

# Improving Focus Targeting in Interactive Fisheye Views

Carl Gutwin

Department of Computer Science, University of Saskatchewan

57 Campus Drive, Saskatoon, SK, S7N 5A9, Canada

+1 306 966 8646

carl.gutwin@usask.ca

## ABSTRACT

Fisheye views allow people to see both a focus region and the surrounding context in the same window. However, the magnification effects of the fisheye lens can cause several problems for users. One of these is focus-targeting, where a user moves the focus to a new location. Magnification makes focus-targeting difficult because objects appear to move as the focus point approaches them. This paper examines how the distortion of a fisheye view affects focus-targeting performance, and present a technique called speed-coupled flattening (SCF) as a way to improve focus targeting in distortion-oriented views. SCF dynamically reduces the distortion level of a fisheye based on pointer velocity and acceleration. In an experiment, the technique resulted in significant reductions in both targeting time and targeting errors. By adjusting distortion based on the user's activity, we can improve usability without requiring the user to manipulate any additional view controls.

## Keywords

Focus+context, distortion-oriented visualization, fisheye views, focus-targeting, Fitts' Law, speed-coupled flattening

## INTRODUCTION

Interactive focus+context techniques show both local detail and global context in the same view (e.g. [4,5]). They provide a user-controlled focus point for indicating which part of the data is to be shown in detail. These techniques are a solution to the space problem in information visualization, and allow many more objects to be displayed than would be possible in an undistorted view.

For two-dimensional data, focus+context techniques include the bifocal lens [5], the dragMag magnifier [15], the radar view [12], and a number of techniques calling themselves fisheye views (e.g. [2,11,13]). Fisheyes, named for the photographer's wide-angle lens, are characterized by their ability to show all of the data in a single view, and by the smooth transition between the high magnification of the focus region and the de-magnification of the context area. An example fisheye system is shown in Figure 1.

Fisheye views use non-linear magnification to achieve the balance between magnification and compression of the

data; depending on where the user's focus point is, different areas of the visualization will be magnified (or de-magnified) by different amounts. Non-linear magnification systems are becoming more and more common as a way to visualize many kinds of data, from file hierarchies [4] to maps [8] to documents [10] to web sites [9]. However, although the non-linear magnification approach effectively addresses the space issue, it also distorts the representation of the data in ways that can cause usability problems.

In particular, this paper is concerned with how the distortion effects of non-linear magnification hinder users' ability to move the focus point to specific parts of the data.

Moving the focus point is something that must be done very often when working with a fisheye view—whenever the user examines a different part of the data, even if that data is only a short distance away, they must re-focus the fisheye on the new target. Therefore, the task of moving the focus to a new target—*focus-targeting*—is one that can greatly affect the usability of a fisheye visualization.

The problem that occurs when focus-targeting in fisheye views is that targets appear to move in the opposite direction to the motion of the magnifying lens. This means that a focus target will move towards an approaching pointer, and away from a retreating one, making it more difficult to precisely position the focus point relative to the underlying visualized data.

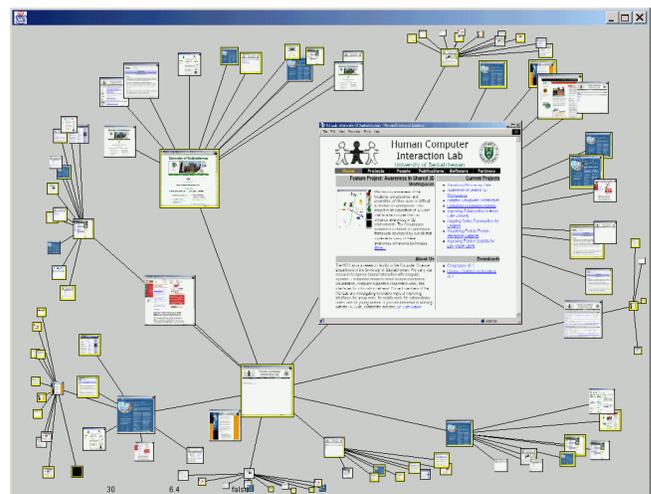


Figure 1. An interactive fisheye view for browsing a web site. The focus point is tied to the mouse cursor, and objects grow larger as the focus approaches them.

Moving targets are always more difficult to hit – but to make matters worse, the gradually increasing magnification of a fisheye lens makes targets move faster and faster the closer the focus comes to them. In fact, targets move at their highest rate of apparent speed at the exact moment that the pointer nears the target, making it difficult for a user to precisely position the pointer over the target.

This paper looks at the effects of distortion and magnification on focus-targeting performance in the well-known Sarkar and Brown (S&B) fisheye [11]. If increasing distortion does make targeting more difficult, then one way to improve targeting in fisheye views is to reduce the distortion level when a user is targeting. This leads to the idea of speed-coupled flattening (SCF), which dynamically adjusts distortion based on the user’s pointer velocity and acceleration. The technique is based on the observation that if a user is moving quickly, they are more likely to be navigating than examining detail.

In an experiment that compared focus targeting in an SCF fisheye and in a normal S&B fisheye, flattening resulted in significant reductions in both targeting time and error rates. This research uncovers one of the usability problems with distortion-oriented visualization techniques—that they make focus targeting difficult—and shows one way to combat the problem. Techniques such as speed-coupled flattening show considerable promise for improving the usability of fisheye views, without requiring users to learn or manipulate any additional view controls.

**FOCUS TARGETING IN FISHEYE VIEWS**

Bringing a target into focus in a fisheye view is similar to ordinary targeting with a pointing device, where the time required to select the target is governed by the distance to the target and the target’s size (that is, Fitts’ Law). Targeting actions can be divided into two phases: *motion*, where the user moves their pointer into the general target region, and *acquisition*, where the user precisely positions their pointer on the target and selects it.

In fisheye views, focus targeting is complicated by fact that the target moves as the focus approaches. The motion is caused by a simple effect of magnification: in any magnification system (including ordinary magnifiers in the real world), moving the lens causes the image of the underlying objects to move as well. For example, with a 2X

magnifying lens, moving the lens by one centimeter will cause the image in the magnifier to move two centimeters, in the opposite direction to the lens’s movement (Figure 2).

Fisheye views demonstrate this movement effect because they use magnification. However, due to the non-linear transformation function of the fisheye lens, (see Figure 3), the magnification and the movement increase as the focus gets closer to the target. As a result, targets move towards the focus more and more rapidly as the focus approaches them. Even though an object also appears to become larger the closer it gets to the lens’s focus, its motion makes acquisition of the target more difficult.

Depending on the magnification of the fisheye (determined in the S&B algorithm by the distortion level  $d$ ), the target may be moving several times faster than the focus point—at exactly the time when the user is attempting precise positioning of the focus. In our experience, this problem can be clearly seen as fisheye users (particularly novices) overshoot their focus targets, “kangarooing” back and forth before correctly positioning the pointer.

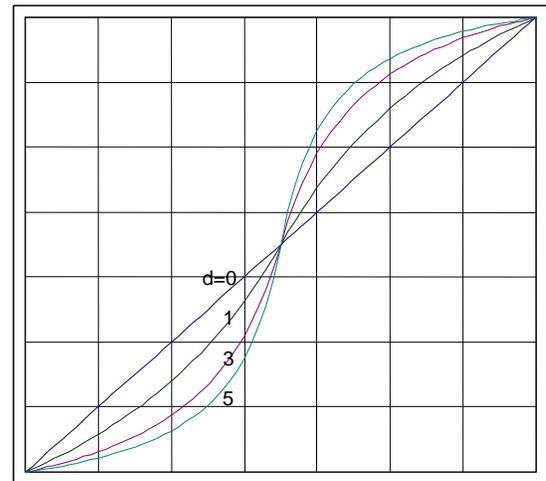


Figure 3. One-dimensional transformation functions [5] for the S&B fisheye at distortion levels  $d=0,1,3,5$ , with the focus set at the midway point. For each function, source data points are projected up from below and reflect left where they meet the curve. Steeper lines imply higher magnification of the source data; a 45-degree line implies no transformation.

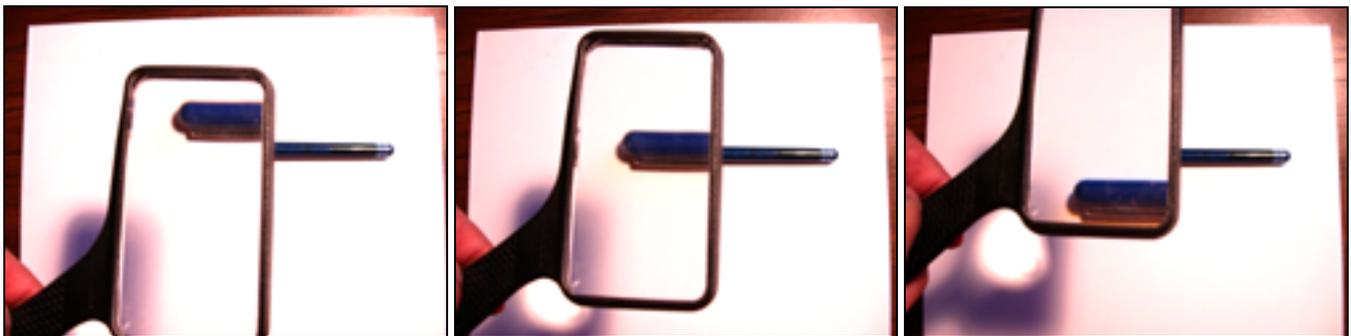


Figure 2. Motion effect of magnification: as the magnifier moves upwards, the magnified image moves down.

## SPEED-COUPLED FLATTENING

If distortion does in fact hinder targeting, then reducing distortion during targeting can reduce the problem. There are several possible ways to do this: for example, by having the user press a mouse or keyboard button to reduce distortion, or by using a dedicated control (such as the mouse wheel) to change distortion level. However, these solutions force users to learn and manipulate additional interface controls to carry out a very frequent task, and may also clash with existing assignments of buttons or other controls in the application.

An alternate approach that does not require any extra effort on the users' part is to automatically recognize when the user is engaged in targeting, based on the motion of their pointer. This is speed-coupled flattening (SCF). The technique uses the pointer's velocity and acceleration to recognize that a targeting action is happening. The underlying idea is that when the focus is moving quickly, or is accelerating, then the user is much more likely to be navigating to a new focus point than to be inspecting details in the data. The approach is similar to that of Tan et al. [14], who couple speed to height during flight through a virtual landscape, allowing the user to get a better idea of the terrain when navigating at higher velocity.

Speed-coupled flattening is designed with the following goals in mind:

- any major changes to the view should occur during the motion phase of targeting, while the user is still far away from the target;
- the view should remain relatively static during the acquisition phase of targeting, in order to simplify precise positioning;
- the changes to the view should be smooth enough that the target is easy to track during the motion phase; and

- the distortion level should be restored smoothly after targeting is complete.

The current implementation (demonstrated in [1]) periodically polls the pointer position and calculates velocity and acceleration. Three-value rolling averages are calculated every 40 ms. Based on the resulting values, the flattening scheme can be in one of four states:

1. *Quick flattening*: pointer acceleration above a threshold, which generally occurs early in targeting, causes a relatively large reduction in distortion.
2. *Moderate flattening*: if the pointer is not accelerating but velocity is above a threshold, distortion is reduced by a static amount (this provides a smooth continued flattening for ongoing movement).
3. *Hold*: when velocity drops below a first threshold, the distortion level is held constant.
4. *Re-magnification*: when velocity drops below a second threshold (or the pointer stops), distortion is gradually increased to its original value.

To test the two hypotheses of whether distortion level affects targeting performance, and whether speed-coupled flattening improves performance, an experiment was designed in which people carried out standard targeting tasks under a variety of distortion conditions in both the S&B and the SCF fisheyes.

## METHOD

### Participants

Ten paid participants (7 male, 3 female) were recruited from the computer science department of a local university. All participants were frequent users of mouse-and-windows based systems (at least 30 hours per week). Six participants of the 10 were familiar with the concept of a fisheye view, and four of these had experience using an interactive fisheye system (at least one hour of use).

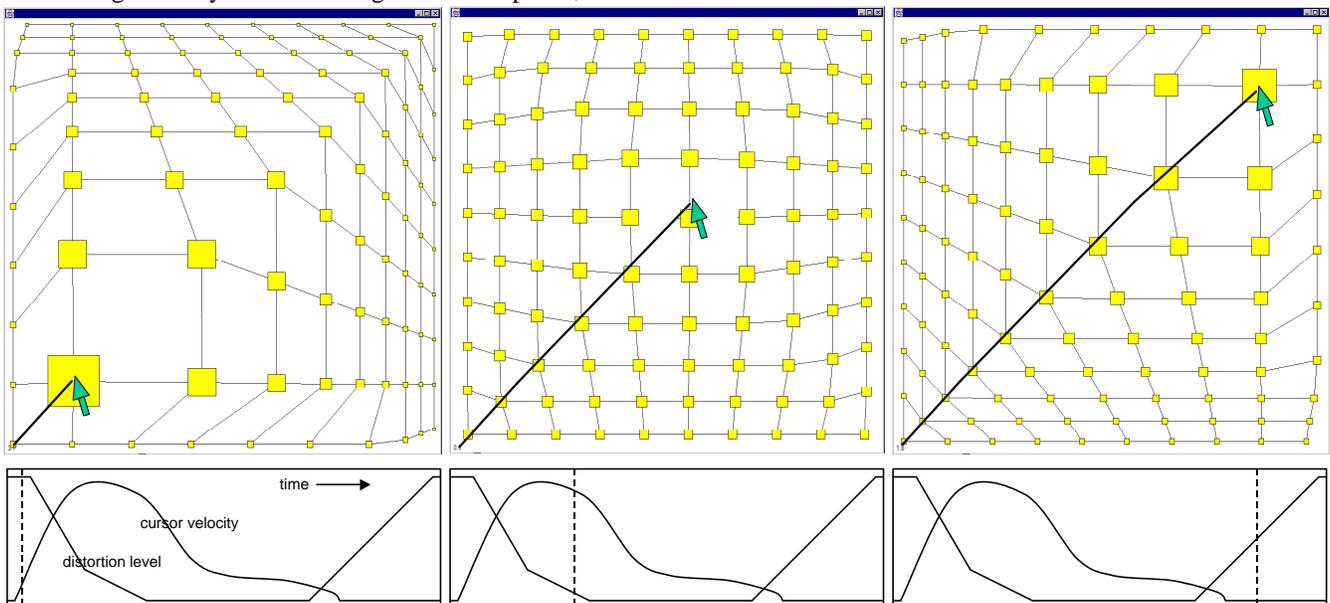


Figure 4. Speed-coupled flattening. Top row shows the fisheye view and pointer path. Bottom row shows a stylized plot of pointer velocity and distortion level. The dotted line indicates the point in time that the corresponding screen was captured.

## Apparatus

The experiment was conducted on a PII Windows 2000 PC running a custom-built Java application. The display was a 21-inch monitor, set to 1280x1024 resolution. The primary targeting task was the multidirectional point-select task described in ISO 9241-9 [ref]. In this task (see Figure 5), 24 circular targets are arranged in a ring, and the participant clicks on each target in succession where the next target is always directly across the ring. The next target to be clicked was coloured green and was marked with a purple cross.

The diameter of the targets was 18 mm and the diameter of the ring was 174 mm. This task has a targeting difficulty of 3.41 bits (using the Shannon formulation given in [6]).

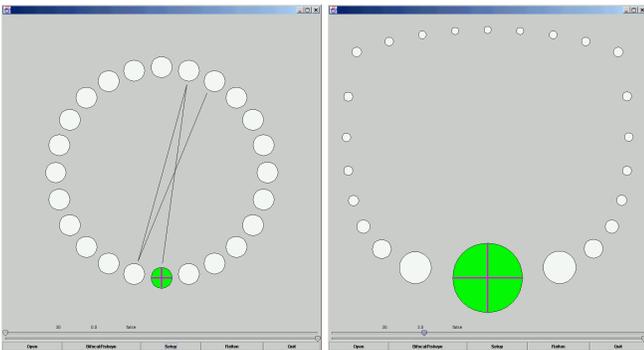


Figure 5. Initial setup of the targeting task with  $d=0$  (left) and  $d=3$  (right). Target path follows the line shown and continues around the ring.

In addition, a second targeting task using more realistic network datasets was given after the ISO task. Three different datasets (called A, B, and C) were used; an example is shown below in Figure 6. The task was carried out in a similar fashion to the ISO task, except that participants did not know beforehand where the next target would appear.

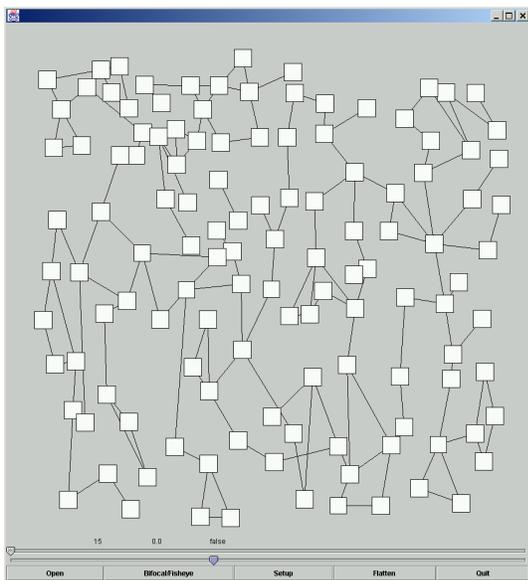


Figure 6. Network dataset B (shown at  $d = 0$ ).

## Fisheye views

The fisheye view used in the experiment was a Java implementation of the Sarkar and Brown algorithm [11], using the polar rather than the Cartesian transformation. The transformation functions for the four levels of distortion used in the experiment are shown in Figure 3.

The SCF version of the fisheye was implemented as described above. However, one change made for the experiment was that the gradual re-magnification that occurs when the mouse cursor slows was cut off at the moment when a target was clicked; at each click, the visualization was immediately returned to the original distortion. Although there was some concern that this sudden change would unnerve participants, it turned out to be easily accommodated and was not a problem for the targeting task.

## Procedure

Participants were randomly assigned to one of two groups: those that used the S&B fisheye first, or those who used the SCF fisheye first. Participants were introduced to the experiment and to the idea of a fisheye view (if they had not seen one before). They were then given a series of practice trials with both types of fisheye, at five different levels of distortion. Participants then completed seven blocks of test trials with the ISO task, one in each of the following conditions:

Block	Distortion	Fisheye type	Trials
1	0	first fisheye	4
2	1	first fisheye	4
3	1	second fisheye	4
4	3	first fisheye	4
5	3	second fisheye	4
6	5	first fisheye	4
7	5	second fisheye	4

Table 1. Experiment conditions for ISO task.

A test trial consisted of 25 clicks, starting with the first click on the first target, and corresponded to one trip around the targeting ring (with one target used twice). For each test condition, participants carried out four trials. Rests were allowed between trials and between conditions. Participants were instructed to click on the targets as quickly and as accurately as possible. After all trials were complete, participants were asked two questions: in which fisheye system did they think they were faster, and which system did they prefer for the targeting task.

After the ISO task was complete, participants also completed a limited set of trials with the more realistic network data. The intention was to get an initial look at how speed-coupled flattening would fare in a more realistic usage situation. Participants completed one trial of 25 target selections in each of three different network datasets (see Figure 6). Participants carried out the trial with both the S&B and the SCF fisheyes, and all trials were carried

out with a distortion level of 5.0. After these trials were complete, participants were again asked the same two questions about speed and preference.

Block	Network	Distortion	Fisheye type	Trials
1	A	5	first	1
2	A	5	second	1
3	B	5	first	1
4	B	5	second	1
5	C	5	first	1
6	C	5	second	1

Table 2. Experiment conditions for network task.

### Experiment design

The primary study was carried out using only the ISO task data, and was a 2 x 3 within-participants factorial design. The factors were:

- Fisheye type: S&B or SCF;
- Distortion level: d-value of 0, 1, 3, or 5.

With ten participants and 25 target selections per trial, there was a total of 4000 target selections in the experiment. Data collected included completion times for each click in each trial, all mouse motion information during the trial (for path analysis), and participant answers to summary questions.

### RESULTS FOR ISO TASKS

The first set of results analysed were those from the standard ISO pointing task described above. Results are organized below by the two main hypotheses: first, that distortion impairs targeting, and second, that flattening reduces the problem. Summary data for targeting time and errors are given in Table 3 and shown in Figures 7 and 8.

#### Effect of distortion on targeting performance

As can be seen in the charts, both targeting time and error rates steadily increase as distortion level increases, in both the S&B fisheye and in the SCF fisheye. At d=0, one targeting action took slightly less than one second; increasing d to 5.0 adds .42 sec. for the S&B fisheye, and .22 sec. for the SCF fisheye. Similarly, people made about one error in 28 trials with d=0; at d=5, people made about one error in 12 (S&B) and one in 20 (SCF).

d	Fisheye	N	Time (ms)		Errors	
			mean	SD	mean	SD
0	---	10	22510	2695	0.88	0.60
1	S&B	10	26603	3361	1.13	0.82
1	SCF	10	24343	2918	0.95	0.79
3	S&B	10	30799	3391	2.02	1.16
3	SCF	10	26021	3024	1.10	0.89
5	S&B	10	33034	4000	2.20	1.67
5	SCF	10	28016	4235	1.27	1.07

Table 3. Means and standard deviations for performance measures for all fisheye types and distortion levels.

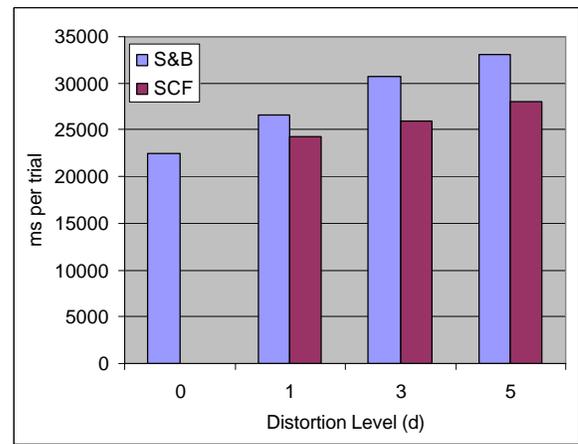


Figure 7. Mean targeting time per 25-target ISO trial.

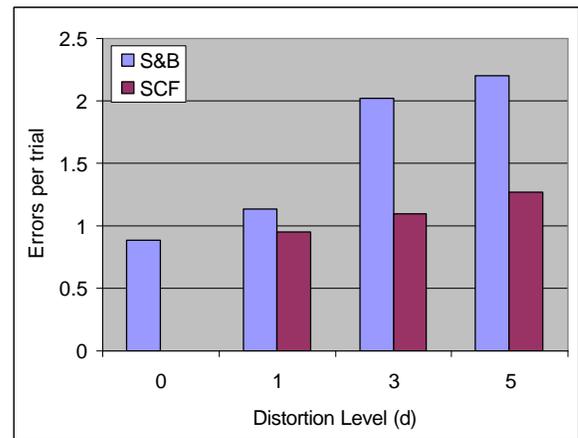


Figure 8. Mean errors per 25-target ISO trial.

Data from the S&B trials were analysed using ANOVA. There were clear main effects of distortion level, both for completion time ( $F_{3,27}=94.4$ ;  $p<.001$ ) and error rate ( $F_{3,27}=8.61$ ;  $p<.001$ ).

#### Targeting path analysis

To explore the reasons behind these differences, the targeting path data were analysed using the measures proposed by MacKenzie et al [7]. Two aspects of the path were considered: first, the divergence from the optimal (i.e. straight) path en route to the target, corresponding to the motion phase of targeting; and second, the pointer movement in the target area, corresponding to the acquisition phase. En-route motion was assessed using movement variability, movement error, and movement offset (see [7] for details). Target acquisition was assessed using orthogonal direction change rate, target re-entry rate, and the total distance traveled after the first target entry.

The data show that distortion did not have an effect on the amount of path divergence. In fact, several participants had paths that were closer to the optimal axis at higher distortion levels. This may be because they were moving more slowly, but it appears clear that the higher targeting times at greater distortion were not caused by extra travel en route to the target.

However, the all of the target-acquisition measures showed significant effects of distortion level: that is, with more distortion people made more direction changes along the task axis, re-entered the target more often, and traveled further between the first target entry and their final acquisition. It is therefore reasonable to assume that part of the increased targeting time and error rate is explained by additional difficulty with the acquisition phase, which is consistent with the earlier observations that overshooting the target is a major problem when focusing in fisheyes.

### Effects of flattening on targeting performance

Given that distortion does affect targeting time and errors, the second hypothesis considered whether speed-coupled flattening made any difference to performance. To test the hypothesis, times and error rates from the S&B and SCF fisheyes were compared for distortion levels 1, 3, and 5.

As can be seen from Figures 7 and 8, there are clear differences between the two fisheye types. At  $d=5$ , using SCF allowed participants to complete each targeting action two-tenths of a second faster, and complete almost twice as many trials between each error. ANOVA showed that there are main effects both for targeting time ( $F_{1,9} = 167.8$ ,  $p < .001$ ) and for error rate ( $F_{1,9} = 23.8$ ,  $p < .001$ ). There were no significant interactions between fisheye type and distortion level.

Performance at the reduced distortion level caused by flattening was about equal to that of the equivalent distortion level in the S&B fisheye. For example, trials at  $d=5$  in the SCF fisheye reduced distortion to about  $d=2$ , and time and errors were between the values for conditions  $d=1$  and  $d=3$  in the S&B fisheye.

These results also falls within original expectations. That is, since increased distortion leads to more problems for focus targeting, it is reasonable that targeting can be improved by reducing distortion when targeting. The study shows that SCF is an effective means of doing this.

### Targeting path analysis

Only the acquisition-phase path analysis was carried out for the comparison of SCF and S&B fisheyes. The same measures were used as above: orthogonal direction change, the rate of target re-entry, and the distance traveled after first entry. Significant differences were found between the two fisheyes for each of these measures, and in all cases, performance was better with the flattening fisheye. Again, this result is not surprising because of the relationship between distortion level and overshooting difficulty, and because the flattening technique reduces the distortion level exactly when users are engaged in precise positioning.

### Preferences

Participants were asked two questions after the trials were complete: in which fisheye did they think that they were faster, and which fisheye did they prefer for doing the targeting task. All ten participants responded that the SCF fisheye was both faster and more preferable than the S&B fisheye. When asked why, most participants made comments indicating that the SCF fisheye had made it

easier for them to operate in the acquisition phase. For example, one person said “the circles don’t jump around as much with the SCF;” another said “I don’t have to be as careful in this one [SCF].”

### Expertise differences

Since four of the ten participants were more experienced with interactive fisheye views, the data were examined to see whether expertise had any effect on overall performance and whether expertise changed the effects of flattening. Although the four experienced participants were considerably faster in all situations, their data showed similar patterns to those reported above both for targeting times and errors. That is, performance was still reduced as distortion increased, and flattening the fisheye still lessened the effect of increased distortion. As an example, targeting times are shown in Figure 9 separated by experience.

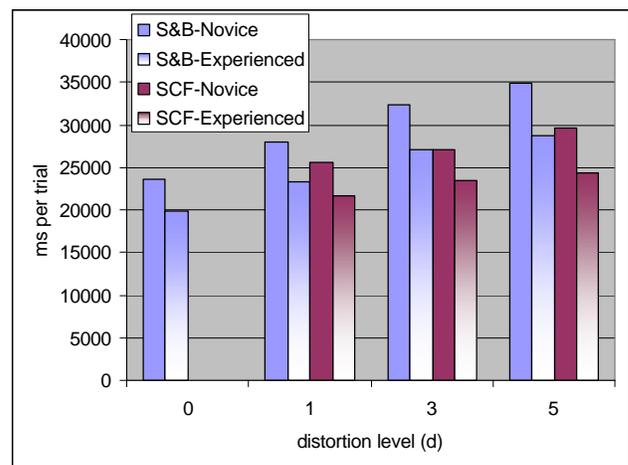


Figure 9. Mean targeting times, showing differences between experienced users (bars with fade fill) and novice users (solid fill) for each fisheye type.

### RESULTS FOR REALISTIC DATASETS

As described earlier, a limited set of trials were carried out with network datasets after the ISO task was complete. These trials were done in order to get an initial look at the effectiveness of the flattening technique with more realistic data. The network datasets (see example in Figure 6 above) were different from the ISO task in three main ways: there were many more objects in the visualization, there was a range of distances to the next target, and the location of the next target was not easily predictable. Data from these trials is shown in Table 4 and Figures 10 and 11.

As in the ISO tasks, targeting times and errors were recorded. Even though only a relatively small sample was collected, significant main effects of fisheye type were found for both targeting time ( $F_{1,9}=17.3$ ;  $p < .005$ ) and error rate ( $F_{1,9}=13.2$ ;  $p < .005$ ).

Participants’ responses to the summary questions, however, were not as one-sided as they had been in the ISO task. Despite the performance differences, people did not strongly prefer the SCF fisheye to the S&B for this task. Four people thought that they were faster using SCF, one

thought that they had been faster using S&B, and five felt that there had been no difference.

The mixed responses for the more realistic task may simply highlight one aspect of speed-coupled flattening: that is it not particularly noticeable in ordinary situations. Since the starting distortion level was fairly high in the network task, and since several of the targets in this were close by, people did not often see an obvious flattening of the data. Nevertheless, the technique did have a significant effect on performance (although the sample size here is small).

Network	Fisheye	Time (ms)		Errors	
		mean	SD	mean	SD
A	SCF	40693	4713	3	2.35
A	S&B	44111	5023	4.9	3.81
B	SCF	40026	4557	1.5	1.26
B	S&B	43718	3790	4.2	2.93
C	SCF	47260	6367	2.4	1.64
C	S&B	50802	6655	4.1	2.99

Table X. Mean targeting times and errors per 25-target trial for network datasets.

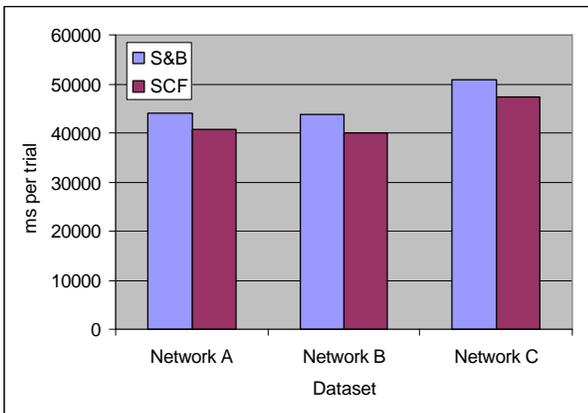


Figure 10. Mean targeting times for network datasets (d=5).

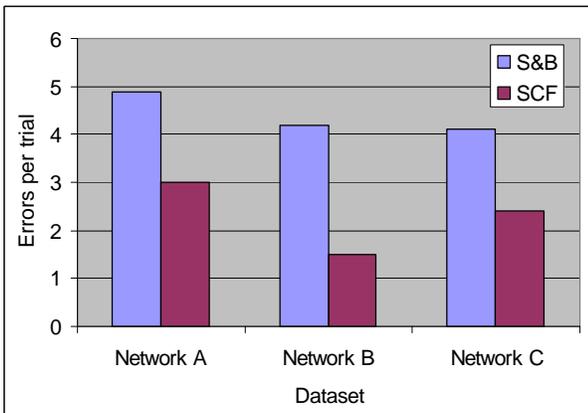


Figure 11. Error rate for network datasets (d=5).

## DISCUSSION

The experimental results support both hypotheses—that distortion level in an S&B fisheye impairs focus-targeting performance, and that speed-coupled flattening can reduce the problem. Flattening appears to work simply because it reduces the distortion level and therefore the motion effects of magnification, particularly during the acquisition phase of targeting where precise positioning is required.

The next sections discuss some of the issues raised by these results. First we consider the tradeoffs inherent in trying to support both focus-based tasks and context-based tasks in the same view, and second, we compare SCF to other techniques for managing this tradeoff.

### Tradeoffs in focus+context views

The problems evident in focus targeting show that there are tradeoffs that must be considered when designing focus+context views. In particular, users carry out different kinds of tasks in the focus region and the context area, but the interaction techniques of most visualization systems do not take these differences into account.

Focus targeting highlights a conflict between navigation (a context-oriented task) and inspection (a focus-oriented task). When a user is navigating, they are unlikely to be inspecting at the same time; yet many systems continue to support the inspection task (through a fully-magnified focus point) at all times. Similarly, the reverse problem often occurs when a user has stopped navigating and is looking closely at a particular object: that not enough screen space is given to the detail region for current purposes, and too much is given to the surrounding context.

Focus+context views would be more useful if they could adjust the support given for focus tasks and context tasks, rather than being stuck at one point on the continuum. This adaptability could be achieved by giving the user greater control over the visualization; however, the example of speed-coupled flattening shows that it may also be possible for systems to adjust themselves automatically to the user's current activity. The study described above suggests that an interface can shift the balance of support towards one type of task over another, based on easily-obtainable user evidence, and make a significant difference in performance.

With adaptive techniques like speed-coupled flattening, the goal is to give users the benefits of both sides of the tradeoff. For example, SCF could allow users to choose a higher initial distortion level to maximize the amount of space devoted to the focus region, without suffering the ill effects of attempting to navigate at high magnification.

### Comparing SCF to other techniques

Coupling distortion level to pointer speed is only one of several possible ways to manage the relative importance of focus and context in a visualization. As discussed above, SCF has advantages of not requiring any additional user input or any additional view controls. In different situations, however, other techniques may also be useful. For example:

- *Click to focus.* Moving the focus only when the user clicks the pointing device is a way to avoid magnification-motion effects entirely. However, the drawback is that exploration is more difficult and may require several clicks to find the correct target.
- *Click to flatten.* Flattening can be tied to an explicit user action such as a right or middle button click rather than to pointer speed. The drawback here is that it is difficult to choose an appropriate flattening rate, something that is automatically determined by SCF.
- *Explicit distortion control.* A separate control such as a mouse wheel can explicitly control the distortion level of the visualization. However, in our experience many users find this setup awkward, even when the control is moved to the non-pointing hand.
- *Lens follows behind.* A technique similar to SCF but without requiring any dynamic adaptation is simply to let the focus point of the magnifying lens trail behind the pointer as if it were attached by a rubber band. High velocity movement results in the lens lagging further behind, which reduces the motion effects; when the pointer slows, the lens catches up again. We are currently comparing this technique to SCF.

## CONCLUSION

Distortion-based focus+context systems allow users to look at and manipulate large datasets in a single view. However, the distorted representation can cause problems for users. In this paper we identified and explored one of these problems—moving the focus of the magnifying lens to a particular target in the data. An experiment with the Sarkar and Brown fisheye showed that the level of distortion has a clear effect on the time people need to select focus targets, and on the number of errors that they make in targeting. Speed-coupled flattening was introduced as a way to reduce the focus-targeting problem by reducing magnification when pointer velocity and acceleration are high. Flattening was able to significantly reduce targeting times and errors compared with the normal S&B fisheye. Problems with focus targeting are one example of the tradeoff between focus and context that all distortion-oriented techniques must manage; SCF shows that these displays can be automatically adjusted to better suit the user's current tasks.

It seems likely that any distortion-oriented visualization in which the user moves a focus point can benefit from some type of flattening scheme (or other means of addressing the tradeoff between focus and context). We are currently expanding our work to look at SCF in other visualization systems and in different kinds of tasks, and to look at other schemes for managing the focus-context tradeoff.

## REFERENCES

1. <author>. Speed-Coupled Flattening for Fisheye Views (video). <http://<address>/downloads/flatten.mov>.
2. Carpendale, M. S. T., Cowperthwaite, D. J., and Fracchia, F. D., Extending Distortion Viewing Techniques from 2D to 3D Data, *IEEE Computer Graphics and Applications*, Vol. 17(4), 42-51.
3. ISO ISO/TC 159/SC4/WG3 N147: *Ergonomic requirements for office work with visual display terminals – Part 9 – Requirements for non-keyboard input devices*, International Organization for Standardization, May 25, 1998.
4. Lamping, J., Rao, R., and Pirolli, P. A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies. *Proceedings of ACM CHI 1995*, 401-408.
5. Leung, Y., and Apperley, M. A Review and Taxonomy of Distortion-Oriented Presentation Techniques *ACM Transactions on Computer-Human Interaction*, 1994 1, 2, 126-160.
6. MacKenzie, I. S., (1992) Movement Time Prediction in Human-Computer Interfaces. *Proc. Graphics Interface '92*, Morgan Kaufmann, 140-150.
7. MacKenzie, I. S., Kauppinen, T., & Silfverberg, M. (2001). Accuracy measures for evaluating computer pointing devices. *Proceedings of ACM CHI 2001*, 9-16.
8. Rauschenbach, U, Weinkauff, T, and Schumann, H, Interactive focus and context display of large raster images, in *International Conference in Central Europe on Computer Graphics, Visualization and Interactive Digital Media*, 2000.
9. Ridsen, K., Czerwinski, M., Munzner, T., and Cook, D. An initial examination of ease of use for 2D and 3D information visualizations of web content. *International Journal of Human Computer Studies*, Vol. 53, No. 5, November 2000, pp 695-714.
10. Robertson, G., and Mackinlay, J. The Document Lens *Proceedings of the ACM Symposium on User Interface Software and Technology 1993*, 101-108
11. Sarkar, M., & Brown, M. (1992) Graphical Fisheye Views of Graphs. In *Proc. ACM CHI'92*, 83-91.
12. Smith, R., Hixon, R., and Horan, B. (1998), Supporting Flexible Roles in a Shared Space, *Proc. ACM CSCW'98*, 197-206.
13. Storey, M.-A.D., Fracchia, F.D., and Müller, H. A., Customizing A Fisheye View Algorithm to Preserve the Mental Map, *Journal of Visual Languages and Computing*, 10, 245-267, 1999.
14. Tan, D.S., Robertson, G.G., & Czerwinski, M. (2001). Exploring 3D Navigation: Combining Speed-coupled Flying with Orbiting. In *Proc. CHI 2001*, 498-505.
15. Ware, C., and Lewis, M. The DragMag Image Magnifier. Video Proceedings of ACM CHI'95 1995, vol. 2, p.407-408.

