Improving Revisitation in Fisheye Views with Visit Wear

Amy Skopik and Carl Gutwin
Computer Science Department, University of Saskatchewan
110 Science Place, Saskatoon, SK, S7N 5C9, Canada
amy.skopik, carl.gutwin @usask.ca

ABSTRACT
The distortion caused by an interactive fisheye lens can make it difficult for people to remember items and locations in the data space. In this paper we introduce the idea of visit wear – a visual representation of the places that the user has previously visited – as a way to improve navigation in spaces affected by distortion. We outline the design dimensions of visit wear, and report on two studies. The first shows that increasing the distortion of a fisheye view does significantly reduce people's ability to remember object locations. The second study looks at the effects of visit wear on performance in revisitation tasks, and shows that both completion time and error rates are significantly improved when visit wear is present. Visit wear works by changing the revisitation problem from one of memory to one of visual search. Although there are limitations to the technique, visit wear has the potential to substantially improve the usability both of fisheye views and of graphical information spaces more generally.

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

ACM Classification Keywords
H.5. Information interfaces and presentation: User interfaces (*Interaction styles, screen design*)

INTRODUCTION
Distortion-based visualization techniques, such as fisheye views, are a solution to the problem of showing both focus and context in a large data space (e.g., [10,15]). These techniques are often interactive, allowing users to explore the data by moving the system's focus point with the mouse. Unfortunately, the distortion effects of interactive focus+context views can cause usability problem, and one problem in particular is that of *memorability*. The memorability of a data space is the degree to which people can remember where things are and where they have been.

Usually, people are good at remembering the locations of objects in a spatial environment [4,14]. In one study, for example, participants were able to remember the locations of web-page thumbnails even several months after they had originally positioned them [4]. However, interactive fisheye views can disrupt this natural spatial ability, because as the magnifying lens moves around the space, it changes object positions, distorts local patterns, and alters the orientation of objects relative to one another.

In this paper we look at ways to address the problem of memorability in fisheye views. We introduce the idea of *visit wear*, a type of read wear [8] that explicitly marks the places that a person has visited in a data space. Using a person's own interaction history as an aid to memory is based on the observation that revisitation is common in many tasks – for example, most of the web pages that people go to have been visited before, and usually fairly recently [21].

In the following sections, we explore the problem of fisheye memorability in more detail, and introduce the concept and design space of visit wear. We also report on two studies, the first done to show how increasing distortion impairs spatial memory, and the second to assess the effect of visit wear on memory tasks in a discrete data set (a graph) and a continuous space (an aerial photo). The second study showed that when visit wear was present, people completed revisitation tasks significantly faster and with significantly fewer errors. Our observations indicate that visit wear turns a memory problem into a perceptual task that is easier to solve. Although there are some limitations to the idea, our results suggest that visit wear could substantially improve the usability of fisheye views, and possibly the usability of other visual representations as well.

FISHEYE VIEWS
A graphical fisheye view is a focus+context technique characterized by a smooth transition between a magnified focus region and a demagnified context area [3,10,15] (see Figures 2,5,7). Most fisheyes used in visualization tasks are interactive, in that they allow the user to vary the level of magnification and to move the focus point with a pointing device such as the mouse. When the user moves the focus over a particular item, they can see details that are not visible when the item is unmagnified.

There are several ways that lenses can vary: in extent, magnification mechanism, and shape [3]. The amount of movement and alteration caused by the fisheye lens, and thus the amount of potential disruption to a user's spatial memory, depends on the lens type and the distortion level. In the studies described below, we used two different lens types: a full-screen lens that shows distortion in all areas of the view, and a constrained round lens that limits the effect to a fixed radius around the focus.

Fisheye views have been tested in several different task contexts (e.g., command and control, web navigation, and menu selection), and have been shown to provide performance benefits over more traditional interaction techniques [1,9,13,16]. However, the distortion of a fisheye has also been shown to cause problems in targeting, steering, and layout tasks [5,6,7]. Perhaps due to these usability problems, fisheye views have not gained wide acceptance in commercial systems aside from a few examples (e.g., www.idelix.com).

SPATIAL MEMORY

In the real world, there are a variety of ways that people remember spatial locations. As people learn a space they remember specific routes, then landmarks, and finally a 'mental map' (*survey knowledge*) of the entire region [11,17]. This knowledge is built up primarily through interaction; that is, people remember locations after having had experience with that location [22].

Previous studies have shown that although abilities can vary widely [18], people are capable of using spatial memory to remember large numbers of items. For example, Robertson and colleagues [14] tested a spatial memory technique (the Data Mountain) against several other methods of recalling web pages. Retrieval of web pages was significantly faster with the spatial technique than with a standard bookmarking system. The spatial memory also persisted over a long time: participants who returned for a follow-up study six months later were able to retrieve items at the same performance level, with only a brief retraining period [4].

However, researchers have also shown that mental maps are difficult to maintain when the data space changes in appearance. Misue and colleagues [12] state that changing the layout of a graph severely compromises people's ability to remember it, even when the topology is not changed (see Figure 1). This is relevant to fisheye memorability since fisheyes change the layout of items in the data space: although qualities such as topology and left-to-right ordering are preserved, fisheyes move items and change the angles formed between objects.

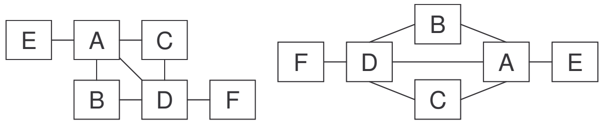

Figure 1. Two layouts of the same graph (from [12]).

READ WEAR AND INTERACTION HISTORIES

Our solution to the problem of memorability is based on Hill and colleagues' work on information physics and interaction histories [8]. Their goal is to "record on computational objects…the events that comprise their use…and display useful graphical abstractions of the accrued histories as part of the objects themselves" ([8], p. 3). Visit wear is essentially a particular type of Hill and colleagues' *read wear*.

It is fundamental to the idea of computational wear that information about the user's actions will be recorded automatically and will be displayed within the item itself. This sets visit wear apart from two well-known revisitation mechanisms: history lists and bookmarking techniques. History lists record interaction history automatically, but are presented separately from the data itself, which forces users to work in two different navigation systems. Manual bookmarking techniques are valuable in some situations (and could be used as an adjunct to visit wear), but there are two reasons to pursue an automatic solution as well: first, explicit marking requires user effort; and second, people often do not know they will want to revisit an item during the initial visit (and so will not think to create a bookmark).

Wear-based visualizations are one way to enrich the fabric of human interaction with computational spaces. Wear can help us see patterns of use, either by ourselves or others, and in a form that is easy to interpret, since people are already familiar with these types of clues [23]. Automatic recording of activity has analogies in the real world, such as footprints on a beach, well-thumbed pages in a book, or paths through a field. There are fewer examples in the computational world, although one is now common: web browsers implement a simple form of visit wear by automatically showing visited links in a different colour.

REVISITATION

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 (for the remainder of this paper, we will assume that the goal is to support revisitation). Different situations have different revisitation patterns, but in most information tasks, many of the items inspected will be revisitations of those already seen.

One well-known example of revisitation is navigation on the WWW: Tauscher and Greenberg [21] found that more than half of pages seen were revisits, and that revisitation occurs mainly to the last few pages visited – the last ten pages seen cover about 85% of revisits. This implies that a recency-based history mechanism can be useful in presenting interaction history and supporting navigation.

With these basics in mind, we now turn to the present research questions: whether distortion impairs spatial memory, how visit wear can be built to assist revisitation, and whether visit wear actually does improve people's ability to find previously-visited features of distorted space.

STUDY 1: EFFECT OF DISTORTION ON MEMORABILITY

Fisheye views move the apparent XY positions of objects in the data space. The location of an object is dependent upon the position of the focus and the amount of magnification, so an object moves to different places depending on where the user is focusing.

This movement can change characteristics of a data set that people use to orient themselves and remember locations. In a graph, for example, even though orthogonal ordering and topology are unchanged, the nodes' positions, the length of edges, and the edge slopes are all altered by the fisheye lens. This means that both the nodes' positions, and many of the memorable shapes formed by nodes and edges, can change as the focus moves.

To investigate whether these differences in appearance would affect people's ability to find nodes in a graph, we carried out the first of two studies.

Method – study 1

Twelve people who were regular users of mouse-and-windows software were recruited from a local university. None of the participants had previous experience with interactive fisheye views. The participants were asked to carry out a number of before-and-after memory tasks in a custom-built fisheye graph browser (see Figure 2). The system used a full-screen fisheye lens implemented with the Sarkar and Brown polar-coordinates algorithm [15].

After an initial introduction to the system and the task, participants carried out tasks in the fisheye system at three different levels of distortion (d = 1, 3, and 5). At each level, participants were first allowed free exploration in the system for two minutes. They then carried out five before-and-after-style memory tasks (two practice and three test).

Each task used two views of the data, a 'before' view and an 'after' view. In the 'before' view, one node in the graph was highlighted with a red border. The participant was asked to memorize the location of the node, using as much time as they needed. The system then showed the participant the 'after' picture, in which the node highlight was removed and the fisheye focus had shifted to another location. The participant was asked to find and click on the target node as quickly and accurately as they could; during this stage of the task, the focus was fixed, and so the graph did not change as the participant found and selected the target node. The before-and-after approach simulates a situation where a user inspects a particular node, then turns their attention elsewhere for a moment, and then needs to find the node again.

The specific nodes (and focus shifts) that were used in the task were chosen such that each participant saw nine different situations, but over the entire experiment all situations were seen equally at each level of distortion. The study used a 1x3 within-subjects factorial design. The single factor was Distortion Level (d=1, 3, or 5). Dependent variables were completion time (time to find the node in the 'after' picture) and number of errors.

We recognize that there is a minor confound in this design, in that the task completion time is made up of a combination of finding time and targeting time. However, the effect of this confound is limited: in our study the focus was fixed during targeting, so that motion effects of the magnifying lens would not occur (as in [6]); this leaves only the smaller target size at higher distortion as a potential difference between low and high distortion conditions. It was clear from our observations that targeting was a relatively minor part of the overall completion time – that is, participants spent most of their time looking for the node, and very little time selecting it.

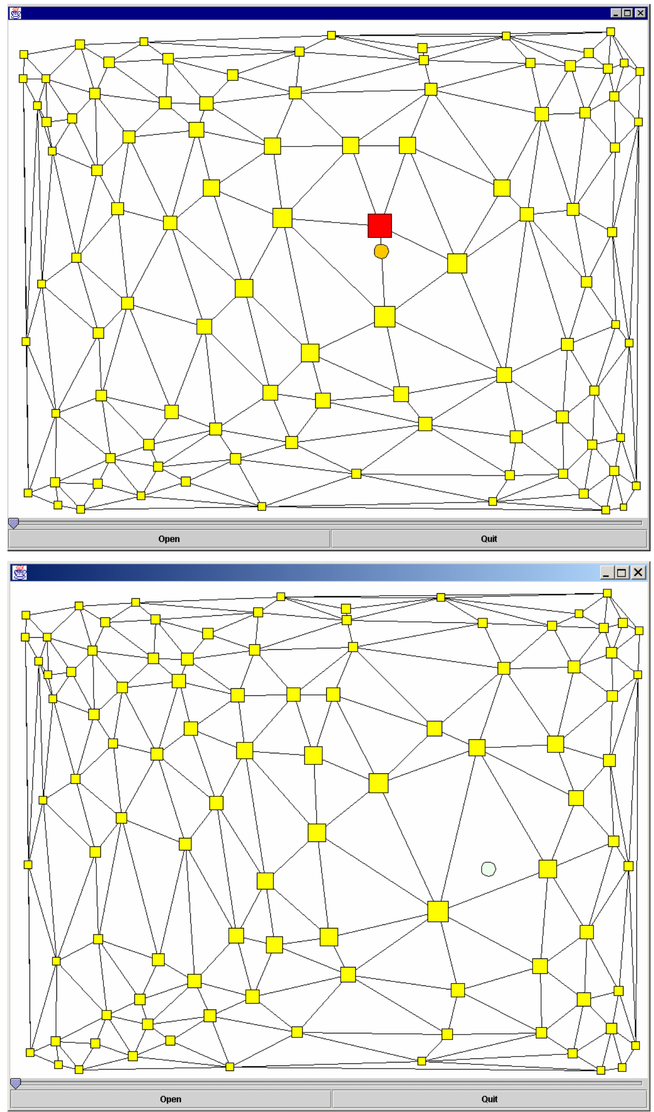

Figure 2. Example task at distortion level 1, with 'before' view above and 'after' view below. The target node is coloured red in the 'before' picture. The filled circle indicates the location of the fisheye focus.

Results – study 1

Increasing the distortion of the fisheye view did affect people's ability to remember and find the target nodes. Task performance is made up of two factors, completion time and error rate. First, a 1x3 ANOVA showed a significant main effect of distortion level on completion time ($F_{2,10}$=11.72, $p<0.001$). As shown in Figure 3, completion times at d=3 were about two seconds longer than at d=1, and were more than six seconds longer at d=5.

Errors were relatively infrequent at all levels of distortion, and no effect of distortion level was found on error rate ($F_{2,10}$=0.20, p=0.82).

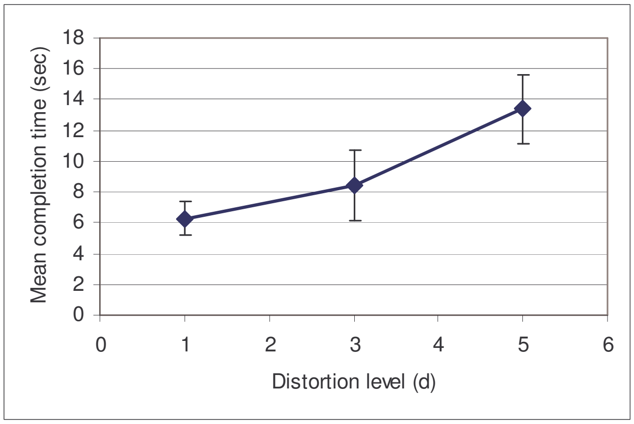

Figure 3. Mean time to find the target node, in seconds, for three levels of distortion. Bars show standard deviation.

Memory Strategies and Landmark Use

People initially attempted to use the strategy of simply remembering the absolute spatial location of the target node. This strategy quickly failed when there was distortion; even during training, people realized that they would have to remember the targets based on information other than absolute position. In particular, people started to use landmarks in the graph – either shapes formed by nodes and edges, or a corner or boundary of the graph.

In fact, some nodes seemed to be quite easy to remember, even at high levels of distortion. Nodes that were at the corners, for example, were always found quickly. Memorability, therefore, seemed to depend on the presence of a useful landmark, and so the memorability of different nodes varied greatly. Even though landmarks overall became more difficult to use as distortion increased, they seemed to be a promising strategy for remembering where things were in fisheye views.

However, it appeared that the landmarks people used were too dependent on the structure of the graph, and were adversely affected by increasing distortion. This observation provided the inspiration to look at adding artificial landmarks to the data, visual indicators that could serve as external memory aids for people as they worked with the data, aids that would be resilient to the changes caused by the fisheye effect. This idea, coupled with the concept of revisitation, leads to visit wear.

VISIT WEAR

Visit wear is a visual effect that adds information to a data space based on the user's activity in the space. Visit wear is a type of read wear that focuses on the idea of visitation – items or locations in the space undergo visual changes once the user has visited that location.

We have identified four main factors that must be considered in the implementation of visit wear: visit recognition, continuity, duration, and appearance.

Visit Recognition

The designer must determine what will constitute a 'visit' to a particular item or location. This determination will be dependent on the task and the underlying data; the designer's goal is to ensure that revisitations are adequately supported without unduly adding visual clutter that distracts the user from their primary task. Criteria for recognizing a visit in a fisheye viewer could include:

- an item is in the focus of the fisheye for a certain amount of time;
- the user carries out an action with an item (e.g., opening a folder, selecting a node, or editing an item);
- the user moves their mouse across an item or location (to be used for techniques such as a cursor trail).

In addition to the initial recognition of a visit, the designer must determine whether visited items can return to the 'unvisited' set for any reason such as the passage of time (see discussion of duration below).

Continuity

Visit wear can be either discrete (restricted to identifiable features) or continuous (recorded and displayed in all parts or pixels of the data space). In most cases continuity is determined by the data and the task: if the data is primarily made up of discrete features on a background (e.g., a graph) then a discrete visit-wear effect is likely to be used. If the data is more dense (e.g., a map or photograph), where any location might be a feature, then a continuous effect may be more appropriate. Note however that for particular tasks, continuous data spaces can be discretized into a set of distinct features (e.g., towns and roads on an aerial photo).

Duration

Duration refers to the time span in which the visit wear effect is visible. Duration may be thought of in terms of elapsed time or visitation sequence, and is also dependent on the task. If the task is to inspect all the items in a set, then visit wear should remain visible throughout the session; if the task involves revisitation, however, then a shorter duration is possible. Shorter durations also limit the number of items that are shown as 'visited' at any one time, reducing clutter.

In many real-world revisitation tasks, an appropriate duration can be calculated automatically. In this scheme, the designer provides a target proportion of revisits that are to be 'caught' with visit wear (i.e., where revisited items are already showing a visit wear effect), and the system dynamically adjusts the duration to meet the target. For example, if a user visits items in the sequence A-B-C-D-A, and the target is to catch all revisits with visit wear, then the duration should be the last four items in the sequence.

The possibility of automatic calculation shows the relationship between visit wear and the revisitation pattern of the task. In tasks where users primarily look back at the last few items (e.g., browsing the WWW), then a short duration is sufficient. In tasks where users need to revisit items much further back, then many more items need to show visit wear at the same time. In some cases, recency may not be a useful measure of revisitation at all; in these situations, other measures such as frequency could possibly be used instead.

Appearance

When considering the appearance of a visit wear effect, the designer must consider where the visualization will be located, what variables will be represented, and how the information will be encoded.

First, there are two main ways to position a visit wear visualization: the most common approach is to add a secondary glyph to a data item (e.g., a border, an overlay, or a separate object); another means is to use the data item itself (although this may alter the meaning of the data).

Second, there are several pieces of information that can be represented in visit wear. We consider two variables to be common to all implementations:

- *membership*: whether an item or location has been visited or not; and
- *age*: the amount of time (or number of visits to other items) since an item was last visited.

Other variables (not considered further here) could include visitation sequence, frequency of visitation, and the identity of the visitor in a multi-user environment.

Third, the information may be mapped to different visual qualities, and may be represented in a variety of ways. Visual properties that can carry a visit wear representation include shape, orientation, colour, texture, black value, size and position of a secondary glyph [2]; in addition, there are other properties such as motion and animation that can encode information in computational spaces.

It is important that the effect remain distinctive even under distortion. Previous work suggests that there are visual properties that are distinctive enough to be landmarks in a fisheye view, and that these are also resilient to distortion. When people were asked how they remembered things in a distorted space, properties like colour and shape were much more likely to be used as landmarks than absolute position or patterns in the graph's topology [20].

From these four basic building blocks, many different visit wear visualizations can be created. As an example, Figure 4 shows three discrete visualizations and indicates how visitation and age have been represented. A more complete catalogue of possible effects is available elsewhere [19].

We used these design principles to develop two visit wear techniques for testing – one version for discrete data sets and one for continuous spaces. Our next step was to investigate whether these techniques work: whether they improve revisitation performance in fisheye views.

Recent visit	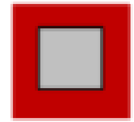	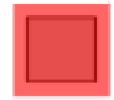	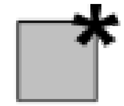
Older visit	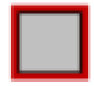	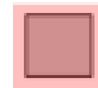	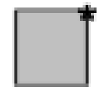
Visitation:	Border highlight	Transparent overlay	Adjacent mark
Age:	Width of border	Transparency level	Size of mark

Figure 4. Examples of discrete visit wear

STUDY 2: EFFECTS OF VISIT WEAR ON REVISITATION

We carried out a second study to investigate the effects of visit wear. The study used two different information spaces, a discrete space (a graph) and a continuous space (an aerial photo). The study used feature-based visit wear for the discrete space and a trail-based implementation for the continuous space. In addition to looking at performance, we also considered whether the visit wear caused problems for the task through clutter or occlusion of the task data.

Method

Participants

Sixteen participants (eight men and eight women) were recruited from a local university. Participants ranged in age from 21 to 45 years and averaged 27 years. All were familiar with mouse-and-windows applications (more than 12 hours/week), and most had experience with computer games (average of 4.5 hours/week). Two participants had seen an interactive fisheye previously.

Apparatus

A custom fisheye viewer was built in C++, using the EPS library [3] and OpenGL. The study was run on a P4 Windows system with an 800x600 screen. Visit wear was implemented by changing the RGB values of the source image – either a predefined border for the discrete data space, or the mouse cursor trail for the continuous space. The application also tracked all items currently affected by visit wear so that the effect could be faded correctly.

Experimental Conditions

In the discrete data space, the visit wear effect appeared as a green halo around the graph node (see Figure 5). The halo deepened in intensity whenever the mouse was over the node, 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 away, the halo started fading and would disappear in one minute.

In the continuous space, visit wear appeared as a green transparent trail 'drawn' by the mouse as it moved (Figure 7). The trail faded completely over a one-minute period (these times were used only for the experimental tasks; real-world effects would persist differently, as discussed above).

Tasks and Datasets

The tasks used in the study represent abstract versions of real-world tasks such as inspection, exploration, and path-finding. Tasks were divided into two phases: an initial visitation phase and a revisitation phase. In those conditions were visit wear was turned on, the visit wear effects were visible in both phases of the task.

Discrete dataset. Participants were asked to move the focus to each of six targets that were labeled one after another at three-second intervals. After the participant had focused on all six nodes, the labels disappeared and the revisitation part of the task began. In this phase, participants were asked to select the nodes that had been labeled, in order, as quickly and accurately as possible.

The task was repeated for six different sets of six nodes. The node sets were selected to have approximately the same total path length, and with similar a priori memorability (due to the nodes' locations in the graph). The node sets were the same for all conditions, although different participants saw them in different orders.

Continuous dataset. In the visitation phase, participants were asked to move between a series of waypoints shown to them one at a time (Figure 8). Participants were asked to observe the path that they took between the waypoints and remember the way that they traveled. After all waypoints had been reached, the participant was asked to revisit their path to find a target marked with a pink two-pixel square dot. The target was placed halfway along the participant's previously-visited path (participants were not told this, and nor was it easy to determine). The target was small enough that it could only be seen with the fisheye's magnification.

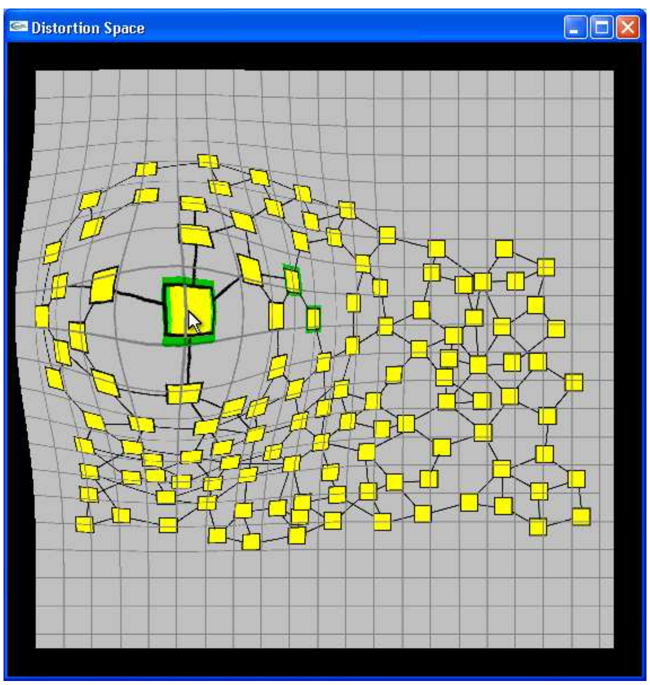

Figure 5: The graph used in the study, showing visit wear on the node under the cursor. Gridlines here illustrate the fisheye effect; they did not appear in the study system.

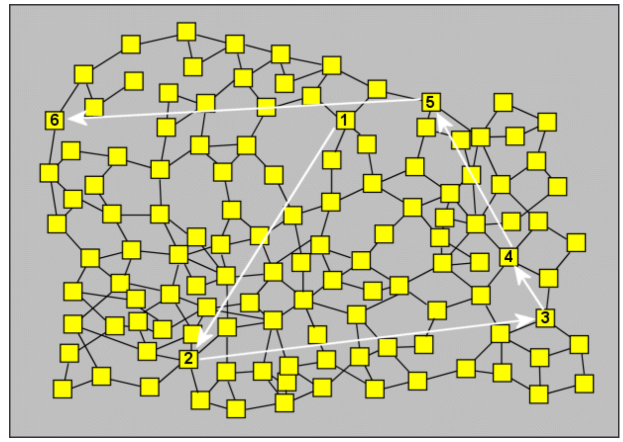

Figure 6: Example node sequence for discrete data set.

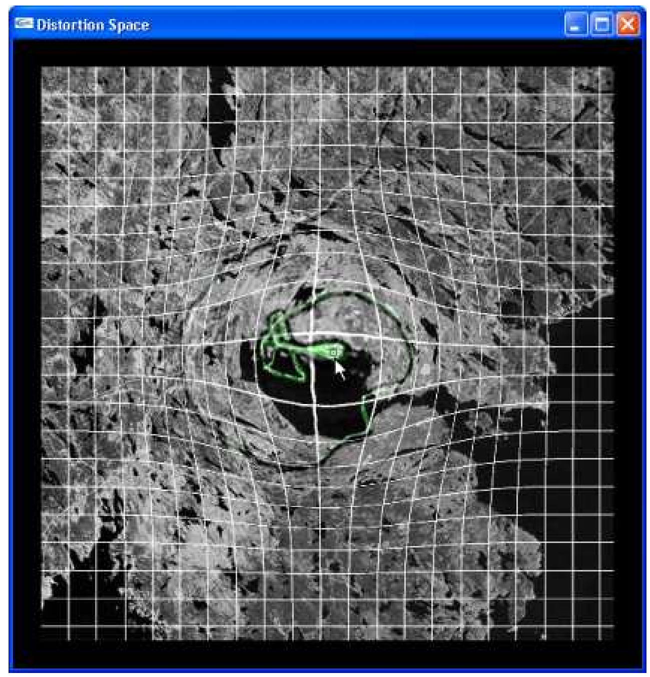

Figure 7: Continuous data space with fisheye lens, visit wear (green trail), and target (pink dot at cursor point). The trail has been intensified here for visibility. Gridlines did not appear in the actual experiment.

This task was again repeated for six different sets of waypoints. The waypoints were selected to give path lengths of roughly the same size (approximately 1240 screen pixels assuming straight lines from waypoint to waypoint). An example task (showing all waypoints at once) is displayed in Figure 8.

All tasks had a 40 second time limit (determined from pilot studies as a reasonable upper bound); if a participant had not finished the task in that time, the system moved on to the next task.

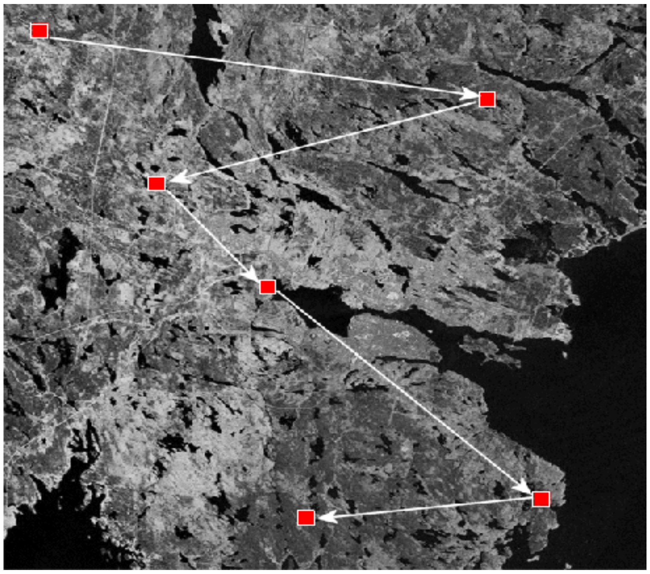

Figure 8: Example task for continuous dataset. In the task, waypoints (red blocks) appeared one at a time.

Procedure

Participants were pre-tested for spatial memory using the Silverman and Eals object location memory test [18], and were divided post-hoc into two groups (above and below the median).

Participants were then randomly placed into an order group and asked to complete practice and test tasks in the different experimental conditions. All participants did all tasks under all conditions. Order was balanced by randomizing each participant's starting point within the task sequence; all conditions were seen in each position in the sequence the same number of times. Participants were allowed to rest between tasks.

After all tasks had been completed, participants filled out a questionnaire based on their experiences. Questions asked which technique (visit wear or no visit wear) they preferred, what strategies people used to remember things, and whether the presence of visit wear distracted them or caused them problems in completing the task.

Study Design

The study used a 2x2x2 mixed factorial design, and analyses were carried out separately for the two different tasks. The within-participants factor was visit wear (present or not present); between-participants factors were sex (male or female) and prior spatial ability (low or high). Dependent variables were completion time and error rate.

The system collected data about mouse movement and button clicks. With six tasks in each condition and sixteen participants, there was a total of 384 data points collected.

Results

Since differences between datasets (discrete or continuous) were expected, the analysis was carried out individually for each task. Tests were carried out to look for effects of visit wear, sex, and prior spatial ability; interactions between these factors were also examined, although the small number of participants per cell (four) reduces the strength of the interaction tests. The following sections provide results for completion time and errors, and then summarize results of the post-task questionnaire.

Completion Time

We analysed the effects of visit wear, sex, and prior spatial ability on task completion time. Figure 9 shows mean completion times for both tasks. For the discrete (graph) dataset, analysis of variance (ANOVA) showed that task completion times were significantly reduced when visit wear was present ($F_{1,15}$=17.00, $p<0.005$). The average time to complete the task with visit wear was 10.2 seconds, compared to 15.1 seconds without visit wear.

For the continuous (aerial photo) dataset, completion time was also significantly lower when visit wear was present ($F_{1,15}$= 11.42, $p<0.005$). The average time with visit wear was 25.1 seconds, and 33.1 seconds without visit wear.

Participants were divided into two groups based on prior spatial ability ('lower' and 'higher'). Analysis of variance found no main effect of spatial ability on completion time (see Table 1); however, there was an interaction between spatial ability and visit wear presence for the discrete dataset. In this task, people in the higher-spatial group improved more with the addition of visit wear than did those in the lower-spatial group. No effects of sex were found on completion time (see Table 1).

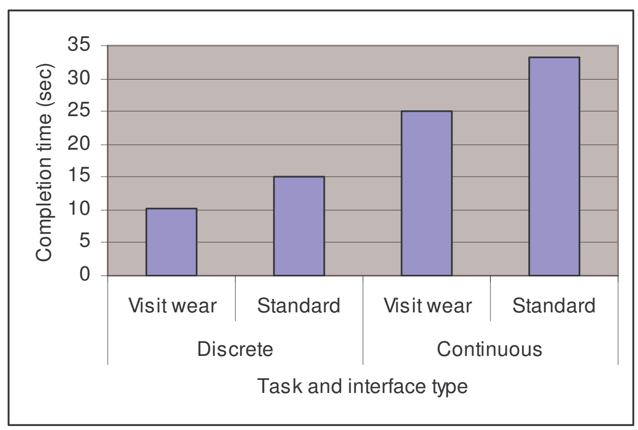

Figure 9. Mean completion times for visit wear and standard interfaces, across discrete and continuous datasets.

	Discrete	Continuous
Visit wear	**$F_{1,15}$=17.0, p<0.005**	**$F_{1,15}$= 11.42, p<0.005**
Sex	$F_{1,15}$=0.66, p=0.43	$F_{1,15}$=0.03, p=0.87
Sex x visit wear	$F_{1,15}$=0.87, p=0.37	$F_{1,15}$=0.02, p=0.88
Prior spatial	$F_{1,15}$=0.22, p=0.65	$F_{1,15}$=0.14, p=0.72
Prior x visit wear	**$F_{1,15}$=6.96, p<0.05**	$F_{1,15}$=0.37, p=0.55

Table 1. ANOVA results for tests of sex and prior spatial ability on completion time (main effects and interaction with visit wear).

Errors

For the discrete dataset, errors were counted in two ways: first, the number of erroneous nodes selected during revisitation (ID errors); and second, the number of our-of-order nodes in the revisitation sequence (order errors). As can be seen in Table 2, ID errors without visit wear were approximately three nodes in every six-node task, but with visit wear, errors were much less than one per task (see also Figure 10).

Analysis using ANOVA showed a significant reduction in ID errors when visit wear was present ($F_{1,15}$=160.69, p<0.001). There was no difference between the conditions in terms of ordering errors.

	No visit wear	Visit wear
ID errors per task	2.99 (50%)	0.27 (5%)
Order errors per task	0.66 (11%)	0.86 (14%)

Table 2. Mean errors in the discrete dataset, per 6-node task. (Order errors are normalized by the participant's number of correct IDs for that task).

For the continuous dataset, error was measured as the number of targets that were not found within the 40-second time period. Again, the presence of visit wear showed a significant reduction in error ($F_{1,15}$=14.70, p<0.005). As can be seen in Figure 10, subjects missed 22% of the targets when visit wear was present; without visit wear, subjects missed 58% of the targets.

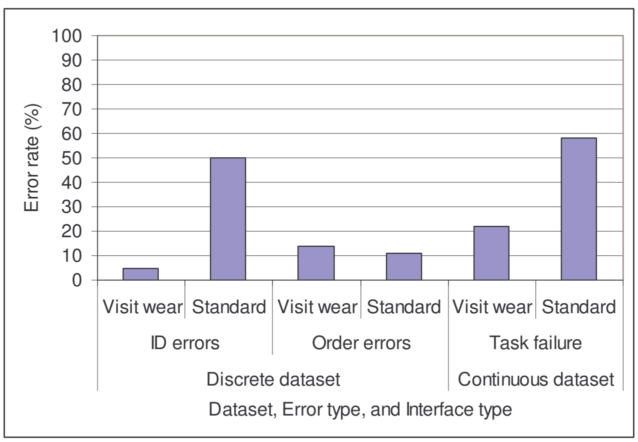

Figure 10: Error rates in discrete and continuous datasets.

No main effects of sex or prior spatial ability were found on error rates for either task (see Table 3), and neither were any interactions with visit wear found.

	Discrete	Continuous
Visit wear	**$F_{1,15}$=160.69, p<0.001**	**$F_{1,15}$ = 14.70, p<0.005**
Sex	$F_{1,15}$=1.19, p=0.29	$F_{1,15}$=0.00, p=1.00
Sex x visit wear	$F_{1,15}$=01.20, p=0.29	$F_{1,15}$=0.44, p=0.52
Prior spatial	$F_{1,15}$=0.001, p=0.97	$F_{1,15}$=0.07, p=0.80
Prior x visit wear	$F_{1,15}$=0.18, p=0.68	$F_{1,15}$=0.80, p=0.39

Table 3. ANOVA results for main effects of sex and prior spatial ability on errors, and for interactions with visit wear.

Preferences

In a post-experiment questionnaire, participants were asked whether they preferred the addition of visit wear in each dataset, 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 to 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 target was easier to see without the coloured mouse trail around it. Even subjects who preferred the visit wear said that they sometimes found themselves concentrating too hard on following the trail and therefore missed the target.

DISCUSSION

Our results show the potential of visit wear for improving people's ability to find items and locations that they have previously inspected. In this section, we suggest explanations for our findings, consider possible limitations to the technique, discuss how the results will generalize to real-world situations, and summarize the lessons that can be taken away by designers of distortion-oriented views.

Explanations of findings

It seemed clear from our observations and from the comments of the participants that visit wear changed the task from a pure memory task to one where memory could be augmented by visual search. With visit wear present in the view, participants did not actually have to remember much information about item locations, since they could find the nodes of interest by inspecting the visual features. This means that visit wear improved revisitation by reducing the need for memory.

Visual markers on visited items are useful because short-term spatial memory is difficult to maintain while carrying out other spatial tasks (that is, new spatial information overwrites the old). Of course, the user still has to remember what the visit wear means, but this type of remembering is apparently not as volatile. The value of these external markers is clear, not just from our study, but from the many real-world examples of people explicitly adding markers – bookmarks, flags, blazes, or breadcrumbs – to help them remember locations. The only difference with visit wear is that the markers are managed automatically; this is valuable because people often forget to place markers, or don't realize that they might want to revisit a location later.

The comments of the participants lend weight to this interpretation. 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 approach (some of which are based in this translation from spatial memory to visual perception), the study showed that visit wear can be extremely effective in some situations.

Design issues

In order for visit wear to be effective in real world visualization tasks, several design issues should be considered: possible occlusion of the task data or of previous visit wear, problems caused by the fading of the effect, the difficulty of determining order, and interpretation costs of the added visual information.

Occlusion of task data or previous visit wear

Visit wear adds visual information to the workspace. Ideally, this information enhances task performance, but it is possible that the added pixels could cause clutter, particularly if the visit wear obscures the task data. In the study system, we tried to avoid this problem by drawing the visit wear effect outside the nodes (in the discrete data set) and by merging the trail with the underlying data in the continuous dataset rather than drawing it on top.

However, in continuous spaces, the colour change caused by adding the trail could become an problem; some of the participants who disliked the visit wear in the continuous space stated that the target was easier to see on the grey photo than it was on the green visit wear. Occlusion will be more of a problem in continuous datasets than in discrete ones. Allowing users to adjust the transparency level of the visit wear effect, or to remove it temporarily, could solve this problem.

It is also possible for a trail-based technique to clutter itself, where a new path occludes an older trail that the user would still like to remember. As one person stated, "the development of new paths in retracing made visit wear confusing." Customization may again help here, as might the dynamic method of determining duration (discussed above). In particular, if revisitation most often occurs to items a few places (or moments) back in the history list, the most recent pieces of a new trail could be hidden until they age sufficiently to be useful to the user.

Problems caused by the effect fading away

If the visit wear effect disappears over time, which is necessary to avoid clutter, situations will arise where users want to revisit items or locations that are no longer marked. If users are relying completely on the visit wear, they may not have recourse to other strategies for finding the items. This problem did occur a few times in the experiment (where the effect duration was arbitrarily determined and fairly short). If participants were slow in moving from target to target, the visit wear effect could have faded from the first node by the time they needed to revisit it. Most participants realized this during training and adjusted their speed so that the visit wear effect did not disappear completely.

Three questions raised here are deserving of further study. The first is the actual frequency of the disappearing-effect problem in real-world use, and whether a duration that is automatically adjusted to the user's revisitation pattern can keep the problem at a low level. The second question is whether visit wear history could be controlled by the user (e.g., a 'history slider') to uncover and show different periods of visitation.

The third question is more general, and asks whether external memory aids like visit wear help or hinder people from forming a longer-term 'mental map' of the data set. Although visit wear could be used too much as a crutch, it is also possible that spatial memory is still exercised when a user revisits items marked by visit wear, so users could still gain survey knowledge of the space.

Determining order

In the discrete dataset, fading over time was used to indicate the age of the last visit. Participants had to use this information to figure out what order they had visited the nodes. However, when the nodes were not side by side, participants found it difficult to accurately determine the order positions of nodes (e.g., the visual difference between nodes visited 20 and 25 seconds ago was slight).

This problem might be solved with a more distinguishable encoding for age (e.g., a dot that moves around the node border), or even an explicit encoding of order (e.g., with numbers). However, an alternative approach is to add a cursor trail (used in the continuous data set) to the discrete data. Although we initially considered that a continuous visit wear effect was superfluous in a discrete space, such a trail could show visitation order better than other representations. This is a case where the user's interaction with the space is continuous (i.e., their path from one node to another), even when the data itself is discrete.

Interpretation of visit wear in visually heterogeneous spaces

The usability of visit wear depends on it having properties that are distinctive with respect to the visual properties of the information around it. Two problems could arise here: the data could already contain the colours or shapes used for visit wear, or other secondary variables could be encoded along with the visit wear using similar visual strategies.

First, the data itself (e.g. web page thumbnails) may already use some of the properties that are available for a visit wear effect. It is possible that image analysis could be used to suggest a set of visual properties that are maximally different from the data; also, allowing the user to customize the effect could solve the problem. Second, some data sets encode other variables with visual properties (e.g., selection or ownership).This can make the visit wear effects less distinctive and requires that users remember the mapping of variables to effects. Again, customizability could assist; we also note that it is rare for current visualization systems to show secondary variables other than selection. Further research is needed to find the upper limit on the number of secondary variables that can be used.

CONCLUSIONS AND FUTURE WORK

In this paper we have shown that memorability – the ability to find items and locations previously visited – is a problem in distortion-oriented visualizations. We introduced the idea of visit wear, a type of read wear, as a technique for improving navigation in tasks where revisitation is common. An empirical study showed that visit wear can significantly improve people's speed and accuracy in finding previously-inspected items and locations, in both discrete and continuous data sets. In addition to these quantitative results, we also discuss several design issues that should be considered when building the technique into a real-world visualization system. These ideas can be used to substantially improve the usability of any visual workspace where revisitation is common, and where memorability is a problem.

We have plans for several further studies in this area. First, we will test visit wear with real-world tasks that have real revisitation patterns. This will allow us to experiment with automatically-determined effect duration, and will show how well visit wear works for different types of revisitation (e.g., revisiting a single item at a time rather than the entire set). Second, we wish to investigate the issues raised above, regarding the effect of visit wear on longer term memory of a space; this can be studied by training people with visit wear and then testing their recall abilities without the effect. Finally, we plan to investigate the use of visit wear in non-distorted spaces. We believe that there are many other situations where memorability can be problematic (e.g., remembering where you've been in a zoom-and-pan system, or finding folders in an file-system window) where revisitation is common and where visualization of interaction history could prove valuable.

ACKNOWLEDGMENTS

Thanks to the anonymous reviewers for their valuable comments. This work was supported by the Natural Sciences and Engineering Research Council of Canada

REFERENCES

1. Bederson, B., Fisheye Menus. *Proc. ACM UIST 2000*, 217-225.
2. Bertin, J., *Semiology of Graphics: Diagrams, Networks, Maps*, W.J. Berg, trans., Madison: Univ. of Wisconsin Press, 1983.
3. Carpendale, S., and Montagnese, C., A Framework for Unifying Presentation Space, *Proc. ACM UIST 2001*, 61-70.
4. Czerwinski, M., van Dantzich, M., Robertson, G., and Hoffman, H., The Contribution of Thumbnail Image, Mouse-over Text and Spatial Location Memory to Web Page Retrieval in 3D Viewing, *Proc. IFIP INTERACT 1999*, 163-170.
5. Gutwin, C., and Skopik, A., Fisheye Views are Good for Large Steering Tasks. *Proc. ACM CHI 2003*, 115-123.
6. Gutwin, C., Improving Focus Targeting in Interactive Fisheye Views, *Proc. ACM CHI 2002*, 267-274.
7. Gutwin, C. and Fedak, Christopher, A Comparison of Fisheye Lenses for Graphical Layout Tasks, *Proc. Graphics Interface 2004*, 9-16.
8. Hill, W., Hollan, J. Wroblewski, D. and McCandless, T., Edit Wear and Read Wear. *Proc. ACM CHI 1992*, 3-9.
9. Lamping, J., Rao, R. and Pirolli, P., A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies, *Proc. ACM CHI 1995*, 401-408.
10. Leung, Y., and Apperley, M., A Review and Taxonomy of Distortion-Oriented Presentation Techniques, *ACM ToCHI* 1 (2), 1994, 126-160.
11. Lynch, K., *The Image of the City*. Cambridge: MIT Press, 1960.
12. Misue, K., Eades, P. Lai, W. and Sugiyama, K., Layout Adjustment and the Mental Map. *Journal of Visual Languages and Computing*, 6 (2), 1995, 183-210.
13. Munzner, T. and Burchard, P., Visualizing the Structure of the World Wide Web in 3D Hyperbolic Space, *Proc. ACM VRML 1995*, 33-38.
14. Robertson, G., Czerwinski, M., Larson, K., Robbins, D., Thiel, D., and van Dantzich, M., Data Mountain: Using Spatial Memory for Document Management, *Proc. ACM UIST 1998*, 153-162.
15. Sarkar, M., and Brown, M., Graphical Fisheye Views of Graphs. *Proc. ACM CHI 1992*, 83-91.
16. Schaffer, D., Zuo, Z., Greenberg, S., Bartram, L., Dill, J., Dubs, S., and Roseman, M., Navigating Hierarchically Clustered Networks through Fisheye and Full-Zoom Methods *ACM ToCHI*, 3(2), 1996, 162-188.
17. Siegel, A., and White, S., The Development of Spatial Representation of Large-Scale Environments. In H.W. Reese ed., *Advances in Child Development and Behavior*, New York: Academic Press, 1975.
18. Silverman, I., and Eals, M., Sex differences in spatial abilities: Evolutionary theory and data. In J. Barkow, L. Cosmides, & J. Tooby, eds., *The Adapted Mind*, New York: Oxford Univ. Press, 1992, 533-549.
19. Skopik, A. *Improving Memorability in Fisheye Views*. Unpublished M.Sc. Thesis, Department of Computer Science, University of Saskatchewan, 2004.
20. Skopik, A. and Gutwin, C., Finding Things in Fisheyes: Memorability in Distorted Spaces. *Proc. Graphics Interface 2003*, 67-75.
21. Tauscher, L. and Greenberg, S., Revisitation Patterns in World Wide Web Navigation. In *Proc. ACM CHI 1997*, 398-406.
22. Thorndyke, P., and Goldin, S., Spatial Learning and Reasoning Skill, in *Spatial Organization: Theory, Research, and Application*, H. Pick and L. Acredolo, eds., New York: Plenum Press, 1983, 195-217.
23. Wexelblat, A., History-rich Tools for Social Navigation. *Proc. ACM CHI 1998*, 359–360.