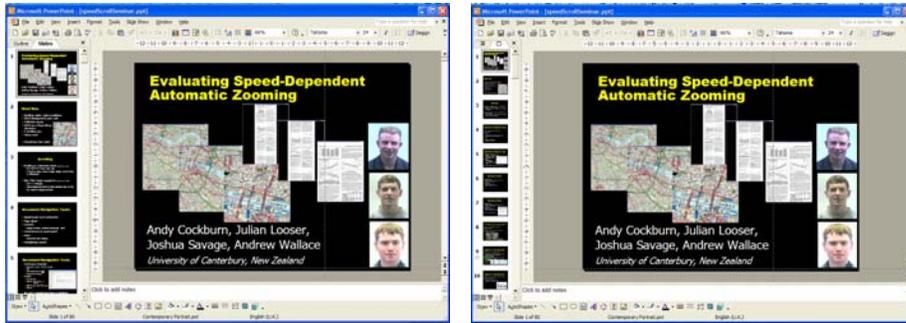
proportion of the information space shown in the detail region; to do so truthfully could cause the thumb to shrink to less than one pixel, creating clear usability problems in acquiring the thumb. For this reason, a minimum thumb-size of approximately ten-pixels is used, which reduces the problems of target acquisition but introduces problems of precise document control because small thumb movements causes large ‘jerky’ document movements. At high scale-ratios, users must find alternative techniques to smoothly move their documents, such as the scrollbar end-arrows, the cursor keys, or rate-based scrolling.

Standard scrollbars do not allow direct configuration of the scale ratio, but there is no reason why they should not. In standard desktop interfaces, length of the scroll-thumb immediately adapts to changes in zoom-level within the detail region. There is no reason why the reverse control could not also be implemented, allowing the zoom-level to be controlled directly by manipulating the length of the scroll-thumb. Range-sliders (Ahlberg and Shneiderman 1994) demonstrate this capability.

Although scrollbars typically encode only spatial information, several researchers have experimented with variants that additionally portray semantic information. *Value bars* (Chimera 1992) embellished the scroll trough with dashed lines, the thickness of which depended on numerical values extracted from columns in formatted data files, to depict the location of ‘interesting’ or outlier data. Hill and Hollan (1992) also used similar techniques to convey the read and edit history of document regions. Byrd (1999) evaluated a related technique with inconclusive results but positive subjective responses from participants. In addition to encoding semantic information about document contents, scroll troughs have also been used to provide awareness of other people’s locations and actions in real-time collaborative tasks (Gutwin, Roseman and Greenberg 1996).

As the level of detail presented in the scroll trough increases, the scrollbar becomes a first-class overview window. Many everyday applications, such as Microsoft PowerPoint (Figure 1) and Adobe Reader, support ‘thumbnail’ document overviews that blur the boundary between scrollbar and overview for navigation.

The effective range of scale-ratios for thumbnail overview is fairly narrow as they provide little benefit in short-cut navigation below 5:1 and they become difficult to distinguish beyond 20:1 (when organised in a single vertical array). Within this range, it is common for interfaces to allow configuration by scaling the proportion of the window allocated to the thumbnail overview (Figure 1 shows Microsoft PowerPoint with two different overview-region sizes).



(a) Five thumbnails in the overview. (b) Ten thumbnails in the overview.

Figure 1: PowerPoint’s overview+detail interface. The scrollable thumbnail overview is on the left-hand side of each window, and the thumbnails can be scaled by configuring the width of the overview region.

Although it is clear that the thumbnails should allow shortcut navigation (clicking on a thumbnail should cause that region of the information space to appear in the detailed view), it is less clear whether the overview and detailed regions should be synchronised so that they continually display corresponding locations. Synchronisation is much less likely to cause disorientation, but it is less powerful as it forbids independent exploration of document regions in the two views. Many variant implementations are possible: for example, Microsoft Powerpoint implements a one-way synchronisation in which the overview is synchronised with the detail view (scrolling the detail causes corresponding movement in the overview), while the detail is unsynchronised with the overview (scrolling the overview has no side-effect on the detail). Given the wide variety of application areas for overview+detail interfaces, it is unlikely that generic guidelines for the ‘right type of synchronisation’ can be provided.

3.2 Gestalt/Radar Views

Gestalt views are commonly used to provide a small overview of the entire information space. The first examples of gestalt/radar views (that we are aware of) are in computer games of the early 1980s, such as Defender (Figure 2). They are sometimes displayed in a separate region displaced on the x,y-coordinates, but they also often overlay a portion of the detail area. As well as providing a miniaturised overview of the entire workspace, they normally also show the region of the workspace displayed in the detail area. Many systems allow the user to control the region portrayed in the detail region by manipulating the window region shown in the overview.

Like thumbnail overviews, the maximum effective scale-ratio is constrained to around 20:1 by the limits of the user’s ability to extract meaningful information from the



Figure 2. An overview+detail interface in the early 1980s game Defender. The white-bordered area at the bottom of the window provides a gestalt/radar overview of the entire game region, with the bounds of the current window depicted in the centre of the overview.

diminished content. Cue-based techniques, such as the Popout-Prism (Section 6), can be used to extend the scale-ratio by providing semantic enhancements to the information shown in the overview.

3.3 Wide field-of-view systems

Although most overview+detail systems present both the overview and the detail regions within a single standard (or mobile) display, several researchers have examined the impact of broadening the display's field of view.

One approach is to use a large projected display. However, these normally achieve their size by simply enlarging each pixel: the resolution remains constant. To support both large size and high resolution, Baudisch, Good, Bellotti and Schraedley (2002) constructed a 'focus+context' display that integrated a small focused display region (1024x768 pixels, flat-panel display) within a large 4x3ft, 1024x768 pixels projected display. The images presented in the two displays were stitched together in software to ensure that panning actions in one caused corresponding updates to the other. Preliminary evaluations showed performance advantages for the focus+context screen in comparison to overview+detail and zoom+pan interfaces. This technology, however, is a stop-gap measure until high-resolution large displays, such as the tiled wall displays by Guimbretiere, Stone and Winograd (2001), are available at low cost.

3.4 Lenses and Z-separation

The systems and techniques described above all separate the overview and detail regions on the x and y coordinates. Several systems, however, separate the views on the z-coordinate, with overviews overlaying or blended with the background detail.

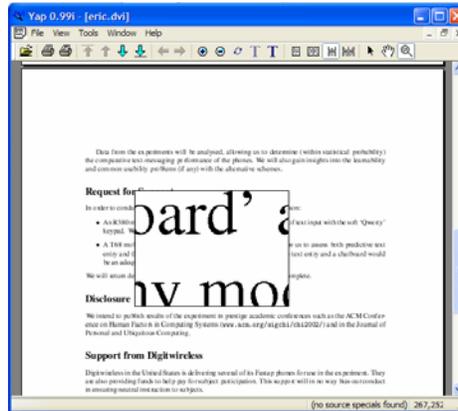


Figure 3: Z-based overview+detail separation. In the “Yap” dvi file previewer a magnified detail region is shown in a lens that follows the user’s cursor.

Lenses are moveable display regions that overlay the default window. Although lenses can be similar to ‘fisheye’ distortions (described later), we categorise them as ‘overview+detail’ because they separate (on the z-plane) the detail and overview.

Lenses have been used as a magnification metaphor in standard desktop environments since the early 1990s: Figure 3 shows a magnification lens in the document previewer ‘Yap’: the magnified region follows the user’s cursor. Bier, Stone, Pier, Buxton, and DeRose (1993) introduced Toolglass widgets and Magic Lenses together under the same design philosophy. Toolglass widgets are resizable see-through windows, normally controlled by the user’s non-dominant hand, that allow users to perform specialized operations on the data space objects over which they are positioned. For example, to change the colour of an object in a drawing application the user would place the colour-selector toolglass over the target using their non-dominant hand, and then click-through the toolglass using the dominant hand. Magic Lenses have the same form as Toolglass widgets, but they transform the visualisation of the underlying objects to allow focused viewing of specific attributes.

Although lenses are normally much smaller than the underlying detailed area, the lens can be the same size as the detailed region, with transparency visually separating the overview and detail ‘layers’. Cox, Chugh, Gutwin, and Greenberg (1998) evaluated one such layered ‘lens’ implementation in a pipeline construction task, with positive results. Display hardware is also available to implement the layered separation. PureDepth[‡] manufactures a dual-layer LCD screen, with a small (1-5cm) separation between front

[‡] www.puredepth.com

and rear LCD layers. Images displayed on the front layer can partially or entirely occlude those displayed on the rear layer, aiding the cognitive separation of displayed layers. An empirical evaluation the effectiveness of PureDepth displays failed to show any advantage over more traditional alpha-blending techniques (Aboelsaadat and Balakrishnan 2004).

4. ZOOMING — TEMPORAL SEPARATION

The second basic category of focus and context interfaces is based on zooming, which involves a temporal separation between the display of focused and contextual views. Naturally, zooming and overview+detail features can be combined, but in this section we focus on the isolated issues of zooming.

Like overview+detail interfaces, some zooming techniques are standard features of desktop user interfaces, but many others remain confined to research systems. Also like overview+detail interfaces, there are few concrete empirical lessons to guide the design of zoomable interfaces. Some of the critical design issues for zoomable interfaces, further discussed in this section, are as follows:

- Discrete versus continuous zoom. Discrete zooming allows users to move between predetermined zoom levels, while continuous zooming allows any zoom level (between the maximum and minimum) to be selected.
- Animation. Zooming involves rapid transitions between display states. These transitions can be abrupt and disorienting unless animation is used to reveal the relationship between the pre- and post-zoom states. The drawback of animation is that it is computationally demanding, which can reduce system performance. In addition, animation takes time which has the potential of slowing down overall performance.
- Zoom-in and zoom-out interface controls. Many zooming interfaces allow the user to select the focal point or region for zooming-in by direct manipulation, but this creates a problem for the inverse operation of zooming-out because the zoom-out region is outside the current view. In addition, many different interaction techniques have been used to let users control the magnification.
- Manual versus automatic zooming. Recent work has investigated automatic zooming in which the zoom-level is automatically adjusted in response to some other facet of user interface control, such as scroll-speed.

- Magnification versus semantic zooming. Most zooming interfaces simply magnify the information as the zoom-level increases, but semantic zooming can alter the representation to maximise the effectiveness of the display-space available at different zoom-levels.

4.1 Standard desktop applications of zooming

Most standard desktop interfaces allow the user to zoom the information space through controls such as a percentage zoom selector in the toolbar or a magnification mode which zooms by a small amount on each mouse-click. These techniques allow users to move between discrete zoom levels. Smooth and continuous zoom controls are also appearing in commercial systems through techniques such as dragging with the middle-mouse button, control-scrollwheel actions (e.g. Microsoft Word), and explicit zoom modes (e.g. Adobe Reader's "Dynamic Zoom"). Region-select zooming is another common control interface, allowing users to drag out a region to be magnified, either in an overview window (e.g. Adobe Reader) or in the main window (many image editing tools). Content-based zooming has also been used to navigate among large image collections (www.photomesa.com).

Anecdotal evidence suggests that controls for zooming actions and their counter-actions should be clearly associated with one another. Undo, the standard mechanism for reversing interface actions, is normally reserved for actions that modify the data-state, so actions that modify the view (such as zoom) cannot be reversed with Undo. Discrete "percentage zoom" controls are unlikely to cause a problem because the control is identical for all zoom-levels, but zoom modes where zoom-in is issued by simple mouse-clicks and zoom-out is controlled by right-, shift-, ALT-, OPT- or FNC- (all used in different commercial systems) introduces a control separation that may not be obvious to the user. Region-select for zoom-in is also risky (but powerful for appropriate users) because there is no equivalent reversing action, unless an overview window is available.

4.2 Zooming toolkits and trajectories

Many zoom-based research systems have been developed to demonstrate interaction and visualization techniques (Card et al. 1999). The Pad system (Perlin and Fox 1993) was the first fully zoomable desktop environment, and it introduced two important concepts: *semantic zooming*, which allows objects to be represented differently at different scales; and *portals*, which allow links between data objects and filters on their representation. Pad prompted extensive further research on these and related topics, and toolkits were developed to ease the implementation of zoomable interfaces, including Pad++

(Bederson, Hollan, Perlin, Meyer, Bacon and Furnas 1996), Jazz (Bederson, Meyer and Good 2000), and Piccolo (Bederson, Grosjean and Meyer 2004). Several application domains have been explored using these toolkits, including drawing tools for children (Druin, Stewart, Proft, Bederson and Hollan 1997), authoring tools (Furnas and Zhang 1998), and image browsers (Bederson 2001).

The experience of building many zoomable applications across multiple domains revealed domain independent design challenges, which Furnas and Bederson (1995) addressed with Space-Scale diagrams, a framework for conceptualizing ZUIs (zoomable user interfaces) and for building efficient navigation trajectories. van Wijk and Nuij (2004) formalised the mathematics involved in generating ‘smooth and efficient animation trajectories’ in pan-zoom space (further discussed below). Another general usability problem associated with zooming is ‘Desert Fog’ (Jul and Furnas 1998), which encapsulates navigations problem caused by feature separation at high zoom levels. To combat Desert Fog, Jul and Furnas introduce ‘critical zones’ (a cue-based technique) to provide a visual demarcation of regions that are guaranteed to yield further information when zoomed.

4.3 Automatic zooming and parallel zoom/pan control

When zooming actions cause a substantial view change, animation is important in helping the user assimilate the relationship between pre- and post-zoom states. Animation causes a brief period of automatic zooming: rather than controlling the animation, the user simply observes it. These softening and spatial-orienting effects are finished within hundreds of milliseconds.

Designing zoom-animations requires finding a suitable transition speed that reveals the relationship between zoom states without slowing the user’s overall interaction. Research suggests that animations should last between 0.3 and 1.0 second (Card, Robertson and Mackinlay 1991; Bederson and Boltman 1999; Klein and Bederson 2005). Other studies emphasise the importance of animation, showing that it improves user performance in scrolling (Klein and Bederson 2005) and in graph navigation (Summers, Goldsmith, Kubica and Caudell 2003). The ideal duration and nature of animation, then, varies with its purpose and nature. For example, most desktop environments allow the user to tailor the animation effect used to display menu items: menus can rapidly ‘scroll’ or ‘fade’ into/out-of view rather than abruptly appearing. The sole purpose of this animation is to ‘soften’ the otherwise abrupt display transition. Windows also display ‘zoom-lines’ as they enter and leave the icon-panel, helping users spatially orient themselves through cues to the items’ location.

Sometimes, however, users need richer cues to maintain comprehension and orientation throughout the animation, demanding longer animated periods and more detailed displays throughout. To this end, van Wijk and Nuij (2004) developed a set of equations for optimising pan/zoom trajectories between two known points in pan-zoom space (from one point in 2D space at one level of zoom to another point at a different level of zoom). Their formulae aim to minimise motion-blur at the user's eye, parameterized on animation velocity (V) and a user-specified pan/zoom trade-off parameter (ρ).

The formulae of van Wijk and Nuij calculate smooth and efficient trajectories for motion between *known* points. Most information browsing, however, involves dynamic modification of paths through scroll/pan (movement) and zoom (scale) actions in response to system feedback. Traditional interfaces impose a serial separation between pan/scroll movements and scale-changing zoom actions, but two recent research threads are investigating concurrent control for motion and zoom.

First, research has shown that target acquisition can be improved by allowing parallel bi-manual (two-handed) control over pan and zoom rather than demanding a serial modal separation of controls (Bourgeois and Guiard 2002). Several studies also show that parallel bimanual input improves over serial input in other tasks such as rectangle editing in drawing tools (Leganchuk, Zhai and Buxton 1998; Casalta, Guiard and Beaudouin-Lafon 1999), in scrolling tasks without zoom (Zhai, Smith and Selker 1997), and in virtual image alignment (Latulipe, Kaplan and Clarke 2005).

Second, several researchers are investigating interfaces that automatically adjust the zoom-level, dependent on some other control such as movement speed. "Depth Modulated Flying" (Ware and Fleet 1997) first demonstrated the automatic coupling of zoom and motion in 3D "fly by" visualisations, with the fly-by motion speed automatically adjusted in response to the user's control of zoom-level. Tan, Robertson, and Czerwinski (2001) inverted the speed/zoom binding in their investigation of "speed-coupled flying with orbiting" in which the user's perspective (effectively zoom-level) is adjusted in response to speed of movement in a 3D world. They showed that the automatic coupling of zoom to speed improved the time to complete navigational tasks by ~10% to ~31%.

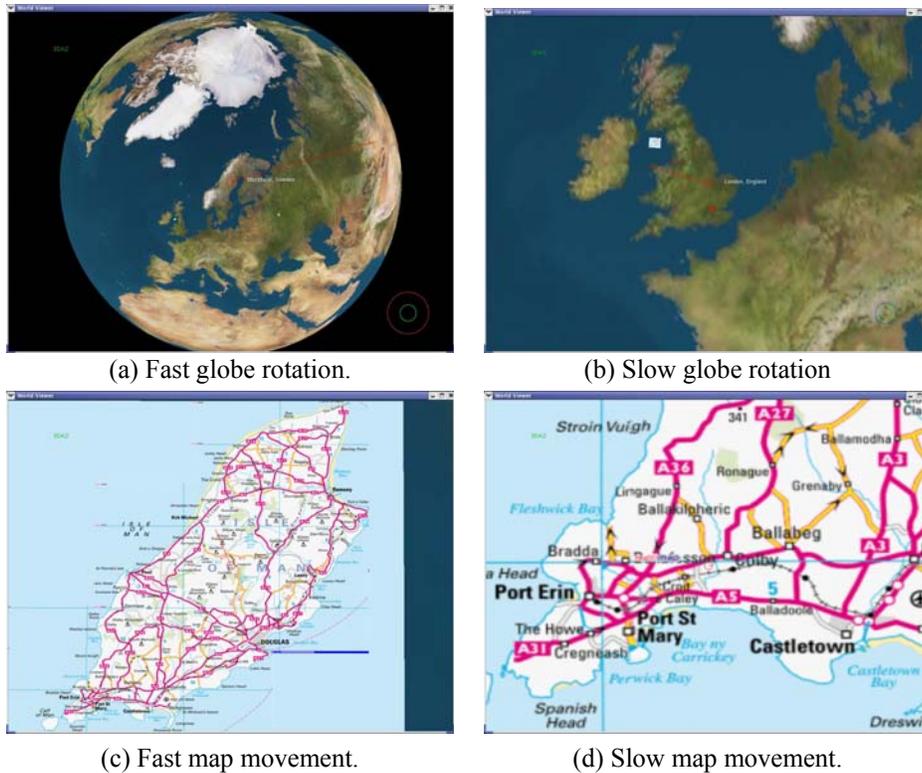


Figure 4. A speed-dependent automatic-zooming ‘globe-browser’ interface (Savage and Cockburn 2005).

In more standard desktop applications Igarashi and Hinckley (2000) implemented several ‘speed dependent automatic zooming’ (SDAZ) systems to overcome the problems of motion-blur when users scroll rapidly. By automatically zooming away from the document as the scroll-rate increases, the pixel-rate of movement is kept within human perceptual limits. Although results from their preliminary evaluations were inconclusive, Cockburn and Savage (2003) and Cockburn, Savage and Wallace (2005) showed that SDAZ allows significant performance gains over traditional panning, scrolling, zooming and using an implementation of van Wijk and Nuij’s formulae. Figure 4 shows one of their systems—a speed-dependent automatic zooming ‘globe-browser’. The globe view (Figures 4a, b) is automatically zoomed out as the rotation speed increases, but when the user slows or stops over a sub-map the system automatically zooms into it (Figures 4c, d) with the same coupling of zoom-level to scroll-speed within the sub-map.

5. FISHEYE DISTORTION — SEAMLESS FOCUS+CONTEXT

The zooming categories discussed so far separate (spatially or temporally) the display of focus and contextual information, leaving the user to assimilate the relationship. The third

approach, which we call “seamless” focus and context, integrates focus and context into a single display where all parts are always visible. The terms “fisheye view” and “focus+context display” are often used to describe this category of visualisation because distortion of the information space is commonly used to enable the focus region to blend into the background context. The aim of “seamless” systems is to enhance the user’s ability to comprehend and manipulate the information due to all regions being continuously present in the display, thus potentially reducing short term memory load since all information is presented at the same time in a single coherent space.

This section provides a brief review and introduction to the wide range of systems, techniques and theories encapsulated in the “seamless focus+context” category. Recurring themes that distinguish the different research contributions include the following:

- Elision versus continual display. Some systems completely remove data items from the display when they are semantically or spatially distant from the focal region. Other systems guarantee that all data items remain visible, even though they may be substantially diminished.
- Distortion function. While some systems distort the information space based on the Euclidian-distance from the focus, others use a structural or semantic function for ‘distance’.
- Smooth-continuous versus discrete-abrupt distortion. Related to the previous point, certain data-types (such as maps and images) are natural candidates for smooth and continuous distortion based on Euclidian-distance, while other data-types such as tables may become hard to interpret if continuous distortion is applied and so use distortions that keep horizontal lines horizontal and vertical lines vertical.
- Single versus multiple foci. While most systems only support a single focal-point, some support an arbitrary number of concurrent focal points.

5.1 Visions and theoretical foundations

Spence and Apperley (1982) described the first seamless focus and context display. Their conceptual “Bifocal Display” used a metaphor of paper stretched across four rollers, with two close rollers giving a rectangular focal region, and two distant rollers on either side giving receding side-plate displays. The user changed their focus by sliding the paper in either direction across the frame. Nearly a decade later Mackinlay, Robertson and Card (1991) implemented the concept in their Perspective Wall, shown in Figure 5.

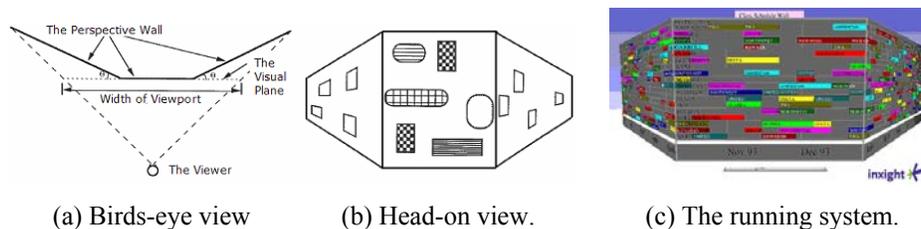


Figure 5. The Perspective Wall (Mackinlay et al. 1991).

The theoretical foundations for seamless focus+content interfaces were first described by Furnas (1986), who described a “generalised fisheye views” formula for calculating the user’s ‘degree of interest’ (DOI) in objects in the data-space: $DOI(x | y) = API(x) - D(x,y)$, where x is a data element, y is the current focus, $API(x)$ is the *a priori* interest in object x , and $D(x,y)$ is the spatial or semantic distance between x and y . Furnas described how his formula could be applied to a variety of information domains, with data objects being elided when they fell below a threshold DOI value for display.

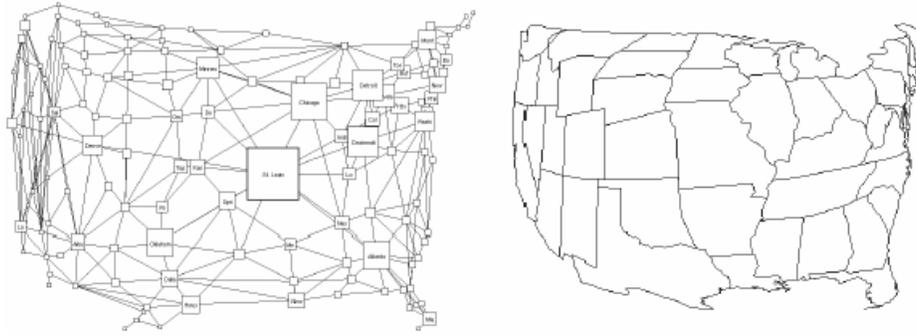
Sarkar and Brown (1992) extended fisheye views to generic graphs and maps by using the Euclidean distance between graph-vertices or map-coordinates. They then produced rich fisheye displays using geometric transformations that determine each point’s size, position, and level of detail in the display (see Figure 6). The Sarkar and Brown algorithm has been heavily used in fisheye visualisations. Lamping, Rao and Pirolli (1995) present an alternative method for seamlessly integrating focus and context based on hyperbolic geometry, which they claim is particularly suitable for layout of hierarchical data.

5.2 Sample applications

Many systems demonstrating seamless focus and context have been developed, but relatively few have been empirically evaluated. This section describes some of the more notable applications and, where possible, it reviews empirical evidence of their success.

5.2.1 Fisheyes for targeting in desktop applications

The MacOS X “Dock” icon-panel incorporates the first large-scale deployment of fisheye-style effects. Items in the icon-panel expand as the user’s cursor moves towards them, providing a dynamic and rich visual effect that many users appreciate. Despite the effect’s popularity, it can frustrate the user’s ability to acquire targets through ‘hunting effects’ that are caused by a separation between the visual location of items and the location of the motor-space that activates them. Details of this effect, and of associated evaluations, are presented in Section 7.1.



(a) Topological map of US cities.

(b) Polar transformation of US States.

Figure 6: Sarkar and Brown (1992) fisheye transformations.

Theoretically similar to the icon-panel, Bederson (2000) developed fisheye menus to help selecting items in long menus: for example, selecting a country-name from a combo-box. Each line of text in the fisheye menu is displayed in an extremely small font, with the font-size increasing as the cursor moves closer. Like other fisheyes, fisheye menus suffer hunting effects, which the design ameliorates by allowing users to ‘lock’ the lens by moving the cursor into a region to the right of the normal menu. Evaluation showed that fisheye menus were faster than traditional scrolling menus, but slower than hierarchical ones. Subjective responses from users also suggest that the dual-mode targeting (fisheye, then ‘locked’) demands higher levels of concentration than traditional menu selections.

5.2.2 *Fisheye documents*

Furnas’s original description of the fisheye technique used the hierarchical structure of computer programs to demonstrate the technique. When the fisheye formula calculated that a program region had a low “degree of interest”, that region was ‘elided’ (or removed) from the display. Many research systems have demonstrated a variety of approaches to program elision, including the Cornell Program Synthesizer (Teitelbaum 1981) and Tioga (Teitelman 1985). Manual controls for eliding program block statements and methods have recently become standard features of programming environments. Everyday word-processors also include structure-based elision capabilities such as the “Outline View” of Microsoft Word, which allows successive folding and unfolding of document sections, and more sophisticated visualisations have been demonstrated by research systems such as the Document Lens (Robertson and Mackinlay 1993), which depicts documents as a 3D pyramid-style perspective visualization.

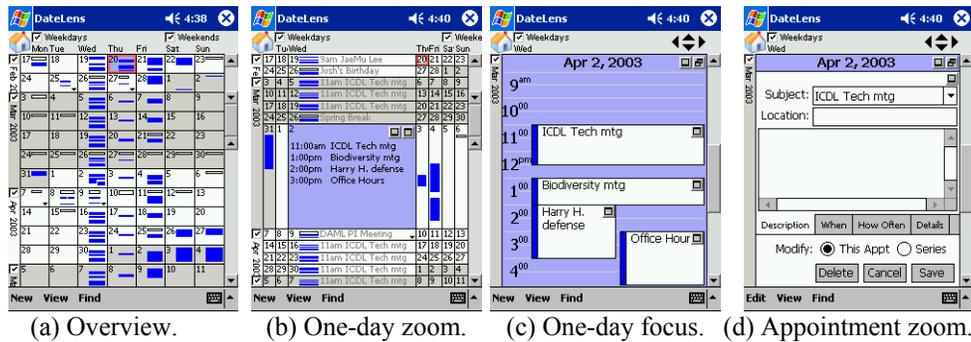


Figure 7. The DateLens interface with the view configured to show 12 weeks at consecutive levels of detail. All transitions between views are animated.

5.2.3 Fisheye tables

Although Furnas’s original paper described how fisheyes could be applied to tabular data such as calendar entries, Table Lens provided the first demonstration of applying fisheye distortion to a general tabular display (Rao and Card 1994). The Table Lens presents a compact overview of large data sets, displaying all rows and columns simultaneously by encoding values as small bars. Fisheye effects are available to selectively expand rows and columns. Expansion was applied independently to rows and columns, allowing multiple focal points, but preserving the familiar rectangular format of cells. Despite its apparent power and intuitive appeal we are unaware of formal empirical evaluation of Table Lens.

Bederson, Clamage, Czerwinski and Robertson (2004) applied concepts from the Table Lens within the DateLens fisheye calendar tool (Figure 7). Designed with the constrained display space of PDAs in mind, DateLens allows powerful and flexible visualisation of different time-spans (days, weeks, months) as well as a variety of search and presentation tools to illuminate patterns and outliers.

6. CUE-BASED TECHNIQUES

The spatial, temporal and seamless approaches described above all modify the size of objects in order to provide focus and context. These scale modifications can be applied purely to the graphical portrayal of objects or semantically so that only objects with certain properties are scaled.

Cue-based techniques, on the other hand, modify how objects are rendered and can introduce proxies for objects that might not be expected to appear in the display at all. They can be used in conjunction with any of the schemes above, and are typically applied in response to some search criteria. Data items satisfying the criteria are then displayed in

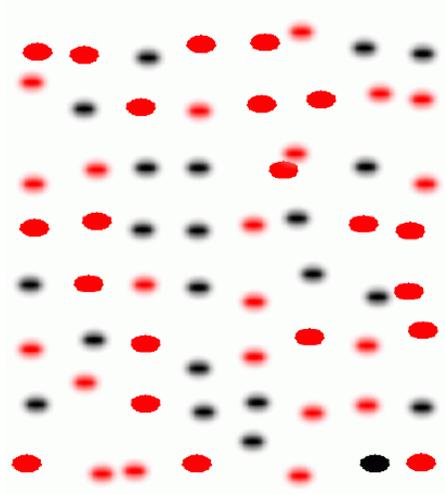


Figure 8: Cue focus+context, based on depth-of-focus ‘blurring’. From (Kosara et al. 2002).

a modified form to alert the user to their presence: making the focal items stand out from their surrounding context.

Given this broad definition of cue-based techniques, it is clear that much of the work on Information Visualization could be included within this category. We confine ourselves to a few examples that are particularly pertinent to the problems of focus and context.

6.1 Cue techniques for highlighting focal objects

Kosara et al. (2002) described a ‘semantic depth of field’ technique that provides a natural interface for drawing the user’s attention to focal items. With this technique items that satisfy search criteria are displayed in focus, while all others are slightly blurred (see Figure 8). An informal study indicated that the technique supports preattentive awareness, allowing focal items to be discriminated in less than 200ms, but it also raised concerns about the degree of visual strain caused by blurring the non-focal items. Baudisch and Gutwin (2004) provide a more detailed exploration of the usability of visual blending and blurring.

Several researchers have examined techniques for adding cues to the presence of search terms in web pages across many types of focus+context displays: normal displays and overviews with the ‘Popout Prism’ (Suh, Woodruff, Rosenholtz and Glass 2002); fisheye views with ‘Fishnet’ (Baudisch et al. 2004); and mobile devices with ‘Summary Thumbnails’ (Lam and Baudisch 2005). Bederson et al. (2004) used visual demarcations

in the scrollbar trough to convey the presence of search matches across several months of appointment data in the PDA calendar DateLens.

6.2 Cue techniques for extending context beyond the window edge

Contextual cues about information lying outside the main display region can be added by ‘decorations’. City Lights, developed by Zellwegger, Mackinlay, Good, Stefik, and Baudisch (2002), used window-edge decorations to indicate the existence, size and/or location of objects that lay beyond the window frame. In Halo, a variation of City Lights, Baudisch and Rosenholtz (2003) explored the use of arcs as decorators, as though each off-screen object were a street lamp just tall enough to cast its circle of light into the screen view-port. Nearby objects cast short, highly curved arcs, while far objects cast long, subtly curved arcs. In this way, object direction and distance are encoded as arc placement and curvature. Evaluations showed that Halo improved user performance.

7. EMPIRICAL EVALUATIONS

Although overview+detail, zooming, and fisheye interfaces have all been deployed in a wide-range of experimental systems since the 1980s, relatively few of the systems have been empirically evaluated. We briefly review this work, grouping the evaluations based on whether they primarily addresses low-level user tasks such as target acquisition, or high-level user tasks such as the ability to comprehend the information space. This summary of evaluative work excludes cue-based techniques: providing a cue to highlight objects that match some semantic criteria will aid users in finding them.

7.1 Low-level evaluations

Most of the low-level evaluations, summarised in Table 1, have each investigated only one type of focus+context method, either comparing performance with and without that method or analysing characteristics of the technique.

Target acquisition, or the user’s ability to select items of different sizes at different distances, is well understood in HCI research, with a well-tested, accurate, and robust model of performance (Fitts 1954). Although Fitts’ Law is traditionally applied to targets that are continuously visible, it also accurately models zoom-based target acquisition across large distances (Guiard et al. 2001; Guiard et al. 2004). In investigating how best to manipulate the concurrent control of panning and zooming, Ramos and Balakrishnan (2005) found that unimanual ‘zlliding’ with a stylus, which used pressure for zoom and dragging for pan, outperformed a variety of techniques that used bimanual separation, while experiments by Savage and Cockburn (2005) showed that an automatic coupling of zoom to panning speed outperformed bimanual explicit control of each parameter.

Table 1. Low-level evaluations of mechanical manipulation and target acquisition.

	Spatial (O+D)	Temporal (Zoom)	Seamless (Fisheye)	Comment
(North and Shneiderman 2000)				Navigating textual census data. O+D interfaces work best when actions in the overview and detail views are coordinated. Coordinated views outperform detail alone by 30-80%
(Beard and Walker 1990)				2D text-tree navigation. Overview allowed faster navigation than scrollbar+resizing. Primarily, a test of O+D versus traditional scrolling.
(Cockburn, Gutwin and Alexander 2006)				Document navigation. Thumbnail-enhanced scrollbars (O+D) are outperformed by a simple zooming interface that presents all pages as a stamp-sheet of page-thumbnail.
(Guiard, Bourgeois, Mottet and Beaudouin-Lafon 2001; Guiard, Beaudouin-Lafon, Bastin, Pasveer and Zhai 2004)				Target acquisition. Careful analysis of zooming as a tool for high index of difficulty pointing tasks.
(Ramos and Balakrishnan 2005)				Target acquisition. Compares parallel input mechanisms for zooming and sliding. Uni-manual stylus 'zsliding' (pressure for zooming, dragging for sliding, beats bi-manual methods).
(Savage and Cockburn 2005)				Document navigation. Calibration of perceptual issues of the relationship between scroll-speed and zoom, and comparative evaluation of automatic zooming versus traditional scrolling.
(Gutwin 2002)				Target acquisition. Describes the fisheye 'hunting effect' problem, and proposes and evaluates 'speed-coupled flattening' which eases the problem.
(McGuffin and Balakrishnan 2002; Zhai, Conversy, Beaudouin-Lafon and Guiard 2003; Cockburn and Brock 2006)				Target acquisition. Targets that expand in motor-space, or purely visually, as the cursor approaches are faster to acquire than static ones.

Gutwin (2002) showed that fisheye views can cause 'hunting effects' that harm target acquisition because the distance to the target change as the cursor approaches. Gutwin showed that the problem can be reduced through "speed-coupled flattening", which reduces or eliminates the fisheye effect as the velocity of the cursor increases. Unfortunately, speed-coupled flattening merely reduces a problem, rather than improving the level of performance attainable without the fisheye. McGuffin & Balakrishnan (2002) and Zhai, Conversy, Beaudouin-Lafon, and Guiard (2003) showed that acquisition times are reduced when discrete targets expand around their centre to fill an enlarged motor-space, even when the expansion starts after 90% of the movement toward the target is complete. Both sets of authors suggest modifications to the fisheye MacOs X Dock that would allow it to maintain the appealing visual effect, but without the adverse effects of target movement. Experiments by Cockburn and Brock (2006) showed that pure visual

expansion (without enlarged motor-spaces) allows most of the performance benefits provided by enlarged motor-spaces.

There is a surprising lack of target acquisition research on overview+detail systems. Guiard, Beaudouin-Lafon, Bastin, Pasveer, and Zhai (2004) research on ‘view pointing’ suggests benefits for overview+detail target acquisition, but their experiments are primarily aimed at examining zooming interfaces. Beard & Walker (1990) and Cockburn, Gutwin and Alexander (2006) compared overview+detail interfaces with a variety of competing interfaces, but again, most of the competitor interfaces include zooming capabilities. Their studies showed that overview+detail techniques out-performed traditional scrolling.

7.2 High-level evaluations

While the low-level evaluations focus on the mechanics of interface manipulation, the high-level evaluations are diverse, examining a wide range of task domains, as summarised in Table 2.

Schaffer, Zuo, Greenberg, Bartram, Dill, Dubs, and Roseman (1996) conducted the first high-level empirical study of focus+context techniques, comparing fisheye and full-zoom interfaces for navigation through 2D graphical networks. The participants’ tasks involved finding and fixing ‘broken network connections’ in hierarchically organised network diagrams. The fisheye view allowed multiple focal and contextual regions, while the zooming condition did not. Participants completed tasks more quickly and with fewer navigation actions when using the fisheye interface, but unfortunately, it remains unclear what caused the improved performance—the fisheye or the multiple foci. An experiment by Plumlee and Ware (2002) highlights these concerns by showing that interfaces supporting multiple overview+detail views outperform zooming interfaces when demands on visual memory are high.

In directly comparing overview+detail and zooming interfaces, Ghosh and Shneiderman (1999) analysed task completion times for tasks that involved extracting medical information from two versions of the Lifelines system (Plaisant et al. 1996). The overview+detail interface allowed tasks to be completed more quickly. It is now common for standard desktop applications to include both overview+detail and zooming capabilities. To better understand the contribution of each of these components to interaction, Hornbaek, Bederson and Plaisant (2002) evaluated user performance in map navigation tasks when using a zooming interface that either had or did not have an additional overview+detail region. Their results showed that overview+detail regions increased task completion times when using a map that allowed semantic zooming, and

Table 2. High-level evaluations of user tasks and comprehension of the information space.

	Spatial (O+D)	Temporal (Zoom)	Seamless (Fisheye)	Comment
(Ghosh and Shneiderman 1999)				Medical histories. The O+D interface to Lifelines (Plaisant, Milash, Rose, Widoff and Shneiderman 1996) allowed tasks to be completed more quickly than a zooming interface.
(Plumlee and Ware 2002)				Abstract graphical task. Evaluation confirms a theory that zooming out-performs multiple O+D views when demands on visual memory are low, but inverse when demands are high.
(Hornbaek, Bederson and Plaisant 2002)				Map navigation. Tests user performance using a zooming interface with and without an overview. Finds some tasks are faster without the overview due to cost of assimilating data.
(Bederson and Boltman 1999)				Spatial memory. Using animation in zooming helps users form a spatial model of the information space.
(Hornbaek and Frokjaer 2003)				Reading comprehension. Compared linear text with O+D and fisheye text interfaces. Comprehension highest with O+D, lowest with fisheye. Fisheye had fastest reading, O+D slowest.
(Baudisch, Lee and Hanna 2004)				Web browsers. Evaluates a cue-enhanced fisheye web browser with and without an overview. Overview and fisheye performed similarly. Overview was popular; fisheye less-so.
(Schaffer, Zuo, Greenberg, Bartram, Dill, Dubs and Roseman 1996)				Graph editing. Compared hierarchical zooming with a continuous fisheye. Fisheye allowed much faster task completion.
(Gutwin and Fedak 2004)				Various tasks on mobile devices. Panning vs two-level zoom vs fisheye. Fisheye fastest for one task, zooming fastest for one task. Panning slowest for all.
(Bederson, Clamage, Czerwinski and Robertson 2004)				Mobile calendars. Fisheye calendar allows complex tasks to be completed more quickly. Fisheye preferred.
(Zanella, Carpendale and Rounding 2002)				Map interpretation. Grids help users interpret the effect of fisheye distortion on spatial layout. Shading is less effective.

they suggest that this cost is due to the overview being made redundant by the rich semantic information in the detail region. When using a map that did not allow semantic zooming (meaning that labels in the map were only legible when zoomed in) there was no difference between performance with overview+detail and zooming interfaces. The participants preferred the interface with the overview despite their slower performance, stating that it helped them orient themselves within the information space. Finally, when the participants were asked to recall elements from the information space, they were better able to do so when they had used the non-overview interface. In associated work on recall, Bederson and Boltman (1999) show that by animating zoom effects, the users' spatial model of their information space is improved. Spatial comprehension remains a

concern for fisheye lenses which distort the information space to integrate both focus and context in the same view. Zanella, Carpendale and Rounding (2002) examined a variety of display enhancement schemes aimed at enhancing users' spatial comprehension of distorted space, concluding that simple parallel grid lines best support comprehension.

Hornbaek and Frokjaer (2003) compared reading patterns and usability issues associated with three forms of electronic document presentation: traditional 'flat' text, an overview+detail interface that enhanced the flat-text view with thumbnails of each page on the left-hand edge of the window, and a fisheye text view that diminished 'less important' document regions. In the fisheye view, the first and last paragraphs of each section remained undistorted, while other paragraphs were diminished unless clicked on by the user. Their evaluation tasks involved reading scientific documents using the different forms of presentation and then either writing short summaries or answering questions about them. Results showed that the fisheye view encouraged faster reading, but that this speed was at the cost of comprehension. The overview interface allowed participants to rapidly navigate through the document, and although they spent longer reading they scored better in comprehension tests. As in other studies, the participants' preference rankings favoured the overview+detail interface. These results are echoed by the findings of Baudisch, Lee and Hanna (2004) who compared three forms of a cue-enhanced web browser (a traditional linear view, a fisheye view, and an overview view). Their participants liked the overview+detail interface, while the fisheye polarised opinions. Their performance data suggested that the benefits of fisheyes are strongly task dependent.

The task specific performance merits of fisheye views are also echoed in two studies on the use of fisheye techniques on mobile devices. First, Gutwin and Fedak (2004) compared a single-focus fisheye interface with two-level zoom and panning interfaces for three tasks: image editing, web browsing, and network monitoring. The fisheye best supported navigation; the zooming interface best supported monitoring; and the traditional panning was slowest for all tasks. The fisheye target-acquisition 'hunting effects' explained the better performance with zooming in the monitoring task, and limitations of the two-level zoom implementation may explain zooming's lack of performance in the navigation task. Second, an evaluation of the fisheye calendar DateLens (Section 5.2.3) in comparison to a standard commercial calendar for PDAs (Microsoft Pocket PC2002) showed that the fisheye allowed improved performance of complex tasks, but little difference for simple ones. Again, like other studies, several users found the fisheye effects disturbing, preferring the standard interface. It remains

unclear whether the relatively poor subjective preference ratings for fisheyes are due to their novelty or due to an enduring distaste for their behaviour.

8. SUMMARY

We have presented four interface approaches to allowing users to work with focused and contextual displays of their information spaces, based on the nature of the separation between display regions and the cues used. Spatial separation systems, typified by overview+detail, allow concurrent view of focus and context but across regions that are spatially distinct in the x, y, or z coordinates. Temporal separation systems, typified by zooming, allow the entire display space to be dedicated to either focused or contextual views by temporally segregating their display. Seamless focus+context systems, typified by fisheye views, present both focus and context at the same time and place by distorting the information space. Finally, cue-based systems modify the display of items to highlight or suppress them, dependent on the user's context or search criterion. Although our review has used this classification to distinguish between these interface types, many interfaces support combinations of the techniques.

None of these approaches is ideal as all three compromise some aspect of interaction with consequent damage to usability. Spatial separation demands that users assimilate the relationship between the concurrent displays of focus and context. Evidence that this assimilation process hinders interaction is apparent in the experiments of (Hornbaek et al. 2002), who note that "switching between the detail and the overview window required mental effort and time". Temporal separation also demands assimilation between pre- and post-zoom states. This difficulty was observed by Cockburn and Savage (2003) in their evaluations, who note that "the abrupt transitions between discrete zooming levels... meant that the participants had to reorient themselves with each zoom action". Animation can ease this problem (Bederson and Boltman 1999), but it cannot be removed completely. Finally, seamless focus+context systems distort the information space, demanding that the user correctly assess the impact of the distortion on the information space. The distortion is likely to damage the user's ability to correctly assess spatial properties such as directions, distances and scales—for example, roads running directly East-West on a map will distort North or South dependent on the focus location. Even in non-spatial domains, evaluations of non-linear distortion techniques have largely failed to provide evidence of performance enhancements, and several studies have shown that fisheyes can damage fundamental components of interaction such as target acquisition.

Regardless of the usability challenges raised by each of the techniques, overview+detail and zooming interfaces are now standard components of many desktop

graphical user interfaces. The usability benefits they can provide outweigh their costs, and several evaluations have demonstrated that access to focus and context functionality outperforms its absence, although sometimes only by a limited amount. For example Hornbaek and Frokjaer (2003) showed that when reading electronic documents both overview+detail and fisheye interfaces offer performance advantages over traditional ‘flat’ text interfaces. Although rare, fisheye views are starting to appear in desktop components such as the MacOSX Dock, but in this case the usability benefits are purely cosmetic ‘eye-candy’, yet many users are happy to trade off efficiency for fashion and ‘coolness’, which is not surprising, since aesthetics and enjoyment are essential aspects of the computer experience.

It therefore seems clear that all three types of focus and context interfaces can improve interaction when compared to interfaces that constrain users to a single view, and thus the approaches are worth considering. The question then becomes which style of interface, or which combination of styles, offers the greatest performance advantages, for what tasks, and under what conditions? The current state of research fails to provide clear guidelines. Relatively few studies have empirically compared alternative focus and context schemes, and when they have, the results are mixed. Furthermore, the studies that do exist always look at a single implementation, and the details of that implementation (such as speed, interface details, and aesthetic design) could strongly influence the results. For example, Hornbaek et al. (2002) showed that fisheye text views allowed documents to be read faster than overview+detail, but that participants understood documents better with overview+detail; Schaffer et al. (1996) showed that hierarchical fisheyes outperformed full-zoom techniques, but it remains unclear whether an alternative implementation of zooming (continuous rather than full) would have defeated fisheyes; Bederson et al. (2004) showed that a table-based fisheye interface for calendars on PDAs outperformed a standard interface for complex tasks, but for simpler tasks the normal interface was more efficient and preferred.

We offer the following concluding comments and recommendations, with the cautionary note that it is important for designers to use their own judgment as to which approach to pursue. We also encourage researchers to continue empirically developing our understanding of the trade-offs between techniques.

Overview+Detail. Studies have consistently found that overview+detail interfaces are popular and commonly preferred to other techniques. For particular domain tasks, such as document comprehension, no alternative has been found to be more effective. Notable disadvantages of the overview+detail are the additional use of screen real estate (which

may be more effectively used for details) and the mental effort and time required to integrate the distinct views. The time-costs of this technique may make it sub-optimal for time-critical tasks.

Zooming. Temporal separation of views can easily create substantial cognitive load for users in assimilating the relationship between pre- and post-zoom states—zooming is easy to do badly. Animating the transition between zoom-levels can dramatically reduce the cognitive load. There is also evidence that interfaces which automatically configure zoom-level in response to other user controls can reduce subjective workload and improve performance, but more empirical work is needed.

Combining Overview+Detail and Zooming. Overview+detail interfaces are commonly supplemented with a variety of zooming capabilities, and in the limit overview+detail and zooming interfaces merge. For example, if the entire window space is dedicated to the overview (for example, a matrix thumbnail overview of a long document) then the interface is best categorized as zooming between page-view and thumbnail overview. In the absence of prior empirical work (to our knowledge), we are currently exploring the efficiency of full-window ‘zooming overviews’.

Fisheyes. The visual effects of fisheye distortion are appealing and many evaluations have revealed increased subjective satisfaction with the ‘cool’ or ‘neat’ effect. Their efficiency, however, is dubious with very few evaluations showing definitive performance advantages, while several reveal performance detriments. They should be used with caution if users need to assess spatial relationships in the data (such as assessing directions and distances). Further work is necessary to determine the tasks and conditions that allow them to succeed.

REFERENCES

- ABOELSAADAT, W. and BALAKRISHNAN, R. (2004). An Empirical Comparison of Transparency on One and Two Layer Displays. *People and Computers XVIII: Proceedings of the British HCI Conference 2004*, 53-67.
- AHLBERG, C. and SHNEIDERMAN, B. (1994). The Alphaslider: A Compact and Rapid Selector. *Proceedings of CHI'94: Conference on Human Factors in Computing Systems*, Boston, Massachusetts, 365-371. ACM Press.
- BAUDISCH, P., GOOD, N., BELLOTTI, V. and SCHRAEDLEY, P. (2002). Keeping Things in Context: A Comparative Evaluation of Focus Plus Context Screens, Overviews, and Zooming. *Proceedings of CHI'2002 Conference on Human Factors in Computing Systems*, Minneapolis, Minnesota, 259--266.
- BAUDISCH, P. and GUTWIN, C. (2004). Multiblending: displaying overlapping windows simultaneously without the drawbacks of alpha blending. *CHI '04: Proceedings of the 2004 conference on Human factors in computing systems*, 367-374.

- BAUDISCH, P., LEE, B. and HANNA, L. (2004). Fishnet, a fisheye web browser with search term popouts: a comparative evaluation with overview and linear view. Proceedings of Advanced Visual Interfaces, AVI04, Gallipoli, Italy, 133-140. ACM Press.
- BAUDISCH, P. and ROSENHOLTZ, R. (2003). Halo: A Technique for Visualizing Off-Screen Locations. Proceedings of CHI'2003 Conference on Human Factors in Computing Systems, Fort Lauderdale, Florida, 481--488.
- BEARD, D. and WALKER, J. (1990). "Navigational Techniques to Improve the Display of Large Two-Dimensional Spaces." Behavior and Information Technology **9**(6): 451-466.
- BEDERSON, B. (2000). Fisheye Menus. Proceedings of UIST'00 Symposium on User Interface Software and Technology, San Diego, California, 217-225. ACM Press.
- BEDERSON, B. (2001). PhotoMesa: A Zoomable Image Browser Using Quantum Treemaps and Bubblemaps. Proceedings of UIST01: ACM Conference on User Interface Software and Technology, Orlando, Florida, 71-80. ACM Press.
- BEDERSON, B. and BOLTMAN, A. (1999). Does Animation Help Users Build Mental Maps of Spatial Information? Proceedings of InfoVis99: IEEE Symposium on Information Visualization, 28-35. IEEE Computer Society.
- BEDERSON, B., GROSJEAN, J. and MEYER, J. (2004). "Toolkit Design for Interactive Structured Graphics." IEEE Transactions on Software Engineering **30**(8): 535-546.
- BEDERSON, B., HOLLAN, J., PERLIN, K., MEYER, J., BACON, D. and FURNAS, G. (1996). "Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics." Journal of Visual Languages and Computing **7**(1): 3-31.
- BEDERSON, B. B., CLAMAGE, A., CZERWINSKI, M. P. and ROBERTSON, G. G. (2004). "DateLens: A fisheye calendar interface for PDAs." ACM Trans. Comput.-Hum. Interact. **11**(1): 90-119.
- BEDERSON, B. B., MEYER, J. and GOOD, L. (2000). Jazz: an extensible zoomable user interface graphics toolkit in Java. Proceedings of the 13th annual ACM symposium on User interface software and technology. San Diego, California, United States, ACM Press: 171-180.
- BIER, E. A., STONE, M. C., PIER, K., BUXTON, W. and DEROSE, T. D. (1993). Toolglass and magic lenses: the see-through interface techniques. Proceedings of the 20th annual conference on Computer graphics and interactive, ACM Press: 73-80.
- BOURGEOIS, F. and GUIARD, Y. (2002). Multiscale Pointing: Facilitating Pan-Zoom Coordination. CHI '02 extended abstracts of CHI2002, Minneapolis, Minnesota, USA, 758-759.
- BYRD, D. (1999). A Scrollbar-based Visualization for Document Navigation. Proceedings of Digital Libraries '99, Berkeley CA, 122-129.
- CARD, S., MACKINLAY, J. and SHNEIDERMAN, B. (1999). Readings in Information Visualization: Using Vision to Think, Morgan-Kaufmann.
- CARD, S. K., ROBERTSON, G. G. and MACKINLAY, J. D. (1991). The information visualizer, an information workspace. Proceedings of CHI91: ACM conference on Human factors in computing systems, New Orleans, Louisiana, United States, 181-186. ACM Press.
- CARPENDALE, M. and MONTAGNESE, C. (2001). A Framework for Unifying Presentation Space. Proceedings of ACM UIST2001: Conference on User Interface Software and Technology, Orlando, Florida, 61-70. ACM Press.

- CASALTA, D., GUIARD, Y. and BEAUDOUIN-LAFON, M. (1999). Evaluating two-handed input techniques: rectangle editing and navigation. CHI '99 extended abstracts on human factors in computing systems, Pittsburgh, Pennsylvania, 236-237. ACM Press.
- CHIMERA, R. (1992). ValueBars: An Information Visualization and Navigation Tool. Proceedings of CHI92: Conference on Human Factors in Computing Systems, Monterey, California, 293-294. ACM Press.
- COCKBURN, A. and BROCK, P. (2006). Human On-Line Response to Visual and Motor Target Expansion. Proceedings of Graphics Interface 2006, Quebec City, Canada Canadian Human-Computer Communications Society.
- COCKBURN, A., GUTWIN, C. and ALEXANDER, J. (2006). Faster Document Navigation with Space-Filling Thumbnails. Proceedings of CHI'06 ACM Conference on Human Factors in Computing Systems, Montreal, Canada, 1-10. ACM Press.
- COCKBURN, A. and SAVAGE, J. (2003). Comparing Speed-Dependent Automatic Zooming with Traditional Scroll, Pan and Zoom Methods. People and Computers XVII (Proceedings of the 2003 British Computer Society Conference on Human-Computer Interaction.), Bath, England, 87-102.
- COCKBURN, A. and SAVAGE, J. (2005). Tuning and Testing Scrolling Interfaces that Automatically Zoom. CHI'05. ACM Conference on Human Factors in Computing Systems, Portland, Oregon ACM Press.
- COX, D., CHUGH, J., GUTWIN, C. and GREENBERG, S. (1998). The Usability of Transparent Overview Layers. Proceedings of CHI'98: Conference on Human Factors in Computing Systems, Los Angeles, 301-302.
- DRUIN, A., STEWART, J., PROFT, D., BEDERSON, B. and HOLLAN, J. (1997). KidPad: A Design Collaboration Between Children, Technologists, and Educators. Proceedings of CHI'97: ACM Conference on Human Factors in Computing Systems, Atlanta, Georgia, 463-470.
- FITTS, P. (1954). "The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement." *Journal of Experimental Psychology* **47**: 381-391.
- FURNAS, G. (1986). Generalized Fisheye Views. Proceedings of the CHI'86 Conference on Human Factors in Computing Systems III, Boston, MA, 16-23. ACM Press.
- FURNAS, G. and ZHANG, X. (1998). MuSE: A Multiscale Editor. Proceedings of the 1998 ACM Conference on User Interface Software and Technology, San Francisco, California, 107--116.
- FURNAS, G. W. and BEDERSON, B. B. (1995). Space-scale diagrams: understanding multiscale interfaces. Proceedings of the CHI95: Conference on Human factors in computing systems, Denver, Colorado, United States, 234-241. ACM Press/Addison-Wesley Publishing Co.
- GHOSH, P. and SHNEIDERMAN, B. (1999). Zoom-Only Vs Overview-Detail Pair: A Study in Browsing TEchniques as Applied to Patient Histories, HCIL, University of Maryland, College Park.
- GUIARD, Y., BEAUDOUIN-LAFON, M., BASTIN, J., PASVEER, D. and ZHAI, S. (2004). View Size and Pointing Difficulty in Multi-Scale Navigation. Proceedings of Advanced Visual Interfaces, AVI04, Gallipoli, Italy, 117-124. ACM Press.
- GUIARD, Y., BOURGEOIS, F., MOTTET, D. and BEAUDOUIN-LAFON, M. (2001). Beyond the 10-bit barrier: Fitts' law in multi-scale electronic worlds. People and Computers XV --- Proceedings of IHM-HCI 2001, Lille, France, 573-587.

- GUIMBRETIERE, F., STONE, M. and WINOGRAD, T. (2001). Fluid interaction with high-resolution wall-size displays. *UIST'01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, Orlando, Florida, 21-30. ACM Press.
- GUTWIN, C. (2002). Improving Focus Targeting in Interactive Fisheye Views. *ACM Conference on Human Factors in Computing Systems (CHI'02)*, Minneapolis, Minnesota, 267-274. ACM Press.
- GUTWIN, C. and FEDAK, C. (2004). Interacting with Big Interfaces on Small Screens: a Comparison of Fisheye, Zoom, and Panning Techniques. *Proceedings of Graphics Interface 2006*. Quebec City, Canada, Canadian Human-Computer Communications Society.
- GUTWIN, C., ROSEMAN, M. and GREENBERG, S. (1996). A Usability Study of Awareness Widgets in a Shared Workspace Groupware System. *Proceedings of CSCW'96: ACM Conference on Computer Supported Cooperative Work*, Boston, Massachusetts, 258-267.
- HERMAN, I., MELANCON, G. and MARSHALL, M. (2000). "Graph Visualization and Navigation in Information Visualization: A Survey." *IEEE Transactions on Visualization and Computer Graphics* **6**(1): 24-43.
- HILL, W. and HOLLAN, J. (1992). Edit Wear and Read Wear. *Proceedings of CHI'92 ACM Conference on Human Factors in Computing Systems*, Monterey, California, 3-9.
- HORNBAEK, K., BEDERSON, B. and PLAISANT, C. (2002). "Navigation Patterns and Usability of Zoomable User Interfaces with and without an Overview." *ACM Transactions on Computer-Human Interaction* **9**(4): 362--389.
- HORNBAEK, K. and FROKJAER, E. (2003). "Reading Patterns and Usability in Visualizations of Electronic Documents." *ACM Transactions on Computer Human Interaction* **10**(2): 119-149.
- IGARASHI, T. and HINCKLEY, K. (2000). Speed-dependent Automatic Zooming for Browsing Large Documents. *Proceedings of UIST'00 ACM Symposium on User Interface Software and Technology*, San Diego, California., 139-148.
- JUL, S. and FURNAS, G. (1998). Critical Zones in Desert Fog: Aids to Multiscale Navigation. *Proceedings of the 1998 ACM Conference on User Interface Software and Technology*, San Francisco, California, 97-106.
- KLEIN, C. and BEDERSON, B. (2005). Benefits of Animated Scrolling. *Extended Abstracts of CHI'05: ACM Conference on Human Factors in Computing Systems*, Portland, Oregon, 1965-1968. ACM Press.
- KOSARA, R., MIKSCH, S. and HAUSER, H. (2002). "Focus+Context Taken Literally." *IEEE Computer Graphics and Applications* **22**(1): 22-29.
- LAM, H. and BAUDISCH, P. (2005). Summary Thumbnails: Readable Overviews for Small Screen Web Browsers. *Proceedings of CHI'05: Conference on Human Factors in Computing Systems*, Portland, Oregon, 681-690.
- LAMPING, J., RAO, R. and PIROLI, P. (1995). A Focus+Context Technique Based on Hyperbolic Geometry for Visualising Large Hierarchies. *Proceedings of CHI'95 Conference on Human Factors in Computing Systems*, Denver, Colorado, 401--408.
- LATULIPE, C., KAPLAN, C. and CLARKE, C. (2005). Bimanual and unimanual image alignment: an evaluation of mouse-based techniques. *Proceedings of UIST'05: ACM Symposium on User Interface Software and Technology*, Seattle, Washington, 123-131. ACM Press.

- LEGANCHUK, A., ZHAI, S. and BUXTON, W. (1998). "Manual and Cognitive Benefits of Two-Handed Input: An Experimental Study." *ACM Transactions on Computer-Human Interaction* **5**(4): 326--359.
- LEUNG, Y. and APPERLEY, M. (1994). "A Review and Taxonomy of Distortion-Oriented Presentation Techniques." *ACM Transactions on Computer Human Interaction* **1**(2): 126--160.
- MACKINLAY, J., ROBERTSON, G. and CARD, S. (1991). Perspective Wall: Detail and Context Smoothly Integrated. Proceedings of CHI'91 Conference on Human Factors in Computing Systems, New Orleans, 173-179.
- MCGUFFIN, M. and BALAKRISHNAN, R. (2002). Acquisition of Expanding Targets. Proceedings of CHI'2002 Conference on Human Factors in Computing Systems Minneapolis, Minnesota, 20--25 April: 57--64.
- MCGUFFIN, M. and BALAKRISHNAN, R. (2002). Acquisition of Expanding Targets. Proceedings of CHI'02 Conference on Human Factors in Computing Systems, Minneapolis, Minnesota, 57-64. ACM Press.
- NORTH, C. and SHNEIDERMAN, B. (2000). Snap-Together Visualization: A User Interface for Coordinating Visualizations via Relational Schemata. Proceedings of Advanced Visual Interfaces, AVI2000, Palermo, Italy, 128-135.
- O'HARA, K. and SELLEN, A. (1997). A comparison of reading paper and on-line documents. Proceedings of CHI'97 ACM Conference on Human factors in Computing Systems, New York, NY, USA, 335-342.
- O'HARA, K., SELLEN, A. and BENTLEY, R. (1999). Supporting Memory for Spatial Location While Reading from Small Displays. Extended Abstracts of CHI'99 Conference on Human Factors in Computing Systems, Pittsburgh, Pennsylvania, 220-221.
- PERLIN, K. and FOX, D. (1993). Pad: an alternative approach to the computer interface. Proceedings of the 20th annual conference on Computer graphics and Interactive techniques, 57-64. ACM Press.
- PLAISANT, C., CARR, D. and SHNEIDERMAN, B. (1995). "Image-Browser Taxonomy and Guidelines for Designers." *IEEE Software* **12**(2): 21-32.
- PLAISANT, C., MILASH, B., ROSE, A., WIDOFF, S. and SHNEIDERMAN, B. (1996). LifeLines: Visualizing Personal Histories. Proceedings of CHI'96 Conference on Human Factors in Computing Systems Vancouver, April 13--18, 221-227.
- PLUMLEE, M. and WARE, C. (2002). Modeling performance for zooming vs multi-window interfaces based on visual working memory. Proceedings of Advanced Visual Interface (AVI02), Trento, Italy, 59-68.
- RAMOS, G. and BALAKRISHNAN, R. (2005). Zliding: Fluid Zooming and Sliding for High Precision Parameter Manipulation. UIST'05: Proceedings of the 18th annual ACM symposium on User interface software and technology, Seattle, Washington, 143-152. ACM Press.
- RAO, R. and CARD, S. K. (1994). The table lens: merging graphical and symbolic representations in an. Proceedings of the CHI'94: ACM Conference on Human factors in computing systems:, Boston, Massachusetts, United States, 318-322. ACM Press.
- ROBERTSON, G. G. and MACKINLAY, J. D. (1993). The document lens. Proceedings of the 6th annual ACM symposium on User interface software and technology. Atlanta, Georgia, United States, ACM Press: 101-108.

- SARKAR, M. and BROWN, M. (1992). Graphical Fisheye Views of Graphs. Proceedings of CHI'92 Conference on Human Factors in Computing Systems, Monterey, CA, 83-91. ACM Press.
- SARKAR, M., SNIBBE, S. and REISS, S. (1993). Stretching the rubber sheet: A metaphor for visualising large structure on small screen. UIST'93: Proceedings of the 16th annual ACM symposium on User interface software and technology, Atlanta, Georgia, 81-91. ACM Press.
- SAVAGE, J. and COCKBURN, A. (2005). Comparing Automatic and Manual Zooning Methods for Acquiring Off-Screen Targets. People and Computers XIX: Proceedings of the 2005 British HCI Conference, Edinburgh, UK, 439-454.
- SCHAFFER, D., ZUO, Z., GREENBERG, S., BARTRAM, L., DILL, J., DUBS, S. and ROSEMAN, M. (1996). "Navigating Hierarchically Clustered Networks through Fisheye and Full-Zoom Methods." ACM Transactions on Computer Human Interaction **3**(2): 162-188.
- SPENCE, R. (2001). Information Visualization, Addison-Wesley.
- SPENCE, R. and APPERLEY, M. (1982). "Database Navigation: An Office Environment for the Professional." Behaviour and Information Technology **1**(1): 43-54.
- SUH, B., WOODRUFF, A., ROSENHOLTZ, R. and GLASS, A. (2002). Popout Prism: Adding Perceptual Principles to Overview+Detail Document Interfaces. Proceedings of CHI'2002 Conference on Human Factors in Computing Systems. CHI Letters 4(1). Minneapolis, Minnesota, 251-258.
- SUMMERS, K., GOLDSMITH, T., KUBICA, S. and CAUDELL, T. (2003). An Experimental Evaluation of Continuous Semantic Zooming in Program Visualization. INFOVIS03: IEEE Symposium on Information Visualization, Seattle, Washington, 155-162. IEEE Computer Society.
- TAN, D., ROBERTSON, G. and CZERWINSKI, M. (2001). Exploring 3D Navigation: Combining Speed-coupled Flying with Orbiting. Proceedings of CHI'2001 Conference on Human Factors in Computing Systems, Seattle, Washington, 418--425.
- TEITELBAUM, T. (1981). "The Cornell Program Synthesizer: A Syntax-Directed Programming Environment." Communications of the ACM **24**(9): 563--573.
- TEITELMAN, W. (1985). "A Tour through Cedar." IEEE Transactions on Software Engineering **11**(3): 285--302.
- TOBLER, W. (1973). "A continuous transformation useful for districting." Annals of New York Academy of Science **219**: 215-220.
- VAN WIJK, J. and NUIJ, W. (2004). "A Model for Smooth Viewing and Navigation of Large 2D Information Spaces." IEEE Transactions on Visualization and Computer Graphics **10**(4): 447-458.
- WARE, C. (2000). Information Visualization: Perception for Design, Morgan Kaufmann.
- WARE, C. and FLEET, D. (1997). Context Sensitive Flying Interface. Proceedings of the 1997 Symposium on Interactive 3D Graphics, Providence, RI, 127--130.
- WOODSON, W. and CONOVER, D. (1964). Human Engineering Guide for Equipment Designers. Berkeley, California, University of California Press.
- ZANELLA, A., CARPENDALE, S. and ROUNDING, M. (2002). On the effects of viewing cues in comprehending distortions. Proceedings of NordCHI'02. , Aarhus, Denmark, 119-128.

ZELLWEGER, P., MACKINLAY, J., GOOD, L., STEFIK, M. and BAUDISCH, P. (2002). City Lights: Contextual Views in Minimal Space. CHI '02 Extended Abstracts on Human factors in computer systems, Minneapolis, Minnesota, USA, 838-839.

ZHAI, S., CONVERSY, S., BEAUDOUIN-LAFON, M. and GUIARD, Y. (2003). Human On-line Response to Target Expansion. Proceedings of CHI'03 Conference on Human Factors in Computing Systems, Fort Lauderdale, Florida, 177-184. ACM Press.

ZHAI, S., SMITH, B. and SELKER, T. (1997). Improving Browsing Performance: A Study of Four Input Devices for Scrolling and Pointing Tasks. Proceedings of INTERACT'97: the sixth IFIP conference on Human Computer Interaction, 286-292.