Improving Memorability in Fisheye Views

A Thesis Submitted to

The College of Graduate Studies and Research

In Partial Fulfillment of the Requirements

For the Degree of Master of Science

In the

Department of Computer Science

University of Saskatchewan

Saskatoon, Canada

By

Amy C. Skopik

Keywords: Fisheye views, focus+context techniques, fisheye usability, spatial memory, memorability, edit wear, visit wear.

Permission to Use

In presenting this thesis in partial fulfillment of the requirements for a Postgraduate degree from the University of Saskatchewan, I agree that the Libraries of this University may make it freely available for inspection. I further agree that permission for copying of this thesis in any manner, in whole or in part, for scholarly purposes may be granted by the professor or professors who supervised my thesis work or, in their absence, by the Head of the Department or the Dean of the College in which my thesis work was done. It is understood that any copying, publication, or use of this thesis or parts thereof for financial gain shall not be allowed without my written permission. It is also understood that due recognition shall be given to me and to the University of Saskatchewan in any scholarly use which may be made of any material in my thesis.

Requests for permission to copy or to make other use of material in this thesis in whole or part should be addressed to:

Head of the Department of Computer Science

University of Saskatchewan

Saskatoon, Saskatchewan S7N 5A9

Abstract

Interactive fisheye views use distortion to show both local detail and global context in the same display space. Although fisheyes allow the presentation and inspection of large data sets, the distortion effects can cause problems for users. One such problem is lack of *memorability* – the ability to find and go back to objects and features in the data. This thesis examines the possibility of improving the memorability of fisheye views by adding historical information to the visualization. The historical information is added visually through *visit wear*, an extension of the concepts of edit wear and read wear. This will answer the question "Where have I been?" through visual instead of cognitive processing by overlaying new visual information on the data to indicate a user's recent interaction history. This thesis describes general principles of visibility in a space that is distorted by a fisheye lens and defines some parameters of the design space of visit wear. Finally, a test system that applied the principles was evaluated, and showed that adding visit wear to a fisheye system improved the memorability of the information space.

Acknowledgements

I would like to thank my thesis advisor, Dr. Carl Gutwin, for his guidance and patience with my research. I would also like to thank my husband, Randy Mosiondz, for supporting my decision to go back to school, and being an all round wonderful person.

Finally, I'd like to thank my parents for providing inspiration in this and every part of my life.

Table of Contents

List of Figures

List of Tables

1.0 Introduction

Information visualization presents information in a graphical form. Visualization is an external aid that amplifies cognition; successful visualizations can reduce the time it takes to get information, can help to make sense out of that information, and can enhance creative thinking (Card, Mackinlay and Schneiderman, 1999). Visualization techniques have been developed for many kinds of data, in particular large data sets which can otherwise be difficult to understand.

There are two common ways to interact with large data sets. Firstly, people often need to see the entire data set at once to get an idea of the scope and limits of the space and to picture where they are within it. Secondly, people also often need to see individual data items in detail. Therefore, the display of the data set must show an overview to provide context, and also be able to focus on details. When the amount of information is too large for the typical screen to present all of it at once in detail, presenting focus and context at the same time is a challenge. If the display space is limited, the trade-off between overview and detail must be carefully considered and managed (Baudisch et al., 2002, Keahey, 1998).

Several methods have been proposed for managing overview and detail in a large information space. These are multiple view methods (Smith et al., 1989), zoom-and-pan methods (Bederson and Hollan, 1994), and focus-plus-context methods (Leung, Spence and Apperly, 1995). Multiple-view methods (e.g. radar views, magnifying lens) use separate screen areas for the overview and the detail view. The zoom-and-pan method allows the user to select a suitable level of magnification interactively, but only one view is shown at a time. Focus-plus-context (F+C) methods integrate the overview and detail into the same area.

In an F+C view, the point of focus is magnified while the rest of the space (the context) is reduced to make room for the enlarged focus area. This reduction requires distortion or obstruction of the space, and this distortion is a hallmark of F+C views. The benefit is that the two views are integrated so that both are always visible and the relationship between the two views (i.e. the location of the focus in relation to the whole) is always clear.

A *fisheye* technique (Figure 1.1) is a type of F+C technique characterised by a smooth transition between the magnified focus region and the de-magnified context area. It is called a fisheye technique because it has a similar effect on the view of the data that an optical fisheye lens has on a photograph. The smooth transition means that neighbouring information is also magnified, though to a lesser extent than the focus, and is therefore almost as legible. This ability to see information about the data close to, but not in, the focus distinguishes the fisheye technique from the other F+C techniques. Seeing neighbouring information is useful in an organized space since in many information sets there is a semantic relationship between data based on proximity (i.e. the closer items are together, the more they have in common). Fisheye views can be either static, where the point of focus does not change, or interactive. In an interactive fisheye view, the location of the focus that provides the maximum magnification can be changed in real time by the user.

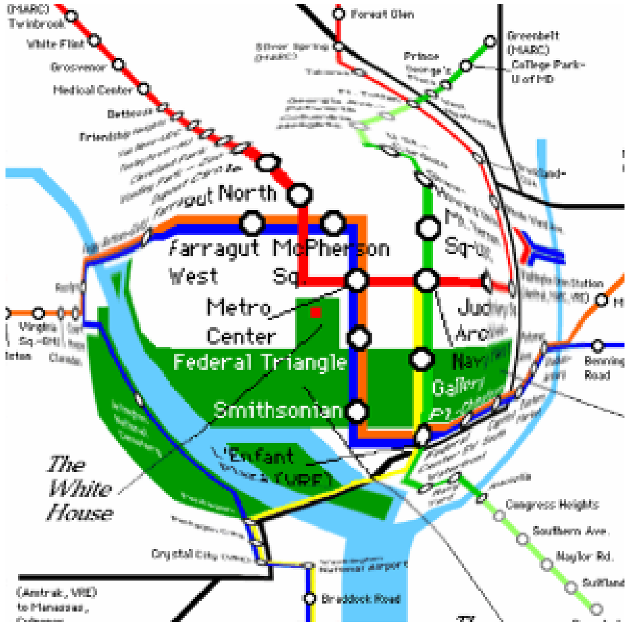

Figure 1.1: Fisheye view of a map of the Washington DC metro system with the focal point centred on the White House (Friendly and Denis, 2004). See also colour figure in Appendix A.

However, along with the advantages of the interactive fisheye method there are also usability problems. The specific problem in interactive fisheye visualizations that will be investigated in this thesis is lack of memorability.

1.1 Problem

The problem to be addressed in this thesis is that people have difficulty remembering where they have been in interactive fisheye views.

In this thesis I define *memorability* as the ability to return to a place where you've already been in a space. Most spaces consist of a collection of sparse features. There are a number of characteristics that make these (or any) features memorable, such as a distinctive property (colour or shape), a distinguishing position within the space, or some semantic content that has significance to the user. A feature possessing these characteristics can become a *landmark* in a space, serving to orient the user and aid in remembering the space around it. Alternatively, there could be characteristics of the environment itself rather than the features that aid in memorability. A *trail* indicating a traveled path is an example of this; if you can look back on your footprints in a snowy field, you can tell where you have been even if the field is featureless.

The effect of these characteristics on spatial memory in the real world carries over to the abstract world of a large data set. Interacting with these data spaces lends itself to spatial analogies like “visiting a feature that you’ve already been to” although no bodily movement has actually taken place. People use the spatial abilities that they use in the real world to remember their interaction with the abstract data set. However, the interactive fisheye method involves a continuously changing distortion of the space, and we are not accustomed to taking distortion into account when using our spatial memories. Figure 1.3 shows two views of the same data when the fisheye focus point has been moved. The node highlighted in Figure 1.2 is difficult to find in Figure 1.3.

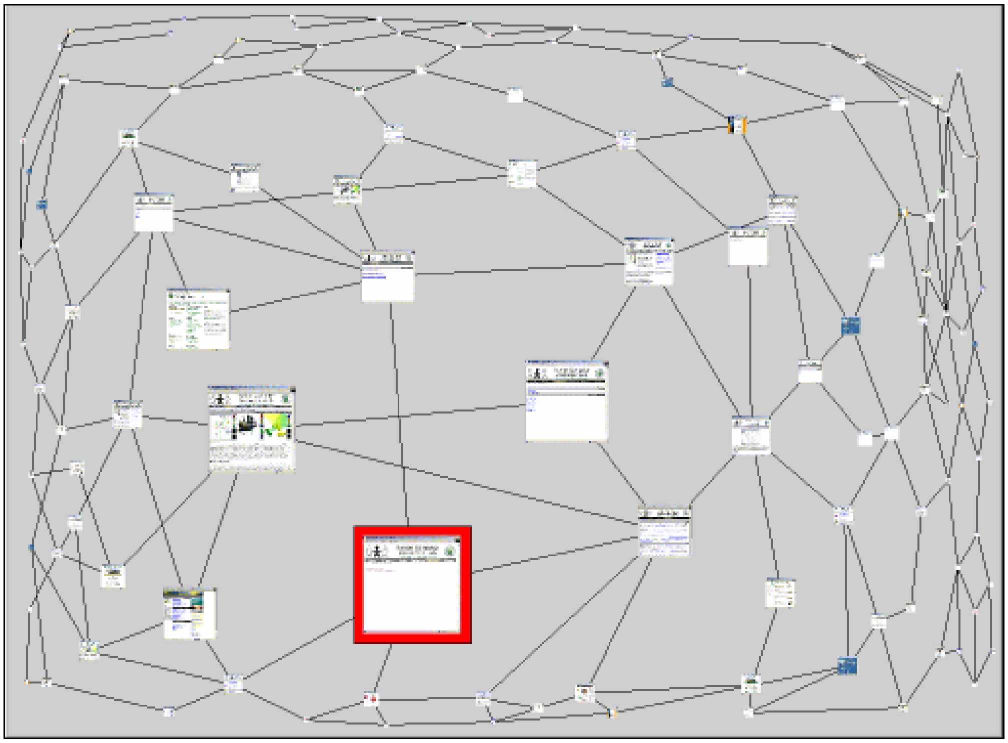

Figure 1.2: Example graph with a fisheye distortion and a random node marked. See also colour figure in Appendix A.

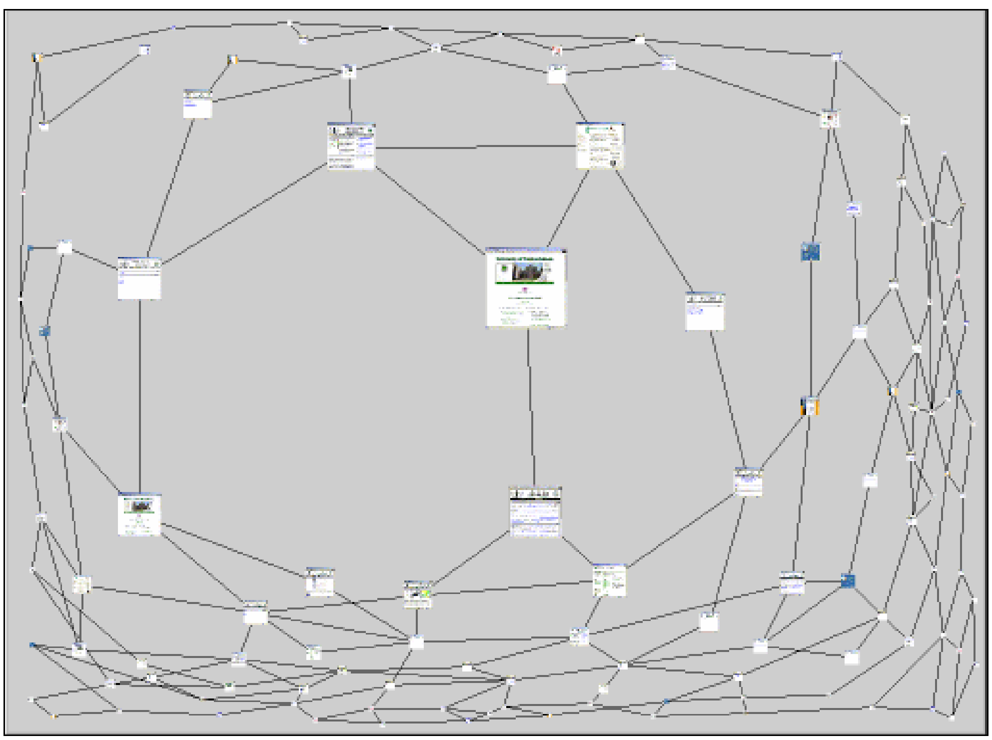

Figure 1.3: The same graph as Figure 1.2 with the focal point in a different place. The node that was marked is difficult to find. See also colour figure in Appendix A.

Users have two main difficulties with memorability in fisheye views: they make more errors when attempting to identify or locate previously visited features in a distorted space, and they take more time to identify and locate previously visited features in a distorted space.

1.2 Motivation

There are two main motivations for this work: improving the usability of fisheye views, and investigating ways of adding memorable features to applications in general.

Improving the usability of fisheye views will make them a practical visualization tool to use in representing large data spaces. Large information spaces have become common, and fisheyes solve several problems of displaying large information sets. Both focus views and context views of data are useful, and fisheyes provide both of these views simultaneously in a space-efficient way.

Improving memorability may also help to improve the usability of applications where the user commonly revisits items that have already been inspected or used. For example, the addition of historical information such as bookmarks and differentiating visited links has been an important part of the usability of Web browsers, where most users revisit previously viewed sites (Tauscher and Greenberg, 1997). Being able to backtrack to the last visited data or last point of decision is important in many tasks involving exploration of large data sets.

1.3 Solution

The problem of memorability in interactive fisheye views can be addressed by adding visual information to represent the user's interaction history.

Representing interaction history visually is called *read wear* (Hill et al., 1992) and is a visual depiction of the history of the user's interaction with the data set. Read wear is presented as part of the data itself, not as a separate screen or menu but as a change in the visual properties of the visited information. This thesis extends this concept to "visit wear," which can make explicit the question of "Where have I been?" and so should solve the problem of memorability by turning it from a memory problem to a visual perception problem, which is easier to solve. Visit wear assumes that the user cares whether or not data has been previously inspected. If the browsing of a data set is completely random (if the user is updating certain information from an external list, for example) then visit wear may not provide any benefits.

Read wear and focus-plus-context techniques seem like a natural combination since in an ordinary data space, the question of exactly where the viewer is looking is difficult to answer without equipment such as an eye tracker. With a fisheye lens, which has a fairly precise magnification area compared to other focus-plus-context techniques, the question "Where is the user's attention?" can most often be answered with "Where the focal point is located" since objects not in the magnification area are typically too small to properly see. The lens movement can pinpoint the user's focus, whereas other explicit actions like opening a file or scrolling cannot.

There are two issues that must be considered when adding visit wear to a fisheye view. Firstly, visit wear should be presented in a way that is minimally affected by the distortion of the lens. If the visit wear information is not distortion-independent, it may be difficult to see and therefore useless as the user moves to a new focal area. If instead the visit wear makes certain elements of the data distinctive no matter where the focus is, the visit wear indication may not only provide historical information, but also act as a potential landmark to aid spatial awareness.

Secondly, however, the visit wear must be added in a way that does not cause new problems through clutter or occlusion. Too much information may distract the user from other tasks. "Too much" here refers to the perceptual weight of the information; the amount of visual attention that the information draws. The displayed history may go far back in time, as long as its representation does not detract from the original data set or the users' interaction with it.

The added information must find a middle ground between too much and too little (Gutwin, 2002). Too much will be cluttered and not distinctive, too little will not be helpful in backtracking and will change too often to be a good landmark.

1.4 Steps in the solution

There were four steps required to carry out this solution. They were as follows.

1. *Determine principles of memorability in a fisheye application by observing user strategies*

Research carried out for this step investigated spatial memory when the space is being distorted with a fisheye lens. There are some visual properties that retain their distinctiveness even when the view is distorted, and these properties are described as being *robust* with respect to distortion. Items with robust properties may be more memorable in a distorted space. I performed a study to test which properties are robust with respect to distortion, in users' opinions. A second study compared users' memorization strategies with performance in actual memory tasks. The result of this step was a discussion of the effect that distortion has on the visual properties of items in the data space, according to user estimation and performance.

2. *Compile a framework of visit wear visualisation*

Research carried out for this step defined the particular type or types of data space that will be used for this experiment (such as graphs or maps) and explored methods of adding visit wear to the space or spaces. Some work has already been done in adding read wear to certain types of spaces (Hill and Hollan, 1993); I examined the appropriateness of existing techniques and developed new ones. This examination did not take into account any distortion of the space. The result of this step was a set of candidate visualisation techniques for visit wear that are appropriate to the chosen data space.

3. *Develop candidate techniques for displaying visit wear in a fisheye visualisation*

To carry out this step, the results of the previous two steps were combined. A list was created of visit wear display techniques that are robust with respect to distortion. The result of this step was a list of methods of adding visit wear that satisfy the criteria for being memorable in a fisheye visualization.

4. *Implement a test application that demonstrates this intersection*

The test application implemented in this step applied a fisheye visualisation to the selected information space and added visit wear to the fisheye visualisation. The test application used selected candidate techniques developed in the previous step, as well as a mode displaying no added historical information (as a control). The result of this step was a test system that allowed measurement of memorability metrics such as target accuracy and time to find the target in a series of memory tasks.

1.5 Evaluation

The hypothesis is that representing interaction history through visit wear in a fisheye visualization will improve the memorability of the data space. This means that users will be able to find previously visited features in less time and with fewer errors.

The hypothesis was tested in two different design and task scenarios, with each scenario comparing users' performance with visit wear to their performance without visit wear. Visit wear was judged to be effective if it let people remember previously visited locations in less time and with fewer errors than when visit wear is not present. In addition, the visit wear should not have added undue clutter or interfere with other tasks.

The evaluation was performed on the developed test application. I recorded and examined the effect of each method on the participants' accuracy and time to complete a series of memory tests. If people had a better performance in memory and browsing tasks with a fisheye lens when visit wear is added, then it can be concluded that fisheye memorability had been improved.

1.6 Contributions

The major contribution of this research is providing empirical evidence that visit wear is a valuable design concept for improving memorability and usability in fisheye views. This will expand the utility of fisheye lenses as a tool for visualizing large data spaces. Minor contributions are:

- An initial definition of the design space of visit wear. Visit wear is a new extension of the concept of read wear, and an exploration of its possible parameters has not yet been done.
- A set of visualisation techniques that can be used for implementing visit wear. The principles of visual cognition can be applied to representing visit wear so that the visit wear is an enhancement to the task rather than a distraction.
- Additional evidence that landmarking strategies are a method of compensating for distortion in a virtual space. The role of landmarks in virtual space has been studied (Darken, Allard and Achille, 1998; Vinson, 1999) but not when combined with distortion of the space.
- Guidelines for automatically calculating the effective duration of displayed history by using revisitation data. Research indicates that many tasks involve revisiting information that has already been seen, and that revisitation patterns can be predicted for many tasks (Tauscher and Greenberg, 1997). This may be able to be applied to the calculation of history list length in visit wear.
- A better understanding of the limits of visit wear. Visit wear is not a suitable solution for all problems, and this thesis discusses issues that ought to be considered when applying visit wear.

1.7 Thesis Outline

The remainder of this thesis is organized as follows:

Chapter 2 outlines the foundation areas that provide a background for the tasks of defining the design space of visit wear, and building and evaluating a system that adds historical information to fisheye views.

Chapter 3 describes two studies carried out to determine which visual properties are effective in aiding memorability in a distorted space. This is done by examining the strategies that people use when trying to remember targets in a distorted environment. The results of the studies determine the principles of memorability in a fisheye space.

Chapter 4 discusses the design space of visit wear, and selects visit wear techniques that also meet the principles of memorability in fisheye space as determined in Chapter 3.

Chapter 5 describes the implementation and evaluation of a visit wear fisheye display, which compares user performance in memory tasks with visit wear to performance without visit wear. The implemented visit wear effects came from the design space discussion in Chapter 4.

Chapter 6 discusses the study results and the implications for design, including an analysis of why the visit wear succeeded when it did, why it failed when it did, the apparent limits of visit wear, and how the results of the experiment can be generalized to a broader class of problem.

Chapter 7 concludes the thesis and describes its contributions and future work that may be done.

2.0 Review of Literature

This chapter reviews previous work in several areas that underlie the proposed research. These areas include visualization of large information spaces, distortion-based visualisations, fisheye views, spatial memory in real and virtual spaces, visual cognition, and techniques for visualising historical information.

2.1 Visualizing Large Information Spaces

Several methods have been proposed for managing overview and detail in a large information space. These include multiple view methods, zoom-and-pan methods, multi-resolution methods, and focus-plus-context (F+C) methods.

Multiple-view methods, for example radar views (Figure 2.1), show the magnified region of the data in a separate display area (Smith et al., 1989). The typical problem with these views is that the user's attention is divided between the detail and the overview, and therefore the view does not preserve a sense of the global context of the magnified data. Putting the magnified view on top of the display results in an effect like the one given by the "Drag-Mag" lens (Figure 2.2; Ware, 1995). While the users' attention is now on the same display, the magnified area obscures part of the unmagnified area.

The zoom-and-pan method allows the user to select a suitable level of magnification interactively, but only one view is shown at a time. A common use is panning through scroll bars in a typical windowing environment, and zooming by selecting a level of magnification through the mouse scroll wheel or through a display menu (e.g. Figure 2.3; Bederson and Hollan, 1994). In these views, however, the sense of context within the space can quickly be lost.

Another technique is to show the focal area in high resolution and the context in a lower resolution (Baudisch et al., 2002). To be effective in displaying large amount of data, however, this needs a large (wall-size) display.

Figure 2.1: A radar view, with the context of the magnified area shown in the small overview to the right (after Smith et al, 1989). See also colour figure in Appendix A.

Figure 2.2: The Drag-Mag lens (after Ware, 1995). See also colour figure in Appendix A.

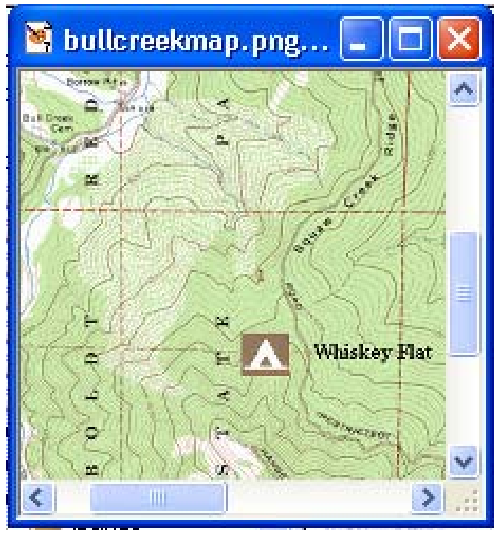

Figure 2.3: A zoom (using the mouse wheel) and pan (using the scroll-bars) interface from Paintshop Pro (Jasc Software). See also colour figure in Appendix A.

Focus-plus-context (F+C) methods integrate overview and detail into the same view. In an F+C view, the point of focus is magnified while the rest of the space (the context) is reduced to make room for the enlarged focus area. In this way, no part of the context area is occluded. However, the context area will be distorted, and this distortion is typical of F+C views (Leung and Apperly, 1994).

2.2 Distortion Based Visualisations

The basic characteristic of distortion-based visualisations is non-occluding in-place magnification that preserves a view of the global context (Leung and Apperley 1994). The distortion of the view is necessary to make room for the magnified portion of the data.

There are two main types of distortion-based visualisations (Keahey, 1998). The first is the type that uses a linear piecewise continuous magnification function. This magnification function can be constant (as in the one-dimensional bifocal view of Figure 2.4 or the two dimensional views of Figure 2.5 and Figure 2.6) or varying (such as the perspective wall of Figure 2.7). The bifocal lenses have only magnified or unmagnified areas; the perspective wall creates the illusion of three-dimensional perspective by applying varying degrees of magnification to the data not in focus.

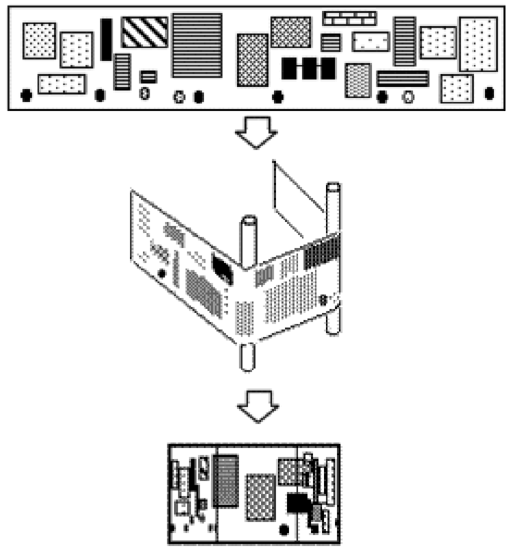

Figure 2.4: Applying the bifocal lens in a single dimension (Spence and Apperley, 1982)

Figure 2.5: Two levels of magnification in two dimensions without distortion. Information is obscured by the magnified area.

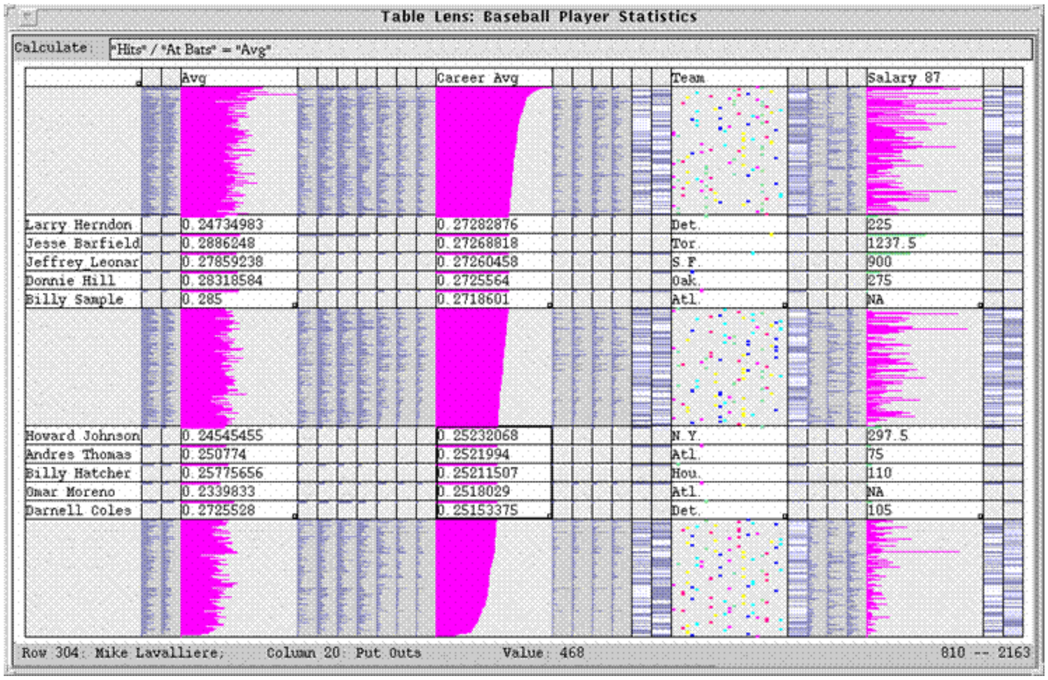

Figure 2.6: The TableLens, a combination of spatial and semantic magnification. (Rao and Card, 1994). See also colour figure in Appendix A.

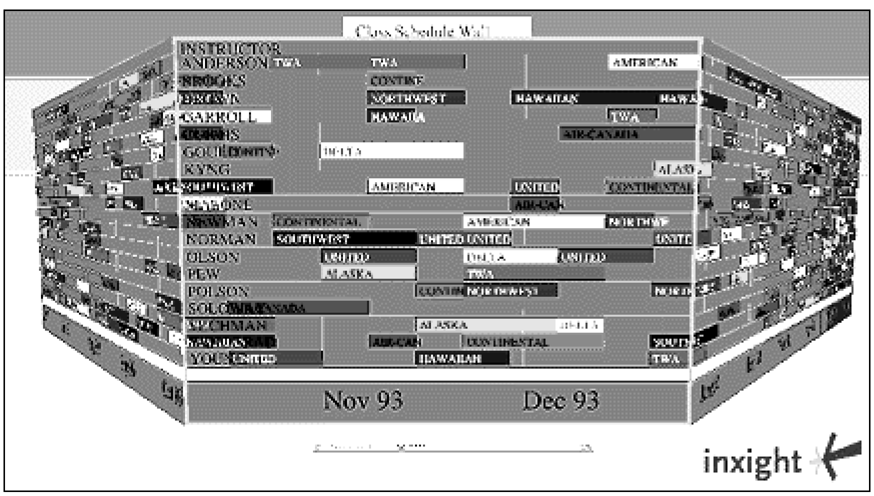

Figure 2.7: The Perspective Wall (Mackinlay, Robertson, and Card, 1991). See also colour figure in Appendix A.

The second type of distortion-based visualisation is the one that uses a continuous magnification function. This is commonly called a fisheye lens, and is discussed in the following section.

2.3 Fisheye Visualisations

A fisheye lens is a distortion-based visualisation that uses a continuous non-linear magnification function (Leung and Apperley, 1994). The idea of in-place magnification of a focus area was first proposed by Furnas and applied to text documents (Furnas, 1981). The display was based on a combination of *degree of interest* (DOI) and proximity to the focus. The DOI was decided prior to inspecting the document. Magnification and reduction in the strictly spatial sense were not involved; rather, an item of text would be highlighted or suppressed depending on its proximity to the current focus and whether its content value was greater than the DOI threshold.

Sarkar and Brown (1992) took the concept of in-place magnification and developed the mathematics to apply it to 2D spatial coordinates. The transformation function T(x) and the magnification function M(x) are given in Equation 1, where d is the *distortion factor* that determines the magnitude of the magnification, and x is the normalized distance from any point under consideration to the point of the focus. A graph of the transformation function with different values of d appears in Figure 2.8.

Equation 1

$$T(x) = \frac{(d + 1)\,x}{(dx + 1)} \qquad \text{and} \qquad M(x) = \frac{(d+1)}{(dx + 1)^2}$$

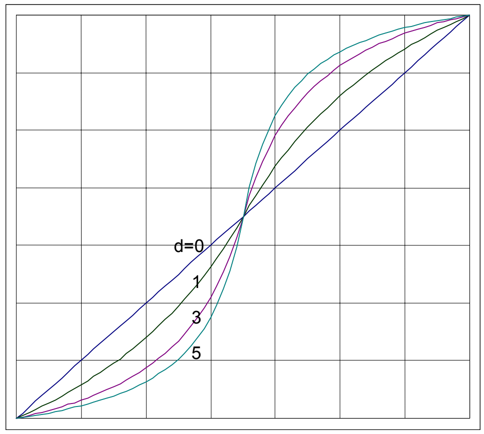

Figure 2.8: Example one-dimensional transformation functions for a Sarkar Brown fisheye at distortion levels d=0, 1, 3, and 5, with the focus set at the midway point. Source data points are projected up from below and reflect left where they meet the curve.

The actual mathematical function can vary (another common one is tanh(x) (Churcher et al., 1997)) but the only requirement is to have a varying slope on the function so that some ranges will have slope greater than one (areas of magnification) and other areas will have slope less than one (areas of reduction). A slope of one (a line at a 45 degree angle) indicates no transformation.

Misue also described similar transformation functions to those of Sarkar and Brown (Misue et al. 1995). In addition to these spatial distortions, Sarker and Brown added Furnas' ideas of DOI to provide both a method for information suppression and enhancement, and also to enhance the illusion of three dimensions due to the perspective-like elimination of details. Most of these graphical fisheyes are interactive, with the focus typically following the cursor to provide the maximum magnification at the point of user interaction.

The magnification function can be applied to the Euclidian X and Y coordinates independently (Figure 2.9), to the Euclidian polar coordinates (Figure 2.10), or by positioning the data in hyperbolic space and mapping it to a Euclidian plane (Figure 2.11). The magnification function can be applied to a two-dimensional space or a three-dimensional one (Carpendale et al., 1997) or the calculations may be done in three dimensions and then mapped back to a two-dimensional representation (Carpendale, 1999). The previous cases have all applied to the entire data space, but the functions can be constrained to part of the space (Figure 2.12).

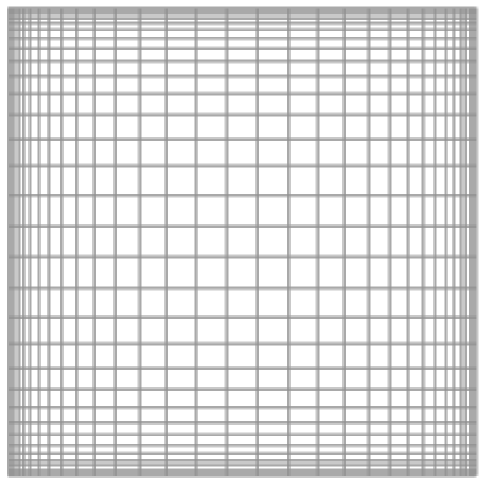

Figure 2.9: Fisheye transformation function applied to x and y coordinates independently. This is the Sarkar and Brown transformation (Sarkar and Brown, 1992)

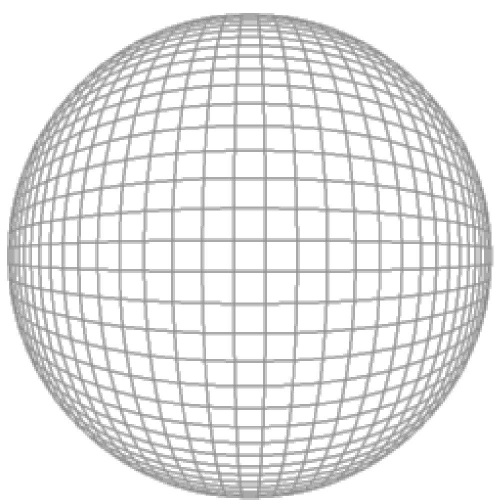

Figure 2.10: Fisheye transformation function applied radially to polar coordinates (Keahey, 1998)

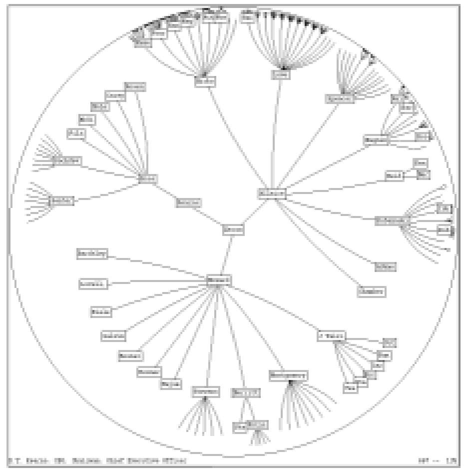

Figure 2.11: Organizational chart displayed in a hyperbolic browser. Components diminish in size as they move outwards, and there is an exponential growth in the number of components as the radius increases linearly (Lamping, Rao and Pirolli, 1995).

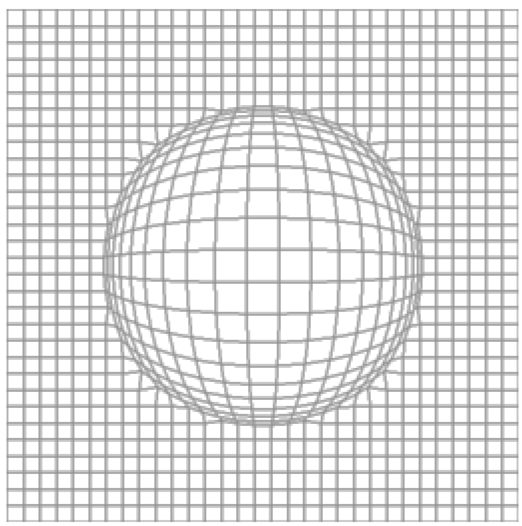

Figure 2.12: Constrained fisheye transformation (Keahey, 1998)

The linear and non-linear types of distortion visualisations can also be combined, so that the area in focus is shown at a constant magnification, while the surroundings are shown with a non-linear magnification (Figure 2.13). This *truncated fisheye* has the advantage of not distorting the area of focus, while preserving the non-occluding view of the context.

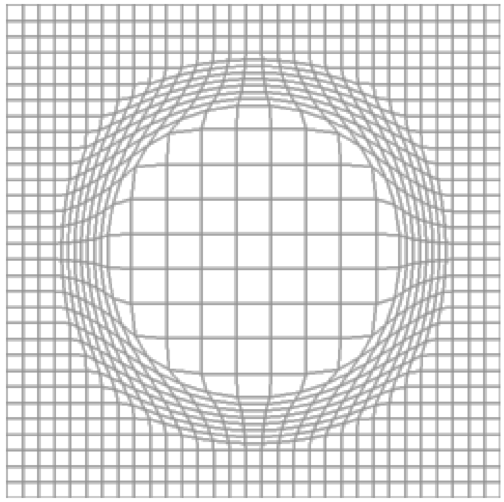

Figure 2.13: A truncated fisheye lens (Keahey, 1998)

The fisheye transformation can also be applied to multiple foci, with the magnification transformations interacting in a predetermined way. This can be especially useful in groupware applications, where each participant has their own fisheye lens. Other participants' lenses can be shown on your screen, but usually at a lower magnification than yours. In this way you know where other participants' attention is directed, which is useful in maintaining awareness of other groupspace participants (Greenberg et al., 1996). Figure 2.14 shows multiple constrained lenses "stacked" on a space.

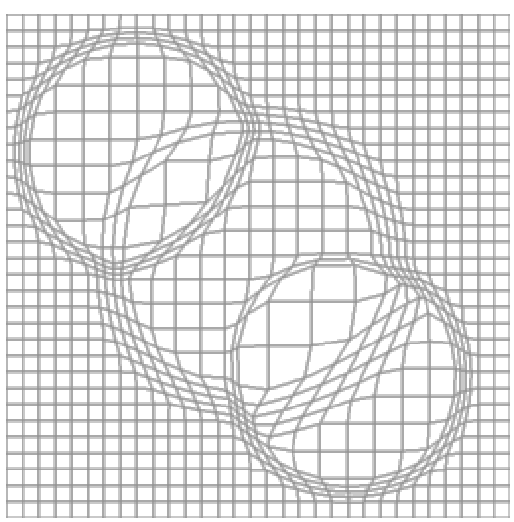

Figure 2.14: Multiple foci (Greenberg et. al, 1996)

Despite fisheye views being an established visualisation technique, they have not made a breakthrough into widespread use. The software architecture tool Rigi (Storey, Fracchia, and Müller, 1999), the Macintosh OS X "dock" (Apple, 2002), and Pliable Display Technology (Idelix, 2003) are the main examples of fisheye use in business applications. Some specific usability problems with fisheyes, such as targeting (Gutwin, 2002), have been identified and solutions suggested, but little research has been done on the effect of the distortion on people's memory of the space.

2.4 Spatial Memory in Real Spaces

The problem of how people remember where they have been in real spaces is one that has been extensively studied by psychologists, geographers, and architects (e.g. Siegal and White, 1975; Lynch, 1960; Sadalla, 1980). Lynch, in his studies on city planning, developed a theory concerning a city's "legibility," or the ease with which its parts can be recognised and organised into a coherent mental model (Lynch, 1960). Though his work concerned large spaces such as cities, it can also be applied to spatial awareness of smaller areas. Siegal and White (1975) generalized Lynch's work to apply to spaces of all sizes. Siegal and White defined the three steps taken to develop spatial awareness as knowledge of *landmarks*, knowledge of *routes* and *survey* knowledge.

In Siegal and White's design approach, knowledge of a space's landmarks is the first step in developing spatial awareness. This involves remembering and using distinctive elements in an environment to orient oneself. Elements can be distinctive through an objective visual property (such as size, shape, colour, location, semantic content) or a subjective one (such as personal memories or associations). Subjective properties usually rely on personal or cultural context, while objective properties are easier to control and more likely to be universally effective.

The next steps in Siegal and White's approach are route knowledge and survey knowledge. Route knowledge is knowing how to get from landmark to landmark. For example, a business traveller to a city may first learn how to drive from the airport to his hotel, probably using a main road. Eventually, he will acquire knowledge of multiple routes. Survey knowledge is knowing the spatial relationship of landmarks without following routes, that is, having a bird's-eye or map view of the area as one's mental model.

Changing the space will disrupt a person's ability to navigate it (Golledge, 1991; Spencer and Schöner, 2000). For example, a person with route knowledge of a city will become disoriented and possibly lost if the route he knows is shut down by construction. A person who uses a construction crane as a landmark to find her way around a campus may become confused once the construction is done and the crane disappears. Spatial memory seems to work best when the absolute positions of features do not change.

2.5 Spatial Memory in Virtual Spaces

In virtual spaces, the problem of spatial memory becomes more acute (Satalich, 1995; Peterson, 1998; Darken, 1998). Virtual spaces offer more control over environment, interaction abilities and perspectives - properties that are difficult or impossible to change in the real world. The designer of a virtual space may not be able to assume, as an architect can, that users will all approach a feature from the same horizontal plane. There are also problems with the typical presentation of a virtual three dimensional space on a small flat monitor. The view onto the virtual space is restricted, and this can lead to decline in performance in many tasks or even "simulator sickness" (Czerwinski, 1997). For example, because of the role of peripheral vision in spatial memory, which virtual spaces are almost never able to provide, users often have trouble forming cognitive maps of a virtual space (Czerwinski, 1997). The controls for moving through the space may also cause problems through unfamiliarity; while few people need to be told how to walk through a park, most cannot initially use the controls of a typical first-person video game without help.

The issue of perspective is the one of interest in this thesis. A virtual space does not necessarily mean a three-dimensional world in which the user has a limited first-person view. The navigation problems in first-person virtual spaces are closer to those of the real world, with complications such as a limited view noted above. In contrast, many data sets are often presented to the user in a bird's eye view, so that the entire space is shown at once. But when a virtual space is shown all at once, Lynch's concepts of route and survey knowledge do not apply. There are no routes to memorise, since they are all displayed (if routes exist at all) and the survey knowledge is likewise already visible. The spatial memory concept applied in this case is called a *mental map* (Misue et al. 1995).

The mental map is the observer's memory of the absolute positions of the features of the space, corresponding to Siegal and White's survey knowledge. The more the observer's mental map differs from the space the more disoriented the user becomes. This has been researched mainly in the area of layout algorithms, where adding a node to a graph, for example, should be done with a minimum of disturbance to the mental map. If adding a node results in every single existing node changing positions, the observer will take a long time to re-familiarise themselves with the space even though its semantic content may have only changed slightly (Figure 2.15).

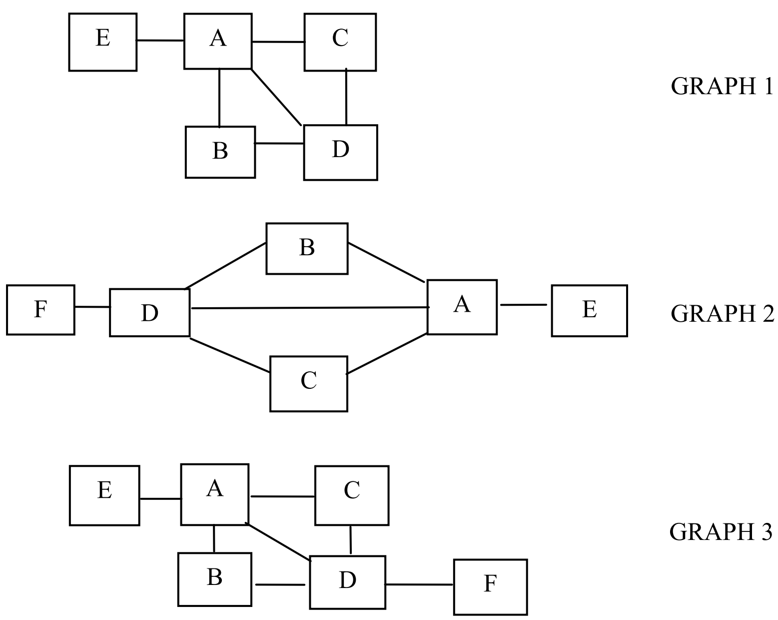

Figure 2.15: When the space is drastically re-ordered, the mental map is drastically disturbed. It is easy to see that Graph 3 is the same as Graph 1 with Node F added to Node D. It is more difficult to see that Graph 2 is the same thing (Misue et al, 1995)

When the space is changed for any reason, Misue stated that three properties must be preserved to maintain the mental map. These properties are orthogonal ordering, clusters, and topology. Orthogonal ordering is the order in which items appear from left to right and top to bottom. Clustering is whether an item is closer to one item than another. Topology is the connections between items. Most algorithms are designed to preserve these properties as much as possible when the space is redrawn (Storey et. al. 1999; Bridgeman and Tamassia, 2000). However in a fisheye distortion, orthogonal ordering can be changed, clustering is almost always changed, and topology is preserved.

2.6 Visual Working Memory

From the earliest studies in human memory there have been two types defined; short term and long term memory. A selected item from the short term memory (which only lasts several seconds) is passed into the long term memory after going through some sort or filter for importance. A model of the memory has been developed which divides the "working memory" into separate areas to handle visual and spatial information as compared to linguistic information. Baddley and Hitch (1974) called this the "visuospatial scratch pad" (Figure 2.16). The visuospatial scratch pad (VSSP) temporarily stores visual information and also acts as a "rehearsal space" for mental construction of spatial images. Information in the VSSP is stored visually, not linguistically, so it is like a "mind's eye" that actually stores images.

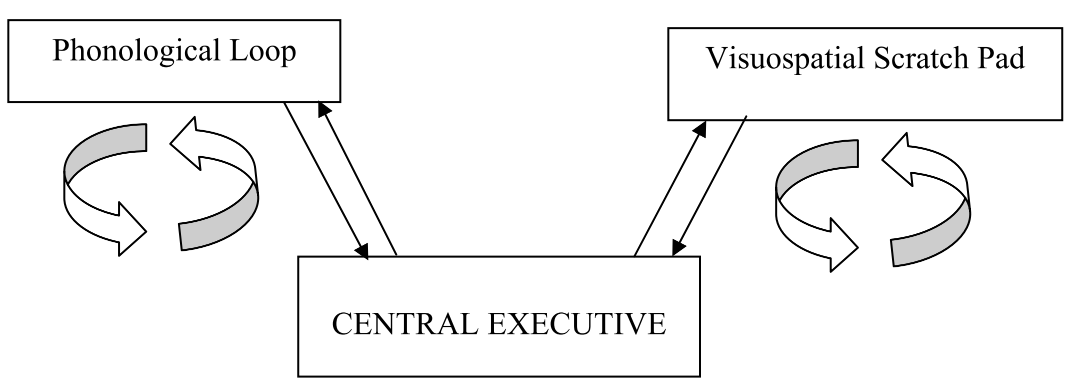

Figure 2.16: Baddley and Hitch's model of working memory (from Wesenick, 2003)

The presence and role of the VSSP has been confirmed by converging research that shows that working memory is not a single general capacity, but instead is made up of several subsystems that complete various types of tasks. Studies of change blindness (Wesenick, 2003) have investigated exactly how much visual information is stored in the VSSP, although this does not diminish the usefulness of the VSSP model. Change blindness is the observation that people actually fail to see changes in visual information when they occur during disruptions such as blinks, eye movements, or during a "cut" in the information flow. This is true even when the changes are large, occur frequently, and the observer is expecting them to occur. Change blindness can be countered by focussed attention and is also reduced by "exogenous cues at the location of the change" such as knowing that the change will always happen where a constantly present arrow is pointing (Rensink, 2000).

In a space that is viewed through a distorted lens, the appearance of the space is changing quickly and constantly. This may not only be disorienting in an abstract sense, but may cause an actual failure to perceive the change in the space.

In addition to the possibility of change blindness, something to consider is the limit of memory in general. The common agreement among psychologists is that the "magic number" for processing information is seven, plus or minus two (Miller, 1956). We can readily remember about seven items (the "span of immediate memory") and readily distinguish between seven levels of a property (the "span of absolute judgement") (Miller, 1956). When we need to remember more information, we use recoding strategies to group information into chunks, so that instead of remembering three phonemes, for example, we remember one word. But in the absence of easy recoding strategies, we cannot remember very many items. External aids, such as lists or markings, take the burden from our memory so that instead of remembering the items themselves, we only have to remember where the list is or how the items are marked.

2.7 Revisitation Patterns in Virtual Space

Visit wear is based on the assumption that users will want to know whether they have seen a data item before, either to revisit it or to avoid revisiting it. In any activity, it is common to repeat actions that have already been done once. We tend to call certain phone numbers over and over, type the same command to compile code, or drive the same route to work. In an information space, a certain percentage of actions are to revisit information that has already been seen once. Browsers that help users navigate the World Wide Web (a very large information space) have history utilities that allow easy revisitation of previously viewed pages. Since the Web has little or no structure (except at the local level) often this history list is the easiest way to find already-visited pages.

Revisitation is common in many interaction tasks. For example, Tauscher and Greenberg discovered that 85% of recurrences are to a page that was in the last ten visited (Tauscher and Greenberg, 1997). Increasing the size of the history (to the last twenty items visited, for example) does not decrease that percentage significantly because the graph of recurrences falls off so sharply (Figure 2.17). Revisitation occurs mainly with the last few pages visited. Users also did not browse in repeatable linear patterns, i.e. set sequences of page traversals.

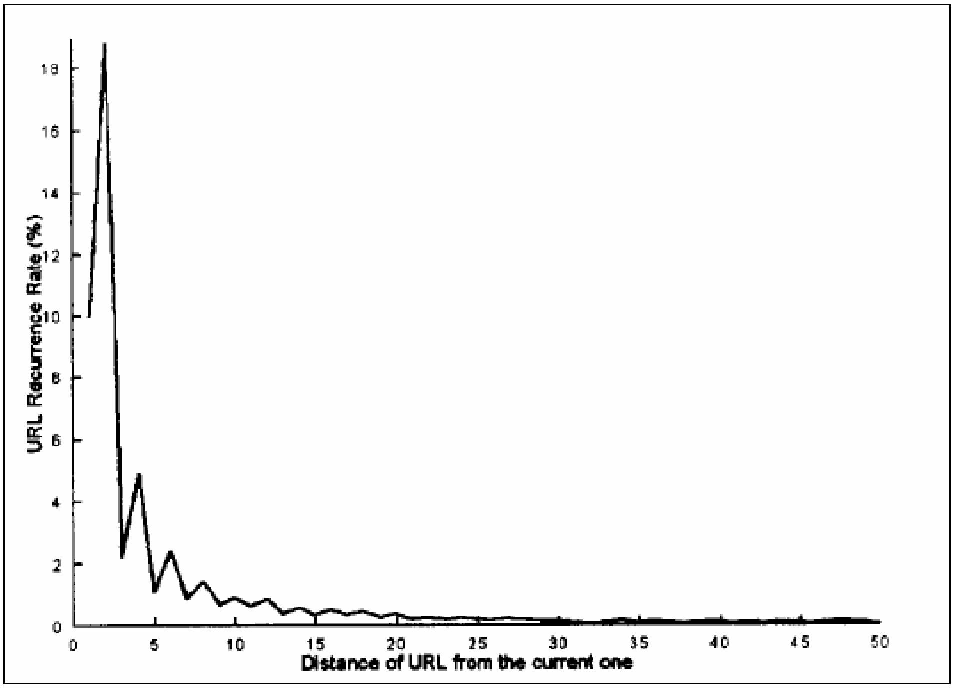

Figure 2.17: URL recurrence rate as a function of distance (Tauscher and Greenberg, 1997)

Research has also shown that a simple recency list of the most recently visited pages, with duplicates eliminated so that only the most recent visit to the page is shown, is an effective way to present interaction history. For example, Cockburn and Jones (1996) showed that the stack-based presentation common in commercial browsers does not necessarily match the mental model that users have of their interaction history, whereas a recency list is a better match. Tauscher and Greenberg (1997) determined that a recency list might be outperformed by other methods that involved organising the history into a hierarchy, but at the cost of an increased cognitive load for the user who must remember the contents of non-visible areas of the hierarchy.

The implication of the research on revisitation patterns is that a simple history of the few most recently inspected items (with duplicates eliminated) will be reasonably certain to contain the next item that the user will want to inspect.

2.8 *Visualising Historical Information*

There are two main ways of adding historical information to an application: manually, in which a utility is provided for the user to add the information themselves, or automatically, in which the information is collected and provided by the system without user interaction. Letting the user explicitly set flags or markings at items of interest has analogies in the real world such as putting slips of paper between pages of a book or marking blazes in orienteering. Automatic recording of activity also has an analogy in the real world, such as leaving footprints on a beach or any other case where our environment shows changes according to our interaction with it, without effort on our part.

2.8.1 Manually Adding Historical Information

As early as 1945, Bush identified the problem of remembering which items in a large space were of interest (Bush, 1945). His solution was to allow the user to add "trails" to the information; records of connections between information items defined by the user as associative indices. Items could belong to as many trails as needed. Although the amount of information available on the Web exceeds even that imagined by Bush, his suggestion of storing trails has not been widely implemented as such. In most commercial browsers, a person would have to use the bookmarking utility and manually group related sites to duplicate Bush's trails. Still, the ability to mark where one left off, or to mark possible sites to revisit, has proven to be important in browsing tasks (Jones and Cockburn, 1996).

Most forms of bookmarking keep the information separate from the main data, and usually do not present this information visually. This is in part because the tasks of the applications are themselves typically not handled visually; for example, most Web browsers present a one dimensional (one page at a time) view of the data. The development of different visualisation methods for the Web, such as a hyperbolic browser, may also result in different methods for displaying bookmarks (Munzner and Burchard, 1995).

2.8.2 Automatically Adding Historical Information

While bookmarking requires conscious participation, the other method of adding historical information requires no effort on the user's part. A browser's history list is maintained automatically by the browser, for example. This list is typically a simple stack, although Cockburn and Jones suggested that usability is improved by "context sensitive Web subspaces" where a cascading menu is developed as different pages are accessed. In this manner, menus tend to be related groups of pages (Cockburn and Jones, 1996). As vital as history mechanisms are (as discussed in the previous section), most browsers display the historical information separately from the information itself, in log files or pulldown menus.

Instead of keeping the historical information separate, it can be shown as a change in a visual property of the information itself. Automatically providing historical information visually is known as *read wear* and mimics the natural way in which everyday objects show traces of their use (Hill and Hollan, 1993). In the real world, wear is a gradual, and usually unavoidable, change due to use (Figure 2.18). Favourite recipes are often identifiable through stains and smudges on their cookbook pages. The shortest distance between two buildings can be guessed by a dirt trail on a lawn where the grass has been worn away. Wear can help us see past patterns of use, either by ourselves or others, and this can give us useful information since we are used to picking up on these clues (Wexelblat, 1998).

Figure 2.18: The path taken by people who have crossed this street in the past is clearly shown. See also colour figure in Appendix A.

In a virtual environment, any changes to the information must be deliberately designed and introduced. Left to itself, a computer document will not yellow with age as a paper document will. Since the changes must be designed, they can follow an alternative physics, where only changes that are of benefit to the effect need be presented (Hill et al 1994). In the real world, for example, everyone who passes will leave footprints on a beach. In the virtual world, however, only selected users or events need leave traces, and even then not under all circumstances (Gutwin, 2002). In a Web browser, for example, the default behaviour is that visited links are automatically shown in a different colour from unvisited links (Figure 2.19) However, this behaviour is completely changeable by the user, from changing the colours to not having it happen at all.

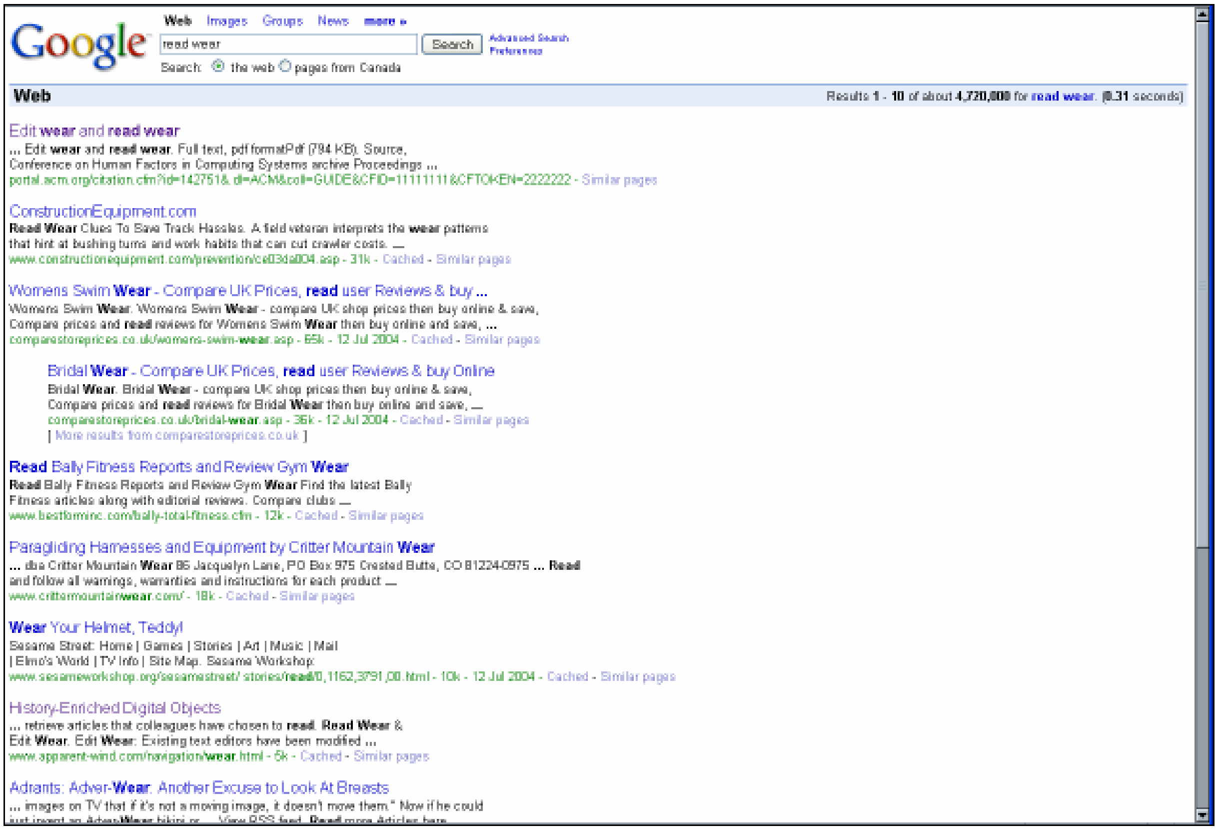

Figure 2.19: A typical Web page showing visited links in a different colour. See also colour figure in Appendix A.

Balancing this freedom to invent new laws of behaviour is the necessity to make the results understandable. For example, a spreadsheet can be enhanced with edit wear so that cells change colour according to how long it has been since they have been changed. Any colours can be chosen, but some will make more sense than others. Figure 2.20 shows two colour schemes; one reflecting a natural and (hopefully) familiar process of yellowing due to age, and one using shades of magenta. With the magenta choice, it might take longer to realize that the more intense and visually arresting shade actually means that the items are older (and therefore possibly of less interest). “Alternative physics,” i.e. the rules of a virtual space, should allow users to exploit their knowledge of the real world while still providing information not normally available (Hill and Hollan, 1992).

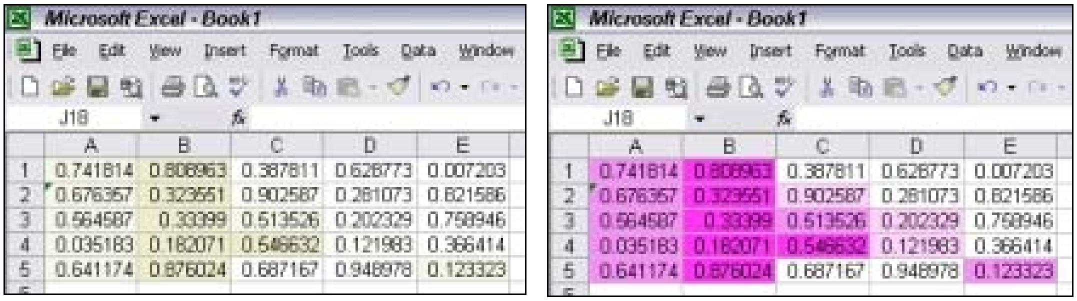

Figure 2.20: Two colour choices for an edit wear enhanced spreadsheet. See also colour figure in Appendix A.

2.9 Visual Cognition and Interaction History

The goal with adding any type of historical information visually is having a user answer the question "Where have I visited?" through visual rather than cognitive processing. The area of visual cognition has been well researched, and there are many factors that can govern the encoding of objects and patterns so that the eye can process them (Card et al., 1999). The main principles that are relevant to the visualization of interaction history are gestalt theory, retinal properties, and pre-attentive processing.

The Gestalt psychologists studied how the brain organizes basic stimuli (Solso, 2001). In particular, the tendency of the human brain to group objects with common characteristics may be useful when adding historical information to a visualisation, by giving visited features a common characteristic. However, not all characteristics identified by the Gestalt psychologists work equally well for the purposes of adding information to an existing data set. For example, one of the characteristics that tend to group objects is proximity (Figure 2.21), but moving visited features on a map together to indicate that they have all recently been visited is not a good solution, since the meaning of map is changed by doing that.

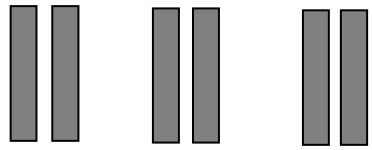

Figure 2.21: Grouping by proximity. Six vertical lines are not seen as six independent lines, but instead as three groups of two.

Bertin (1983) defined "retinal properties" of graphical elements, as shown in Figure 2.22. The human eye is sensitive to these properties, independent of the position of the objects. However, Bertin's properties were restricted to those that can be recorded with paper and ink. Other researchers have expanded this list to include properties that can be conveyed on a computer such as direction of illumination, direction of motion, and 3D depth cues (Healy, Booth and Enns, 1995).

In the field of cognitive psychology, Treisman and Gelade (1980) conducted studies that showed that searching for a target distinguished by a conjunction of features (e.g., colour and shape) demanded attention, whereas a search for a target distinguished by one feature (e.g., colour) could be done pre-attentively. Pre-attentive processing means that the time it takes to pick out these targets remains constant even if the size of the data set increases.

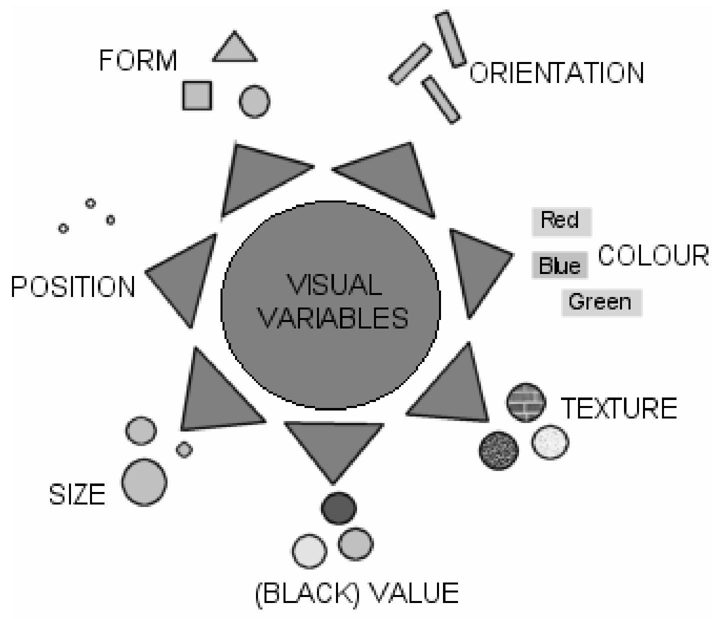

Figure 2.22: Retinal properties (Bertin, 1983)

Obviously, having historical information represented by something with pre-attentive properties would effectively answer the “Where have I been?” question. The fact that pre-attentive processing does not increase with set size is also of benefit. According to Triesman and Gelade, though, the target must be distinguished by a single property. Choosing a single property that makes an item stand out from the other data may be a challenge if the data are varied in appearance. However, many researchers currently favour the idea that information can be selected from different levels of representation, such as a conjunction of features or grouping, depending on the particular demands of the task (Baylis and Driver, 1993).

The addition of visual historical information should be done in such a way that little or no mental effort is needed to identify and understand the added information. The principles of gestalt theory, retinal characteristics, and pre-attentive processing can be used to make the added information easier to identify.

3.0 Memorability in Fisheye Applications

3.1 Introduction

As described in the previous section, efficiently navigating a data space relies on the user's memory of it. If the space changes through the distortion of a fisheye lens, the user can become confused. Some properties such as absolute position can change dramatically, some properties such as colour, texture and semantic content can become invisible if the object is not in focus, and some, such as proximity, change relatively little. Some visual properties as described in Sections 2.5 and 2.9 may retain their distinctiveness even when the view is distorted, and I will describe these properties as being *robust* with respect to distortion. Items with robust properties may be more memorable in a distorted space, both in and of themselves, and because they can become *landmarks* that aid in remembering nearby targets as described in Section 2.4.

To test which properties are robust with respect to distortion, two studies were performed. The first study gathered information about which properties people thought would be effective in remembering particular screen objects at several increasing levels of distortion. The second study looked in detail at the strategies used to actually re-find an object in a distorted view. If the strategy used was mostly based on the visual properties of the target or of landmarks, then the properties were robust and useful.

The results of the studies are in Table 4, which lists visual properties identified by the participants and a description of the effect of distortion on these properties. In addition, Table 5 predicts the effects that distortion will have on the visual properties that were not tested.

3.2 Experimental Setup

Both studies used the same two-dimensional graph for the data space (see Figure 3.1, Figure 3.2, Figure 3.3 and Figure 3.4, noting that the focal point is in a different position in each figure). The graph was always displayed using a Sarkar-Brown fisheye visualisation at one of four different distortion levels (d=0 (no distortion), d=1, d=3 and d=5, where d is the distortion factor). The nodes were repeated thumbnail representations of Web pages, with several colour schemes, patterns and semantic content. There were varying numbers of edges between nodes, and the topology of the graph created several unique node-and-edge shapes, referred to below as “constellations.”

This information space was deemed to have realistic visual properties, since data objects seldom have only one of the properties mentioned by Bertin or Misue. The properties that I expected participants to identify were grouped into five categories:

1) Absolute position on the screen
2) Node colour or pattern (visual properties described in Section 2.9)
3) Orthogonal ordering (remembering the position of nodes relative to each other, as described in Section 2.5, especially edges and corners)
4) Constellations (remembering the shapes that the nodes and edges made, a combination of topology and clustering as described in Section 2.5)
5) Semantic content of the nodes.

The Sarker-Brown lens was chosen because it affects every part of the space. When the focal point of the lens is moved, every item in the space is magnified or de-magnified to some extent to make room for the magnification at the focal point. Any distortion effects are therefore likely to be maximized and thus the effect of distortion is easier to observe. Other types of fisheye lens, such as constrained lenses, may show fewer effects from distortion since the distortion does not affect the entire screen.

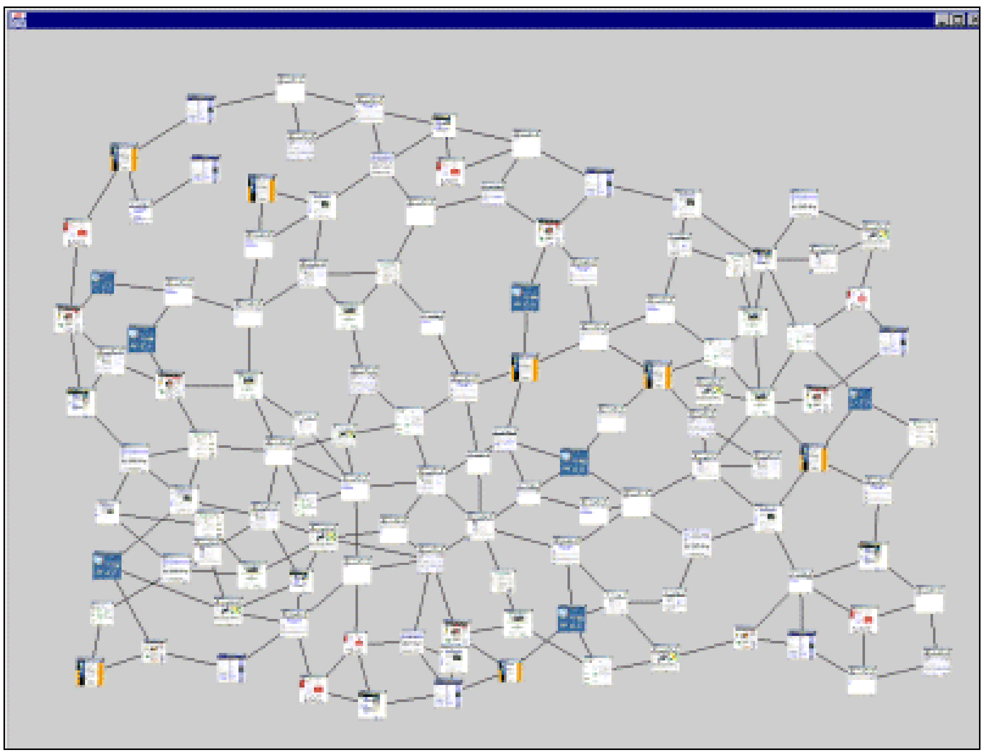

Figure 3.1: Graph used in the study, showing distortion level 0 (d = 0). See also colour figure in Appendix A.

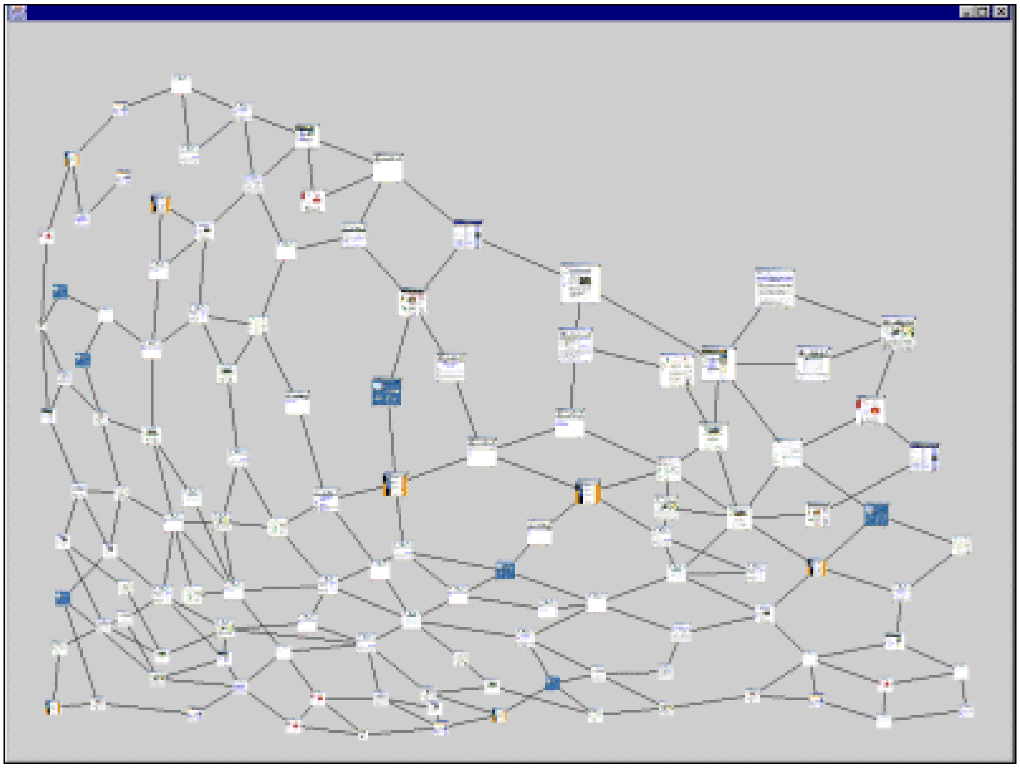

Figure 3.2: Graph used in the study, showing distortion level 1 (d = 1). See also colour figure in Appendix A.

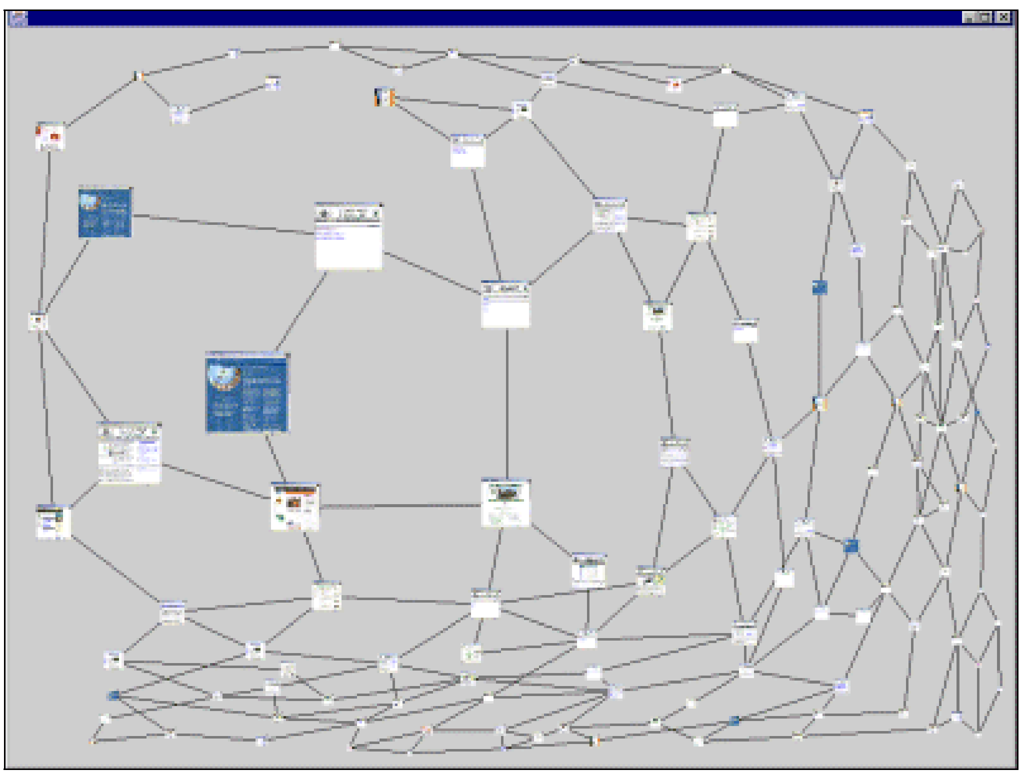

Figure 3.3: Graph used in the study, showing distortion level 3 (d = 3). See also colour figure in Appendix A.

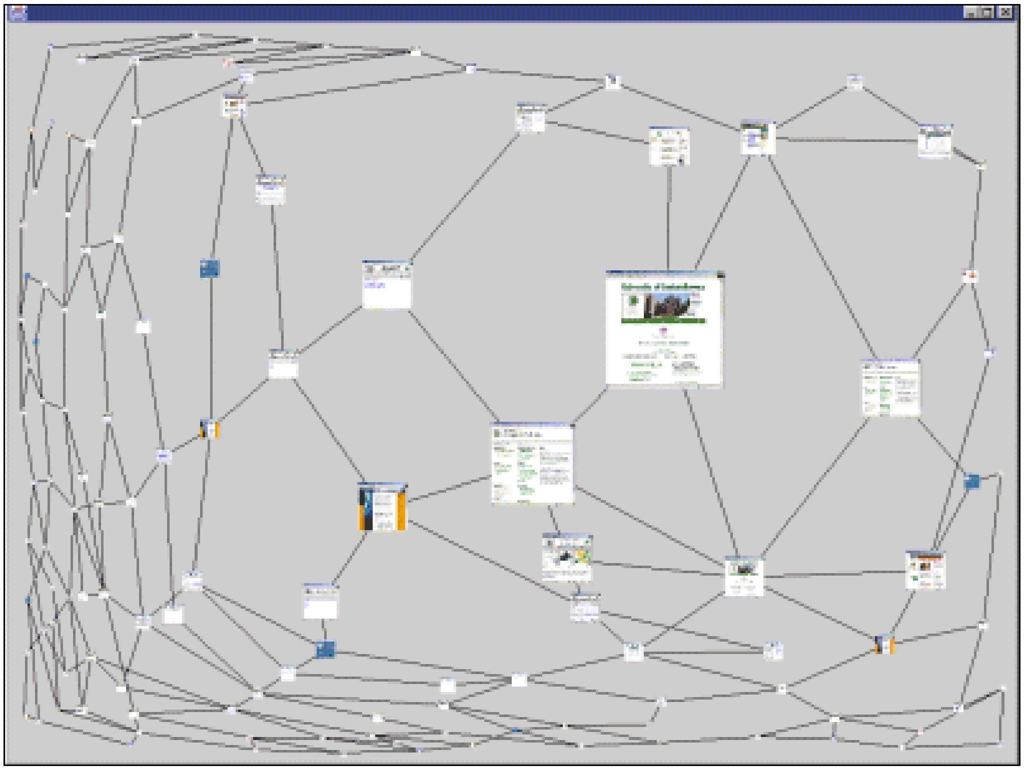

Figure 3.4: Graph used in the study, showing distortion level 5 (d = 5). See also colour figure in Appendix A.

3.3 Identification of Memorable Properties

The first study gathered information about which properties people thought would be effective in remembering particular screen objects at several increasing levels of distortion. I expected participants to use the five properties described in the previous section (absolute position, colour and pattern, orthogonal ordering, constellations and semantic content). Of course, any properties not in the above list were also noted as they were mentioned by the participants.

3.3.1 Participants

The subject group for the first study was seventeen fourth-year Computer Science students at the University of Saskatchewan. Their participation was rewarded by a bonus mark in a class that they were all taking. All participants were frequent computer users (at least twelve hours a week) and only one had ever used a software application with a fisheye view before. Although this is not a demographically varied participant pool, the purpose of this study was to do a general exploration of the problem of re-finding targets in a distorted space, and to provide direction for a more rigorous study later.

3.3.2 Apparatus

The experiment was run on a PII Windows NT PC. The graph was shown on the PC with a custom built Java application to provide the fisheye distortion. The focal point of the fisheye lens was linked to the mouse cursor, and participants were able to freely move the focal point to investigate the effects of the distortion on the graph.

3.3.3 Procedure

Participants were asked to complete an informed consent form and a short demographic questionnaire, as given in Appendix B. They were then introduced to the system and to the fisheye representation. They were allowed to interact with the system and move the focus point for several minutes, in order to familiarize themselves with the way that the distortion effect worked. Participants were then asked to describe verbally how they would remember particular nodes in the graph, for each of several distortion levels and focus positions. With each screen, participants were given time to understand the current distortion effect by freely moving the mouse for one minute.

If participants did not mention any particular memorable properties, they were asked if there were any objects that stood out with the current distortion level and also if there were any areas of the graph that would be particularly difficult to remember.

3.3.4 Results

I expected participants to mainly identify screen position and constellation as memorable properties at low levels of distortion, since those are properties defined as important in the mental map and at low levels of distortion, the users' mental map would be unchallenged. At high levels of distortion I expected greater use of colour and semantic content, since these properties do not vary with distortion. However, contrary to expectations, participants chose colour or proximity to a colour as the most distinctive property at all levels of distortion except d=5 (Figure 3.5). At this level of distortion, the majority opinion was that out-of-focus nodes were too small to see their colour or pattern easily. Being part of, or near, a constellation was the second most popular property, and actually became the first choice at the maximum distortion level despite the distortion's effect on the constellations' shapes. Absolute screen position remained secondary to the other properties, but at higher levels of distortion it became a popular way to specify a neighbourhood so that the focus magnification would show enough detail to use other properties as landmarks.

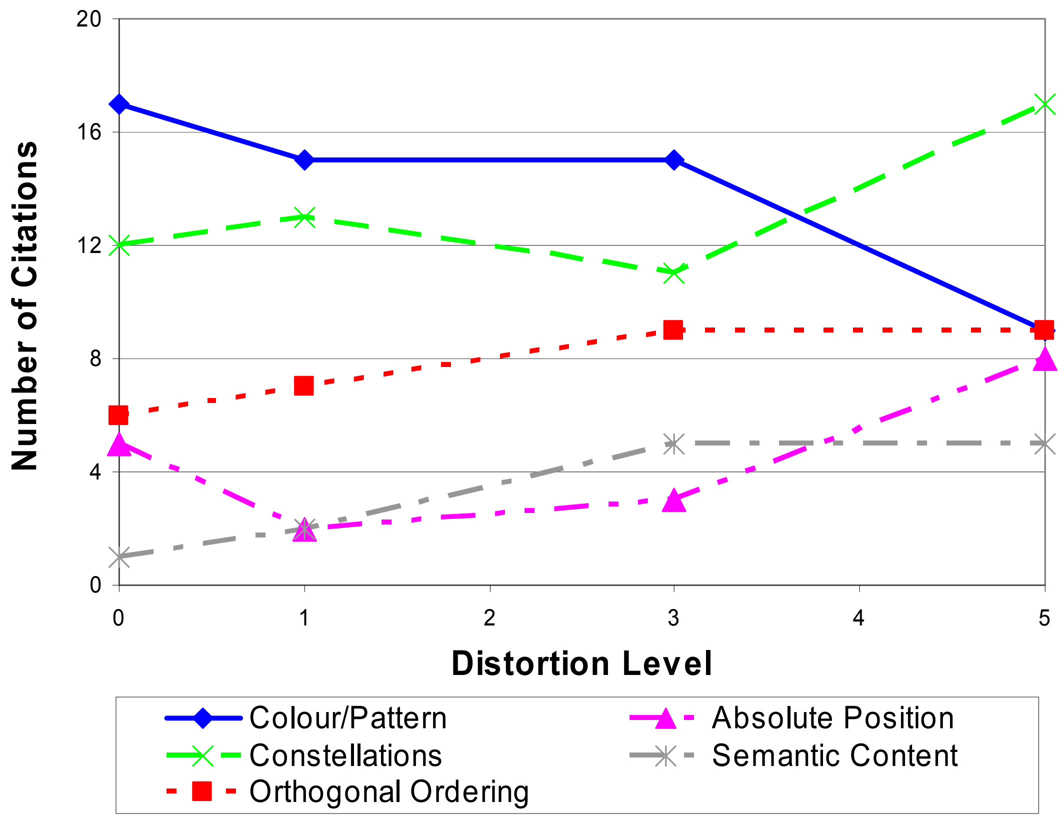

Figure 3.5: Memorable visual property choices (number of times mentioned over all tests) for each distortion level. See also colour figure in Appendix A.

Certain screen objects were identified by most participants as having memorable properties, and these objects were commonly used as landmarks to aid in remembering nearby objects as well. Figure 3.6, Figure 3.7, Figure 3.8 and Figure 3.9 show the most commonly chosen landmarks circled for each distortion level. The weight of the circle is roughly proportional to the number of participants who identified that landmark, with the lightest weight indicating about three identifications, and the heaviest weight indicating eight or more identifications.

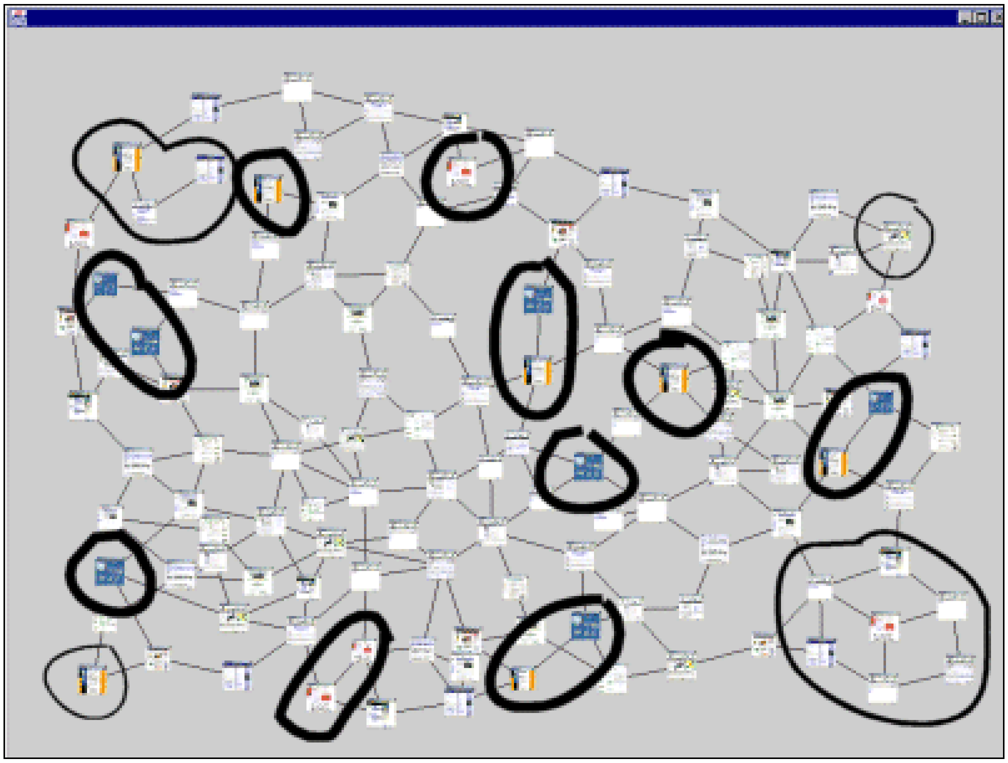

Figure 3.6: Landmarks chosen at d = 0 (no distortion) See also colour figure in Appendix A.

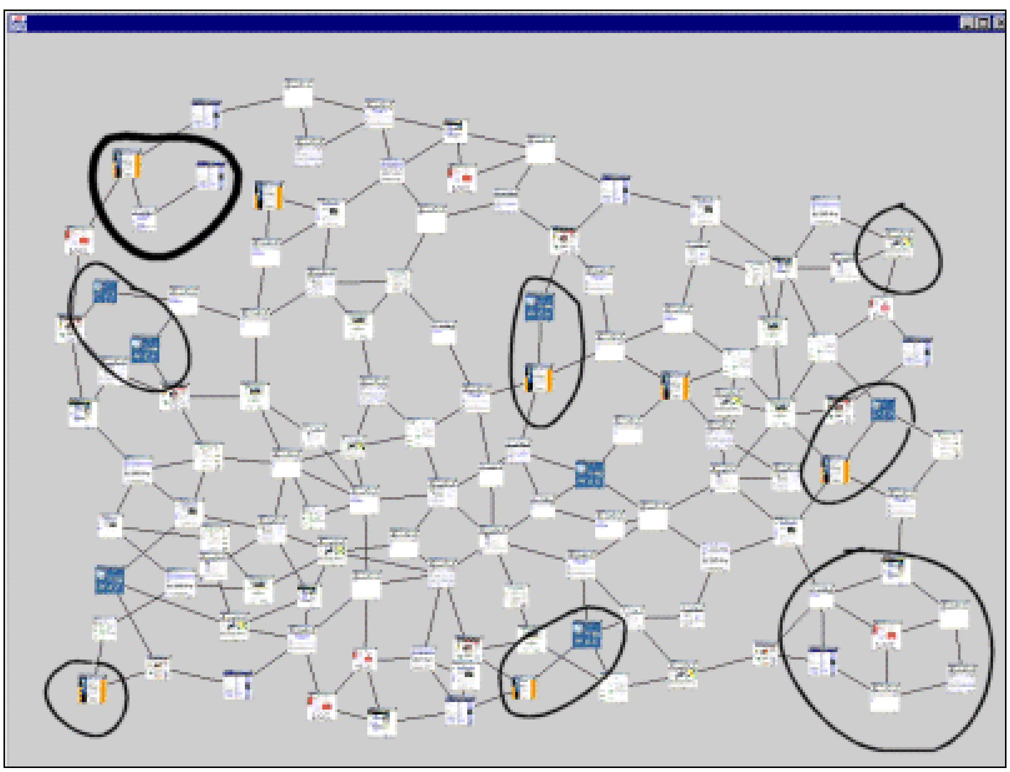

Figure 3.7: Landmarks chosen at d = 1 (note that distortion is not shown) See also colour figure in Appendix A.

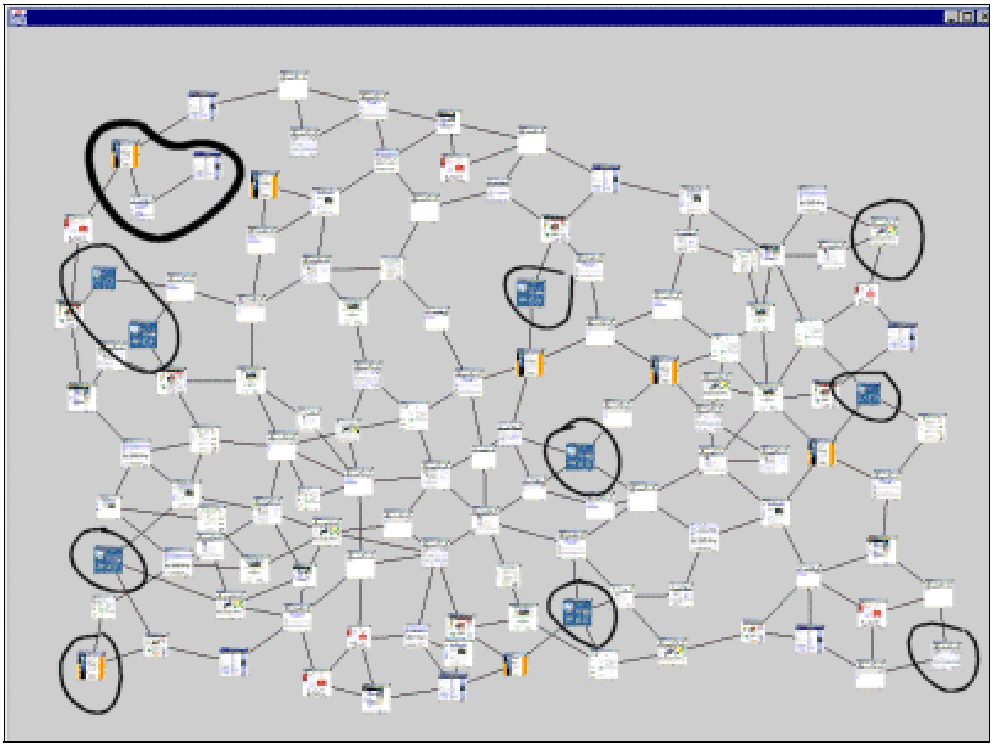

Figure 3.8: Landmarks chosen at d = 3 (note that distortion is not shown). See also colour figure in Appendix A.

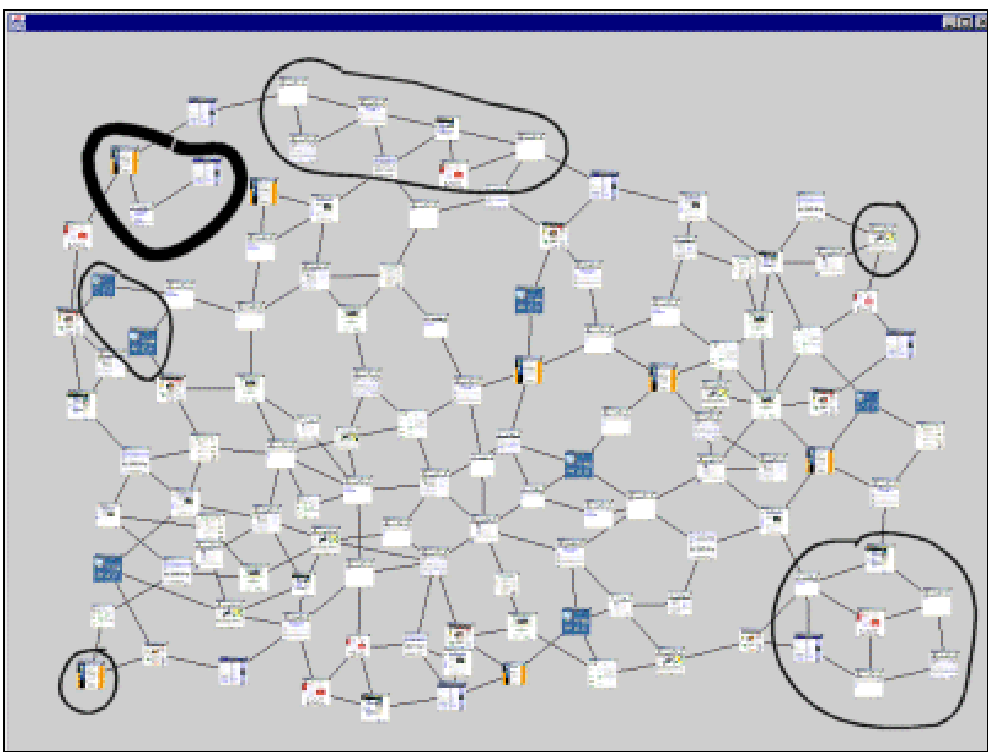

Figure 3.9: Landmarks chosen at d = 5 (note that distortion is not shown). See also colour figure in Appendix A.

At zero distortion (d = 0), all participants chose colour or proximity to a distinctive colour as the way that they would remember any screen object. All seventeen participants mentioned colour as the first property that made an area distinctive. Pairs of distinctively coloured nodes were identified as being especially useful by five subjects. When constellations or graph edges and corners (orthogonal ordering) were mentioned, they were always secondary to colour.

When distortion was added to the graph (d = 1), colour was still the most commonly mentioned property. However, nine participants said that now the pairs of coloured nodes were more useful than the single nodes. Also, as shown in Figure 3.2, fewer participants used absolute screen position and more relied on the edges and corners of the graph itself. Six participants, however, said that this level of distortion did not affect their memory strategies at all.

As distortion increased (d = 3), the effect of the fisheye lens became noticeable to all the participants. Seven said that only dark blue was an effective colour now, since the nodes became too small when not magnified to see any of the other colours. However, eight people also said that the constellations became too distorted to be effective aids to memory. The overall effect was described as "very uncomfortable" by one subject, and several said that nothing was useful as a landmark unless the focus was already in the neighbourhood. Absolute position was mentioned only in the context of getting "into the neighbourhood" and then using other landmarks. The use of edges and corners, in default of any other feature, was more frequently mentioned. The semantic content of the nodes, while mentioned by a few participants at no or low distortion, was also more frequently mentioned. This was because the magnification of the focus was finally large enough for the details of the node to become easily visible.

At the maximum level of distortion (d = 5), most participants found that even the dark blue nodes became too small to distinguish when they were out of the focus. Almost all participants that still used colour as a property did so only after using absolute screen position to put the focus in the neighbourhood first. At this distortion level, however, the greater apparent motion of the nodes as the fisheye focus moved seemed to help several participants see the constellations. One subject explicitly commented "I just now noticed that sawtooth shape!" and several said that the cube shape in the lower right corner was actually easier to see as a unit as it was being moved around by the changing focus. Semantic content was mentioned, like colour, as being useful when the focus was in the neighbourhood.

After the distortion sessions, participants were also asked to identify areas of the graph they thought would be especially difficult to remember. All participants identified the lower left of the graph as being difficult, since it is a dense and homogenous area where no objects have very distinctive properties.

3.4 Effects of Properties on Memory Tasks

The second study examined users' actual performance in memory tasks. The purpose was to investigate any correlation between the ability to remember the location of a target object with the object's possession of a memorable property mentioned in the first study. The properties were absolute position, colour and pattern, orthogonal ordering, constellations and semantic content. If the target object itself did not possess any memorable properties, then the user was expected to either try to remember the target's absolute position on the screen (a mental map strategy), or try and remember the target's relationship with a nearby object that did have a memorable property (a landmarking strategy). I was interested in seeing whether the users' ability to recognize memorable properties changed with distortion level in the same way that the first study predicted.

3.4.1 Participants and Apparatus

The subject group for the second part of the experiment was seven participants recruited from the first study. The second study used the same experimental setup as the first, and used screenshots from the same application using the same two-dimensional graph.

3.4.2 Procedure

The second study tested the participants' memory of the experimental space through a set of memory tasks. For each distortion level, five memory tasks were given. For each of these, the graph was shown with the focus fixed in a certain location, and a node in the graph (not necessarily at the focus) was outlined in red. Participants were asked to memorise the location of the red target node. A second static image of the graph was then shown but with the focus at a different point and without the red outline on the target. Again, the mouse cursor did not affect the focal point. The participants were then asked to click on the target node.

The study used a 4x1 within-participants factorial design. The factor was distortion level (d=0, 1, 3, 5). With seven participants and five trials, there were 140 tasks recorded in total. The study system recorded the participant's selections and the time required to find each target in the second view.

3.4.3 Results

In each task, two factors determined how much the graph changed between the first and second view: the distortion level, and the distance that the focal point changed moved (*focal point delta*). These factors affected the absolute position of each node of the graph to some extent. However, participants' performance was also dependent on the proximity of the target node to a distinctive landmark. The following sections discuss how distortion, focal point delta, and proximity to landmarks affected performance.

Effect of Distortion Level on Performance

In general, I expected participants' accuracy in the memory tasks to decrease as the distortion level increased, and the average time to find the target to increase with the distortion level. In general, this did occur (see Figure 3.10 and Figure 3.11); however, there is considerable variation within each distortion level. Some tasks took longer to complete at a lower distortion level than ones at a higher level, and some tasks at low distortion had lower accuracy than tasks at a higher distortion.

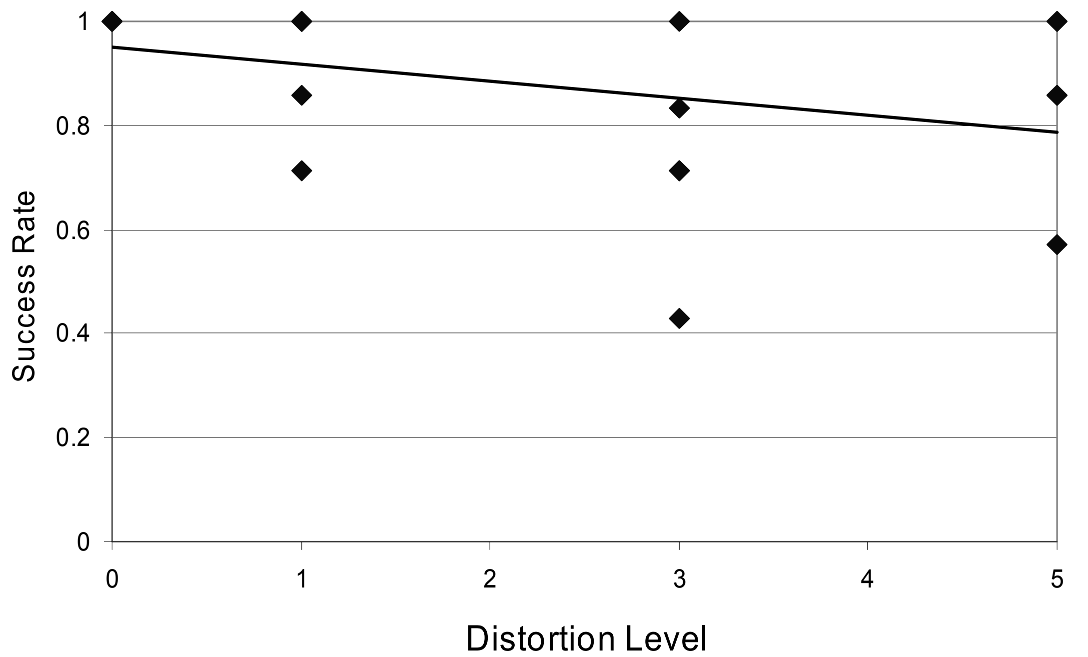

Figure 3.10: Accuracy in each memory task as a function of distortion level. The line represents the trend of the mean; each mark represents the mean success rate of an individual trial. There are five trials per distortion level; some marks are superimposed.

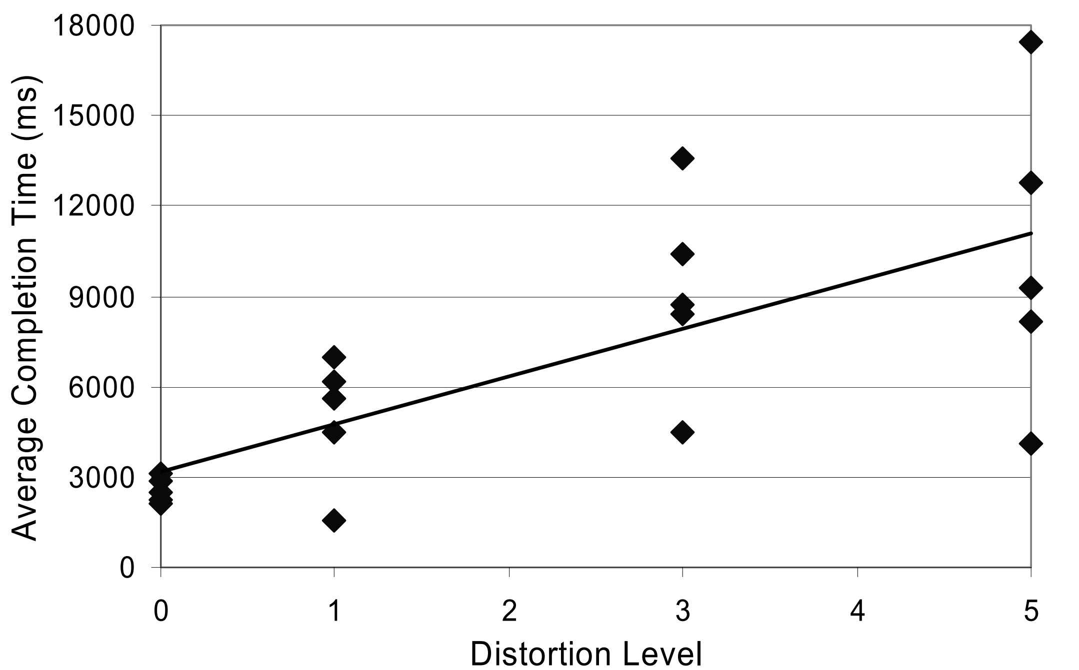

Figure 3.11: Completion time in each memory task as a function of distortion level. The line represents the trend of the mean; each mark represents the mean completion time of an individual trial. There are five trials per distortion level; some marks are superimposed.

There is only a very weak relationship between accuracy and distortion level ($r^2 = 0.037$ $p < 0.005$), and a slightly stronger one between completion time and distortion level ($r^2 = 0.18$, $p < 0.001$).

Effect of Focal Point Delta on Performance

Similar to the effect of the distortion level, I expected participants' accuracy in the memory tasks to decrease if the focal point delta increased, and the average time to find the target to increase if the focal delta increased. Again, this occurs in general, but with large variations (Figure 3.12 and Figure 3.13). Completion time and accuracy are not well predicted by focal point delta. For accuracy, $r^2 = 0.0072$, $p = 0.32$ and for completion time $r^2 = 0.0001$, $p = 0.93$.

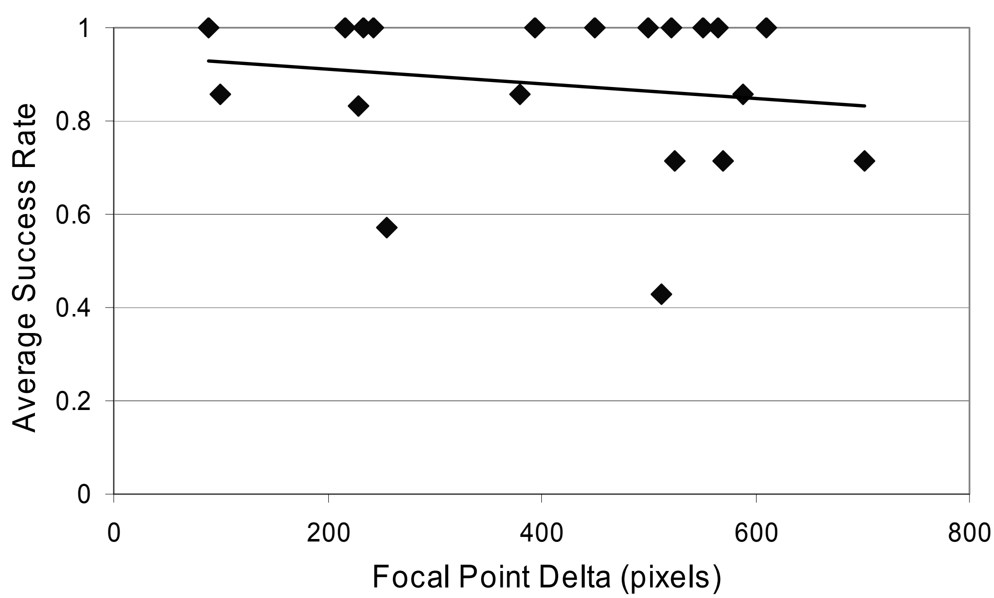

Figure 3.12: Accuracy in each memory task as a function of focal point delta. The line represents the trend of the mean; each mark represents the mean success rate of an individual trial (four distortion levels with five trials each, some marks are superimposed).

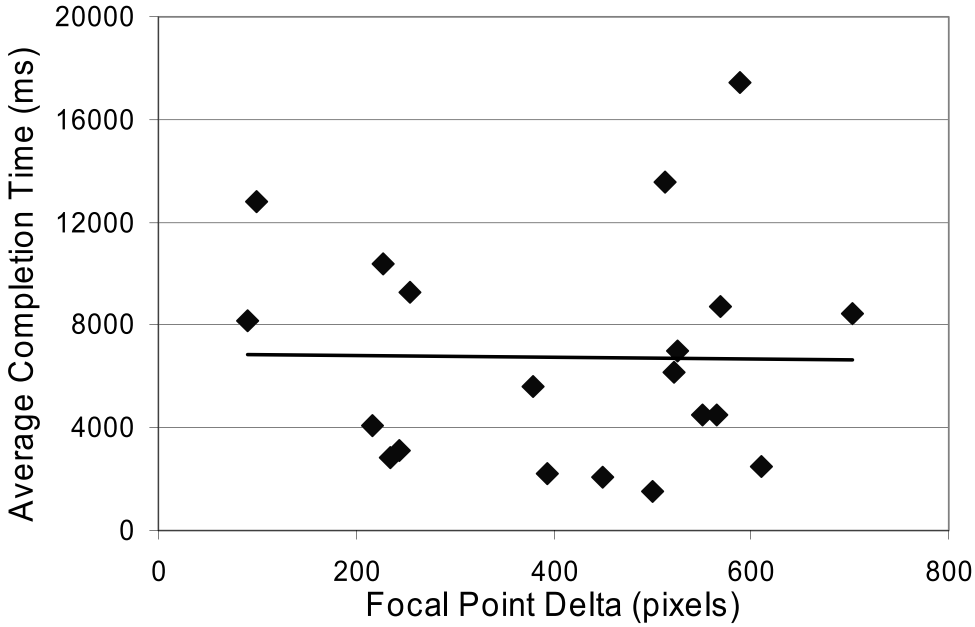

Figure 3.13: Completion time in each memory task as a function of focal point delta. The line represents the trend of the mean; each mark represents the mean completion time of an individual trial (four distortion levels with five trials each; some marks are superimposed)

Effect of Landmark Proximity on Performance

In the first study, participants used objects with distinctive visual properties as landmarks. We therefore investigated the relationship between performance and landmark proximity. These effects are shown in Figure 3.14 and Figure 3.15. The distance to the landmark is expressed in steps - the graph path length between the target and the nearest landmark. Landmark locations were defined by the participants in the first study, and counting steps in the graph was the method used by all participants when relating the target to a landmark.

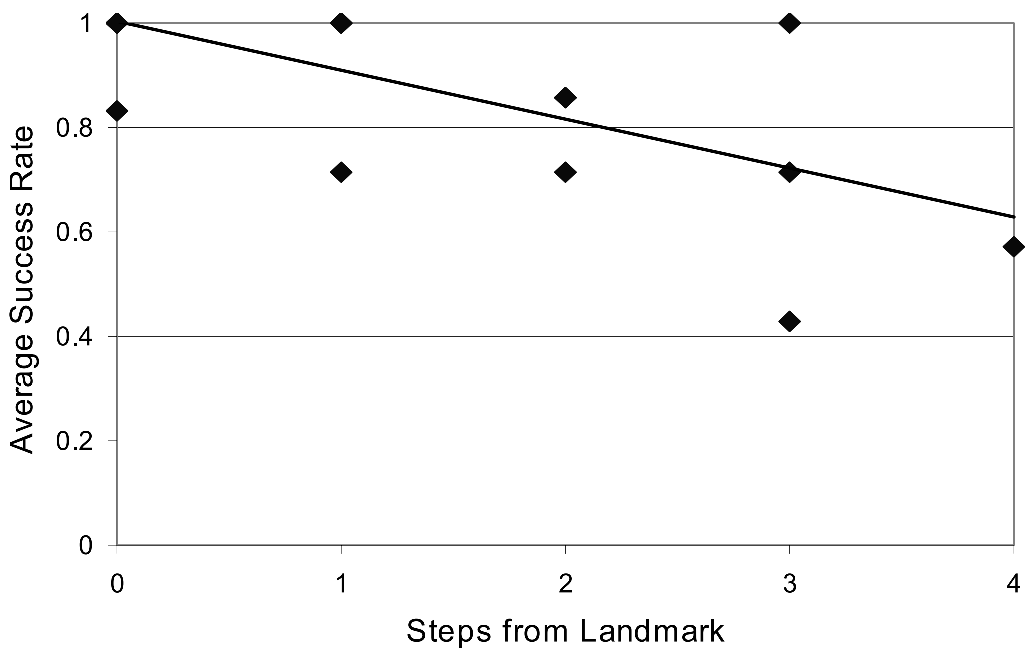

Figure 3.14: Accuracy in each memory task as a function of steps from a landmark. Each mark represents the mean success rate of an individual trial (four distortion levels with five trials each; some marks are superimposed).

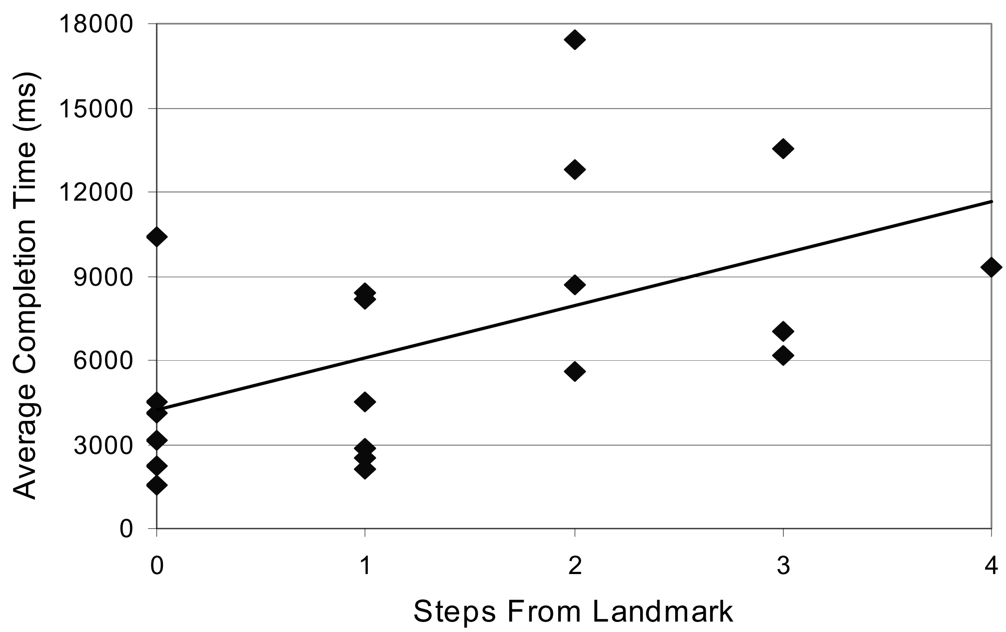

Figure 3.15: Completion time in each memory task as a function of steps to landmark. Each mark represents the mean completion time for an individual trial (four distortion levels with five trials each; some marks are superimposed)

Completion time and accuracy are strongly predicted by landmark proximity (for accuracy, $r^2 = 0.12$, $p < 0.001$; for completion time, $r^2 = 0.089$, $p < 0.001$). They are as strongly predicted by landmark proximity as they are by distortion and focal delta When all three factors are combined, the prediction is stronger, but still explains only a small fraction of the overall variance: for accuracy, $r^2 = 0.14$, $p < 0.001$; for completion time, $r^2 = 0.22$, $p < 0.001$).

Case Studies

Several of the tasks provided interesting case studies of how distinctive properties of objects interacted with the performance measures. In these cases, tasks at the same distortion level and with comparable (within 15%) focal point deltas had very different results for accuracy and completion time. Three in particular are considered, one at d=1, one at d=3, and one at d=5.

Tasks 8 and 9, both occurring at d=1, are compared below in Table 1 (actual target locations are shown in Figure 3.16). In Task 8, the target is on an edge, and connected to a coloured corner node (one step away from a commonly identified landmark with a distinctive property). The target in Task 9, in contrast, is in the middle of the homogenous area that was identified as difficult by all participants. It is three steps away from the closest identified distinctive landmark, (a blue pair on the left edge).

Table 1: Comparing Tasks 8 and 9

	Task 8	**Task 9**
Avg. Accuracy	100%	72%
Avg. Completion Time (ms)	1556	7005
Focal Delta (pixels)	500	524
Landmark Distance	0	3

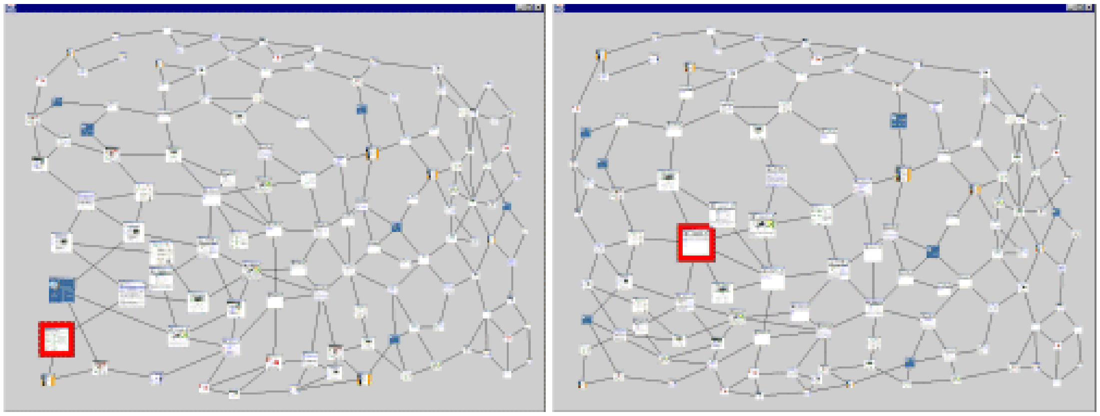

Figure 3.16: Target locations for Tasks 8 (left) and 9. See also colour figure in Appendix A.

Tasks 12 and 15, both occurring at d=3, show a similar difference in Table 2 (target locations shown in Figure 3.17). The Task 12 target is again located in the homogenous area, three steps away from the nearest commonly identified distinctive object (the blue node on the lower left edge), while the Task 15 target is only one step away from the dark blue node on the upper right edge.

Table 2: Comparing Tasks 12 and 15

	Task 12	Task 15
Avg Accuracy	42%	100%
Avg Completion Time (ms)	13,553	4514
Focal Delta (pixels)	512	564
Landmark Distance	3	1

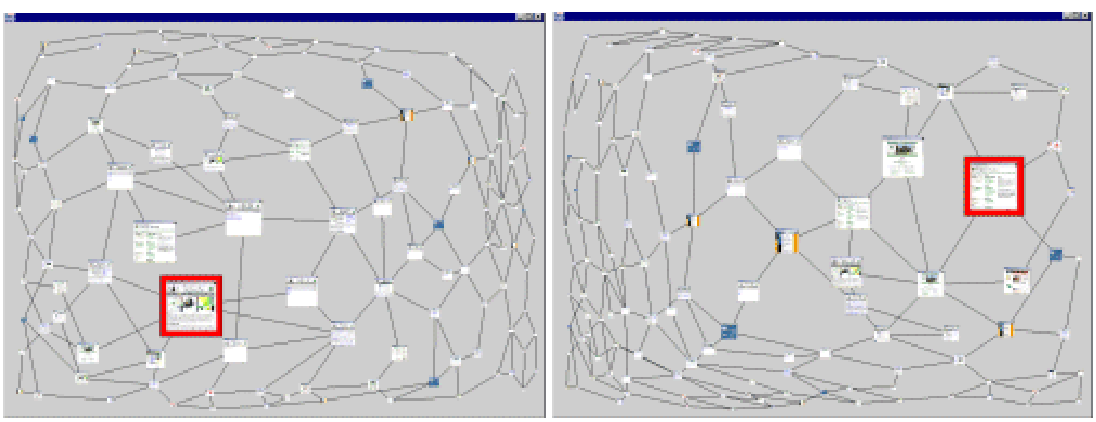

Figure 3.17: Target locations for Tasks 12 (left) and 15. See also colour figure in Appendix A.

Finally, compare Tasks 18 and 19 (at d=5) in Table 3 (target locations shown in Figure 3.18). Again, the target of Task 18 was located in the homogenous area, while that of Task 19 was on the graph edge and next to a very commonly identified topological formation (the “hook” shape).

Table 3: Comparing Tasks 18 and 19

	Task 18	Task 19
Avg Accuracy	57%	100%
Avg Completion Time (ms)	9299	4090
Focal Delta (pixels)	254	216
Landmark Distance	4	0

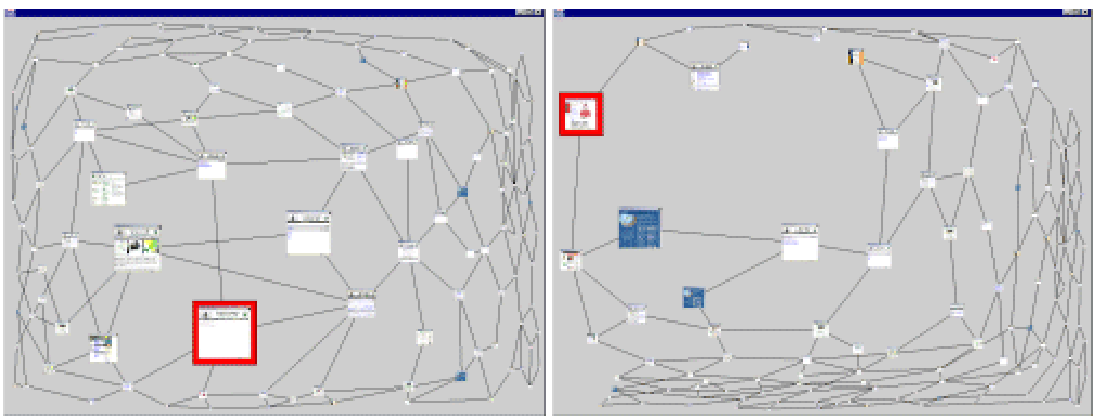

Figure 3.18: Target locations for Tasks 18 (left) and 19. See also colour figure in Appendix A.

Task 18 target was near one potential landmark; the notch in the edge, but only one subject identified that as a landmark. Most instead counted from the left edge, and one subject said they had no strategy at all for that task.

In all of these cases, the accuracy in the memory tasks could be predicted only by considering the focal delta and the distortion level in light of target location in relation to the identified landmarks. Distinctive visual properties aided in remembering the targets.

3.4.4 Discussion

The first study showed that people tend to use visual properties of objects rather than their absolute locations to describe their location in a distorted space. Distortion-based presentations, such as fisheye lenses, affect the absolute location of objects to such an extent that users are encouraged to find additional methods to remember the location of objects. Instead of saying "the top left corner" they will say things like "to the right of the blue one in the hook shape."

The second study confirmed that some visual properties are robust with respect to distortion and aid in remembering where objects are in a space. If an object itself did not have a robust visual property, then a nearby object that did was used as a landmark to remember the target object. It is clear from our studies that success in remembering a target depends considerably on whether the target has, or is near something with, distinctive visual properties. The experimental data space contained many features that could be used as landmarks, and some features were identified by all the subjects as distinctive. Targets that were near these identified landmarks were easier to remember. Distortion level and focal point delta alone did not predict the participants' accuracy in memory tasks.

3.5 Principles of Memorability in a Distorted Space

If objects have visually distinctive properties that are robust with respect to distortion they become easier to remember, and can act as landmarks to re-find other objects that do not have these properties.

The most distinctive properties tested in the study were colour and pattern. Contrasting colours were identified as memorable over any other property by all participants, for as long as colour was visible. Constellations were the next most popular feature since the way that the graph nodes were connected to each other was not affected by the distortion. If three nodes are connected linearly (A to B to C) then distortion will not cause A to be connected to C, and the fisheye will not move the nodes very far away from each other. Similarly, the orthogonal ordering of the objects was not changed much by the distortion; a node on the edge of the graph with no nodes to its right will always have no neighbours to its right.

A table summarizing the tested properties and the effect that distortion had on each one appears below (Table 4). In addition, the effects of using a constrained lens (Figure 2.12) instead of the Sarkar-Brown lens are predicted.

Table 4: Tested visual properties and distortion effects

Property	**Effect of Distortion on Memorability**
Colour/Pattern	Extreme distortion meant that colours were difficult to see and patterns were impossible when objects were not in focus (context features were only a few pixels wide). Until then, this was the preferred property of the users. Even at maximum distortion, any visible colours or patterns were used in preference to any other property. With a constrained fisheye lens, the demagnification of context information is reduced. Colour and pattern in a constrained fisheye would be even more robust than they were with the Sarkar-Brown lens because they would be more visible.
Constellations (Topology and Clustering)	Distinctive shapes formed by screen elements remained recognizable even at maximum distortion. Distortion did not affect connections between nodes and edges. Some users reported that distortion actually let them see the shapes more clearly. With a constrained fisheye lens, the topology and clustering of the context area remains mostly undistorted. Relative distances and angles remain unchanged. If the constrained lens is a truncated lens (Figure 2.13) then the constellations are even less distorted and would be more robust.

Property	Effect of Distortion on Memorability
Orthogonal Ordering	Distortion did not change relative positions of nodes as long as they were at roughly the same magnification. For example, being to the left of a blue node still was a valid mnemonic. Orthogonal ordering was most valuable when nodes had no neighbours in one or more directions; nodes on the border and corners were easy to remember because the border of the graph remained the border throughout all distortion levels. This situation should be the same with a constrained lens, as neighbouring nodes are still at the same level of magnification.
Absolute Position	With the Sarkar-Brown fisheye lens used, increased distortion had a large effect on absolute position of individual screen objects. People did not tend to use absolute position by itself when distortion levels were high, but used it only to establish where the focal point should be to magnify the correct neighbourhood. A constrained lens might make absolute position an even less robust property. With the Sarkar-Brown lens, a screen object is moved in the same direction as its neighbours, just to a greater or lesser degree according to its distance from the focus. With a constrained lens, only the objects under the lens are moved, which may result in a greater relative change in position to the unmagnified objects.
Semantic Content	Like colour and texture, extreme distortion made semantic content almost impossible to see in context objects. At low magnification, users looked for the visual pattern that the node information made rather than at the actual semantic content With a constrained fisheye lens, the demagnification of context information is reduced. Semantic content in a constrained fisheye would be more robust than it was with the Sarkar-Brown lens.

The following table discusses the extrapolated effect that distortion would have on the visual properties that were not tested by this study (Table 5).

Table 5: Untested visual properties and extrapolated effects of distortion

Property	Predicted Effect of Distortion on Memorability
Value	The black value of an object (how dark it is) will be affected by distortion in the same manner as the colour of the object. Users will likely see objects with different values as actually having different colours (dark blue versus medium blue versus light blue). Therefore, like colour, the value will probably be used until the demagnification of the object makes seeing the value difficult. With a constrained fisheye lens, the demagnification of context information is reduced. Black value in a constrained fisheye would be more robust than it would be with the Sarkar-Brown lens.
Size	A fisheye lens alters the size of objects by magnifying or de-magnifying them, but a property must be distinctive to be an effective memorability aid. If the fisheye effect makes too many objects have the same size (some inherently, some because of the fisheye effect) then size ceases to be distinctive. Size is not robust with respect to distortion, either in a Sarkar-Brown or a constrained lens.
Form (Shape)	In the lens used in the studies in this chapter, the magnification of the objects was done orthogonally; i.e. the shape of the objects was not changed by the distortion. Each node remained a rectangle with a constant aspect ratio, no matter its dimensions. In this situation, form could be a robust visual property that would only cease to be useful if the demagnification made it too small to see. In other types of lens, the shape of objects is distorted. This may make it more difficult to determine the true shape of the item, in the same way that it is difficult to tell if a shape is a true circle when one looks at it in perspective. On the other hand, the shapes formed by the topology of the graph remained distinctive even at high levels of distortion. If the form of an object is distinctive enough (a circle in a field of squares) then form could be robust until demagnification of the object makes seeing the form difficult. With a constrained lens, form would be more robust since it is not distorted in the context.

Property	Predicted Effect of Distortion on Memorability
Orientation	It is likely that a distinctive orientation would be a robust property until the demagnification of the object makes seeing the orientation difficult. This property would therefore be more robust in a constrained lens, where since it is not de-magnified as much in the context.
Clustering (without topography)	Clustering and topography were very similar properties in the study since a connected graph was used. In the absence of actual topography (as in a collection of unconnected objects) the constellations used as landmarks would likely be formed by clustering instead, much as literal constellations are made from stars. More experimentation would be needed to see if the implied connections in a cluster constellation would be as robust as the visible connections inside a topologically connected constellation.

The studies performed in this chapter established what users tend to see as distinctive properties of objects in distorted space. However, it may not be possible to manipulate all of these properties in a given information space when adding visit wear. As mentioned in Section 2.9, some of these properties may encode other information that should not be altered. Connecting more edges to a node in a graph may make it more memorable, but it would also change topology and therefore the meaning of the graph. Additionally, the added information should not unduly increase the cognitive load of the user. The next chapter looks at how to choose visual properties for visit wear.

4.0 The Design Space of Visit Wear

4.1 Introduction

The principles of visit wear have been described generally in Section 2.8. Visit wear means that information in a virtual space should automatically undergo a change in visual properties when the user has visited the information. This chapter sets out the distinctions and dimensions in the design space for visit wear. Since the appearance of visit wear effects is not limited by physical reality, we have almost complete freedom in the representation. However, we ought to consider several factors so that historical information becomes an enhancement to the users' task rather than a replacement for it. As Wexleblat (1998) puts it, "If interacting with the history information becomes a central focus, we have probably changed the task too much." (p. 3). We need to consider the principles of visibility, cognition, and realism. If we do not consider what people consider natural or obvious from their experience with the rest of the world, the representation may not convey what we are intending, and if we do not consider the principles of visibility, the representation may not be noticed at all. Issues such as how much history to show, and how long the history remains shown, are important because too much historical information can clutter the representation rather than enhance it.

The design space of visit wear is made of several independent parameters, as described below in Table 6.

Table 6: Visit wear design space parameters

Parameter	Questions
Data space type	Does the data consist of discrete features separated by an information-less background (like a graph), or is it a continuous spread of information (like a map)? Will the visit wear be shown as part of the features of the data or as part of the environment?
Duration	What time span should the visit wear cover? Should the representation change with time? How much visit wear is possible before the space becomes cluttered?
Appearance	Should the visit wear be represented as a change in the data or by an additional glyph? What visual properties should the visit wear possess to balance distinctiveness with distraction?

In addition, visit wear can be used to represent other information, such as who the visitor was in a multiple-user environment. There is also the possibility of making visit wear manual instead of automatic, in other words, having the visual changes only appear on the user's command (like bookmarking a Web page). This chapter concentrates on the parameters in Table 6, since they will be common to all visit wear designs.

4.2 Data Type

If data is sparse, such as in a graph, the areas of information are usually separated by a background which does not contain information. This can be called a *discrete* space. In a discrete space, a user visiting the background (which contains no information) is not as important as a user visiting a data item. It is therefore natural that only the data items are affected by visit wear. Figure 4.1 shows an example of visit wear in a discrete space. The questions are how many items to affect (i.e. how long a history list to maintain) and how to differentiate recent visits from older visits.

These are discussed in later in this chapter, as is the question of exactly how to draw the visit wear.

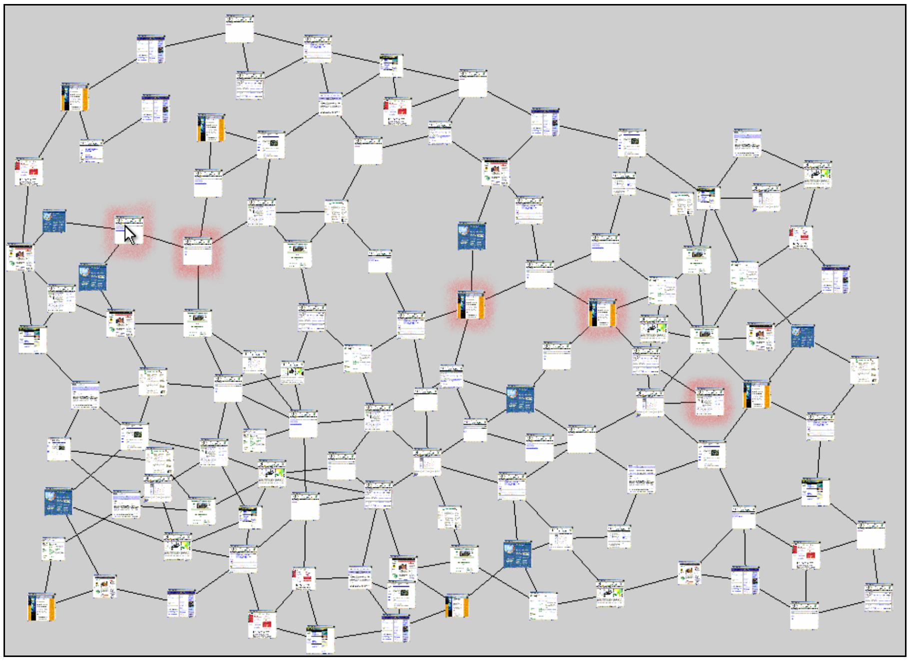

Figure 4.1: Feature-based visit wear in a discrete space. The last five visited nodes are marked with a red halo, though these specifics ("five" and "red halo") are arbitrary. See also colour figure in Appendix A.

However, the visit wear can also be implemented as a change in the data space itself rather than as a change in the features. In this case, a visit to any part of the space is recorded, whether the part is a feature or not. This results in a trail formed of recorded mouse locations, as shown in Figure 4.2. The same interaction with the data space occurred, but there is no distinction made between visiting a feature and visiting the background.

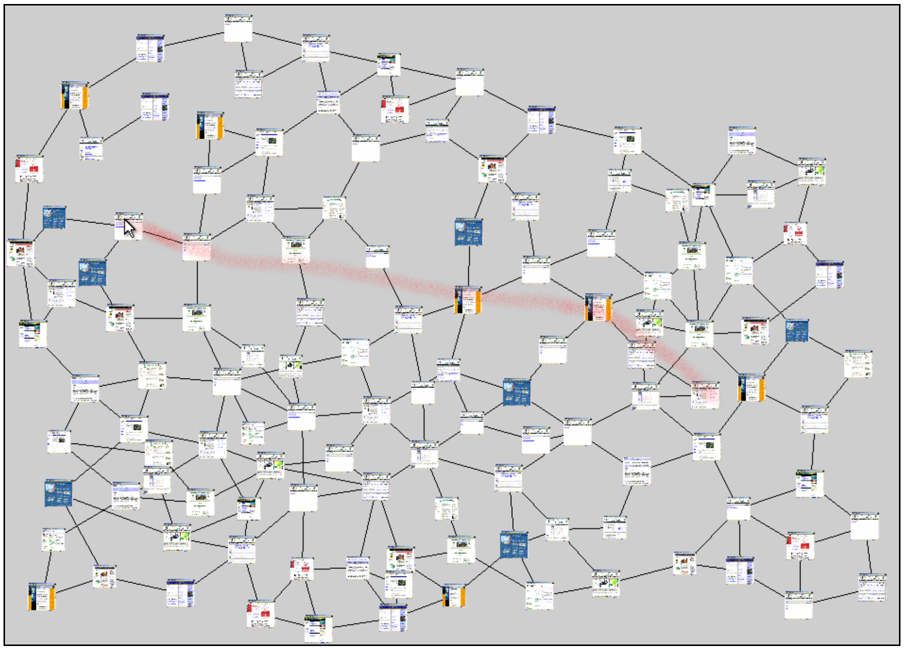

Figure 4.2: Environment based visit wear in a discrete space. The location of the mouse cursor in the recent past is shown, whether it was over a feature or not. The trail is shown in red, but this choice and the choice of the length of the trail are arbitrary. See also colour figure in Appendix A.

In a space where the data is dense and continuous, such as a map, defining features is more difficult since the areas of interesting information may depend on the task. In a detailed map, for example, the areas of interest might be different for a geologist (rock outcroppings), a civil engineer (roads and bridges) and a hydrologist (lakes and rivers). Defining features as areas of the space that will display visit wear must be done *a priori*, and might not match the user's idea of a feature. Figure 4.3 shows a map where "features" has been taken to mean "campsites." This might be acceptable for campers, but a park employee might also be interested in visiting the water towers and bridges, for example.

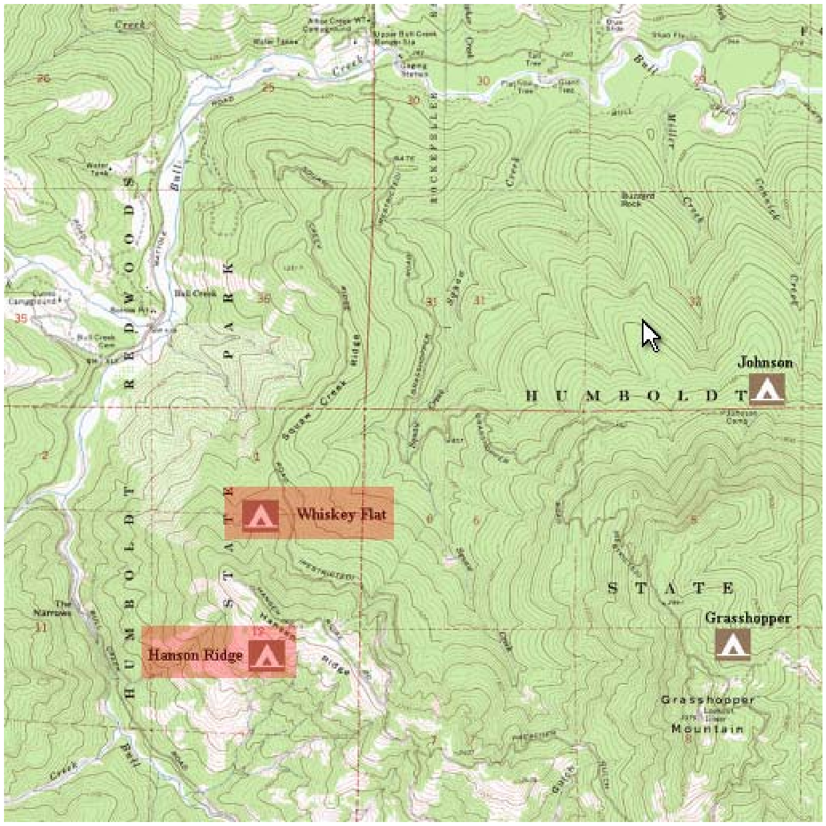

Figure 4.3: Feature-based visit wear in a continuous space. The features have been predefined to be the campsites; no other area of the space will show visit wear. See also colour figure in Appendix A.

In a continuous space, the fact that every pixel contains information means that potentially everything is a feature. In this rich environment, the path between two points might be just as important as the points themselves. Showing the visit wear as change in the data space itself rather than only in parts of the space means that none of the richness is lost. Figure 4.4 shows the same interaction as Figure 4.3, but because the interaction is displayed as a change in the environment rather than as a change in a predefined feature, we can see the parts of the interaction (visiting the top of the map) that was lost in Figure 4.3.

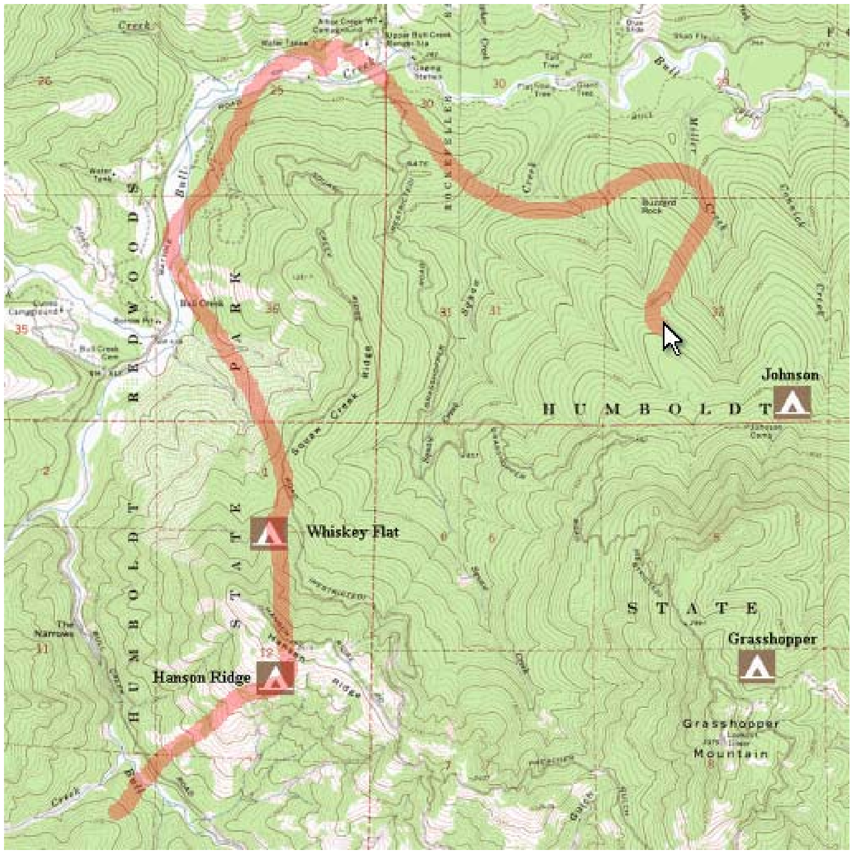

Figure 4.4: Environment based visit wear in a continuous space. No matter what the user considers important in the space, the visiting of it has been shown. See also colour figure in Appendix A.

4.3 Duration

Because visit wear is shown as a change in the data on the screen, it cannot accumulate without limit as it can in a pull-down menu or a log file. If the history is not limited, then any lengthy interaction with a space could result in almost everything being marked as visited. The history can be limited either by size (the number of pixels on the screen that are affected by any visit wear) or by time (having the visit wear effect disappear after a certain period of time) or both.

The advantage of having the visit wear change with time is that it gives a sense of sequence to the representation. Older items will be marked differently from the newer items. This makes it easier to backtrack along the visiting history, and the visit wear need not disappear completely if persistence is desired.

One way to calculate the length of the history list in feature-based visit wear can come from the revisitation research by Tauscher and Greenberg described in Section 2.7 (Tauscher and Greenberg, 1997). A visit wear application could keep a running total of all re-visits and their position in the current set of visited features. If the application wants to be able to satisfy 50% of a user's revisitation requests, for example, then the set of visited features is marked with visit wear, starting from the most recent, until the visit wear set includes 50% of the recurring features. For environment-based visit wear, a similar calculation could be done but with revisiting sections of the space instead of features. Figure 4.5 shows this concept. The sequence of visits to the nodes is B-C-D-B-A-D-E. Keeping track of revisits, we find that there are two; to B and to D. Therefore, if we want to include 100% of the revisits, we need to make sure B and D are in the set of nodes marked with visit wear. Marking the last four nodes is the minimum history size that will achieve this.

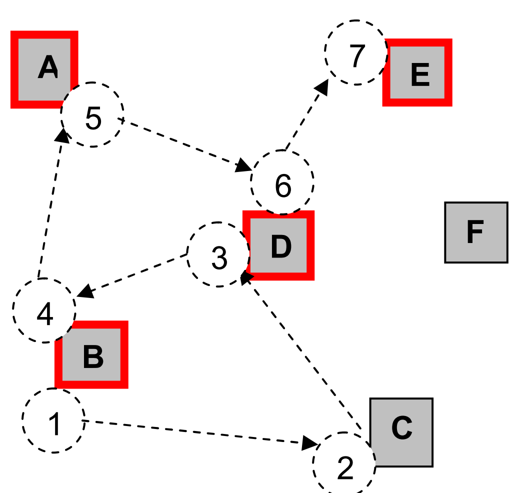

Figure 4.5: An example of automatically calculating visit wear duration. See also colour figure in Appendix A.

In Tauscher and Greenberg's research, they found that most of the revisits occurred within a certain number of the last places visited. The graph of the recurrence pattern (Figure 2.17) showed a "knee" that meant that there were diminishing returns to making the history list longer. The last 30 items did not include significantly more revisits than the last fifteen items. The list of features marked with visit wear should have a maximum limit to avoid clutter in the display.

4.4 Appearance

There are three aspects of visit wear that we may want to convey through appearance. These are 1) whether a screen item has been visited or not, 2) The age of the visit and 3) how often it has been visited. Each of these aspects will be discussed below, with some possible techniques to represent them.

As mentioned in Section 2.9, an item can have a visual property that lets a user pick it out immediately from other objects that do not have this property. Bertin defined these properties as form (shape), orientation, colour, texture, black value, size and position (Bertin 1967). In addition, there are properties that cannot easily be shown on paper, such as motion. If the visit wear representation has some of these characteristics, it will be distinguishable to the user with a minimum amount of attention.

Visit wear can be implemented by either changing the appearance of a data item or by adding a secondary glyph. Changing the appearance of the data item requires care so as to not change the meaning of the data. This can be the case if information is already encoded visually. Maps and scale drawings should not have their data's position changed, for example, and neither should the colour of text that indicates markup information in a word processing document. An application that moves information around so that recently visited information is always at the top of the screen is possible but research into the mental map model (as described in Section 2.5) indicates that users could have trouble remembering the information space.

4.4.1 Visitation

Representing visitation means showing whether an item is in the current history list or not. For feature-based visit wear, an item would be a feature. For environment-based visit wear, the item would be a screen area, possibly as small as a pixel.

Visitation can be shown either by changing the appearance of the data itself or by adding a secondary glyph. Changing the appearance of the data itself could mean changing its size, colour or black value, since other properties are more likely to change the data's meaning (like position) or be visually confusing (like texture).

Changing the size of visited items would create an effect like a series of miniature fisheye lenses that remain stationary on the screen after the mouse has passed, even when there is no fisheye effect on the space. Figure 4.6 shows a possible appearance of this effect.

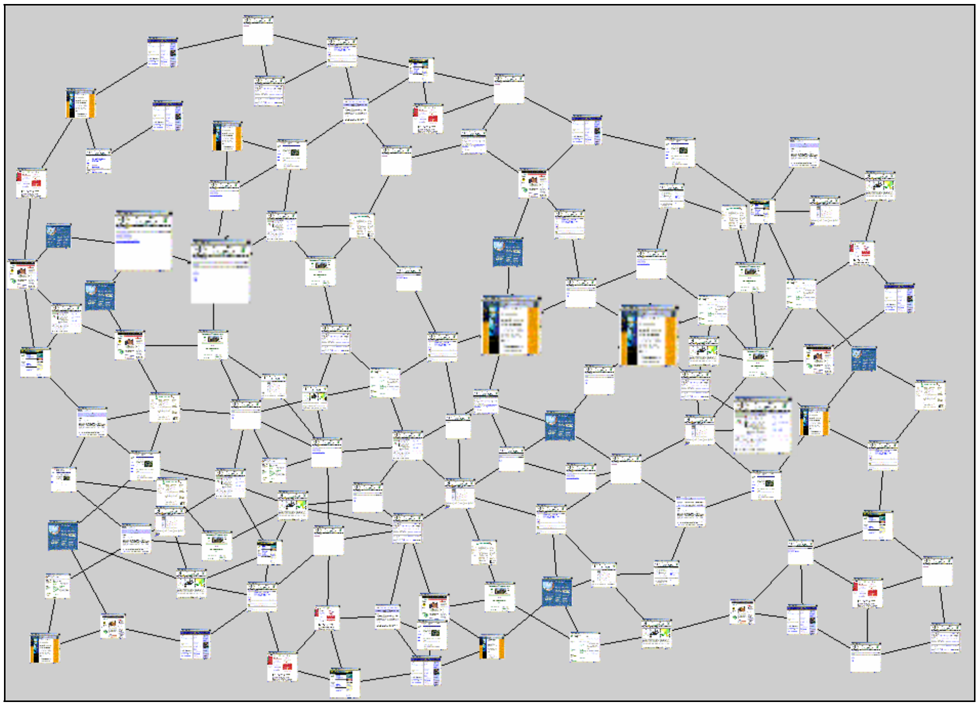

Figure 4.6: Representing visit wear by changing the size of a visited node. See also colour figure in Appendix A.

Changing the colour or black value of the data itself still has the potential to change the meaning of the data. Figure 4.7 shows what changing these properties can look like, using a node from the graph in Figure 4.1 as an example. Note, however, that a user looking for a “white, orange and blue page” might not recognize the versions in B or C as the target since the colours of the page have been changed.

Figure 4.7: Changing the properties of the data to indicate visitation. A) The original data B) Changing colour C) Changing black value. See also colour figure in Appendix A.

A safer way to add visitation information to the data is to add a secondary glyph. A border around the visited item could have its properties change instead of the data itself, for example. Figure 4.8 shows some different ways to add a secondary glyph, using a node from the graph in Figure 4.1 as an example.

Figure 4.8: Adding a secondary glyph and changing its properties to indicate visitation. A), the original node B) Adding a border and altering its colour C) Adding a border and altering its texture D) Adding a secondary text glyph. See also colour figure in Appendix A.

A concern when adding secondary glyphs is that there is a limited amount of screen space available for any representation. A data space might already have indicators for such things as “marked for deletion” or showing an intentional bookmark set by a user.

Most of these effects can be applied to environment-based visit wear too, as shown in Figure 4.9. The secondary glyph added in environment-based visit wear is the trail itself

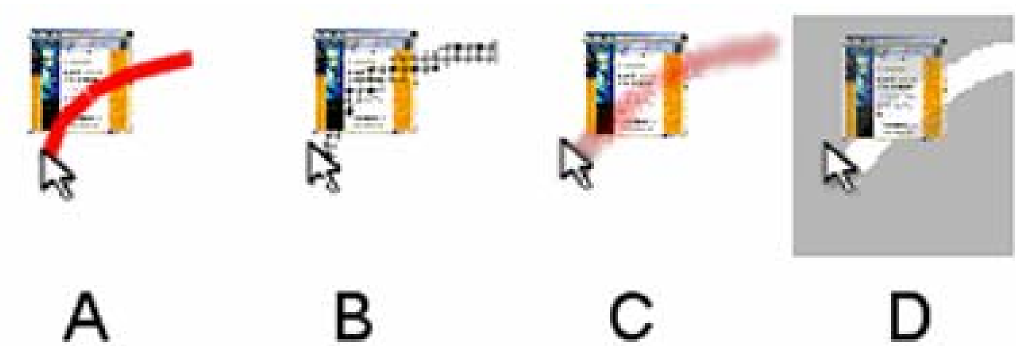

Figure 4.9: Methods of representing visitation by changing A) the colour B), texture C) colour and texture, and D) black value of the mouse trail. See also colour figure in Appendix A.

Choosing a specific colour or texture to represent visitation is a challenge, since the best colour depends on the colour of the data. The red used in Figure 4.7

and Figure 4.8 would not be effective if it did not stand out against the white background and the colours of the data. On the other hand, the red might be *too* effective; it might actually distract the user from the data rather than provide an enhancement. Similarly, the texture trail stands out against a plain white background, but might get lost when it is superimposed over a detailed topographical map. The ideal effect is almost subliminal. Customisability of the visit wear effect would be one solution to make sure the visit wear remains an enhancement rather than a distraction.

Another characteristic that is possible to change in virtual spaces might be motion. In the real world, items often are left moving slightly after we've used them. A swing sways back and forth after the child using it has gone, and curtains move slightly after they've been closed. Having the viewed items "vibrate" slightly might logically indicate that they have just been used. The trouble is that motion probably draws our eye too successfully; that is, "Wiggly stuff is distracting" (Human Factors International promotional button). The effect would have to be very subtly done to avoid distraction.

4.4.2 Age of Visit

With visit wear, there will be a list of the items or areas that have been visited. In one extreme, the history list could grow infinitely, and in the other extreme, the list could always be only the last few nodes visited. Between these extremes, Section 4.3 discussed how the length of the visit wear history list could be calculated from the revisitation pattern of the user. In all cases, however, having the visit wear representation change with time will keep the sense of the order in which the information was visited.

If visitation is indicated with a colour change (either of the information or through a coloured secondary glyph) then the visit wear effect could increase its transparency over time. If the history list is not infinite, then the visit wear would become completely transparent (i.e. disappear) when the item was no longer on the history list. Figure 4.10 shows an example of this effect.

Figure 4.10: Visit wear transparency indicates age of visit. Full opacity (far left) indicates a recent visit, full transparency (far right) means the node either has not been visited or was visited long enough ago that it is not on the history list. See also colour figure in Appendix A.

If visitation is indicated though a secondary glyph, then the size of the glyph could shrink over time. Again, if the history list is not infinite then the glyph would completely disappear when the item was no longer on the history list. Figure 4.11 shows an example of this effect.

Figure 4.11: Size of visit wear secondary glyph indicates age of visit. Full size (to the far left) indicates a recent visit, zero size (far right) means the node either has not been visited or was visited long enough ago that it is not on the history list. See also colour figure in Appendix A.

If the history list is infinite, then items that have been visited will always show an indication of the visitation. In this case, a transition between two colours could be used to indicate the age, since visited nodes will always have some sort of marking and the visit wear effect will never fade completely. Using a Web standard as an example, Figure 4.12 shows an example of this effect. Recent visits are shown in purple, which changes to blue as the visit gets older.

Figure 4.12: Visit wear colour indicates age of visit. Purple (far left) indicates a recent visit, blue (far right) indicates long ago visit. See also colour figure in Appendix A.

For a more abstract representation, another glyph could be added to indicate the age of the visit. A glyph resembling stopwatch hands, for example, could move to indicate the approximate time since the last visit (Figure 4.13). The speed of the "hands" would depend on the length of the history list.

Figure 4.13: Additional glyph indicates age of visit. Glyph moves like clock hands, starting at "12:00" when t = 0 (to the far left) and moving clockwise with age. See also colour figure in Appendix A.

The intention here is to keep the visit wear representation visual rather than cognitive; adding a glyph that said "2 minutes 34 seconds" would require too much cognitive interpretation. With any age representation, too much information should not be encoded, or the memory burden for the user will become too great.

4.4.3 Frequency

Another variable of visit wear that could be represented visually is the frequency of visits. Information that is revisited often is probably of more importance to the user than information that is revisited only once.

A change in a property of the secondary glyph could be used, like the techniques to represent the age of the visit described in Section 4.4.2. Transparency, size and colour could all be changed, as long as the same property was not used for age. A border could become more transparent with age, for example, but thicker with the number of visits. Alternatively, a secondary glyph could be repeated to indicate the number of visits (Figure 4.14).

Figure 4.14: Change in secondary glyph indicates number of visits (shown below the example node). One visit is shown with a one-pixel wide border, or with one extra glyph. Ten visits are shown with a ten pixel wide border or ten extra glyphs. See also colour figure in Appendix A.

Combining frequency representation with age representation may increase the user's memory burden unacceptably. Trying to remember whether thickness of a border is proportional to number of visits or age of visits might take too much cognitive effort and distract the user from the real tasks. The frequency information should be important to the user and the application before it is represented at all.

4.5 Visit Wear in a Fisheye Environment

The design parameters of visit wear discussed in this chapter were the data space, the duration, and the appearance of the visit wear. When the visit wear information is presented through a fisheye lens, an additional parameter – robustness with respect to the fisheye distortion – is also added to the design space.

The data space, whether the information is sparse or continuous, is affected in the same way by the fisheye view. The question of the length of the history list and the fade time are also independent of the distortion effect. The appearance of the visit wear is the main aspect of the design space that should be reassessed with respect to robustness.

Chapter 3 investigated visual properties that were relatively unaffected by distortion. The experiment showed that colour and pattern (texture) are very distinctive properties, even at high levels of distortion. It was extrapolated that value, being a similar property to colour, would also be robust with respect to distortion.

Some other robust properties from Chapter 3 were the constellations formed by the information and the orthogonal ordering of the information. However, visit wear would require changing these properties, to distinguish visited information from non-visited information. As discussed in Section 4.4, changing position or order of the information could alter the information or disrupt the users' mental map, which is already being affected by the fisheye lens.

Changing the size of the information itself, as discussed in Section 4.4.1, might not combine well with the magnifying effect of the fisheye lens. Changing the size of the secondary glyph (as discussed in Sections 4.4.2 and 4.4.3 to indicate age or frequency of visits) might also be difficult to distinguish when parts of the space are magnified and parts are not.

Table 7 shows a summary of the results of Chapters 3 and 4, with a list of visual properties, their robustness with respect to distortion, and whether they will be suitable to represent visit wear. Texture (pattern), colour, black value and orientation are the best candidates for representing visit wear in a fisheye lens. Form might be suitable, but as discussed in Table 5 it depends on the type of fisheye lens.

Table 7: The intersection of robustness with respect to distortion and suitability for representing visit wear

Property	Robust with Respect to Distortion	Suitable for Representing Visit Wear
Colour/Pattern	Yes	Yes
Constellations	Yes	No
Orthogonal Ordering	Yes (edges)	No
Absolute Position	No	No
Semantic Content	Yes	No
Black Value (Darkness)	Yes	Yes
Size	No	Yes
Form (Shape)	Possibly	Yes (added glyph)
Orientation	Yes	Yes (added glyph)
Clustering	Unknown	No

4.6 Choosing Properties to Test

As Section 4.5 concluded, texture, colour, value and orientation are the best properties for representing visit wear in a fisheye lens. In order to keep the data itself as unaffected as possible, a secondary glyph that represents visit wear should be added, and have its texture, colour, value or orientation changed.

The effects of visit wear on memorability in a space distorted by a fisheye lens have not been tested before, and the design space of visit wear has not been previously defined. The experiment described in the next chapter, therefore, isolated the visit wear effects and tested them to determine if visit wear has any effect on memorability under the most basic conditions.

This initial experiment will choose one visual characteristic to represent visit wear, and test this visual characteristic both in a space with sparse information (a discrete space) and a space with dense information (a continuous space). In both cases, a secondary glyph will be added to represent the visit wear; a halo around individual items in the discrete space and a trail behind the mouse in the continuous space.

The discrete space was chosen to be the same graph used in the experiment described in Chapter 3, since that study has already identified the landmarks and memorable properties of the graph. As was discussed, node colour and pattern were distinctive properties even at high levels of distortion. Therefore to isolate visit wear as a memorability factor, the contents of the graph nodes were removed and all nodes were coloured yellow.

The continuous space was a detailed greyscale satellite map of the Halifax harbour area, chosen because of the density of the information. Because the map is a continuous space (every pixel contains information) it is more difficult to predict what parts of the space the user will be inspecting. With the graph, we can safely say that the user will be interested in the contents of the nodes. With the continuous map, we cannot necessarily say whether the user will be interested in the lakes, the roads, areas above a certain elevation, etc. For this space, therefore, the visit wear is a continuous trail that does not depend on the data being inspected.

A secondary glyph will represent the visit wear. The visual characteristic that changes on this glyph will be colour. Pilot tests showed that changing the black value (darkness) of the mouse trail over the greyscale map was not visible enough because of the varied darkness values of the information itself, and texture would likely have had the same problem. A semitransparent green was chosen after piloting as being sufficiently distinctive (Figure 4.15). Changing the glyph's orientation is not easily applicable to the continuous visit wear representation.

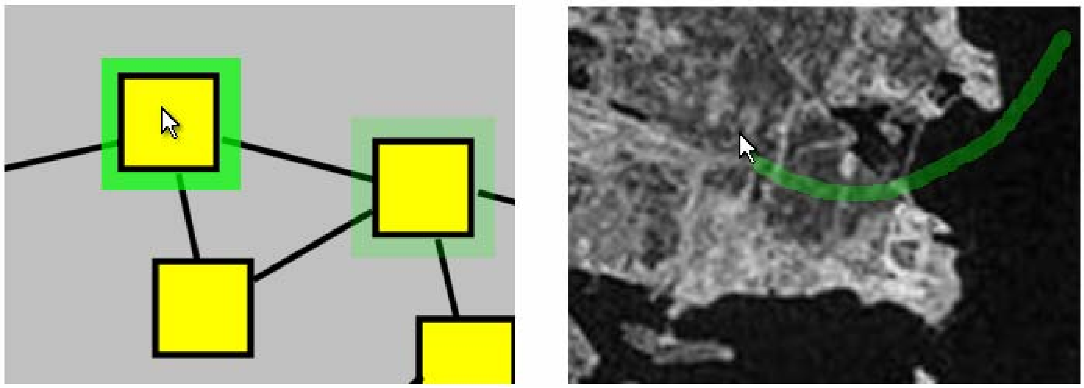

Figure 4.15: Example visit wear effects to be experimentally tested in a discrete space (left) and a continuous space (right). See also colour figure in Appendix A.

The duration of the visit wear will be a simple timer instead of the revisitation analysis technique described in Section 4.3. The visit wear will start increasing in transparency as soon as it is drawn, and after one minute, will have attained complete transparency (i.e. disappeared). This is because the experiment will be simple memory tasks to test memorability, and revisitation analysis requires a longer and more realistic interaction with the data set.

Frequency of visits was not tested at this time, since the main hypothesis of the thesis is that visit wear improves memorability. Frequency of visits, while probably useful in a real world application, is not directly related to memorability.

The next chapter describes the experiment that tests the described visit wear implementation.

5.0 Experiment and Evaluation

Previous chapters have discussed the principles behind visit wear and fisheye views, and the design considerations to be taken into account when implementing visit wear in a fisheye view. The experiment described in this chapter uses these principles and considerations to investigate whether adding visit wear really does allow people to complete memory tasks in a fisheye view more quickly and more accurately than when visit wear is not present. Because adding visual information to a space can cause more problems than it solves, an important secondary question is whether the added information affected the users' perception of the difficulty of the task through clutter or occlusion.

5.1 Method

The study used two types of information spaces, a discrete space (in which the data are sparse and separated by a background of non-information) and a continuous space (each pixel of the space representing information). As discussed in Section 4.1, different types of visit wear are suitable for each of these types of space. Therefore, the study used a feature-based visit wear implementation for the discrete space and a trail-based implementation for the continuous space. Section 4.6 discussed the selection of visit wear design parameters for the experiment. There was one memory task (slightly different for each space), repeated with and without visit wear.

5.1.1 Experimental Apparatus

The application was built in C++ specifically for the experiment, using the EPS library (Carpendale, 1999) and OpenGL to implement the fisheye lens and the visit wear. The system on which the experiment ran was a Dell Dimension 8300 with a Pentium 4 processor at 2.66 GHz and 512 MB RAM. The screen was an 800x600 LCD monitor. The experiment monitor was set up as the secondary monitor on the experiment computer, so that the experiment window was the only thing visible on the monitor shown to the participants.

The application was a single window that texture-maps any bitmap image onto a planar surface defined in OpenGL. The EPS library adds the fisheye distortion effect to the planar surface by remapping the points of the surface in three-dimensional space and then projecting the points back onto the plane. Since the image is texture-mapped onto the distorted points, the image is also distorted.

To add the feature-based visit wear, the area of each feature was defined in a data file, and the application determined whether or not the mouse cursor was currently over a feature by comparing the current mouse position to the feature data file information. If it was, the application used the defined area of the feature to alter the RGB values of the pixels in the texture image that bordered the feature. By altering the texture image itself, the visit wear was also affected by the distortion of the fisheye lens. The application kept a list of all features currently affected by visit wear so that the visit wear could be faded correctly.

The trail-based visit wear was also added directly to the texture map by altering the RGB values of the affected pixels. Every time the mouse position changed, the system drew a line between it and the previous recorded mouse location using Bresenham's algorithm (Bresenham, 1965). The system kept all points in the trail in a list so that the trail could be faded appropriately.

5.1.2 Experimental Conditions

The appearance of the visit wear was decided as detailed in Section 4.6. If visit wear was enabled and the data space was defined as discrete, the system expected a list of the coordinates of all the features in the space. Given these coordinates, a green halo deepened in intensity around a feature when the mouse was over it, reaching full intensity in three seconds. The halo remained at full intensity for as long as the mouse was over the feature. When the mouse moved off the feature, the halo started fading. After one minute, the halo was completely faded (Figure 5.1).

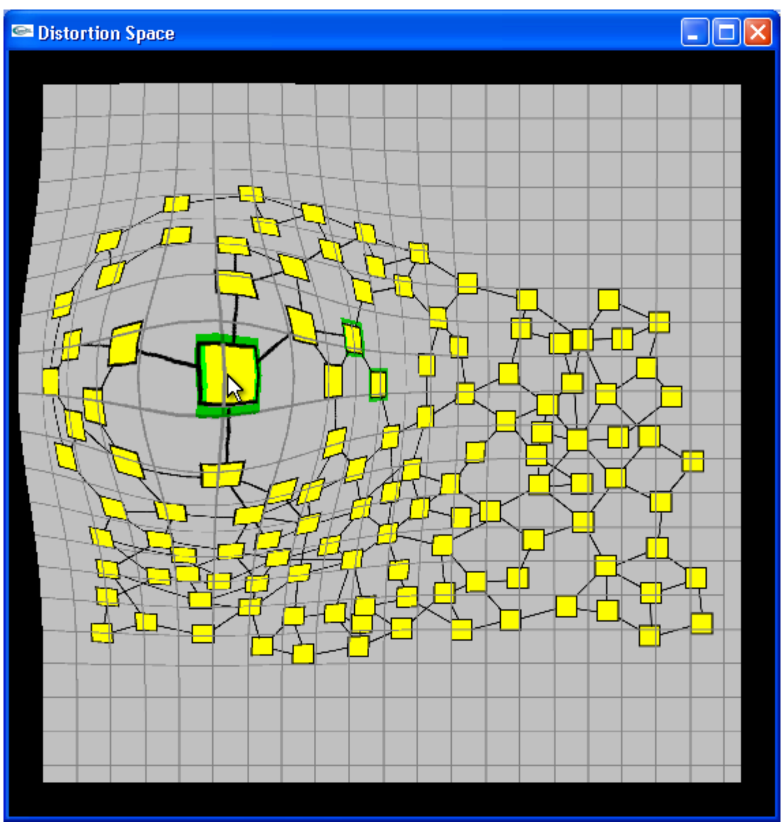

Figure 5.1: The experimental discrete data space, the fisheye lens, and the visit wear effect used in the experiment. The grid is shown to make the fisheye effect clear; it did not appear in the actual experiment. See also colour figure in Appendix A.

If visit wear was enabled and the data space was defined as continuous, then a green trail followed the mouse as it moved, fading over time (Figure 5.2). One minute after the mouse trail was drawn, it had completely disappeared. Figure 5.2 also shows the magenta target (see Section 5.1.3) in the center of the window, at the focal point of the fisheye lens.

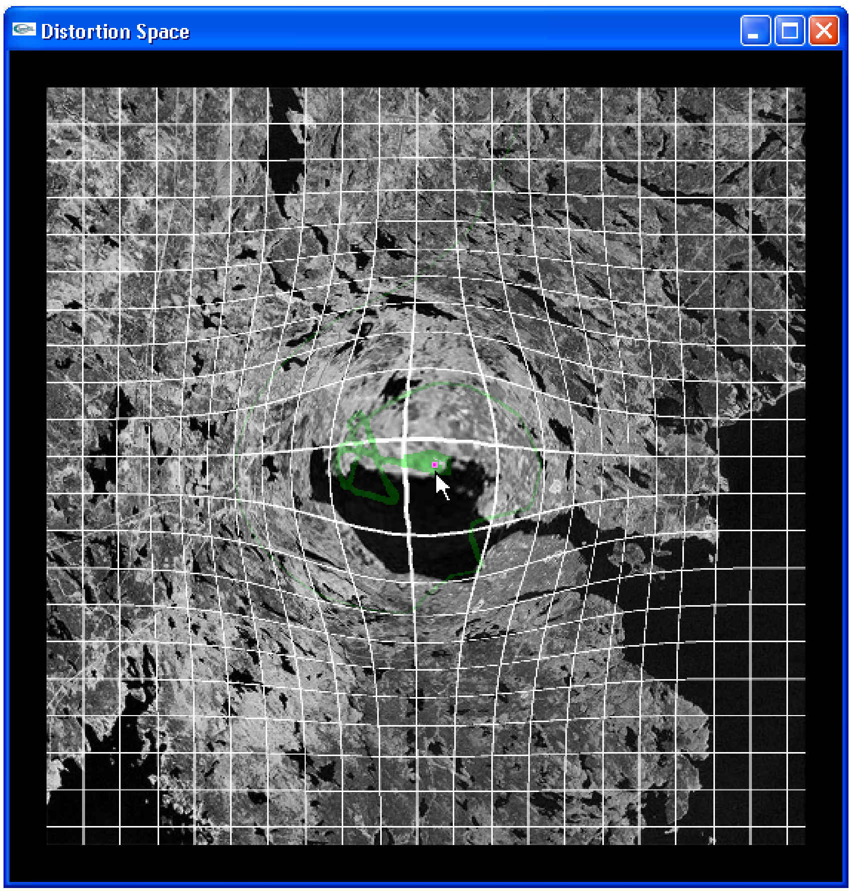

Figure 5.2: The experimental continuous data space, the fisheye lens, and visit wear effect used in the experiment. The grid is shown to make the fisheye effect clear; it did not appear in the actual experiment. See also colour figure in Appendix A.

In all cases, the visit wear was affected by the distortion of the fisheye lens. The halos around the features enlarged as the feature became magnified, and the mouse trail appeared larger as it entered the lens.

5.1.3 Tasks

The tasks chosen represent abstract versions of real-world tasks. Because subjective association can be an important factor in memory (Siegal and White, 1975), the abstraction was used to isolate visit wear as a factor.

For the discrete space, six target labels (“1” to “6”) appeared sequentially on certain nodes in the graph. The next target appeared when the subject had spent enough time on the current target for its visit wear to reach maximum visibility (whether visit wear was visible or not). When all six labels were visible, all the labels disappeared and the subject was asked to click on the nodes that had been labelled, in the same order. This same task was repeated for six sets of six nodes. The node sets were selected to all be approximately the same total path length. However, the node sets did differ in how easy they were to remember. Each node was rated as 1 (easy) 2 (medium) or 3 (difficult), indicating the memorability of the node as judged by subjects in the study in Chapter 3. An example node set is given in Figure 5.3. In this example, the node difficulty rating is 1-1-1-2-1. The next node in the sequence was labelled after the subject had spent a few seconds on the current node (to represent examination time), and node labels were not erased until the last node had been examined.

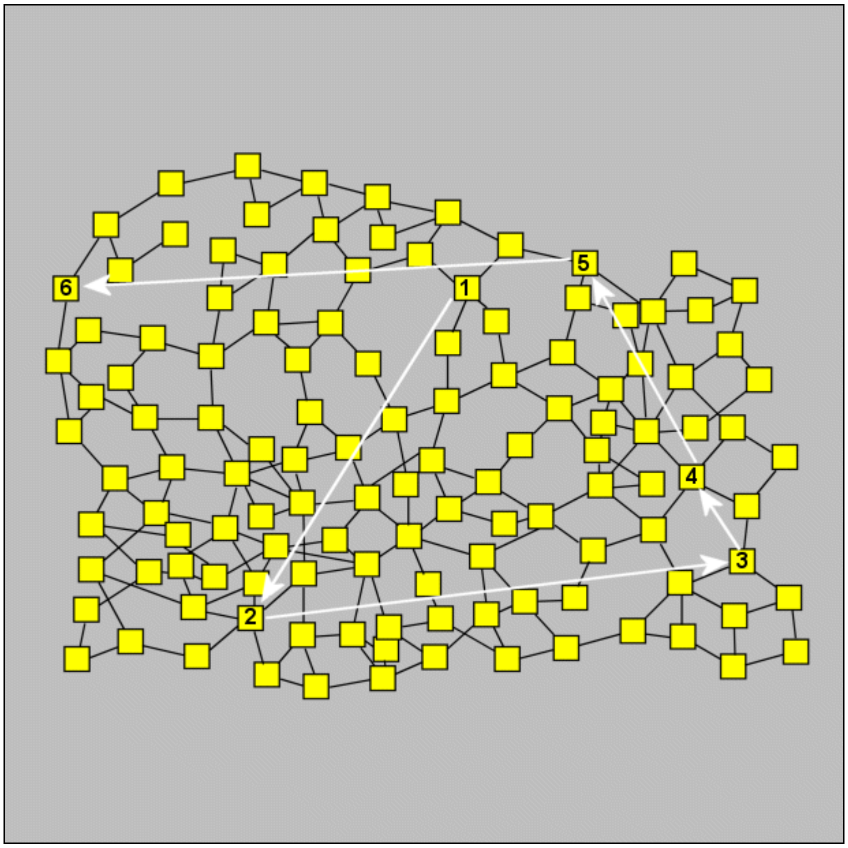

Figure 5.3: Example task sequence, discrete space. White arrows show the subject's expected path. See also colour figure in Appendix A.

For the continuous space a red 50-pixel square was shown. As soon as the subject's mouse cursor was over this red square, it disappeared and was redrawn in another location. The subject would move the mouse to the new location, and the red square would again be redrawn in a new location. The subject therefore drew a path roughly of a given shape, although the route that the subject's cursor took to get from one location to the next was not guided or constrained. After the sixth red square was reached and had disappeared, the subjects were told that to imagine that they had dropped their keys somewhere along the path that they had just traveled. The "keys" were represented by a 2-pixel square magenta target, which was small enough to be only visible with the magnification of the fisheye. The exact position of the "keys" was not calculated until the subject had visited all six targets. At that time, the magenta target was drawn halfway along the path between the first red square position and the sixth red square position. The point was halfway so as not to favour starting from the beginning or the end of the trail.

This same task was repeated for six sets of six red square positions. The positions were selected to give path lengths of roughly the same size (approximately 1240 screen pixels assuming a straight line from target to target). An example set of red square positions is given in Figure 5.4; however, note that in the actual trials, only one red square was visible at a time.

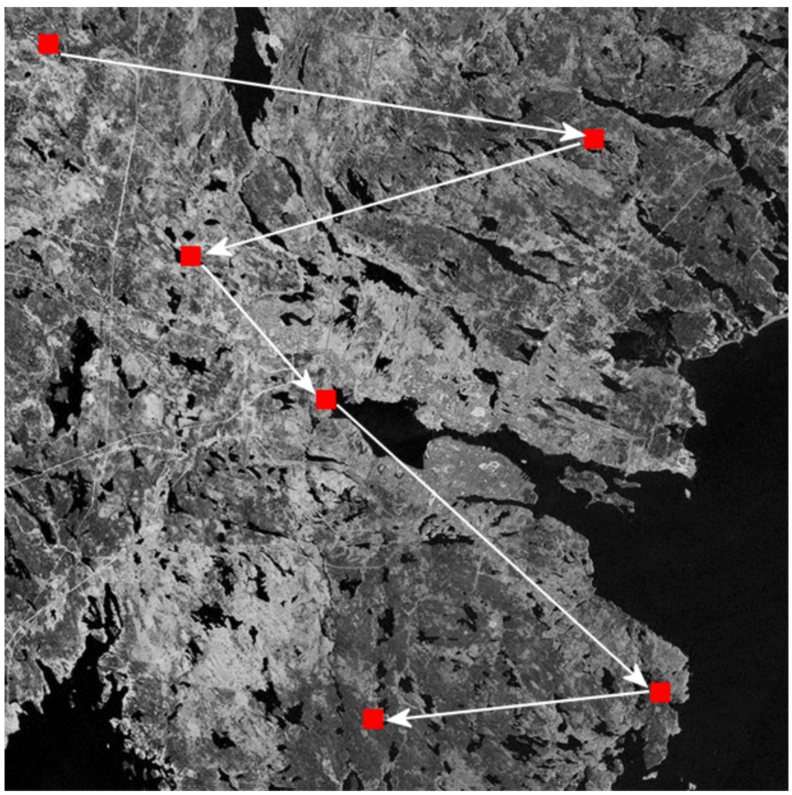

Figure 5.4: Example task for continuous space. Red squares appeared one at a time; white arrows show subject's expected path between squares. See also colour figure in Appendix A.

The system recorded all mouse movement and mouse button activity during the tasks, with time stamps. All tasks had a 40 second time limit; after the time had expired the screen went blank and the system was ready to start the next task. When the screen was blank the system was not recording data, so the subjects could take as long a break as they wanted between tasks.

5.1.4 Study Design

The study used one main independent variable (visit wear presence) in a within-subjects design. All subjects did all tasks under all conditions. The two sets of six tasks were presented in the same sequence to each subject, but each subject's starting point within the sequence was randomly determined and balanced so that all conditions were seen in each position in the sequence the same number of times.

Another independent variable, spatial ability of the subjects, was tested using an object location memory test developed by Eals and Silverman (1992). This "E+S" test was designed to measure differences in location memory. The test consists of an array of twenty-seven "familiar" objects (a cup, a paintbrush, a suitcase, etc.) and a second array where some of the objects are switched from their original positions. The participant identified which objects have moved and which have not. Scoring was done by counting how many objects the participant has correctly identified as having moved or not moved; the maximum score is therefore twenty-seven. The test sheet is reproduced in Appendix B.

For the discrete space, two dependent variables (completion time and accuracy) were measured to test the hypothesis that adding visit wear will improve memorability in a fisheye view. For the continuous space, three dependent variables were measured within each cell (completion time, accuracy, and distance traveled during the search). These variables are discussed in Section 5.1.5. The starting conditions (with or without visit wear, discrete or continuous space) were balanced.

5.1.5 Measures of Memorability

The definition of memorability in this thesis is the ability to locate a previously inspected area of a space. It is impossible to observe someone's thought processes directly but we can observe how people act based on their memories, and that is how the variables used in the experiment were chosen.

1. *Completion time* assumes that if someone has accurately remembered a target, they will identify the target more quickly than if they did not remember the target.
2. *Accuracy* assumes that if someone has remembered a target, they will be able to identify it from among other candidates. If they did not remember it, they will make more mistakes in identifying it.
3. *Distance traveled* was used only in the continuous space trials, to indicate the search strategy used when the subjects were trying to find the target. If the subjects did not remember the target's location, they could try a "sweep" strategy and hope to find the target by chance. This sweeping would likely result in more cursor travel than if the target's location was remembered. Measuring the distance traveled shows when the subjects were using a sweep strategy instead of memory.

The participants were also asked if they preferred visit wear or no visit wear in the trials, and if they had any other comments about the implementation of the visit wear. These questions were intended to determine if the visit wear caused clutter and occlusion, since those can be subjective. Participants were also asked what strategies they used to remember things when there was no visit wear. This was to identify any other memorability cues apart from visit wear that might have affected the experiment.

5.1.6 Participants

The experiment tested sixteen participants recruited from the student and staff community at the University of Saskatchewan. Participants were paid $10 for taking part in the study. Criteria for inclusion were that the participants be experienced with mouse-and-window based applications and have normal (or corrected-to-normal) vision.

There were 8 men and 8 women in the study, ranging in age from 21 to 45 years and averaging 27 years. All worked with computers for over 12 hours per week, and most had experience with computer games (average of 4.5 hours/week) and some experience with graphical programs (average of less than 1 hr/week). The game and graphics programs experience was asked to determine familiarity with targeting and clicking actions. Only two participants had any experience with fisheye systems, and those were in previous departmental experiments rather than commercial applications.

5.1.7 Procedure

The experiment was carried out in several stages. Participants were first asked to sign consent forms and fill out a short demographic questionnaire, included in Appendix B. Then they were asked to complete the Eals and Silverman spatial memory test described in Section 5.1.4 and Appendix B.

After the spatial memory test, participants were introduced to the application using standard instructions (see Appendix B). For each task (discrete and continuous) the participant was allowed to do two practice trials, one with visit wear and one without. After the practice trials, the participant completed twelve trials for each task, six with visit wear and six without. The starting conditions (with or without visit wear, discrete or continuous space) were balanced and the tasks within each trial set started at a random location in the sequence for each participant.

After all tasks were completed, participants filled out a final questionnaire related to their preferences (see Appendix B).

5.2 *Results*

In the following sections the results from each task for each dependent variable are reported, and then the results of the participants' preferences are presented. In all cases tests were Analysis of Variance (ANOVA) except for the analysis of the E+S scores, which was done through a regression analysis.

5.2.1 Completion Time

Figure 5.5 shows the mean results over all participants of the completion time for all trials within the discrete task, for trials with visit wear and trials without visit wear.

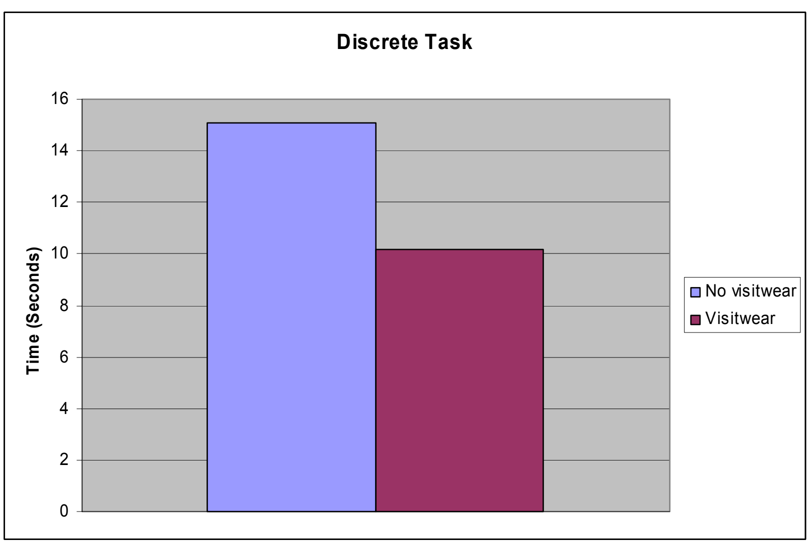

Figure 5.5: Comparison of completion time for the discrete task, with and without visit wear

For the discrete task, the addition of visit wear showed a significant improvement in time to complete the task ($F(1,15) = 16.994$, $p < 0.005$). The average time to complete the task with visit wear was 10.2 seconds, compared to 15.1 seconds without visit wear

Figure 5.6 shows the mean results over all participants of the completion time for all trials within the continuous task, for trials with visit wear and trials without visit wear.

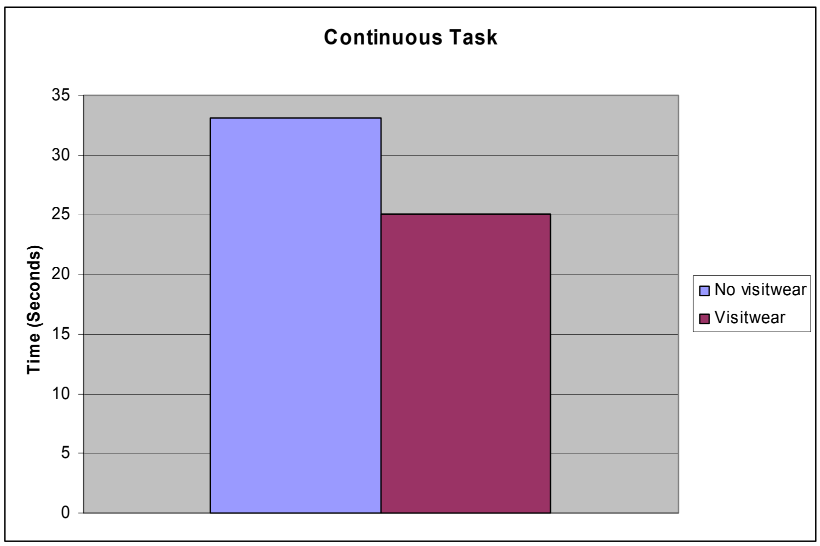

Figure 5.6: Comparison of completion time for the continuous task with and without visit wear.

For the continuous task, the addition of visit wear showed a significant improvement in time to complete the task ($F(1,15) = 11.416$, $p < 0.005$). The average time with the visit wear was 25.1 seconds, without the visit wear was 33.1 seconds.

No effect of spatial ability, as measured by the Eals and Silverman test score groups in a regression analysis, was found on completion time for either the discrete task ($F(1,190) = 0.1226$, $p = 0.7266$) or the continuous task ($F(1,190) = 0.2395$, $p = 0.6251$). No effect of gender, using an ANOVA, was found on completion time for either the discrete task or the continuous task ($F(1,14) = 0.659$, $p = 0.431$) or the continuous task ($F(1,14)=0.027$, $p = 0.872$). For the discrete task, the difficulty of the sequence had a small effect on completion time, using an ANOVA, which was expected ($F(5,75)=2.516$, $p = 0.037$).

5.2.2 Accuracy

The task in the discrete space was to click on the nodes that were labelled, in the order of the labels. To determine the accuracy of the answers, I assigned each node in the answer two possible scoring points. One point came from clicking on the correct node, the second from the node being in the correct place in the sequence. For the six node sequence, therefore, the perfect score was 12.

Figure 5.7 shows the mean results over all participants of the accuracy for all trials within the discrete task, for trials with visit wear and trials without visit wear.

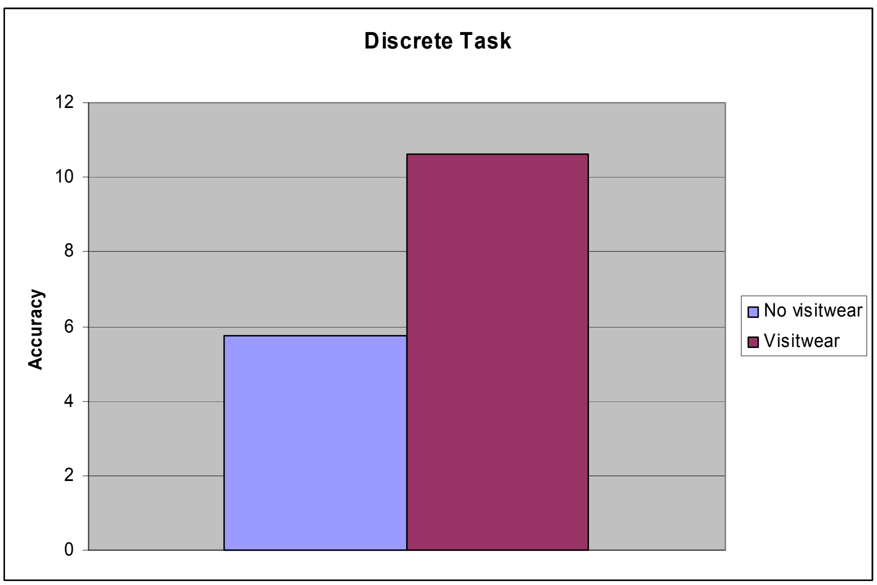

Figure 5.7: Comparison of accuracy in the discrete task with and without visit wear.

For the discrete task, the addition of visit wear showed a significant improvement in accuracy ($F(1,15) = 117.874$, $p < 0.005$). Out of a perfect score of 12, the average score with visit wear was 10.6, compared to 5.7 without visit wear.

Furthermore, the errors made with the visit wear were mostly ordering errors (the correct nodes were identified, but in the wrong order).

The continuous task was to find a very small target along a path that the mouse had previously taken. For this task, the accuracy was measured to be either yes or no (0 or 1) depending on whether or not the target was found. Figure 5.8 shows the mean results over all participants of the accuracy for all trials within the continuous task, for trials with visit wear and trials without visit wear.

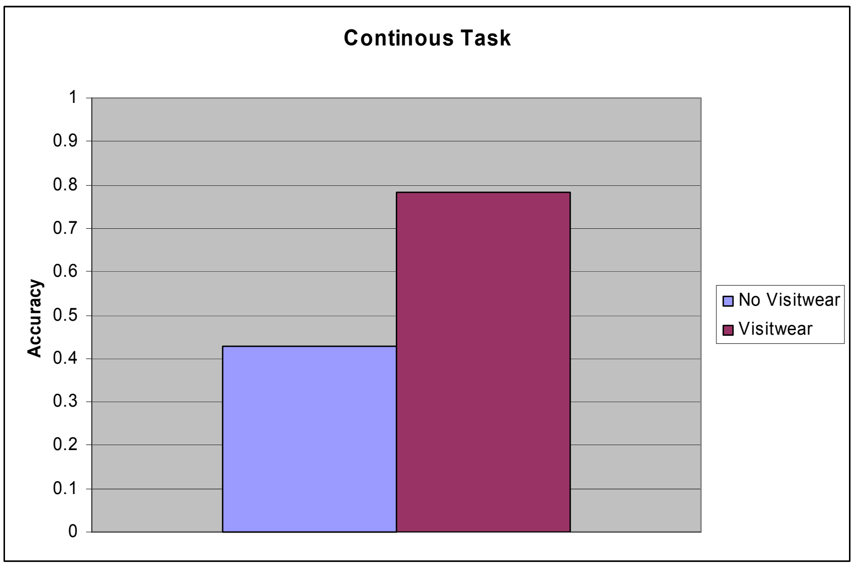

Figure 5.8: Comparison of accuracy in the continuous task with and without visit wear.

For the continuous task, the addition of visit wear showed a significant improvement in accuracy ($F(1,15) = 14.695$, $p < 0.005$). With visit wear, subjects found the target 78% of the time, and without visit wear subjects found the target 42% of the time.

No effect of spatial ability, as measured by the Eals and Silverman test score groups and using a regression analysis, was found on accuracy for either the discrete task ($F(1,190) = 0.1890$, $p = 0.6643$) or the continuous task ($F(1,190) = 0.0336$, $p = 0.8547$). No effect of gender was found on accuracy using an ANOVA, for either the discrete task (F(1,14)=1.190, p = 0.294) or the continuous task (F(1,14)=0.000, p = 1.000). For the discrete task, the difficulty of the sequence had no effect on accuracy using an ANOVA (F(5,75)=0.835, p = 0.529).

5.2.3 Distance Travelled

This measure was only analysed for the continuous task, since the "sweeping" strategy described in Section 5.1.4 was not useful in the discrete task. Figure 5.9 shows the mean results over all participants of the mouse distance traveled when looking for the target, for all trials within the continuous task, for trials with visit wear and trials without visit wear.

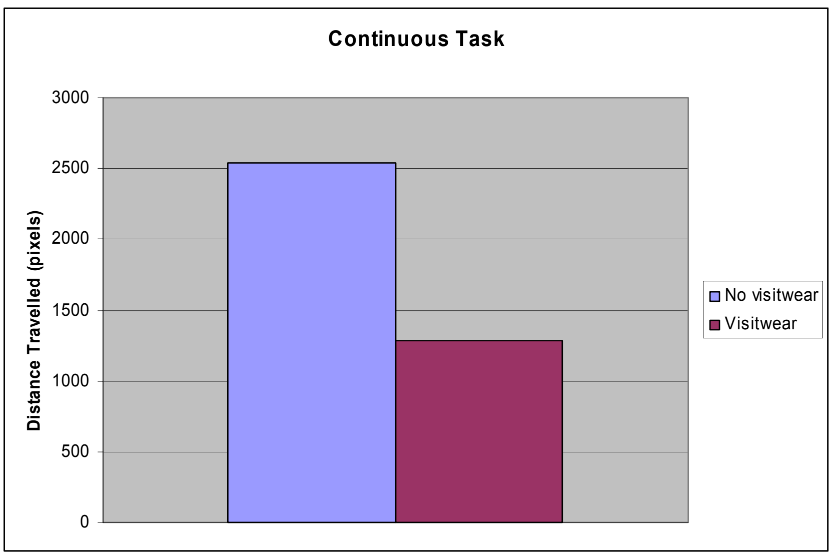

Figure 5.9: Comparison of mouse distance traveled in the continuous task with and without visit wear.

For the continuous task, the addition of visit wear showed a significant reduction in the amount the mouse traveled before the target was found (F(1,15) = 25.294, p < 0.005). The average distance with the visit wear was 1282 pixels. The target was always halfway along the subjects' mouse path, and the most direct path through all the targets was always approximately 1240 pixels long, therefore the most direct path to the target was approximately 620 pixels. Without the visit wear, the average distance traveled was 2536 pixels.

Given that the most direct path was 620 pixels, the average distance with visit wear of 1282 pixels seems surprisingly large. Breaking the average results down further into successful trials and unsuccessful trials with and without visit wear is shown below in Figure 5.10.

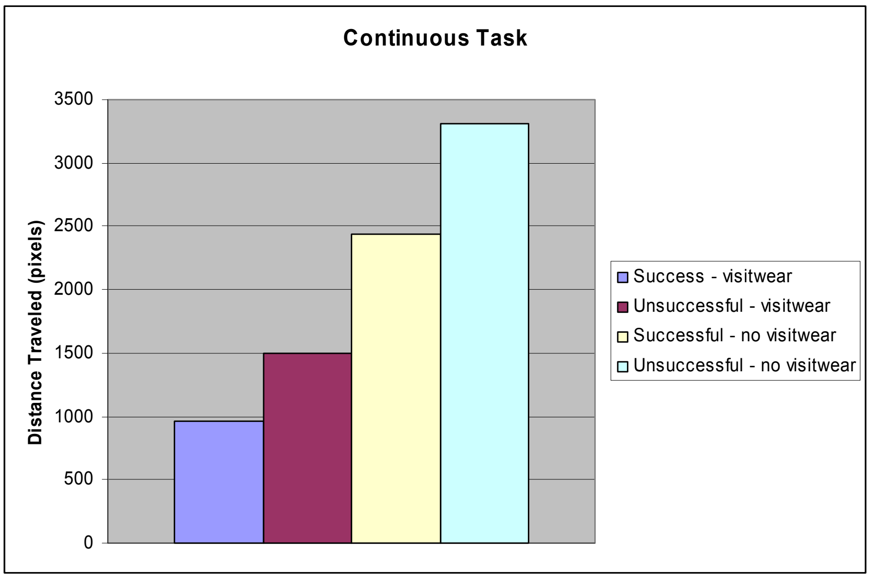

Figure 5.10: Comparison of mouse distance traveled in the continuous task with and without visit wear.

Figure 5.10 shows that if the subjects found the target, the average distance traveled with visit wear was 961 pixels, which is closer to the theoretical shortest possible distance. The remainder of the difference may be because the direct path is measured starting from the first target, while the mouse distance traveled included the distance that the subject covered to get to the start of the trail from wherever their current mouse position was. If the subjects didn't find the target with the visit wear, the average distance traveled was 1495 pixels, which is approximately what one needs to traverse the trail completely. When no visit wear was present, the average distance traveled was 2432 pixels for a successful trial and 3313 pixels for an unsuccessful one. These averages seem to indicate that without visit wear even successful trials used a sweeping strategy to find the target rather than a memory strategy.

No effect of spatial ability, as measured by the Eals and Silverman test score groups using a regression analysis, was found on distance traveled ($F(1,190) = 1.8282$, $p = 0.1779$). No effect of gender using an ANOVA was found on distance traveled ($F(1,14) = 1.212$, $p = 0.290$).

5.2.4 Participants' Preferences

In a post-experiment questionnaire, participants were asked whether they preferred the addition of visit wear in each task, and why or why not. All sixteen participants preferred visit wear in the discrete task. One subject said "Without visit wear I could only remember the general region, not the specific node." Another said "With visit wear, I didn't have to guess." Another subject called the graph without visit wear "featureless" and said that visit wear was the only way s/he could identify the target nodes at all.

Twelve out of the sixteen participants preferred the visit wear in the continuous task. Typical comments were "It narrows down the area that I have to search" and "Without the visit wear, it was impossible to remember exactly where I'd been." The dissenting four, however, all said that the visit wear trail made the space too cluttered, that they got confused by the overlapping trails, and that the "keys" were easier to see without the coloured mouse trail around them. Even subjects who preferred the visit wear said that they sometimes found themselves concentrating too hard on following the trail and therefore missed the "keys."

6.0 Discussion

6.1 The Effect of Visit wear

As Section 5.2 showed, the addition of visit wear to a data space distorted with a fisheye lens let people find previously inspected things more quickly and more accurately. For a task in a discrete space that involved targeting six previously identified nodes in a graph in the correct order, the average time to complete the task with visit wear was 10.2 seconds, compared to 15.1 seconds without visit wear. Out of a perfect score of 12 (all six nodes correctly identified in the correct order), the average score with visit wear was 10.6, compared to 5.7 without visit wear. For a task in a continuous space involving re-visiting a randomly chosen area that had been visited once already, the average time with the visit wear was 25.1 seconds, without the visit wear was 33.1 seconds (these averages include the times where the task was 'completed' because it timed out). With visit wear, subjects found the target 78% of the time, and without visit wear subjects found the target 42% of the time.

The results appeared to not be dependent at all on the participants' spatial ability. Any effects seemed to be due solely to the visit wear.

In both cases visit wear improved average performance, and from the comments of the participants, it was because the visit wear turned both tasks from memory tasks into visual tasks. Participants had to remember a minimal amount of information, which in turn led to a preference for the visit wear. All participants preferred visit wear in the discrete space, while twelve out of sixteen preferred visit wear in the continuous space. In the absence of visit wear, participants attempted to remember targets by landmarks of the data space or by the absolute screen position of the target. With visit wear, the strategies to remember targets were simpler and mainly concerned with filling in the gaps that visit wear did not provide. The success of these strategies, and the reasons for the four who preferred no visit wear, are discussed more fully in the remainder of this chapter.

6.2 Reasons for Visit wear Effects

Information visualisation in general is concerned with providing external aids to memory, thought and reasoning or, as Bertin says, "using vision to think" (Bertin, 1967/1993). Humans usually need external aids to cognition; not many of us can multiply large number in our heads, but most of us are able to do it using a pen and paper. Similarly, external aids to memory are common. We make shopping lists, use bookmarks and highlight passages in books, put sticky notes on monitors, and tie strings around our fingers. All we have to remember is the aid, not what it is referring to. And if the aid is distinctive enough, we can even forget about the aid and count on noticing it again later (my mother used to put a shoe on the oven to remind herself to turn it off; she could concentrate completely on something else until she would ask herself "what's that shoe doing on the oven?" and have her memory jogged). In the same manner, adding a visual effect to something that you are supposed to remember means that you do not have to really remember it. In terms of the visuospatial scratch pad (VSSP) described in Section 2.6, when tasks are easier there is less load on the VSSP.

That visit wear provided this external aid is supported by the improved performance in the memory tasks, as already discussed, but also by the comments of the participants. One participant commented "The visit wear is like cheating; I don't even have to try and remember where the nodes are." Similarly, another said that with visit wear "I can see where I've been, I don't have to remember." Although there are limits to the visit wear approach (as discussed in Section 6.5 below) the study showed that the technique can be extremely effective in some situations.

The visit wear representations were chosen according to the principles of spatial memory in virtual spaces (see Section 2.5) and the principles of visual cognition (see Section 2.9). The improved performance in the memory tasks showed that the visit wear provided the external aid to memory in a visually effective way. The retinal properties that make an item visually distinctive were effective in visit wear.

In Section 2.6, the possibility was mentioned that disorientation in a fisheye view might be caused by change blindness. Change blindness is when people actually fail to see even large changes if they occur during a cut in the information flow. In an interactive fisheye view, the appearance of the space is always changing and a "jump" caused by change blindness could cause a user to miss a visual target. While this was not explicitly tested, one of the ways to minimize change blindness is to provide visual cues that do not change. If visit wear is visible even in the context area of the fisheye space, it may serve as such a cue.

6.3 When Visit wear Failed

With the visit wear, participants made fewer errors than without it, but they still did make errors. This section examines and analyses the circumstances of the errors.

In the discrete space, there were three main categories of errors:

1) *Participants took too long and the visit wear effects around the nodes completely faded before the nodes were revisited*

The appearance of the visit wear did not change between the initial visitation phase and the revisitation phase of the experiment. The visit wear effect always darkened around any visited node and always started fading as soon as the mouse cursor left the node. The visit wear faded completely within one minute. In order for the visit wear to still be visible when the participants started to revisit the nodes, therefore, the participants had to click on all six target nodes within a minute.

Participants all realized, usually in the practice trial, that speed was important. Even so, sometimes the participants took longer than the fade time, sometimes because of targeting difficulties (the cursor had to be on the current node in order for the next one to appear) and sometimes because they were taking time to memorize the target nodes' surroundings. When this happened, the first (and sometimes second) nodes in the target sequence had no visit wear effect remaining and the participants had to remember or guess as to the location.

An increased fade time might have helped this problem, but also might have cause too much visual clutter in a more realistic task (one with extended interaction with the space). Also, it might have made the problem of non-target nodes being affected by the visit wear worse (see Problem 3 below)

2) *The degrees of fading in the visit wear effect were too subtle to be easily seen*

The visit wear started fading as soon as the mouse was off the node and faded in a minute smoothly from full intensity to nothing. A node that had been visited 55 seconds ago therefore had a much fainter effect than one that had been visited only five seconds ago. However, the difference between 30 seconds and 33 seconds was much harder to distinguish, since the shades of green were very similar and the nodes were not viewed side by side. Because participants tended to move between nodes fairly quickly (to keep from losing the visit wear – see Problem 1 above) there tended to be only a few seconds between the visit times for any two nodes.

Ordering errors were the most common error in the visit wear trials. Participants would identify all the target nodes correctly, but typically with a pair of nodes reversed. This most often happened with the third and fourth or fourth and fifth nodes, which matches the tendency of people to remember the first and last things in a set. Several participants commented that with visit wear, the only thing they had to remember was the order of the target nodes; they were not relying on the visit wear to reliably provide that information.

This problem might be solved with a more sophisticated fading effect; one that not only fades in transparency but also changes from one colour to another as discussed in Section 4.4.2. This might make comparison of two non-adjacent shades less difficult. In a task with more realistic data, however, choosing a colour effect that is visually distinctive is difficult enough with one colour. An alternate solution, again as discussed in Section 4.4.2, would be to use a visual indicator that is easier to compare, such as changing the position or orientation of a secondary glyph.

3) *Non-target nodes were affected by the visit wear*

As stated, appearance of the visit wear did not change between the initial visitation phase and the revisitation phase of the experiment. Any time the participants dwelled on a node, the visit wear effect started to darken around it. The participants almost always figured this out in the practice trial, and either avoided non-target nodes altogether or moved the mouse quickly enough over them that the visit wear effect did not appear. Also, the visit wear took three seconds to achieve full intensity and the participants usually realized their mistake before this. The visit wear effect on the mistake nodes therefore disappeared more quickly than the targets because the fading was starting from a lighter point. While the mistake nodes had some effect in terms of distraction, very few of the participants actually chose a mistake node instead of a target node during any of the trials.

The choice of three seconds for the time that it took for the visit wear to achieve full intensity was fairly arbitrary. Increasing this time, so that one has to inspect a node for longer before the visit wear reaches full intensity, is certainly possible. In a real task where the user's revisitation patterns could be monitored, the size of the visit wear history list could be automatically calculated as discussed in Section 4.3.

In the continuous trial, there were two main categories of error:

1) Participants passed over the target without seeing it

As in the discrete task, the visit wear effect was present in both the visitation and revisitation phases of the task. The mouse continued to leave a trail even when the participants were looking for the target. Most participants used a strategy of backtracking (traveling over the trail in reverse order instead of moving to the start of the trail) to minimize clutter, but because the trail faded with time, even the participants that moved to the start of the trail rarely got confused as to whether they were following their original trail or a newer one. This was possibly because the "new trail" colour was always available behind the mouse for purposes of side by side comparison.

Following the correct trail was therefore not a problem for most participants, but what did seem to happen was that they focused too much on the trail. Four subjects specifically commented that it was easy to concentrate too much on tracing the trail, as if that were the purpose of the task, instead of looking for the target. This happened even though they knew that they did not have to be exactly over the target to see it, and had completed two practice sessions. For almost every participant, there was at least one instance where I saw the target as the participant was moving over it, but the participant did not.

Four participants also said that the target was easier to see (had a higher contrast) when the trail was absent. The magenta target for them stood out more on the grey map than it did on the green visit wear. This might have been a problem with the specific implementation of the visit wear.

2) *Participants did not reach the target*

Whether they were trying to trace the trail exactly, or just taking time to thoroughly inspect the focal area, some participants moved the mouse cursor too slowly for some sets and therefore did not reach the target in time.

This seems to be a problem with the task rather than with the visit wear implementation. The forty second timeout was chosen just to keep the experiment from taking an unreasonable amount of time. In a real world task, of course, a time limit of this nature is less likely. Because the participants probably would have seen the target given more time, this is not so much a problem with the visit wear as with the experimental conditions.

6.4 The Effect of Strategy

For the discrete task with visit wear, the participants almost all stated that they did not need to use a strategy as such. The visit wear was effective enough that they could just use visual information. Some participants mentioned, however, that the degrees of fading in the visit wear effect were too subtle to be easily seen (as discussed in the previous section). Participants therefore remembered the order of the targets on their own, rather than relying on the visit wear.

Without the visit wear, the most common memory strategy in the discrete task was remembering the topology of the surrounding nodes; the "constellations" referred to in Chapter 3. Other strategies were counting from the edges, remembering the shape created by the six target nodes, and remembering the absolute screen position of the target nodes. Most participants stated that none of these strategies were very effective, and this is supported by the difference in accuracy between the data with the added visit wear and without. The fisheye distortion made the absolute positioning unreliable and distorted the shapes, and counting was often too high a memory load.

For the continuous task with visit wear, as mentioned in the previous section, most participants used a strategy of backtracking (traveling over the trail in reverse order instead of moving to the start of the trail) to avoid cluttering the space with overlapping trails. The participants also realized that if they moved too quickly over the target, they might not see it since the target was only visible when it was in the fisheye magnification area. Moving slowly was the obvious strategy, but the time limit of forty seconds meant that moving slowly enough to recognize the target had to be balanced with moving quickly enough to have time to inspect the whole trail.

Unlike the discrete space, in the continuous space the participants knew whether they had succeeded in the task or not and therefore had a chance to refine their search strategies. Participants who moved too slowly realized this and corrected their strategy, figuring that two quick sweeps of the entire trail was more likely to find the target than a slow inspection of half of the trail. Because the target was always halfway along the trail (though they did not know this), this was a better strategy.

Without the visit wear, participants used two main strategies to remember their paths. Like the discrete task, they remembered the shape that the six red squares made. However, because all six squares were never shown at once (unlike the discrete task), this was not as successful. Participants also tried to remember landmarks in the map, though again the fisheye distortion made this difficult as is predicted by the study in Chapter 3. For example, the following conversation was recorded during the experiment:

> Subject: "This is hard because I don't have any landmarks like I would in the real world"
>
> Researcher: "You don't?"
>
> S: "Well, none that aren't changing too much to use"

Many participants also stated that they just used random sweeping and no real strategy at all. The participants were not pre-tested for visual recognition and hand-eye coordination, but these abilities may have influenced success on the non-visit wear continuous task more than spatial memory.

6.5 Limits to Visit Wear

Although visit wear is a valuable technique, it has limits. The following issues must be considered when adding visit wear to an application.

- *Selection of visual properties*

The visibility of visit wear depends on it having properties that are distinctive with respect to the visual properties of the information around it. For heterogeneous data sets choosing a distinctive colour or texture, for example, might be difficult. Users' abilities to distinguish colours or patterns can also vary. Although none of the participants in the experiment mentioned difficulty in seeing the visit wear, they represented a fairly narrow demographic sample in terms of age and computer experience.

To address this problem, the visit wear information could be coded in several simultaneous ways, such as by the colour and shape of a secondary glyph, although as mentioned in Section 2.9 combining visual characteristics can reduce their pre-attentive qualities. Alternatively, the visit wear appearance could be customizable so that users can adjust it to their applications and needs.

- *Noise*

Visit wear, by definition, adds visual information to the screen. Ideally, this information enhances the user's performance, but it is easy to go from enhancement to clutter. All participants preferred visit wear to no visit wear when data was sparse. In this situation, the visit wear was an enhancement. A typical comment was "I could see where I'd been and didn't have to remember it." However, with the continuous space (which is richer in information) four out of the sixteen participants preferred no visit wear, and it did not improve their performances. A typical comment of this group was "The development of new paths in retracing made visit wear confusing." For a fairly large minority, adding visit wear to an already visually dense space was noise rather than an enhancement.

Having the visit wear appearance be customizable may help this problem, as might a more dynamic method of determining an appropriate history list length (see Section 4.3). Another possibility, although departing from the strict definition of visit wear, would be to allow users to turn off the creation of new visit wear when they are backtracking or revisiting. As an alternative to allowing users to turn off the visit wear creation manually, this could be done automatically by determining how far back in the history list the user typically revisits, then not showing visit wear for more recent visits than that.

- *Duration*

If the visit wear effect disappears over time, which is an obvious way to avoid clutter, then there is the possibility that the visit wear effect will disappear while users still want to revisit the information. The users could also become lost when the visual cues disappear because they were relying on these cues rather than memory of the space.

This did occur in the experiment. In the discrete space, the visit wear effect started fading from the time participants moved the mouse off the graph node. This meant that if they were slow in moving to the next target, the first visit wear effect could have faded by the time they needed to remember it. However, since the next target did not appear until the current target had reached maximum colour opacity, which took a few seconds, there was a limit to how quickly they could complete the target-visiting sequence. Most participants realized this during the trial sessions and adjusted their speed so that the visit wear effect did not have time to disappear completely, despite the built-in “waiting” time.

In the continuous space, the trail fading was less of a problem. This was partly because there was no limit to how quickly the participant could move to the next target; unlike the discrete space where the next target did not appear until after a certain time, the next target appeared immediately in the continuous space. The trail fading might have been less of a problem also because the trail continued to be drawn when participants would retrace the original route. Unless they were moving very slowly, the route would be visible because of the visit wear laid down during the retracing when the original faded.

Solving the problem of the premature disappearance of visit wear must be balanced with the possibility of creating clutter or occluding data. If the visual cues never fade, then the information marked with visit wear will always be visible. One possibility is to have a "time snapshot" option, where the user could see what was marked with visit wear in the past. Viewing past states of the information could let users find visit wear that had faded, as long as they knew the approximate time that it was last visible.

- *Occlusion*

Related to the noise problem, there is the possibility that the visual effects of visit wear could actually obscure the information that the user is trying to see if the visit wear effects are too large or too opaque or if the task involves observing fine detail. The experiment tried to avoid this problem in the discrete task by keeping the visit wear effect over the background, and in the continuous task by always drawing the magenta target on top of the visit wear effect. Even this did not succeed all the time, though, as the participants who disliked the visit wear in the continuous space pointed out that the magenta target was easier to see on the grey map than it was on the green visit wear. One participant commented "I concentrate too much on the green."

Occlusion will be more of a problem in continuous spaces than in sparsely populated discrete spaces, since the continuous spaces have no background that visit wear can be safely drawn over. Even discrete spaces can have occlusion problems, though, as the density of the information increases. Being able to customize the transparency of the visit wear effect might help, as would making visit wear be an overlay on the data that could be called up only when necessary.

- *Information overload*

There is also the possibility in an information space that visual characteristics are already being used to encode information (such as markup information or ownership), and too much visually encoded meta-information could confuse the user as to what is representing what. This was not tested in the experiment, where visit wear was the only encoded information, but could be a problem in other information spaces. Making the appearance of the visit wear customizable might help this, depending on what else is encoded.

6.6 Generalizing the Results

The experiment was deliberately limited to isolate and explore the effects of visit wear on memorability. The intent was to determine if the techniques could be valuable in real world applications with real world tasks. The experiment showed that visit wear improved memorability, and there are reasons to expect similar results in the real world despite the experiment's limitations.

The fisheye lens was the same as would be used in a real world task; the EPS software is used in commercial applications (Idelix, 2003). However, only one style of fisheye lens was used, and it is not clear whether the results from the constrained lens could be generalized to predict results from other styles of lens. A Sarkar-Brown lens should give similar results since the visual properties of the visit wear were derived from test with a Sarker-Brown lens. A truncated lens might let different visual properties successfully represent visit wear, as mentioned in Table 4 and Table 5.

The continuous data space was an actual satellite map, while the discrete graph was small but similar in structure to real graphs. The biggest simplification in the data was the elimination of data from the graph so that all the nodes in the graph were yellow squares. This was done to isolate the visit wear effects, however if the fisheye lens included semantic zooming (data only becoming visible at certain magnifications) then information could appear as a blank node when it is in the context area of the visualization. As will be discussed, a more realistic task will also involve more realistic data.

The participants were experienced in visual systems, and their use of the system was not artificially limited except by the task constraints. The participants' reactions and use of the visit wear was reasonably realistic.

The basic aspects of the experimental tasks (inspecting the data space, moving the focus of the fisheye lens, finding information that had already been visited) were the same as they would be in the real world, but the tasks themselves were very simplified. In more realistic tasks, participants would work with more detailed semantic content over a longer period of time. The tasks in this experiment were very short and uncomplicated, unlike most real world interactions with large data sets. Real world tasks more often involve recalling a specific piece of information after a reasonable length of time, rather than recalling information in a certain order after a short period of time.

This thesis is the first study to have investigated the effects of visit wear, and was necessarily limited. To be able to achieve more general results, more realistic tasks with real data should be performed. This will illustrate the potential problems with other visual data that might interfere with the use of the visit wear. Also, more realistic tasks would give the opportunity to examine revisitation patterns and their use in calculating a dynamic history length for visit wear applications.

7.0 Conclusion

The problem addressed in this thesis was that people have difficulty remembering where they have been in interactive fisheye views. The main motivation for solving this problem was that fisheye views are a good method of providing a focus-plus-context view of a large information space. Solving fisheye usability problems will make them a more useful tool for visualising large data spaces. Another motivation is that studying general methods of improving memorability may improve the usability of any application where the user commonly revisits items that have already been inspected or used.

The solution explored in this thesis was to develop the concept of "visit wear" and add it to the fisheye visualisation. Visit wear is the addition of visual information to the display to represent the user's interaction history, and is an extension of the existing techniques of read wear and edit wear. To develop the idea of visit wear, this thesis examined the principles of visibility in a distorted space, defined the parameters of the visit wear design space, and finally implemented a visit wear application to test its effect on users' performance in memory tasks.

The test application showed that the addition of visit wear to a fisheye visualization let users complete simple memory tasks faster and with fewer errors than when the tasks were done without visit wear.

7.1 Summary of Research

The research in this thesis consisted of an experiment to determine the basic principles of memorability in a fisheye distortion, the development of a design space for visit wear in general and for distorted spaces in particular, and finally a second experiment to test the effectiveness of visit wear in a distorted space.

The first experiment determined the principles of memorability in a fisheye application by observing user strategies when they were asked to remember a specific target in a distorted space. The purpose was to see which visual properties of a screen object retained their distinctiveness even when they were distorted by a fisheye lens. These properties were called *robust* with respect to distortion and were colour, pattern, topology, orthogonal ordering, semantic content, black value and orientation. Absolute position, size and form were not robust with respect to distortion.

The main parameters of the visit wear design space were defined and explored. The data space type, the duration of the visit wear and appearance of the visit wear were identified as the three core parameters. Different values and methods for applying each of these parameters were presented. Finally the intersection of the visit wear design parameters with the distortion-robust properties already determined was established, resulting in a list of visit wear implementation methods that would work well with a fisheye lens. Texture, colour, black value and orientation are the best properties for representing visit wear in a fisheye lens.

The second experiment implemented two possibilities from the visit wear design space to test their effect on memorability in a fisheye view. In simple memory tasks with visit wear, the accuracy, completion time and mouse distance traveled were measured and compared to the same measurements performing the same tasks without visit wear. Visit wear improved the accuracy and decreased the completion time of the memory tasks. Visit wear also decreased the average mouse distance traveled while completing the tasks, even when the tasks were successfully completed.

7.2 Contributions

The major contribution of this research is providing empirical evidence that visit wear is a valuable design concept for improving memorability and usability in fisheye views. This will expand the value of fisheye lenses as a tool for visualizing large data spaces. Minor contributions are:

- An initial definition of the design space of visit wear. Visit wear is a new extension of the concept of read wear, and an exploration of its possible parameters has not yet been done.
- A set of visualisation techniques that can be used for implementing visit wear. The principles of visual cognition can be applied to representing visit wear so that the visit wear is an enhancement to the task rather than a distraction.
- Additional evidence that landmarking strategies are a method of compensating for distortion in a virtual space. The role of landmarks in virtual space has been studied but not when combined with distortion of the space.
- Guidelines for automatically calculating the effective duration of displayed history by using revisitation data. Research indicates that many tasks involve revisiting information that has already been seen, and that revisitation patterns can be predicted for many tasks. This may be able to be applied to the calculation of history list length in visit wear.

- A better understanding of the limits of visit wear. Visit wear is not a suitable solution for all problems, and this thesis discusses issues that ought to be considered when applying visit wear.

7.3 Future Work

This thesis is the first study to have investigated the effects of visit wear, and only the basic principles and simplest cases were studied. There are many aspects of the visit wear design space that bear further exploration, especially the idea of using revisitation patterns to automatically determine a good history list length. It is also possible that a more quantifiable set of guidelines could be developed for designing the appearance of visit wear; for example, given a certain background information density, any textures used to indicate visit wear should have a certain granularity.

The design space of visit wear described in Chapter 4.0 is based on work done on both visual cognition and spatial memory in virtual spaces. The experiment described in Chapter 5 showed that visit wear aided memorability in the given tasks, but did not prove that it was because of having pre-attentive visual properties, compensating for change blindness, or some other reason. Further investigation into isolating and comparing visit wear representations could better prove that the design space of visit wear is based on valid principles. Performing the experiment in Chapter 3 in a more rigorous manner could also strengthen the foundations of the design space.

With regards to the fisheye lens, the effect of different lens styles on memorability is another area that needs further exploration. A constrained lens was used to test the visit wear in this thesis, but studies should also be done on the effects of the technique with other types of lens, such as a Sarkar-Brown or a constrained truncated lens. Of course, with the many ways to represent visit wear, much research could be done on the possible permutations of lens type and visit wear style.

More realistic tasks with more realistic data will allow the use of revisitation data to calculate history length. More realistic data will also give a better idea of whether other visual data will interfere with the use of the visit wear. Tasks of a longer duration would also better explore the possibility that visit wear adds too much clutter to the visualisation.

This thesis represents the first steps in defining and analysing visit wear. It is a promising technique for improving the usability of fisheye views with potential to improve the usability of other applications as well.

References

Apple Software (2002). The Macintosh OS X Operating system. http://www.apple.com/macosx/theater/dock.html

Baddeley, A. D. and Hitch, G. (1974). Working Memory. The Psychology of Learning and Motivation: Advances in Research and Theory, G. H. Bower (Ed.), NY Academic Press, Vol. 8, pp. 47-89

Baudisch, P., Grod, N., Bellotti, V. and Schraedley, P. (2002). Keeping Things in Context: A Comparative Evaluation of Focus Plus Context Screens, Overviews and Zooming. In Proc. ACM CHI 2002, pp. 259-266

Baylis, G. C., & Driver, J. (1993). Visual attention and objects: Evidence for a hierarchical coding of location. Journal of Experimental Psychology: Human Perception and Performance, Vol. 19, pp. 451-470

Bederson B. and Hollan J. (1994). "Pad++: A zooming graphical interface for exploring alternate interface physics". In Proc. ACM User Interface Software and Technology, pp. 17-26

Bertin, J. (1967/1983). Semiology of Graphics: Diagrams, Networks, Maps (W.J. Berg, trans.) Madison WI: University of Wisconsin Press

Bresenham, J. E. (1965) Algorithm for Computer Control of a Digital Plotter*,* IBM Systems Journal, 4(1), pp. 25-30

Bridgeman, S. and Tamassia, R. (2000). Difference Metrics for Interactive Orthogonal Graph Drawing Algorithms. Journal of Graph Algorithms and Applications, Vol. 4 No. 3, pp. 47-74

Bush, V. (1945). As We May Think. Atlantic Monthly, 176(1), pp. 101-108.

Card, S., Mackinley, J. and Shneiderman, B. (Eds.) (1999). Readings in Information Visualization, San Francisco CA: Morgan Kaufmann Publishers

Carpendale, M.S.T. (1999). A Framework for Elastic Presentation Space. Ph. D. Thesis, Department of Computing Science, Simon Fraser University

Carpendale, M.S.T., Cowperthwaite, D. and Fracchia, F. (1997). Extending distortion viewing from 2D to 3D. Computer Graphics and Applications, pp. 42-51

Churcher, N. Prachuabmoh, P. and Churcher, C. (1997). Visualization techniques for collaborative GIS browsers. In International Conference on GeoComputation, August issue

Cockburn, A. and Jones, S. (1996). Which Way Now? Analysing and easing inadequacies in WWW navigation. Int. J. Human Computer Studies, 45(1), pp 105-129

Czerwinski, M. (1997). Spatial Ability and Visual Navigation: An Empirical Study. In The New Review for Hypertext and Multimedia, Volume 3, pp. 40-66

Darken, R.P., Allard, T., & Achille, L. (1998). Spatial Orientation and Wayfinding in Large-Scale Virtual Spaces: An Introduction. Presence: Teleoperators and Virtual Environments, Vol. 7, No. 2, pp. 101-107

Friendly, M and Denis, D. (2004) Milestones in the History of Thematic Cartography, Statistical Graphics, and Data Visualization http://www.math.yorku.ca/SCS/Gallery/milestone

Furnas, G.W. (1981) The FISHEYE View: a new look at structured files, Bell Laboratories Technical Memorandum #81-11221-9

Golledge, R. G. (1991) Cognition of physical and built environments. In T. R. Garling & W. Evans (Eds.), Environment, cognition, and action: An integrated approach. pp. 35-62. New York: Oxford University Press

Golledge, R.G. (1999) Wayfinding Behavior: Cognitive Mapping and Other Spatial Processes, Johns Hopkins University Press

Greenberg, S., Gutwin, C. and Cockburn, A. (1996) Using distortion-oriented displays to support workspace awareness. In People and Computers XI (Proceedings of the HCI'96), August issue

Gutwin, C. (2002) Improving Focus Targeting in Fisheye Views. In Proc. ACM CHI 2002, pp. 267-274

Gutwin, C. (2002) Traces: Visualizing the Immediate Past to Improve Group Interaction. Proceedings of GI 2002, pp. 43-50

Hill, W. and Hollan, J. (1993) History-Enriched Digital Objects. Proceedings of Computers, Freedom and Privacy (CFP'93) http://www.cpsr.org/conferences/cfp93/hill-hollan.html

Hill, W., Hollan, J. Wroblewski, D. and McCandless, T. (1992) Edit Wear and Read Wear. In Proc. CHI 92, pp. 3-9

Idelix Corporation. (2003) Pliable Display Technology, www.idelix.com.

Jones, S. and Cockburn, A. (1996) A study of navigational support provided by two World Wide Web browsing applications. In Proc. 7th ACM Conference on Hypertext, pp 161-169

Keahey, T. A. (1998) The Generalized Detail-In-Context Problem. Proceedings of the IEEE Symposium on Information Visualization, IEEE Visualization, October

Lamping, J. and Rao, R. (1996) The Hyperbolic Browser: A Focus+Context Technique for Visualising Large Heirarchies. Journal of Visual Languages and Computing, Vol. 7 No. 1, pp 33-55

Lamping, J., Rao, R. and Pirolli, P. (1995) A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies. In Proc ACM CHI 1995, pp. 401-408

Leung, Y. and Apperley, M. (1994) "A review and taxonomy of distortion-oriented presentation techniques", ACM Transactions on Computer-Human Interaction, Vol. 1 No. 2, June 1994, pp. 126-160

Leung, Y.K., Spence, R., and Apperley, M.D. (1995) Applying bifocal displays to topological maps. International Journal of Human-Computer Interaction, Vol 7, No. 1 pp. 79-98

Lynch, K., (1960). The Image of the City. Cambridge, MA: MIT Press

Mackinlay, J. D., Robertson, G. G. and Card, S. K. (1991) The Perspective Wall: Detail and context smoothly integrated. In Proc. of CHI'91, ACM Press, pp. 173-179

Miller, G.A. (1956) The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. The Psychological Review, Vol. 63, pp. 81-97

Misue, K., Eades, P. Lai, W. and Sugiyama, K. (1995) Layout Adjustment and the Mental Map. Journal of Visual Languages and Computing, Vol. 6 No. 2, pp. 183-210

Munzner, T. (1998) Exploring Large Graphs in 3D Hyperbolic Space. IEEE Computer Graphics and Applications, Vol. 18, No. 4, pp. 18-23

Munzner, T. and Burchard, P. (1995) Visualizing the structure of the World Wide Web in 3d hyperbolic space. In Proc. VRML '95, pp 33--38. ACM Press

Peterson, Barry. (1998) The Influence of Whole-Body Interaction on Wayfinding in Virtual Reality. Masters' Thesis, University of Washington

Rao, R. and Card, K. (1994) Table lens: Merging graphical and symbolic representations in an interactive focus plus context visualization for tabular information. In Proc. CHI'94, ACM Press, pp 318-322

Rensink, Ronald A. (2000) When Good Observers Go Bad: Change Blindness, Inattentional Blindness, and Visual Experience. Psyche 6(9)

Sadalla, E. K., Burroughs, W. J., and Staplin, L. J. (1980) Reference points in spatial cognition. Journal of Experimental Psychology: Human Learning and Memory, Vol. 6 No. 5, pp. 516-528

Sarkar, M. and Brown, M. (1992) Graphical Fisheye Views of Graphs. In Proc. ACM CHI'92, pp. 83-91

Satalich, G. (1995). Navigation And Wayfinding In Virtual Reality: Finding The Proper Tools And Cues To Enhance Navigational Awareness. Master's Thesis, University of Washington, Seattle. http://www.hitl.washington.edu/publications/satalich

Siegel, A.W., and White, S.H., (1975). The Development of Spatial Representation of Large-Scale Environments. In H.W. Reese (Ed.), Advances in Child Development and Behavior. New York: Academic Press

Silverman, I., & Eals, M. (1992). Sex differences in spatial abilities: Evolutionary theory and data. In J.H.Barkow, L. Cosmides, & J. Tooby (Eds.), The Adapted Mind (pp.533-549). New York: Oxford

Skopik, A. and Gutwin, C. (2003). Finding Things in Fisheyes: Memorability in Distorted Spaces. Proceedings of GI 2003, in press

Smith, R.B., O'Shea, T., O'Malley, C., Scanlon, E. & Taylor, J. (1989) Preliminary experiments with a distributed, multimedia problem solving environment. CITE

Report 86. Institute of Educational Technology, Open University, Milton Keynes, UK

Solso, Robert L. (2001) Cognition and the Visual Arts, Cambridge MA, MIT Press

Spence, R. & Apperley, M.D. (1982) Data Base Navigation: An office environment for the professional, Behaviour and Information Technology, 1, 1, pp.43-54

Spencer, J. and Schöner, G. (2000). A Dynamic Field Model of Location Memory. In Proc. CogSci2000, http://www.cis.upenn.edu/~ircs/cogsci2000/PRCDNGS/SPRCDNGS/posters/spe_sch.pdf

Storey, M., Fracchia, F., and Müller, H. (1999). Customizing a Fisheye View Algorithm to Preserve the Mental Map. Journal of Visual Languages and Computing, Vol. 10 No. 3, 245-267

Tauscher, L. and Greenberg, S. (1997) Revisitation Patterns in World Wide Web Navigation. In Proc. ACM CHI'97, pp. 398 - 406

Treisman A. and Gelade G. (1980) A feature integration theory of attention, Cognitive Psychology, Volume 12, pp. 97-136

Vinson, N.G. (1999) Design Guidelines to Support Navigation in Virtual Environments. In Proc. CHI 99, ACM Press, pp. 278–285

Ware, C. and Lewis, M. (1995) The DragMag image magnifier. In Proc. CHI 95, ACM Press, pp. 407-408

Wexelblat, A. (1998). History-rich tools for social navigation. In Proc. CHI 98, ACM Press, pp. 359–360

Wesenick, M-B. (2003) Limitations of Human Visual Working Memory. Doctoral Thesis, Ludwig Maximilians University, Munich

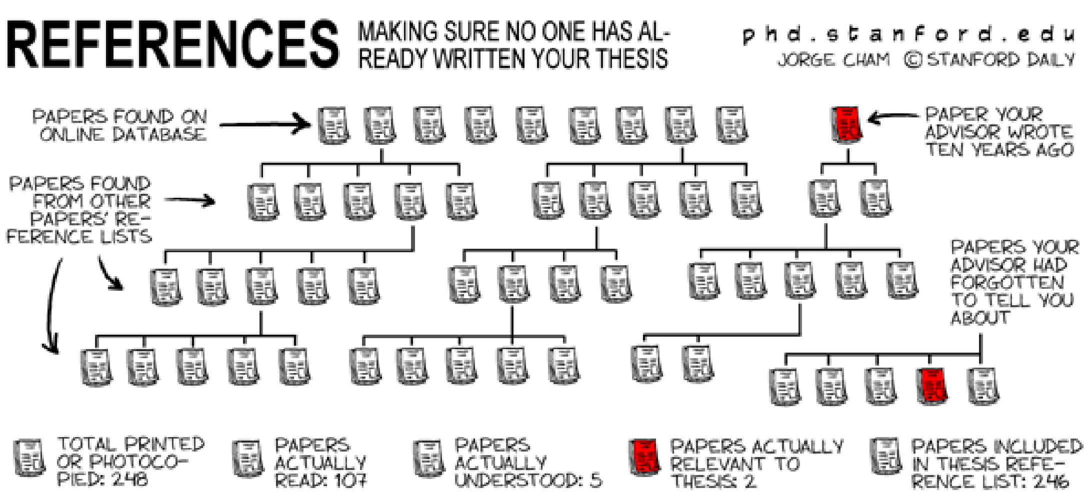

Appendix A - Colour Figures

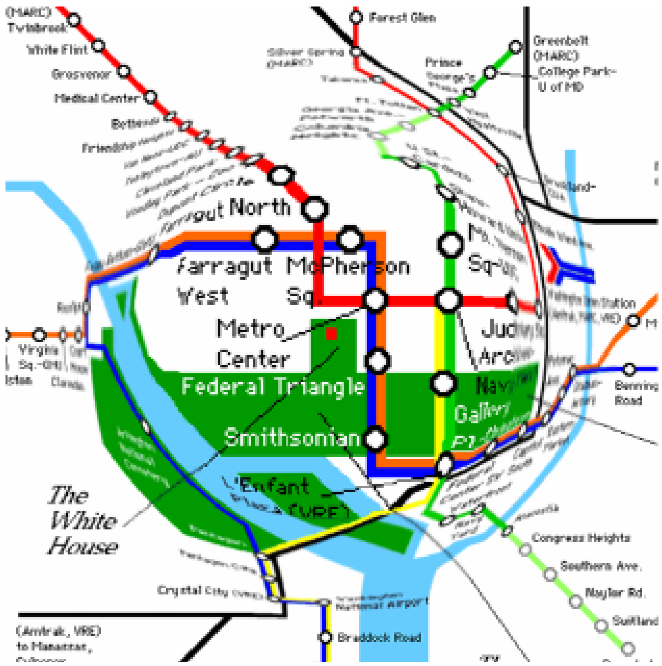

Figure 1.1: Fisheye view of a map of the Washington DC metro system with the focal point centred on the White House (Friendly and Denis, 2004).

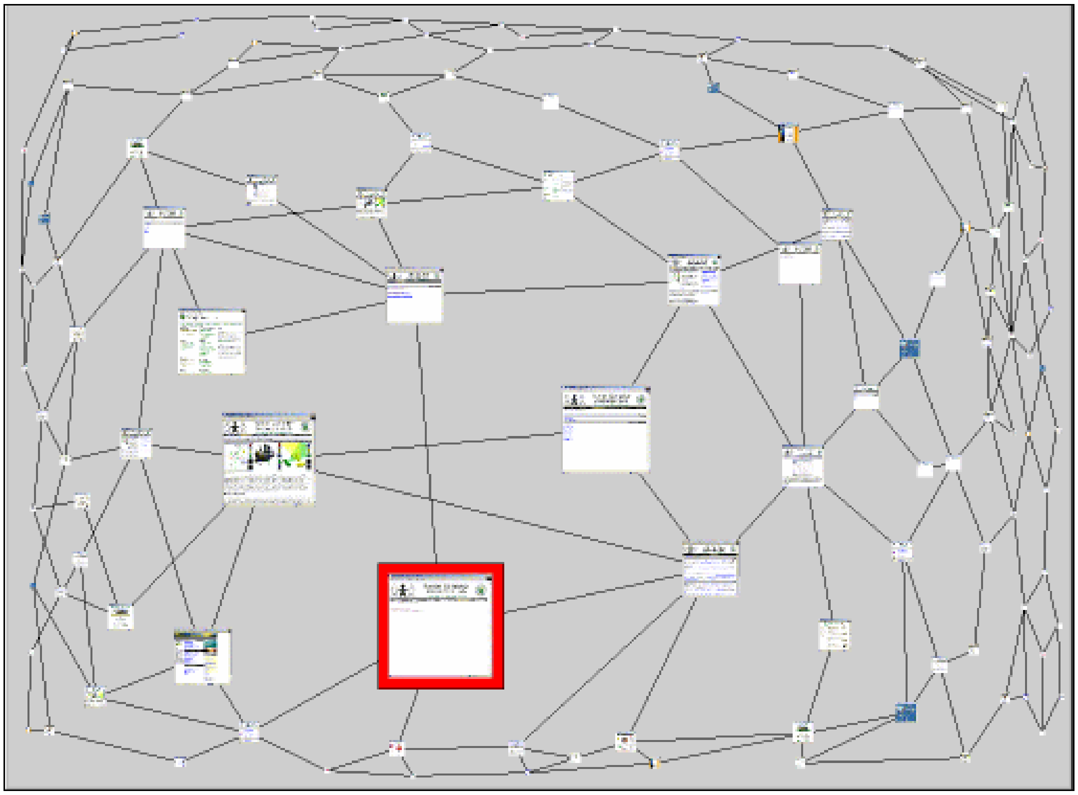

Figure 1.2: Example graph with a fisheye distortion and a random node marked

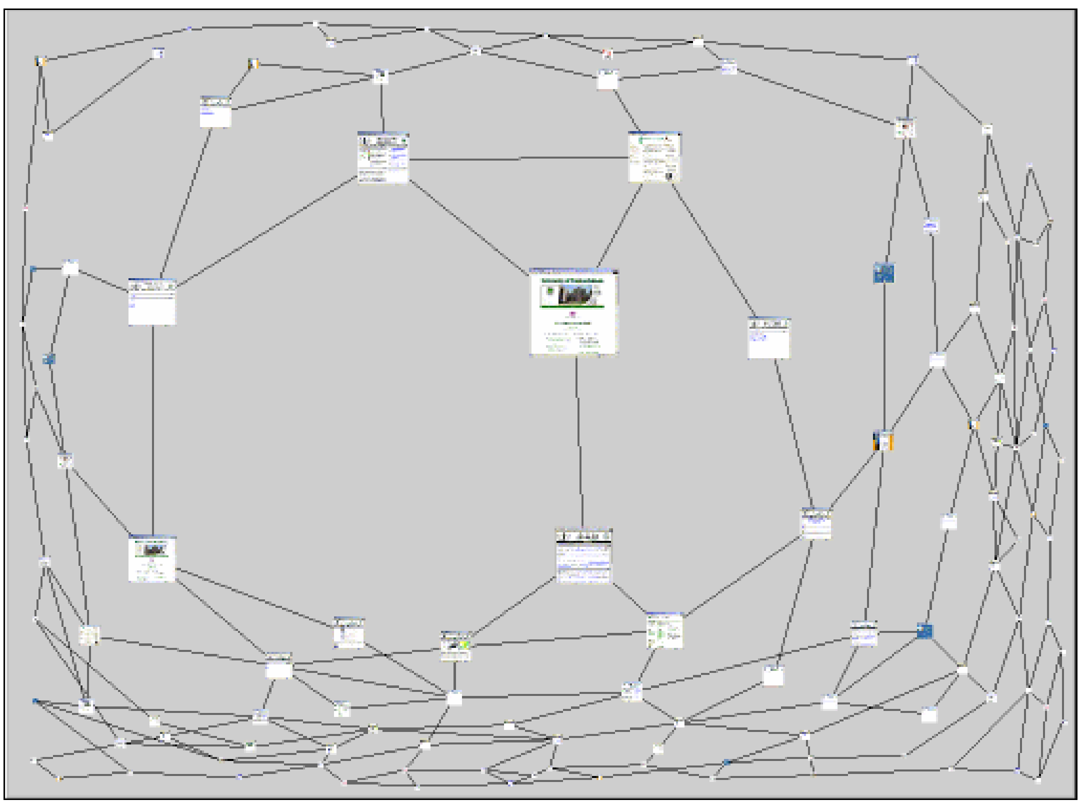

Figure 1.3: The same graph as Figure 1.2 with the focal point in a different place. The node that was marked is difficult to find.

Figure 2.1: A radar view, with the context of the magnified area shown in the small overview to the right (after Smith et al, 1989)

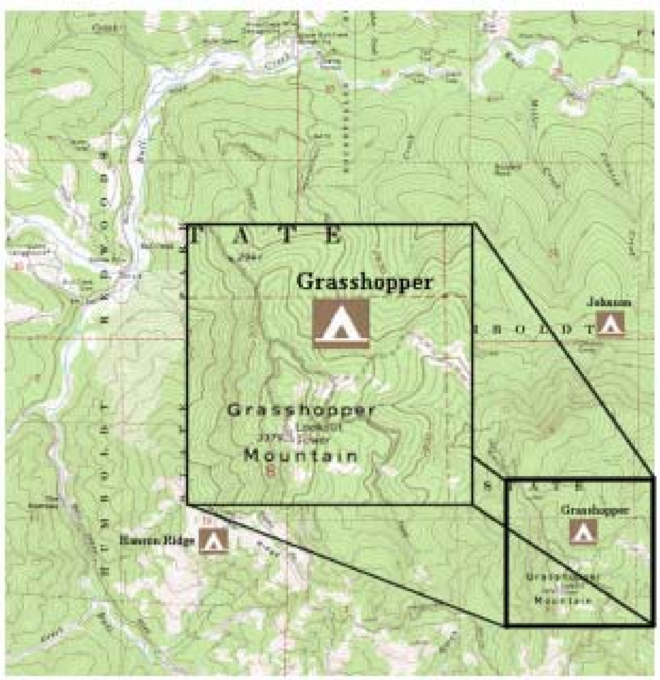

Figure 2.2: The Drag-Mag lens (after Ware, 1995)

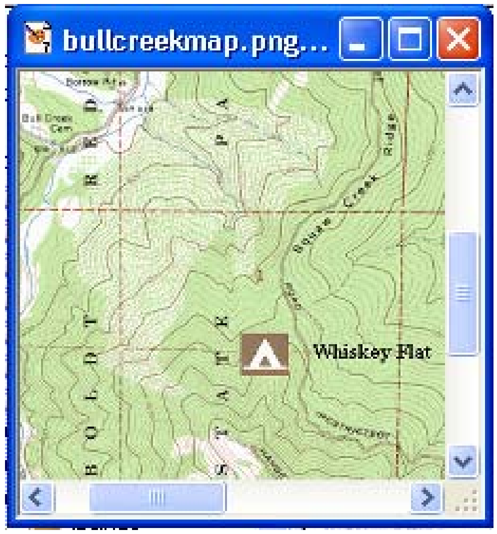

Figure 2.3: A zoom (using the mouse wheel) and pan (using the scroll-bars) interface from Paintshop Pro (Jasc Software)

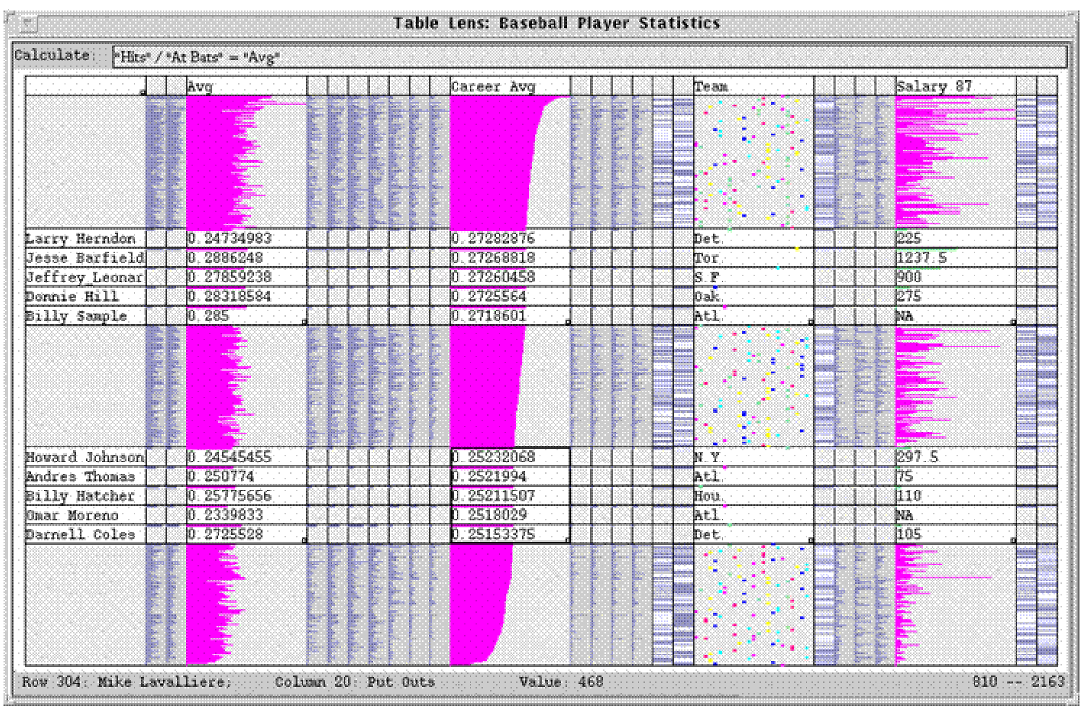

Figure 2.6: The TableLens, a combination of spatial and semantic magnification. (Rao and Card, 1994)

Figure 2.18: The path taken by people who have crossed this street in the past is clearly shown.

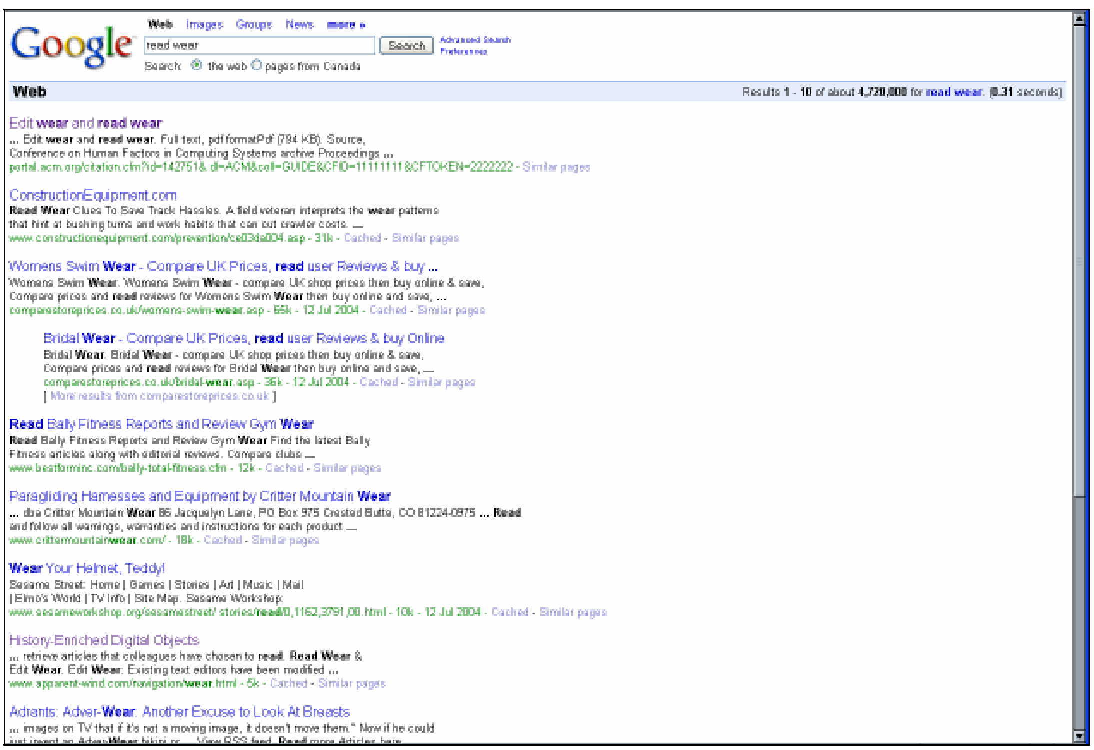

Figure 2.19: A typical Web page showing visited links in a different colour.

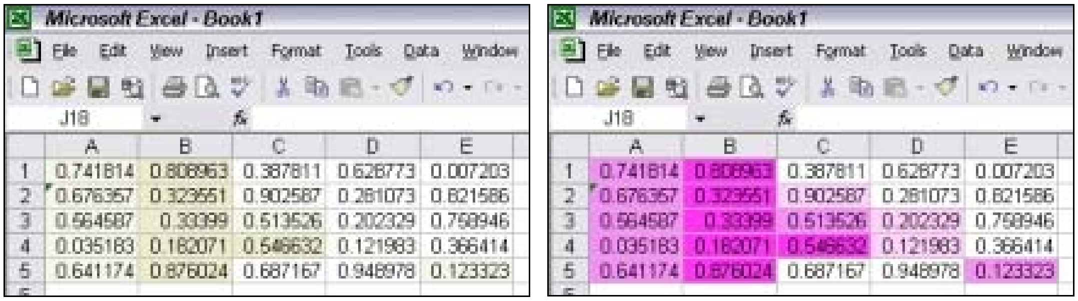

Figure 2.20: Two colour choices for an edit wear enhanced spreadsheet

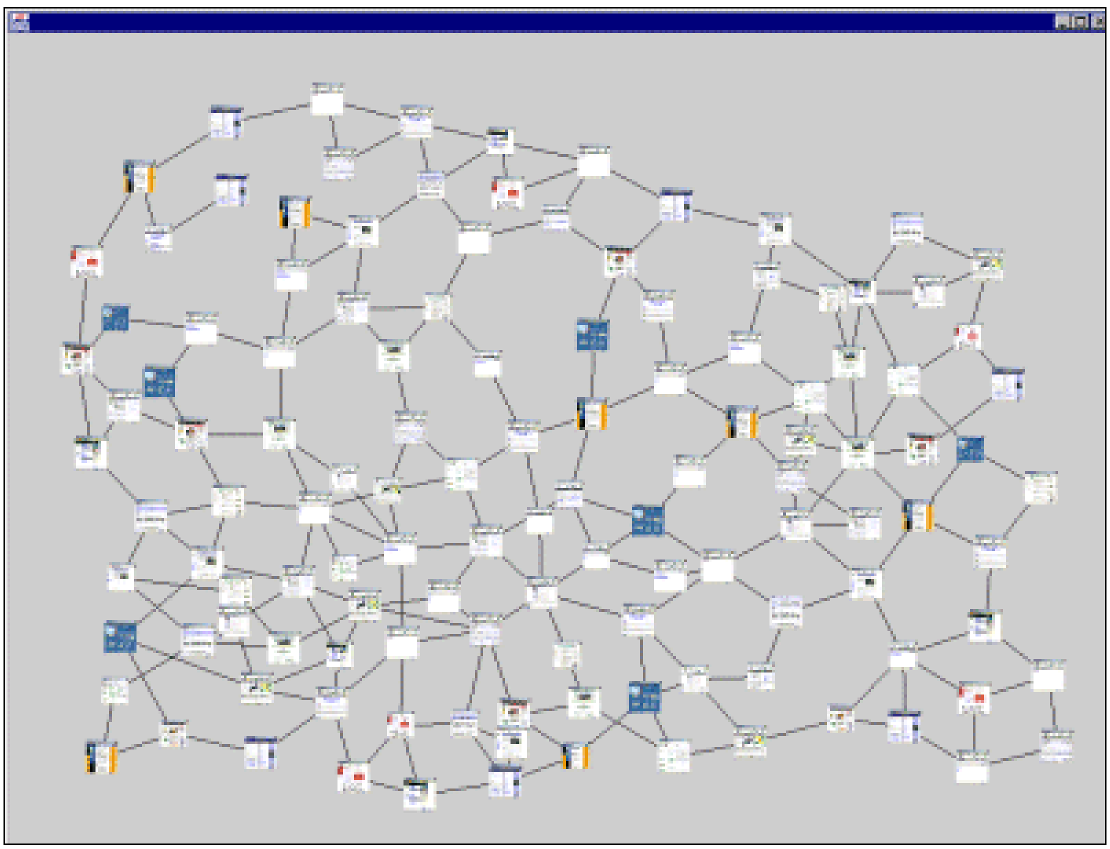

Figure 3.1: Graph used in the study, showing distortion level 0.

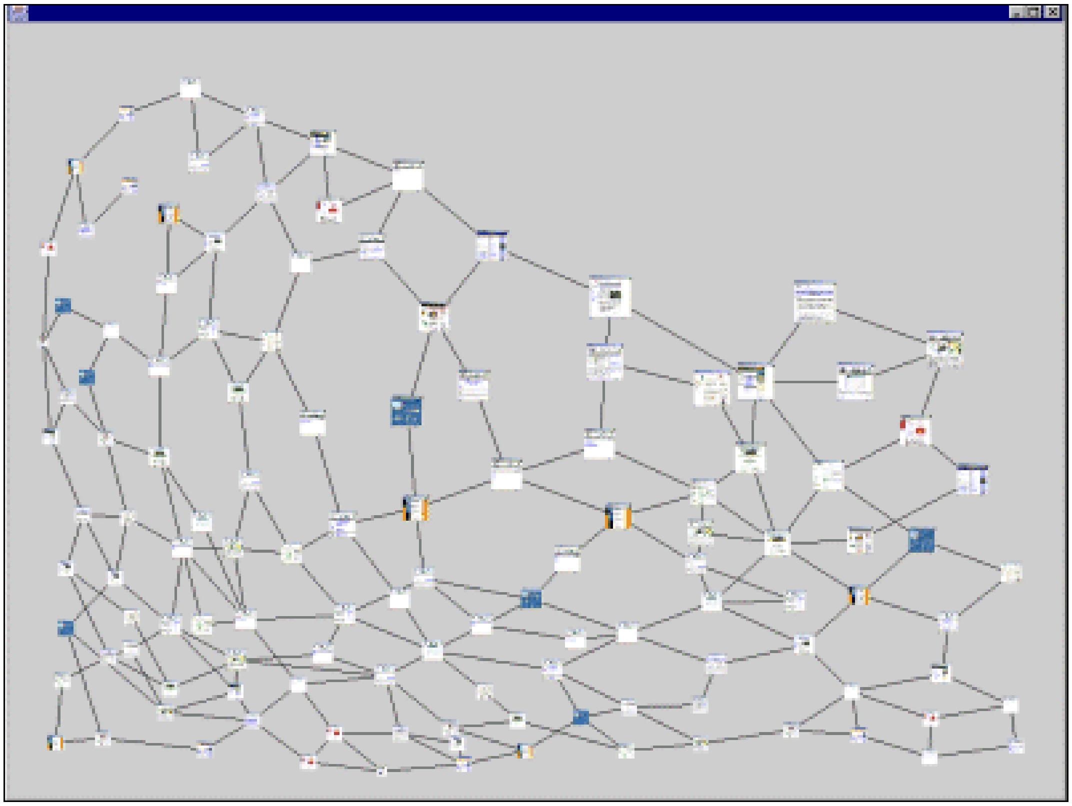

Figure 3.2: Graph used in the study, showing distortion level 1.

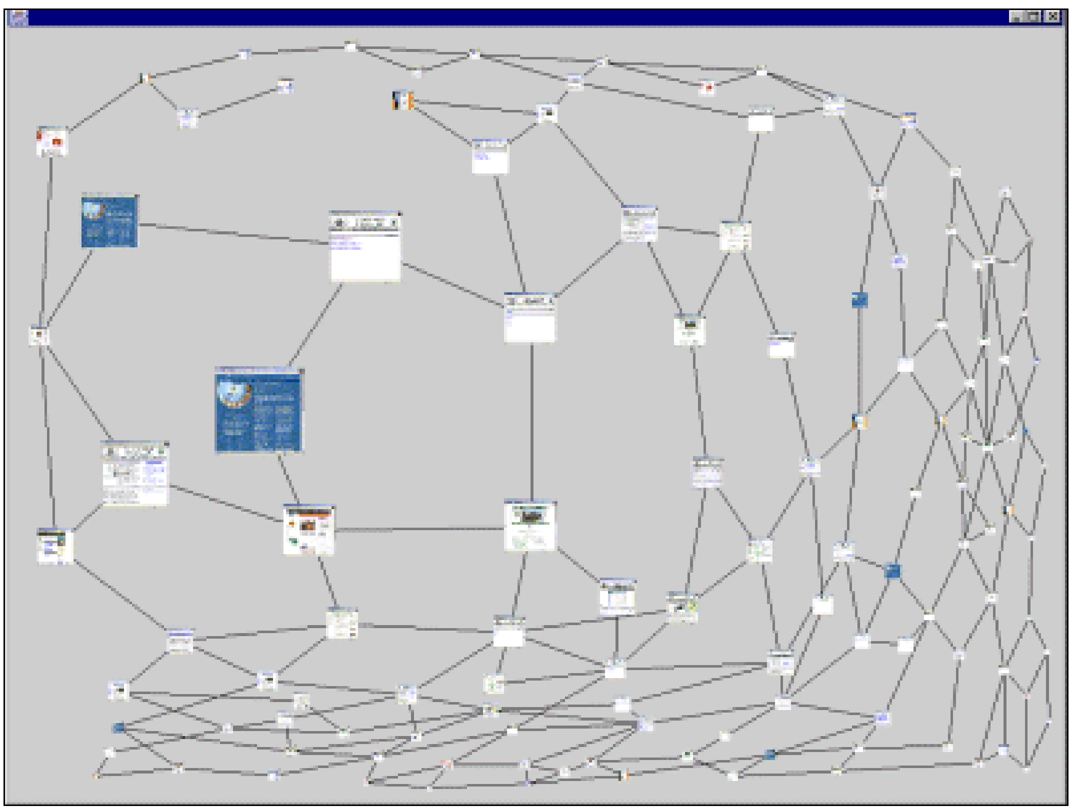

Figure 3.3: Graph used in the study, showing distortion level 3.

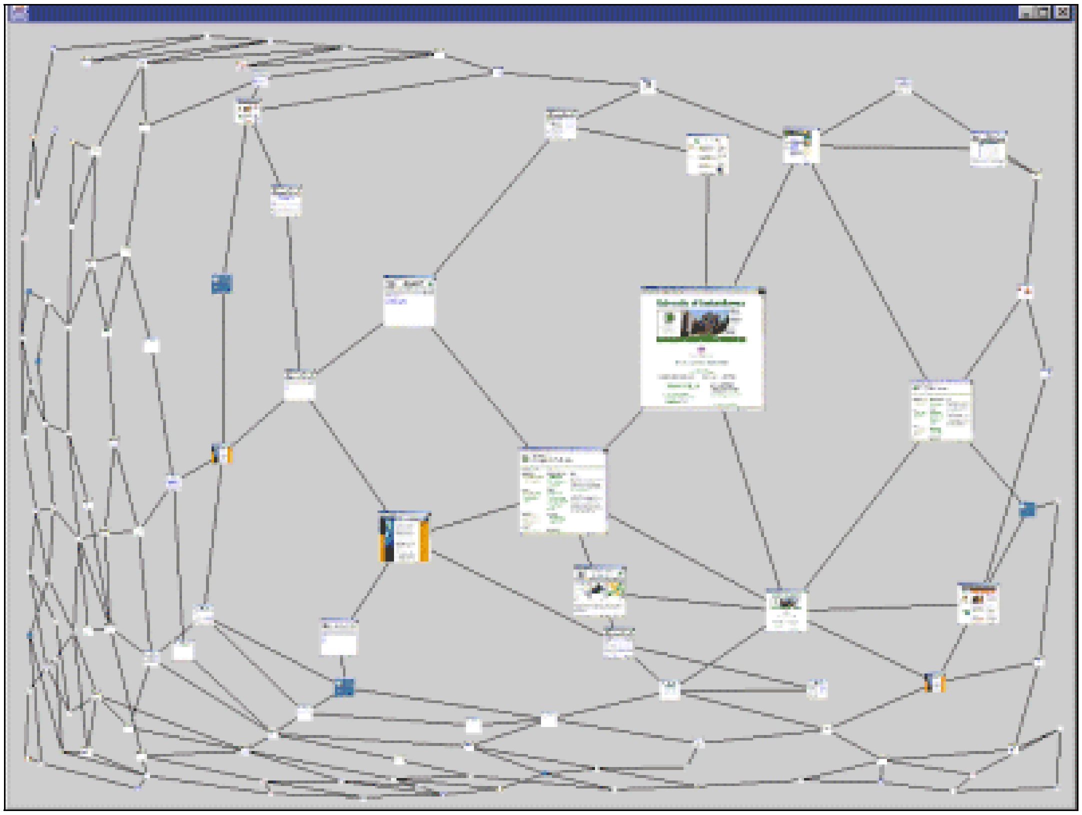

Figure 3.4: Graph used in the study, showing distortion level 5.

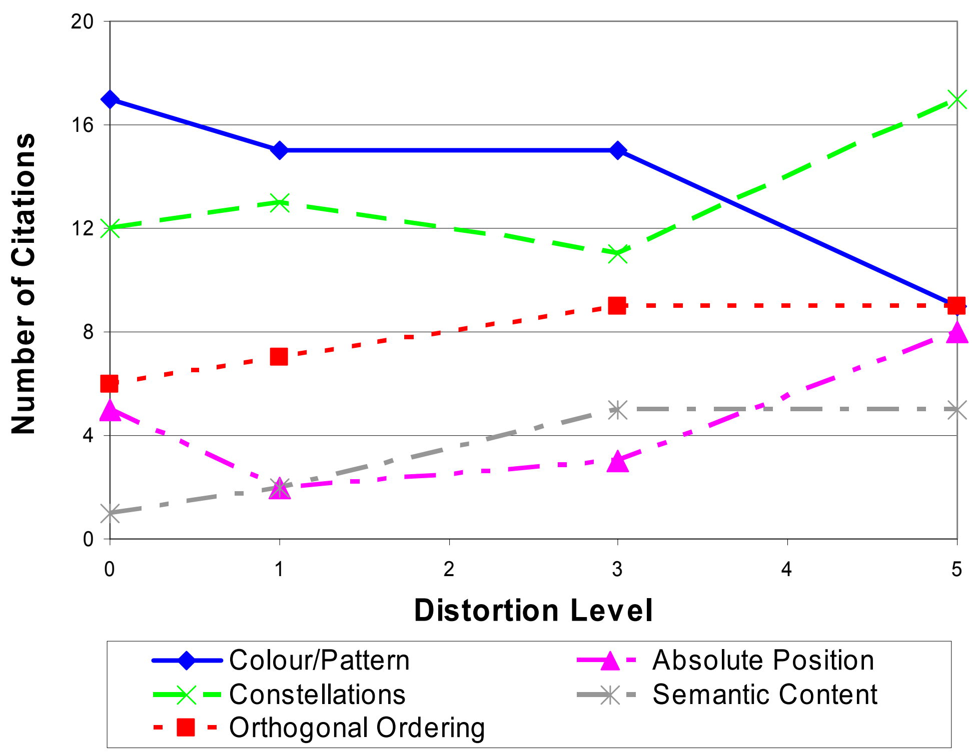

Figure 3.5: Memorable visual property choices (number of times mentioned over all tests) for each distortion level.

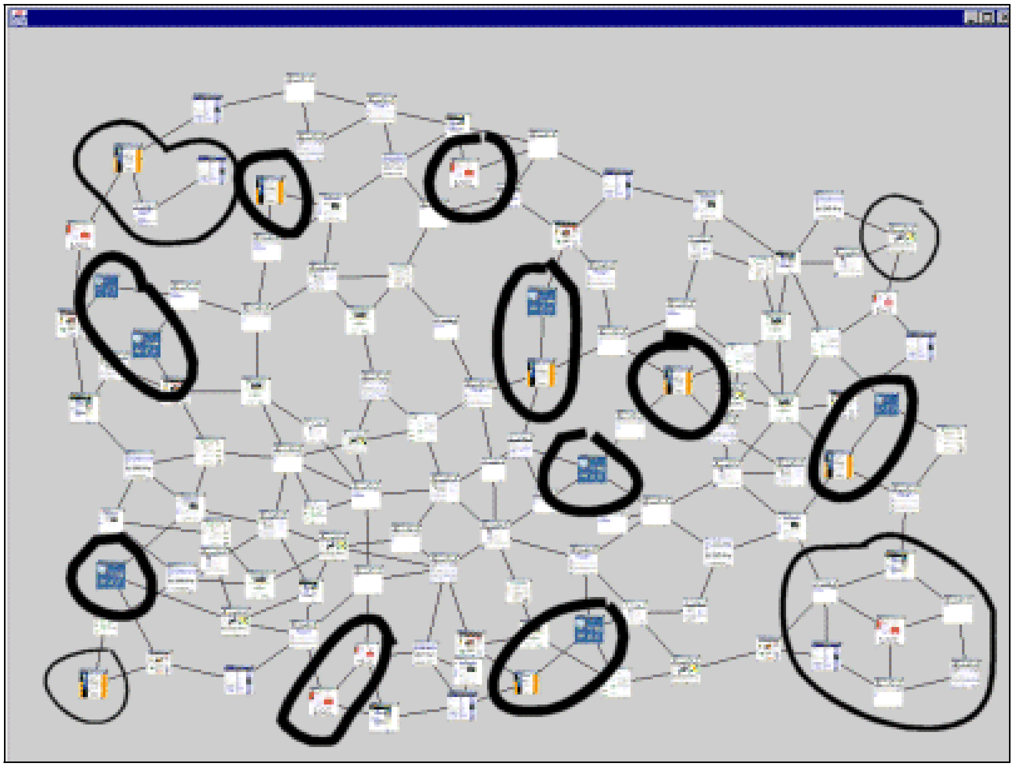
Figure 3.6: Landmarks chosen at d = 0 (no distortion)

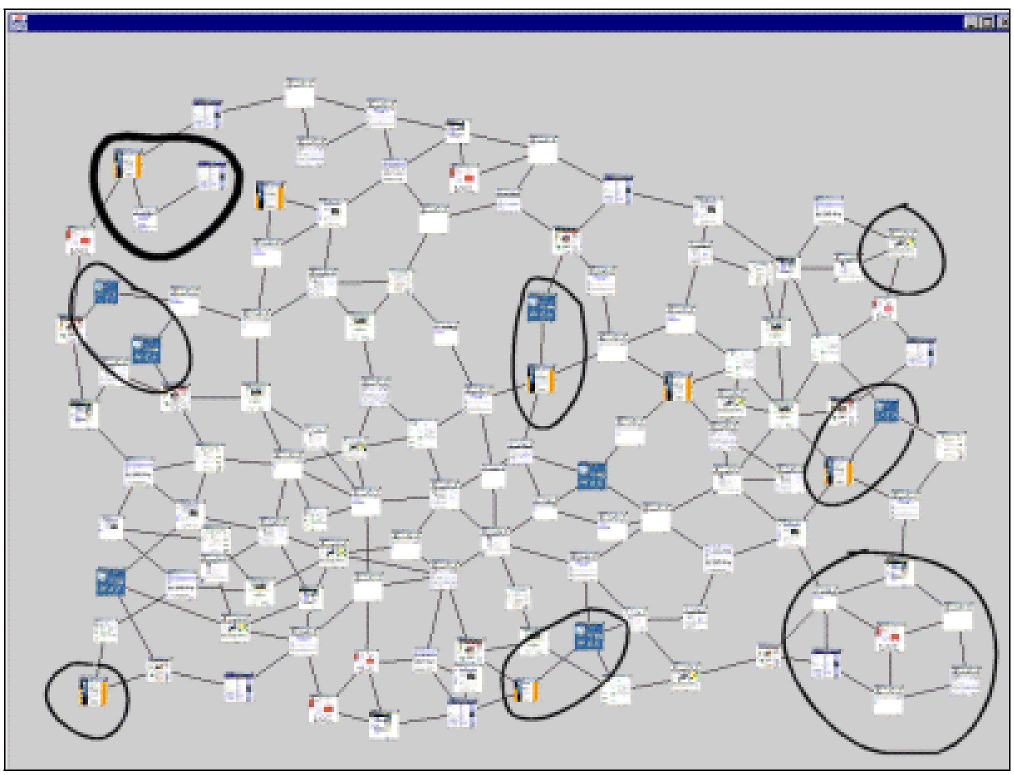
Figure 3.7: Landmarks chosen at d = 1 (note that distortion is not shown)

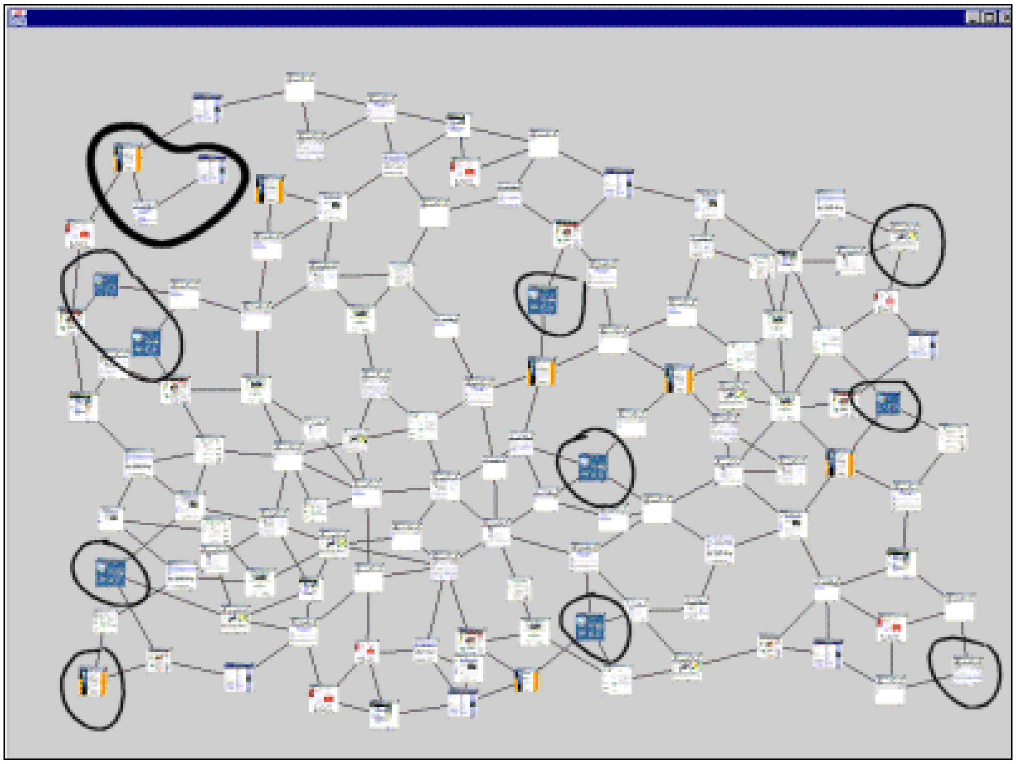

Figure 3.8: Landmarks chosen at d = 3 (note that distortion is not shown)

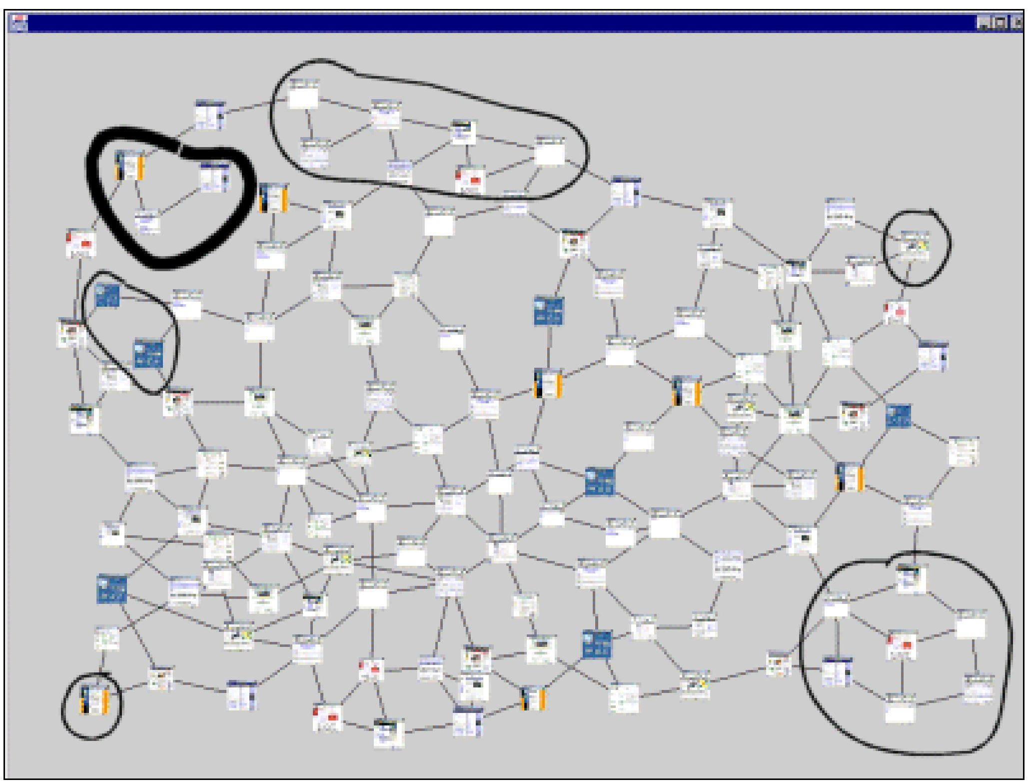

Figure 3.9: Landmarks chosen at d = 5 (note that distortion is not shown)

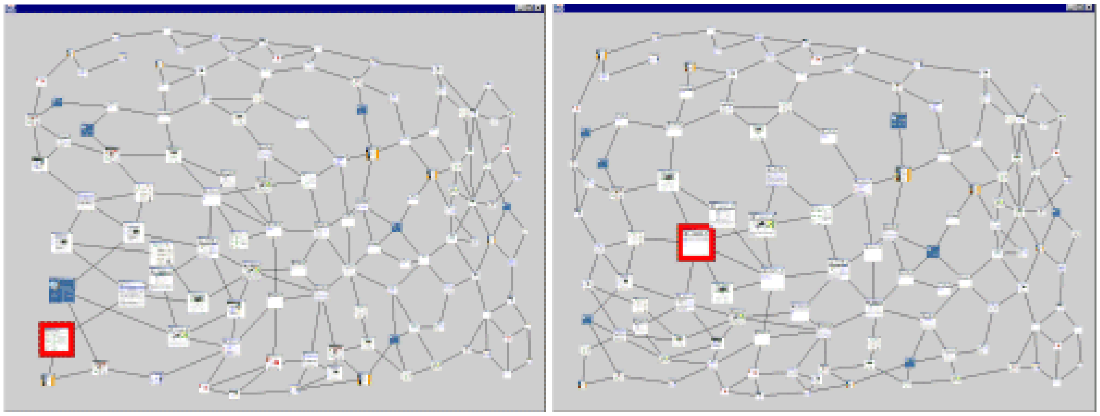

Figure 3.16: Target locations for Tasks 8 (left) and 9

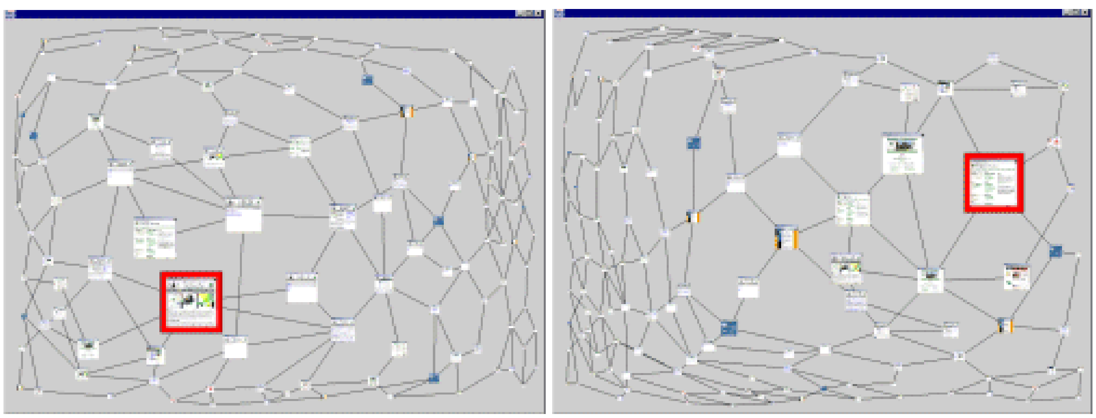

Figure 3.17: Target locations for Tasks 12 (left) and 15

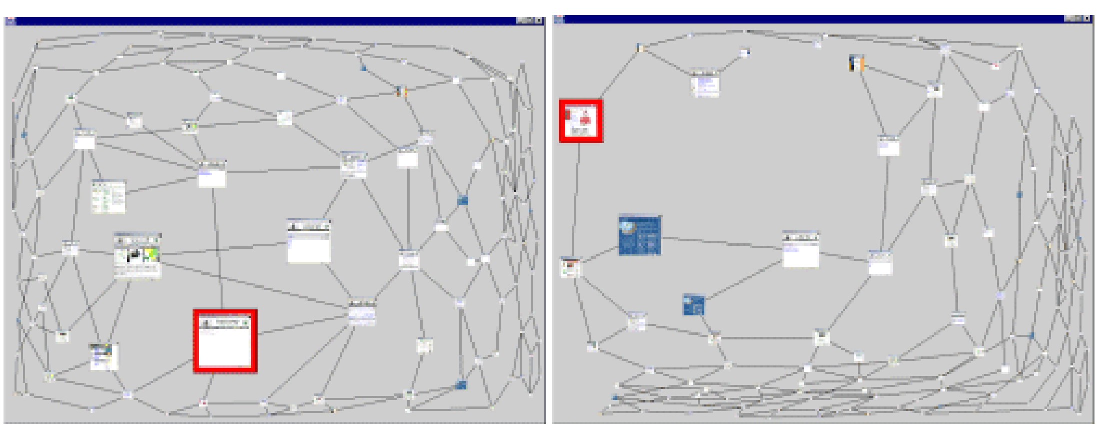

Figure 3.18: Target locations for Tasks 18 (left) and 19

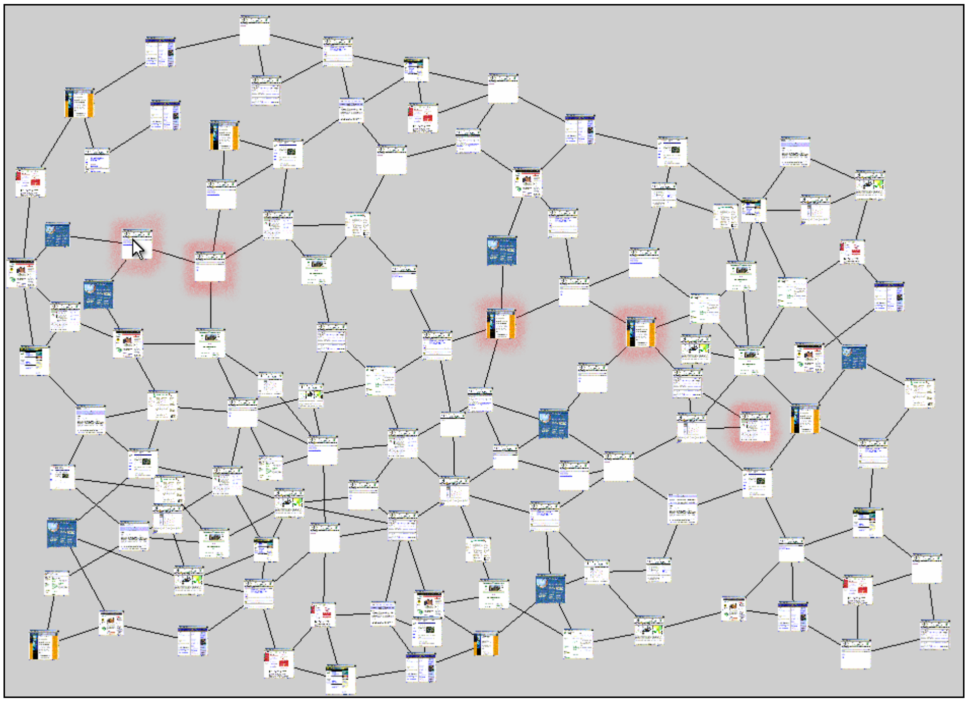

Figure 4.1: Feature-based visit wear in a discrete space. The last five visited nodes are marked with a red halo, though these specifics (“five” and “red halo”) are arbitrary.

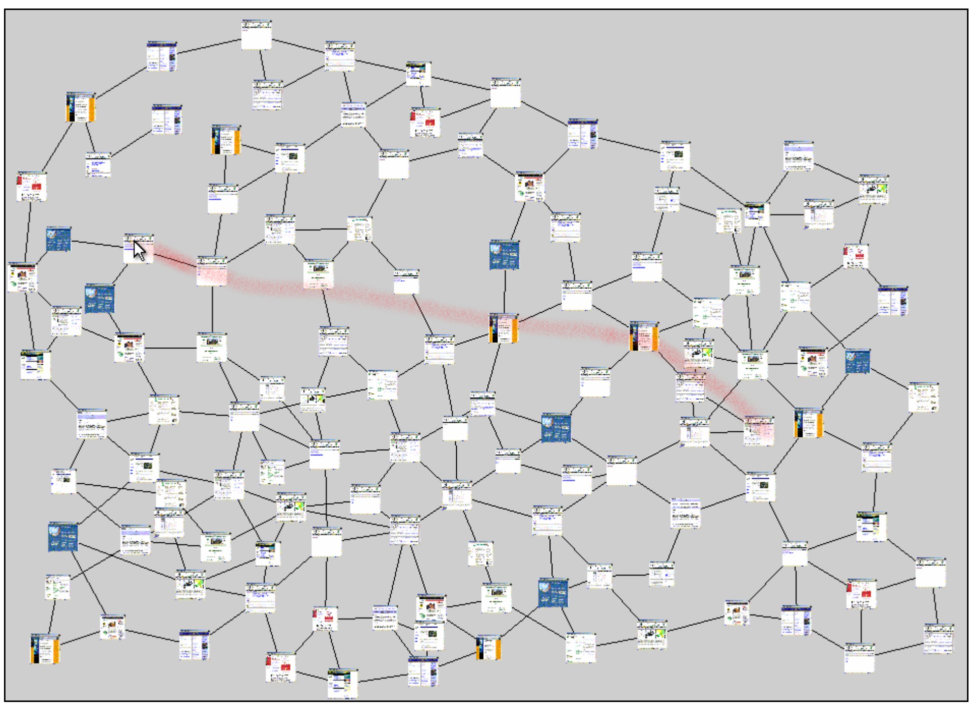

Figure 4.2: Environment based visit wear in a discrete space. The location of the mouse cursor in the recent past is shown, whether it was over a feature or not. The trail is shown in red, but this choice and the choice of the length of the trail are arbitrary.

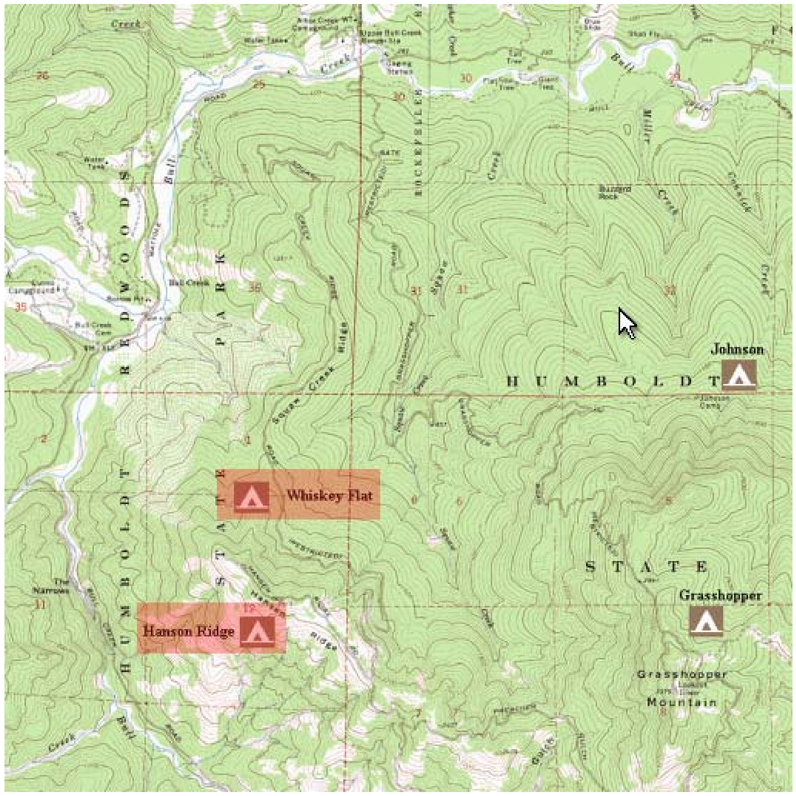

Figure 4.3: Feature-based visit wear in a continuous space. The features have been predefined to be the campsites; no other area of the space will show visit wear

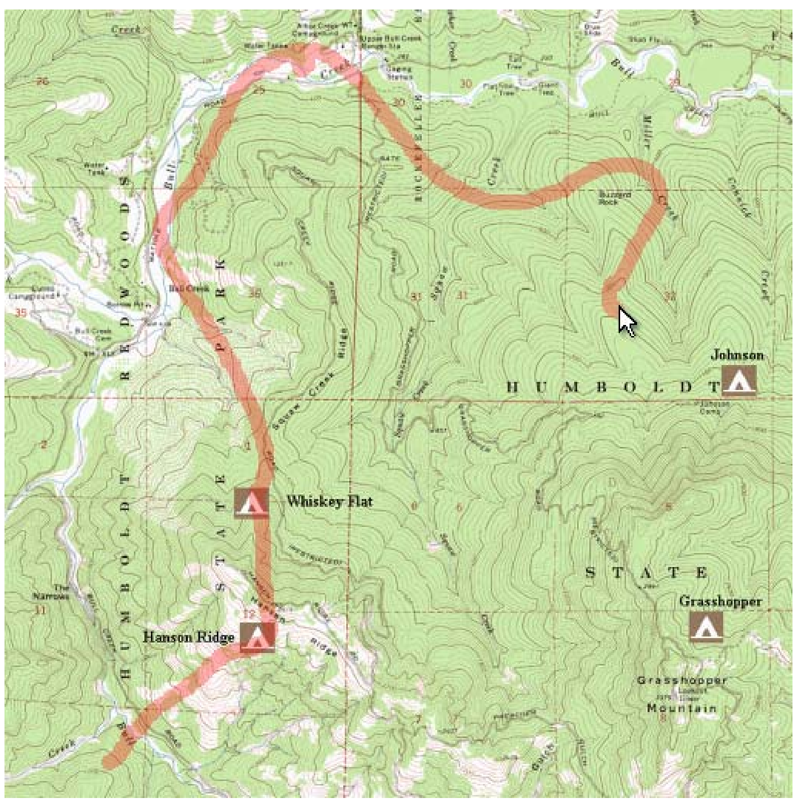

Figure 4.4: Environment based visit wear in a continuous space. No matter what the user considers important in the space, the visiting of it has been shown.

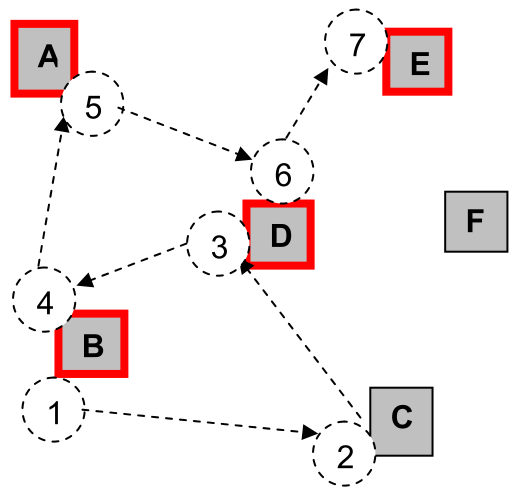

Figure 4.5: An example of automatically calculating visit wear duration.

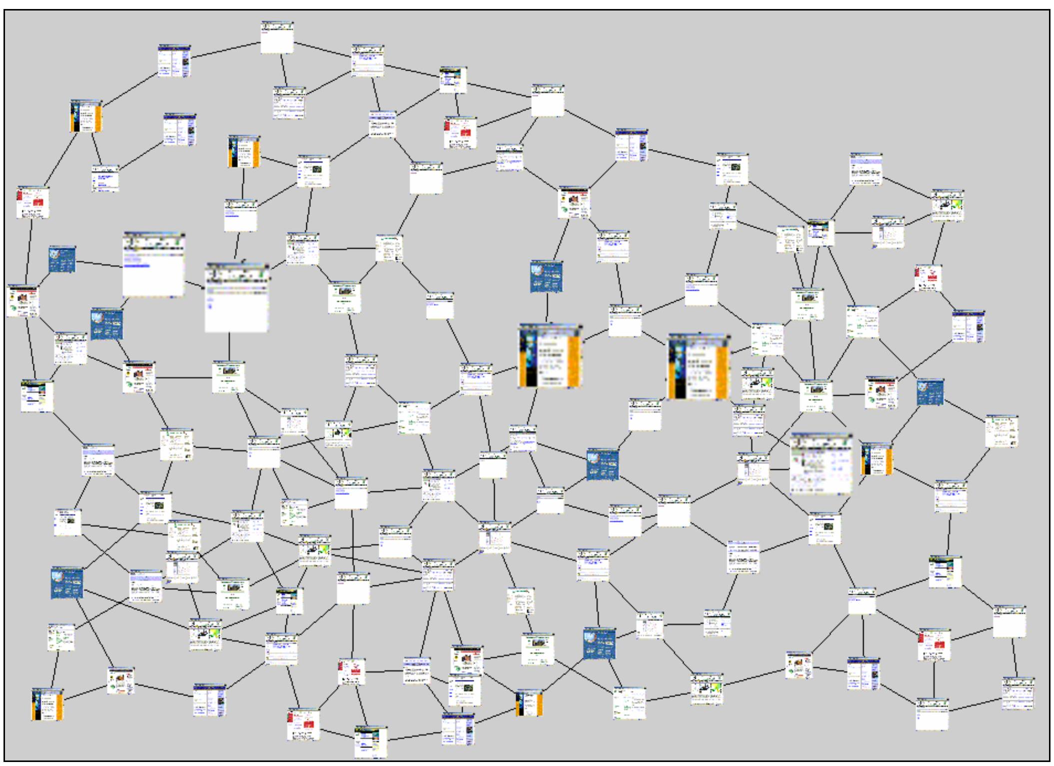

Figure 4.6: Representing visit wear by changing the size of a visited node.

Figure 4.7: Changing the properties of the data to indicate visitation. A) The original data B) Changing colour C) Changing black value

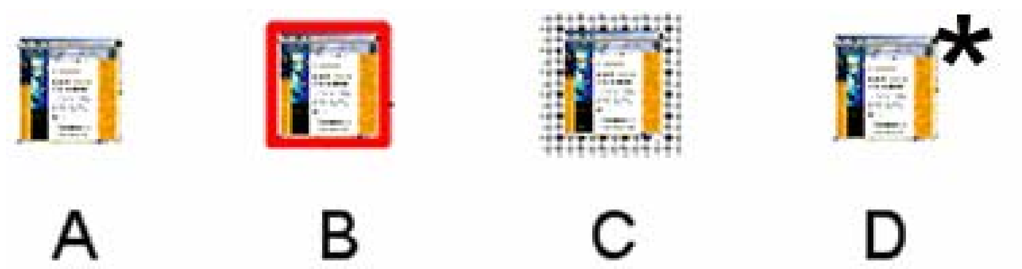

Figure 4.8: Adding a secondary glyph and changing its properties to indicate visitation. A), the original node B) Adding a border and altering its colour C) Adding a border and altering its texture D) Adding a secondary text glyph

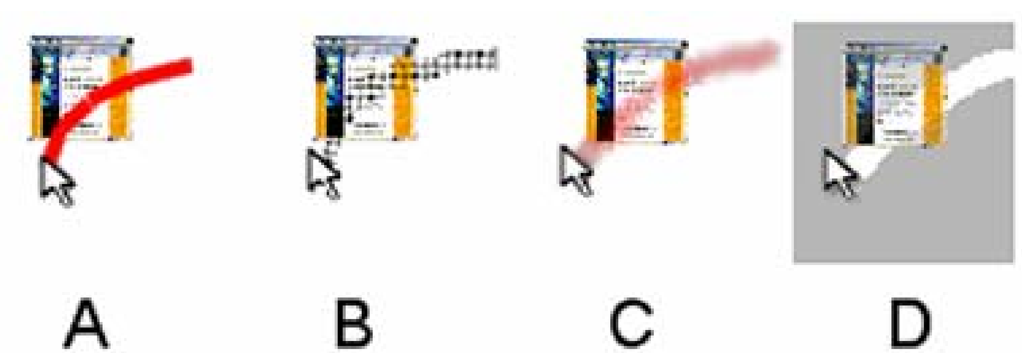

Figure 4.9: Methods of representing visitation by changing A) the colour B), texture C) colour and texture, and D) black value of the mouse trail.

Figure 4.10: Visit wear transparency indicates age of visit. Full opacity (far left) indicates a recent visit, full transparency (far right) means the node either has not been visited or was visited long enough ago that it is not on the history list.

Figure 4.11: Size of visit wear secondary glyph indicates age of visit. Full size (to the far left) indicates a recent visit, zero size (far right) means the node either has not been visited or was visited long enough ago that it is not on the history list.

Figure 4.12: Visit wear colour indicates age of visit. Purple (far left) indicates a recent visit, blue (far right) indicates long ago visit.

Figure 4.13: Additional glyph indicates age of visit. Glyph moves like clock hands, starting at t = 0 to the far left and moving clockwise with age.

Figure 4.14: Change in secondary glyph indicates number of visits (shown below the example node). One visit is shown with a one-pixel wide border, or with one extra glyph. Ten visits are shown with a ten pixel wide border or ten extra glyphs.

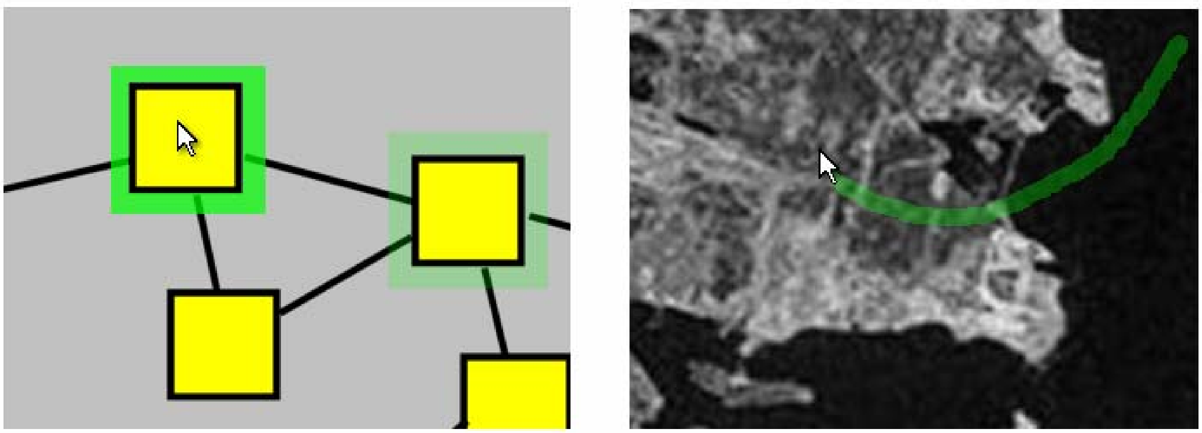

Figure 4.15: Example visit wear effects to be experimentally tested in a discrete space (left) and a continuous space (right)

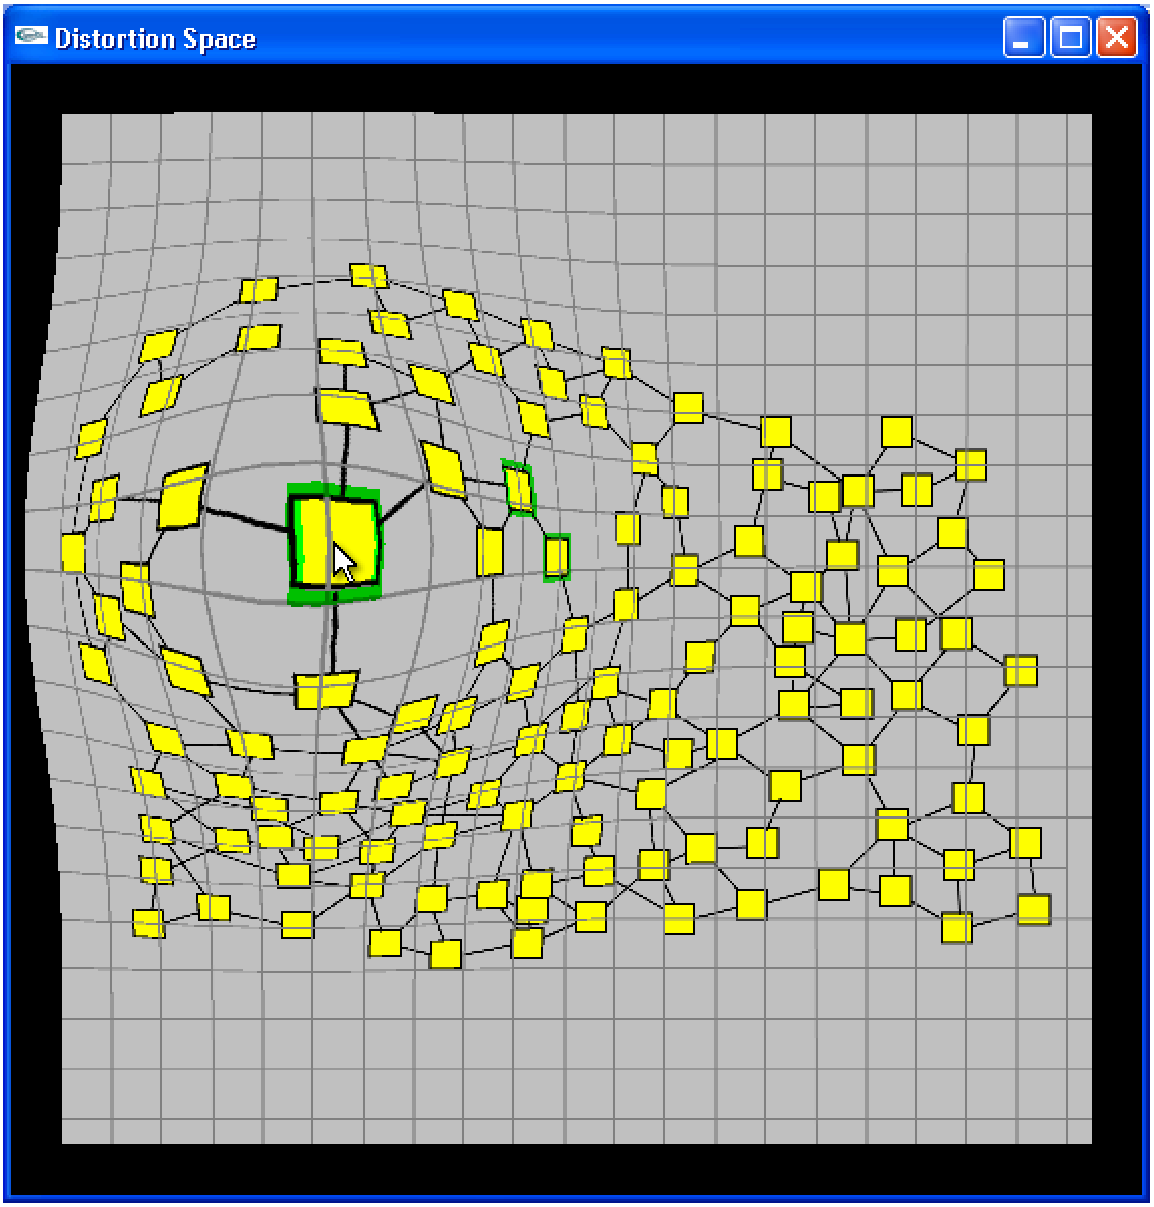

Figure 5.1: The experimental discrete data space, the fisheye lens, and the visit wear effect used in the experiment. The grid is shown to make the fisheye effect clear; it did not appear in the actual experiment.

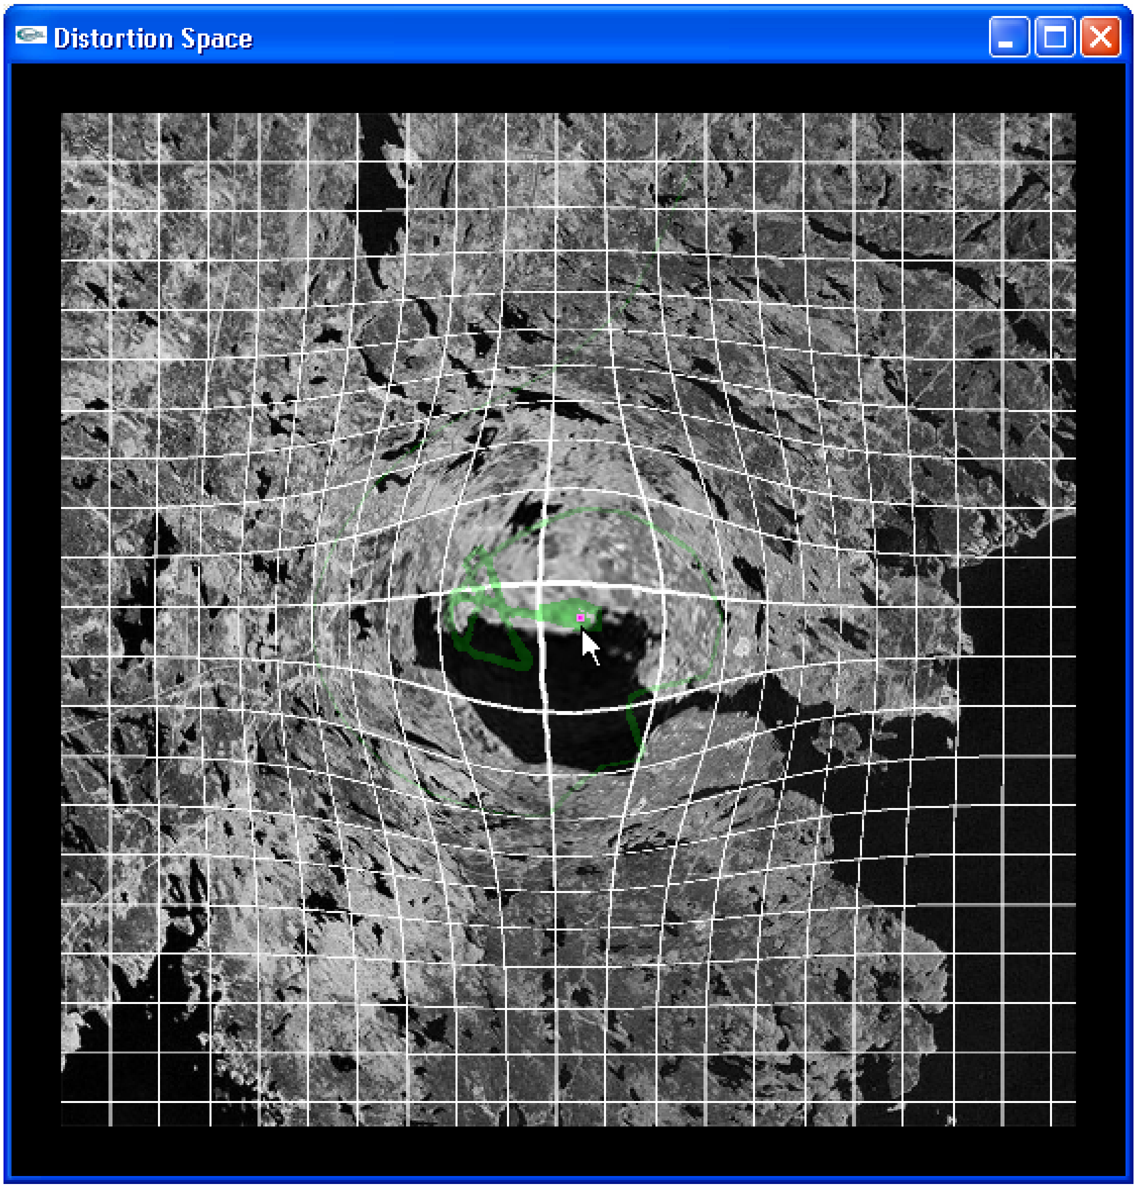

Figure 5.2: The experimental continuous data space, the fisheye lens, and visit wear effect used in the experiment. The grid is shown to make the fisheye effect clear; it did not appear in the actual experiment.

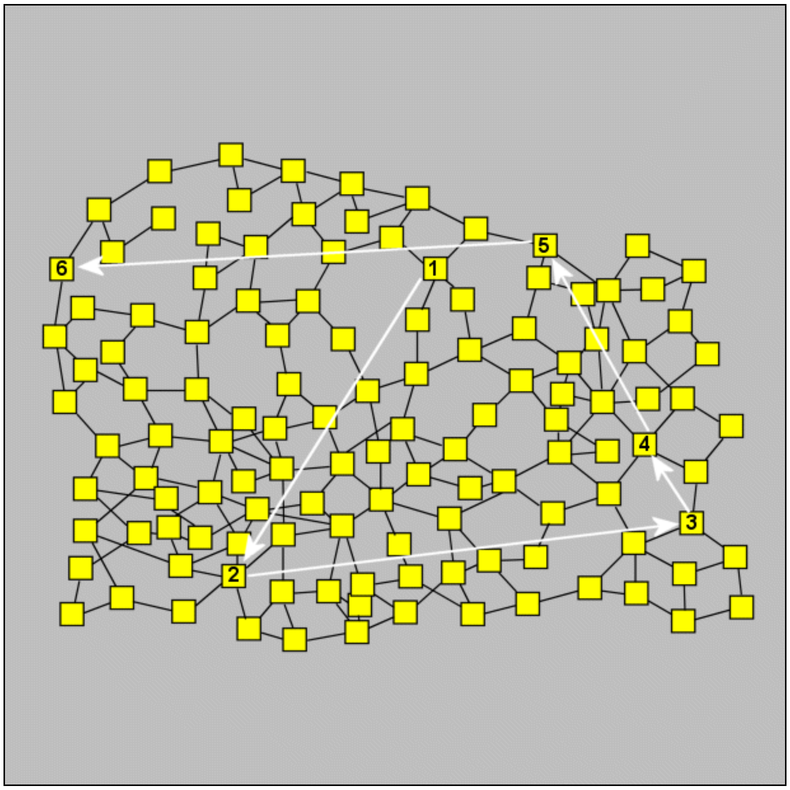

Figure 5.3: Example task sequence, discrete space. White arrows show the subject's expected path.

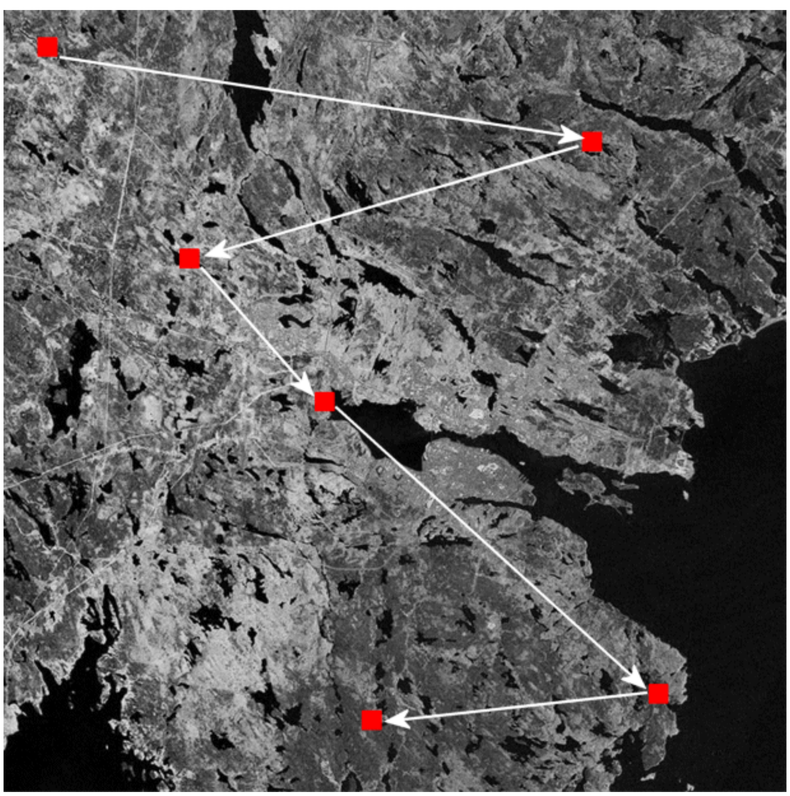

Figure 5.4: Example task for continuous space. Red target squares appeared one at a time; white arrows show subject's expected path between targets

Appendix B – Evaluation Materials

Experiment 1

1. Informed Consent Form
2. Demographic Survey
3. Instructions to Participants

Experiment 2

1. Informed Consent Form
2. Demographic Survey
3. Eals and Silverman Test
4. Instructions to Participants
5. Post-study Questionnaire

Experiment 1 – Informed Consent Form

Research Project:	Fisheye Navigation
Investigators:	Amy Skopik, Department of Computer Science (966-6593)
	Dr. Carl Gutwin, Department of Computer Science (966-8646)

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you a basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, please ask. Please take the time to read this form carefully and to understand any accompanying information.

This study is concerned with evaluating the effects of a fisheye distortion on the navigation of an abstract data space. Your experiences and comments during this study will be used to determine the direction of further experiments.

The session will require about one hour, during which you will be asked to familiarize yourself with the effects of a distortion lens, and carry out some search tasks in the space. At the end of the session, you will be given more information about the purpose and goals of the study, and there will be time for you to ask questions about the research.

The data collected from this study will be used in articles for publication in journals and conference proceedings.

As one way of thanking you for your time, we will be pleased to make available to you a summary of the results of this study once they have been compiled (probably in about three months). This summary will outline the research and discuss our findings and recommendations. This summary will be available on the web at: hci.usask.ca/.

All of the information we collect from you (data logged by the computer, observations made by the experimenters, and your questionnaire responses) will be stored so that your name is not associated with it (using an arbitrary participant number). Any writeups of the data will not include any information that can be linked directly to you. The research materials will be stored with complete security throughout the entire investigation. Do you have any questions about this aspect of the study?

You are free to withdraw from the study at any time without penalty and without losing any promised benefits. Withdrawal from the study will not affect your academic status or your access to services at the university. If you withdraw, your data will be deleted from the study and destroyed. In addition, you are free to not answer specific items or questions on questionnaires. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a participant. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. If you have further questions about this study or your rights as a participant, please contact one of the following:

Dr. Carl Gutwin, Assistant Professor
Department of Computer Science
(306) 966-8646
gutwin@cs.usask.ca

Office of Research Services
University of Saskatchewan
(306) 966-4053

Participant's signature:____

Date:_____________________

Investigator's signature:____

Date:_____________________

A copy of this consent form has been given to you to keep for your records and reference. This research has the ethical approval of the Office of Research Services at the University of Saskatchewan.

Experiment 1 – Demographic Survey

1. Personal Information:

Gender	❑ Male	❑ Female		
Age Range	❑ 16 - 20	❑ 21 - 25	❑ 26 - 30	❑ 31 - 35
	❑ 36 - 40	❑ 41 - 45	❑ 46 - 50	❑ 50+

2. How many hours a week, on average, do you spend working with computers?

   ❑ 0 - 4 ❑ 4 - 12 ❑ 12+

3. Have you ever used a computer program that used a fisheye distortion technique before?

   ❑ Yes ❑ No

4. How many hours a week, on average, do you spend playing computer games?

   ❑ 0 - 4 ❑ 4 - 12 ❑ 12+

5. If any, what type of games? (e.g. first person shooter, sport, strategy, etc):

   ____
   ____
   ____

6. How many hours a week, on average, do you spend working with graphics programs?

   ❑ 0 - 4 ❑ 4 - 12 ❑ 12+

7. If any, what type of programs (e.g. CAD, draw/paint, rendering, etc):

   ____
   ____
   ____

Experiment 1 – Instructions to Participants

1) Setup; get window to starting point, get snapshot of starting screen

2) Consent form (two copies), describe experiment, reiterate “stop at any time”

3) Questionnaire; write trial number

4) Exploration; landmarks are features of graph that are easy to find again, easy to describe to others, points of reference for less distinguished parts of graph. Look at graph for 30 secs – 1 minute, describe five landmarks and why they are landmarks. Describe areas that are difficult to navigate and why.

   Then explore; click on green nodes to turn them normal.

   Describe fisheye distortion, note that what’s under the mouse looks the biggest, father away is smaller. Does this affect landmark judgement? Repeat for three distortion levels, new landmarks each time

5) Describe search tasks, node will turn red, then new view of graph will be shown. Click on the node that was red. Understand? Estimate ease of finding node on red screen, give confidence level afterwards.

6) Save file in fisheye program.

7) Thank subject. Any questions?

Experiment 2 - Informed Consent Form

Research Project: Improving Memorability in Fisheye Visualizations
Investigators: Amy Skopik, Department of Computer Science (966-6593)
Dr. Carl Gutwin, Department of Computer Science (966-8646)

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you a basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, please ask. Please take the time to read this form carefully and to understand any accompanying information.

This study is concerned with evaluating the effects of adding historical information to a fisheye visualization of an abstract data space. Your experiences and comments during this study will be used to determine the direction of further experiments.

The session will require about one hour, during which you will be asked to familiarize yourself with the effects of a distortion lens, and carry out some search tasks in the space. At the end of the session, you will be given more information about the purpose and goals of the study, and there will be time for you to ask questions about the research.

The data collected from this study will be used in a Masters' thesis, as well as in articles for publication in journals and conference proceedings. The data will be stored for five years and then destroyed.

As one way of thanking you for your time, we will be pleased to make available to you a summary of the results of this study once they have been compiled (probably in about three months). This summary will outline the research and discuss our findings and recommendations. This summary will be available on the web at: hci.usask.ca/.

All of the information we collect from you (data logged by the computer, observations made by the experimenters, and your questionnaire responses) will be stored so that your name is not associated with it (using an arbitrary participant number). Any writeups of the data will not include any information that can be linked directly to you. The research materials will be stored with complete security throughout the entire investigation. Do you have any questions about this aspect of the study?

You are free to withdraw from the study at any time without penalty and without losing any promised benefits, including the $10.00. Withdrawal from the study will not affect your academic status or your access to services at the university. If you withdraw, your data will be deleted from the study and destroyed. In addition, you are free to not answer specific items or questions on questionnaires. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a participant. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. If you have further questions about this study or your rights as a participant, please contact one of the following:

Dr. Carl Gutwin, Assistant Professor
Department of Computer Science
(306) 966-8646
gutwin@cs.usask.ca

Office of Research Services
University of Saskatchewan
(306) 966-4053

Participant's signature:____

Date:____________________

Investigator's signature:____

Date: ____________________

A copy of this consent form has been given to you to keep for your records and reference. This study has been approved on ethical grounds by the University of Saskatchewan Behavioural Sciences Research Ethics Board.

Experiment 2 - Demographic Survey

2. Personal Information:

Gender	❑ Male	❑ Female		
Age Range	❑ 16 - 20	❑ 21 - 25	❑ 26 - 30	❑ 31 - 35
	❑ 36 - 40	❑ 41 - 45	❑ 46 - 50	❑ 50+

2. How many hours a week, on average, do you spend working with computers?

❑ 0 - 4 ❑ 4 - 12 ❑ 12+

3. Have you ever used a computer program that used a fisheye distortion technique before?

❑ Yes ❑ No

4. How many hours a week, on average, do you spend playing computer games?

❑ 0 - 4 ❑ 4 - 12 ❑ 12+

5. If any, what type of games? (e.g. first person shooter, sport, strategy, etc):

____
____
____

6. How many hours a week, on average, do you spend working with graphics programs?

❑ 0 - 4 ❑ 4 - 12 ❑ 12+

7. If any, what type of programs (e.g. CAD, draw/paint, rendering, etc):

____
____
____

Experiment 2 - Eals and Silverman Test

This test is designed to measure differences in location memory. The test consists of an array of twenty-seven “familiar” objects and a second array where some of the objects are switched from their original positions. The image arrays are printed on 11”x17” paper.

The participant is shown one of the arrays first; it does not matter which one. After one minute of inspection, the first array is removed and the second one is shown to the participant. The participant then has as long as he or she needs to identify which objects have moved and which have not. Scoring is done by counting how many objects the participant has correctly identified as having moved or not moved; the maximum score is therefore twenty-seven.

The two image arrays are shown on the following page.

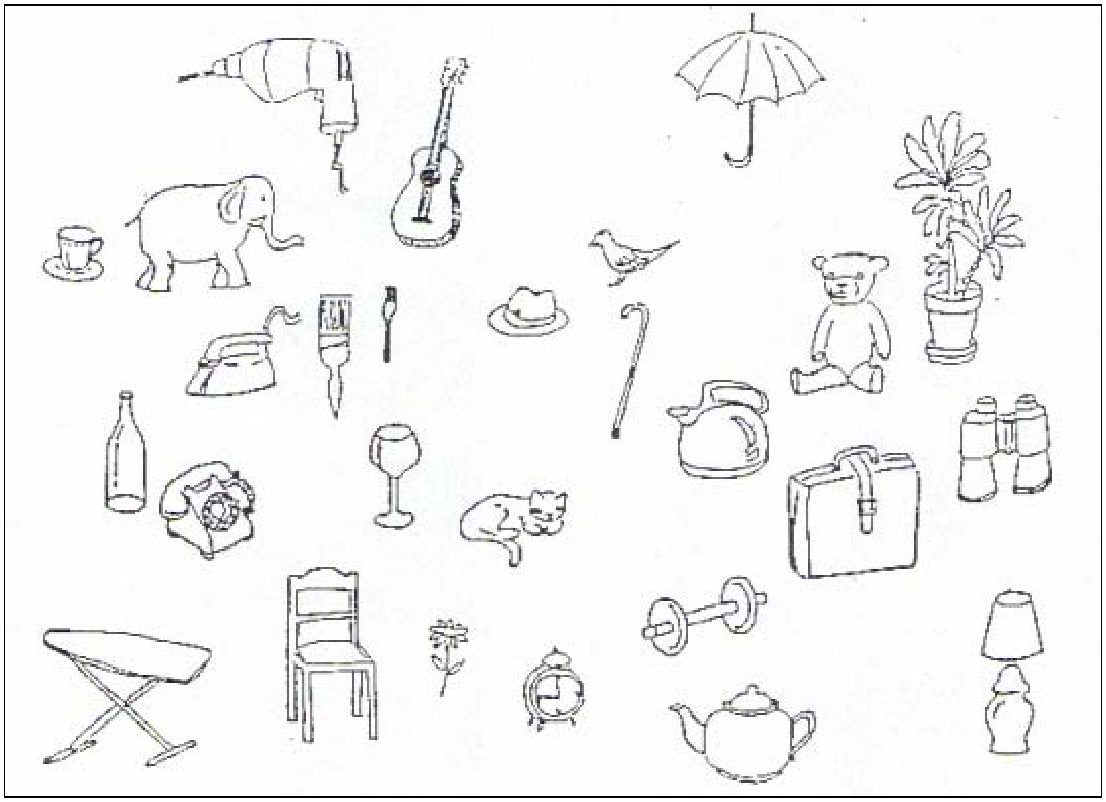

Experiment 2 - Instructions to Participants

1. First, I'd like you to read and sign this consent form, and ask me if you have any questions about it. The consent form assures you that your data will be stored anonymously and securely, and that you can quit the experiment at any time if you're at all uncomfortable.
2. Now, please fill out this short demographic questionnaire. Ask if you have any questions. A fisheye lens is a view of information that distorts the presentation so that what's in focus is shown larger than what's not in focus, with a smooth transition between.
3. Now you'll do a spatial memory task that cognitive psychologists have used for years. I'll give you a piece of paper with pictures of some objects on it. You'll have one minute to study the paper. At the end of the minute, I'll give you another piece of paper with the same objects on it, but some of them will have changed position. On this second piece of paper, I'd like you to circle the objects that have *NOT* changed, and put an 'x' through the objects that *HAVE* changed. You'll have as much time as you like to do this, and when you're done every object should either be circled or x'd.
4. Now you'll do some similar tasks but on a computer and with a fisheye lens. As I said, a fisheye lens is a view of information that distorts the presentation so that what's in focus is shown larger than what's not in focus, with a smooth transition between. We're looking at what happens when we add "visit-wear" to a fisheye view. Visit-wear ("wear" as in "wear out" or "wear down") is adding visual information automatically to a presentation to show how you've interact with it. In the real world, things wear out when they're handled or visited a lot, so we're seeing what happens when computer things do the same thing.

    a. This information space, a graph, is called "discrete" because the information is in certain areas (the yellow squares) while the grey background doesn't contain any information. The visit wear in this case is highlighting graph nodes that you visit. When you spend time on a node, the green halo darkens around it. When you move off the node, its halo starts fading. The fading happens more slowly than the darkening, and if you move quickly enough over a node, the darkening doesn't happen at all. Feel free to get familiar with how the fisheye and visit wear work

       The memory task will start when you're ready. There will be two practice trials first, one with visit wear and one without. In both cases the task is the same. One of the nodes in the graph will have the number '1' appear on it. Move the mouse to that node. When you've spent a few seconds on that node, another node will have the number 2 appear on it. Move to that node. This will continue until six nodes are numbered 1 to 6. After a few seconds on Node 6, all the labels will disappear. Then I want you to click on those same six nodes, in the same order that they were shown to you, 1 to 6.

You'll have forty seconds to click on the six nodes, or I can skip to the next trial when you're ready. In between trials, the screen will be blank and purple. The trials won't continue until you're ready, so if you need to take a break, the blank purple screen is the best place to do it. Are you ready to start the practice trials?

b. This information space, a map, is called "continuous" because every pixel represents something and contains information. The visit wear in this case is drawing a trail that shows where your mouse has been. The trail starts fading the instant that it's drawn, so that if you leave the mouse still for long enough, the trail disappears completely. Feel free to get familiar with how the fisheye and visit wear work

The memory task will start when you're ready. There will be two practice trials first, one with visit wear and one without. In both cases the task is the same. A red square will appear on the screen. Move the mouse to the red square, and it will then move to another location on the screen. You can move to the square however you like. After the square has appeared in six places, I want you to imagine that the path that you have just traveled between the squares was a path that you walked. And you've just discovered that you dropped your car keys during the walk. You know that they keys are somewhere that you have been. They will be in a pixel that your mouse has passed over. The keys are a small magenta square; note how you can only see them when you are fairly close, but you do not have to be right over top of them. When you find your car keys, click on them (or in their area, don't' bother trying to get exactly on them).

You'll have forty seconds to find the car keys, or I can skip to the next trial when you're ready. In between trials, the screen will be blank and purple. The trials won't continue until you're ready, so if you need to take a break, the blank purple screen is the best place to do it. Are you ready to start the practice trials?

5. Now there is a brief post-task questionnaire. "Visit wear" means the green halo or mouse trail, the "graph" was the yellow boxes, and the "map" was the find-the-keys task.
6. Thank you for participating! Print and sign your name, fill in the date, and here's your $10.

Experiment 2 - Post-Experiment Questionnaire

1. In the graph, which method (visit wear or no visit wear) do you think let you carry out the tasks best?

Main reasons why:

2. In the map, which method (visit wear or no visit wear) do you think let you carry out the tasks best?

Main reasons why:

3. What strategies did you use to remember things if there was no visit wear?

4. Other comments: