Perceptual Invariance of Nonlinear Focus+Context Transformations

Keith Lau * Ronald A. Rensink Tamara Munzner

University of British Columbia

Abstract

Focus+Context techniques are commonly used in visualization systems to simultaneously provide both the details and the context of a particular dataset. This paper proposes a new methodology to empirically investigate the effect of various Focus+Context transformations on human perception. This methodology is based on the *shaker paradigm*, which tests performance for a visual task on an image that is rapidly alternated with a transformed version of itself. An important aspect of this technique is that it can determine two different kinds of perceptual cost: (i) the effect on the perception of a static transformed image, and (ii) the effect of the dynamics of the transformation itself. This technique has been successfully applied to determine the extent to which human perception is invariant to scaling and rotation [Rensink 2004]. In this paper, we extend this approach to examine nonlinear fisheye transformations of the type typically used in a Focus+Context system. We show that there exists a *no-cost zone* where performance is unaffected by an abrupt, noticeable fisheye transformation, and that its extent can be determined. The lack of perceptual cost in regards to these sudden changes contradicts the belief that they are necessarily detrimental to performance, and suggests that smoothly animated transformations between visual states are not always necessary. We show that this technique also can map out *low-cost zones* where transformations result in only a slight degradation of performance. Finally, we show that rectangular grids have no positive effect on performance, acting only as a form of visual clutter. These results therefore demonstrate that the perceptual costs of nonlinear transformations can be successfully quantified. Interestingly, they show that some kinds of sudden transformation can be experienced with minimal or no perceptual cost. This contradicts the belief that sudden changes are necessarily detrimental to performance, and suggests that smoothly animated transformations between visual states are not always necessary.

CR Categories: H.1.2 [Models and Principles]: User/Machine Systems—Human information processing H.5.2 [Information Interfaces and Presentation]: User Interfaces—Evaluation/methodology

Keywords: information visualization, visual search, fisheye transformations, Focus+Context, visual representation

1 Introduction

One of the major challenges in the field of information visualization is to keep users from getting lost when navigating through visual representations of large complex datasets. Many visualization systems use Focus+Context techniques, in which a nonlinear transformation is applied to the image to provide both the focus and the context in one integrated image [Furnas 1986; Leung and Apperley 1994; Munzner et al. 2003]. Another definition from Keahey characterizes these transformations as “non-occluding, in-place magnification which preserves a view of the global context” [1997].

The literature on evaluating the efficacy of Focus+Context interfaces contains mixed results. Studies have shown that fisheye or other distortion approaches are beneficial for tasks such as interaction [Gutwin and Skopik 2003], navigation [Gutwin and Fedak 2004b; Risden et al. 2000; Schaffer et al. 1996], and even calendar use [Bederson et al. 2004]. However, other studies found that distortion can impair performance for tasks such as layout [Gutwin and Fedak 2004a] and scanning large local areas [Kobsa 2003]. The costs of Focus+Context interfaces have not yet been quantified, and their effect on visual perception is of particular interest. As such, it is difficult to determine what kinds of difficulties may be incurred by the use of such systems, and what particular kinds of transformations might minimize such difficulties.

In this paper, we propose a way of quantifying the perceptual cost of nonlinear transformations of the type generally used in Focus+Context systems. We show that the shaker paradigm [Rensink 2004] can be adapted to measure the perceptual cost of these transformations, both in terms of their static aspects – the perception of the static transformed image – and their dynamic aspects – the effort of shifting from an old visual representation to a new one. We examine the way that these costs vary with the degree of transformation applied, and determine whether adding visual cues, like grid lines, to the background has any effect on the task.

1.1 Focus+Context Systems

A common problem in the exploration of large datasets via a single window is that, due to the lack of explicit information about anything else other than what is currently on the screen, focusing on a particular portion of the dataset leads to a loss of context. Furthermore, the cognitive overhead of maintaining a mental model of navigation history is high [Zhang 1991]. Without any additional aid, people must often backtrack to remember where they have been.

One popular approach is to provide an small overview window that always shows the current location relative to the overall view, as evaluated by Hornbaek et al [2002]. This approach works by explicitly providing the context and reducing the cognitive load of the user. However, one disadvantage is that switching between the two views might be quite distracting for the mind, as the user constantly tries to relate the small overview window with the main one. In addition, since this approach provides no support for navigational history, the solution of creating a single integrated Focus+Context view showing the details surrounded by context is appealing. If done correctly, the complete context could be preserved, without the need for supporting features such as animated transitions or multiple views.

The risk of this approach is that a distortion that is too extreme would render the neighbouring context unrecognizable. Nonetheless, this method of nonlinear image transformation has

*{klau,rensink,tmm}@cs.ubc.ca

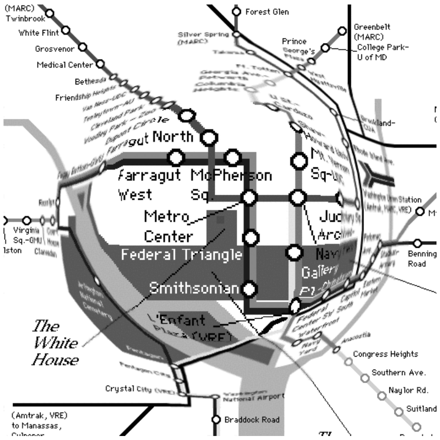

Figure 1: Example of radial fisheye transformation, courtesy of T. Alan Keahey [Keahey and Robertson 1996].

been researched extensively under various names, including: Focus+Context [Rao and Card 1994], fisheye views [Furnas 1986; Sarkar and Brown 1994] distortion-oriented presentation techniques [Leung and Apperley 1994], nonlinear distortion [Keahey and Robertson 1997], pliable surfaces [Carpendale et al. 1995], and elastic presentation spaces [Carpendale and Montagnese 2001]. Figure 1 shows an example of a radial fisheye transformation. Hyperbolic tree and graph browsers [Lamping et al. 1995; Munzner 1997] are examples that use a nonlinear transformation variant that has a single radial focus, magnifying the node of interest and shrinking other nodes of the tree. The metaphor of a stretchable rubber sheet with a rectangular lens used for distortion was introduced by Sarkar et al [1993]. This approach has been used in many systems, ranging from the early Document Lens [Robertson and Mackinlay 1993] to the recent TreeJuxtaposer system for fast structural comparison of large trees [Munzner et al. 2003].

While systems employing nonlinear transformations are often an appropriate solution for this problem, several key issues remain unresolved. We would like to understand what kinds of perceptual costs are incurred by both static and dynamic perception of transformations, and how to minimize them. Static perception may be affected by effects such as the clutter or crowding of the transformed image. Dynamic aspects of perception include the remappings that may need to take place as the viewer moves around. Despite recent work on comparing the efficacy of different distortion functions [Gutwin and Fedak 2004a; Gutwin 2002], an optimal design for nonlinear transformations has not yet been discovered.

1.2 Transformational Invariance

The effects of geometric transformation on visual perception have received only a limited amount of investigation over the years. Among the earliest studies were those that measured how well observers could recognize pictures of a three-dimensional object that was rotated by some amount in three-dimensional space. It was found that the time to determine this was proportional to the difference in angle, implying that observers could mentally rotate the representation of the object when needed [Shepard and Metzler 1971; Tarr and Pinker 1989]. Later, size scaling was also found to function similarly [Bundesen and Larson 1975; Bennett 2002], as was position [Bennett 2002]. These results suggest that rotation, scaling and translation are *natural* transformations, in that if sufficient time is given, they can be compensated for by the visual system.

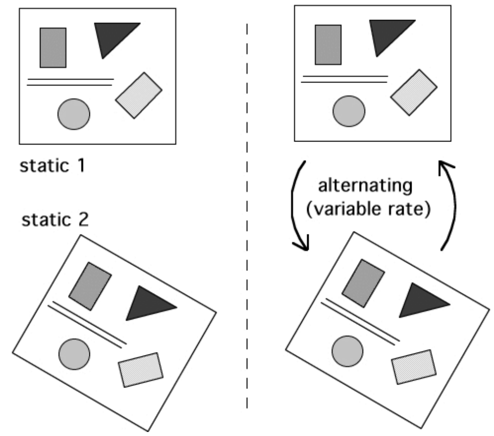

Figure 2: Shaker paradigm. The static cost of the transformation is measured by comparing performance for static image 1 (original image) against static image 2 (transformed image). The dynamic cost of the transformation itself is measured by comparing the average performance for images 1 and 2 against the alternating case, where the image appears to *shake* back and forth. The alternating case has an equal amount of time for both static images, while also having a transitional component that depends on the alternation rate.

Another line of research originated from the observation that pictures in movies do not appear distorted, even for audience members located far from the viewpoint for which the projection would be correct. In particular, Cutting [1987; 1991] investigated why rigid objects still appear to be rigid, even though their distorted projections should make them appear to be nonrigid. Cutting's results indicate that the visual system does not compensate for the distortions caused by off-axis viewing; rather, it simply appears to use local information that is not greatly affected by such distortions.

However, neither type of study is helpful for the purpose of understand visualization systems, because both types involve transformation of a single object rather than the spatial layout of a collection of objects. Furthermore, they do not involve an active scanning of space, a key component of information visualization; therefore, other methods must be used.

1.3 Shaker Paradigm

A more direct method of examining the effect of geometric transformations on visual perception is the shaker paradigm [Rensink 2004]. Here, performance of a visual task, such as memorization or scanning, is compared under three different conditions: (i) an original image, (ii) a transformed static image, and (iii) the original and transformed images rapidly alternating with each other, as in Figure 2. The extent to which the static aspect of the transformation affects perception is measured by comparing performance on the static conditions (i) and (ii). The extent to which the dynamic aspect of the transformation affects perception is measured by comparing the average of the two static performance measures with the performance on the rapidly alternating sequence.

To measure the effect of a given transformation on the scanning of space, a natural task is *visual search*, in which the observer must search a large array of items for the presence or absence of a given target figure [Treisman 1985; Treisman and Gormican 1988]. Performance here is typically measured by the response time (RT) required to detect the presence or absence of the target, with accuracy

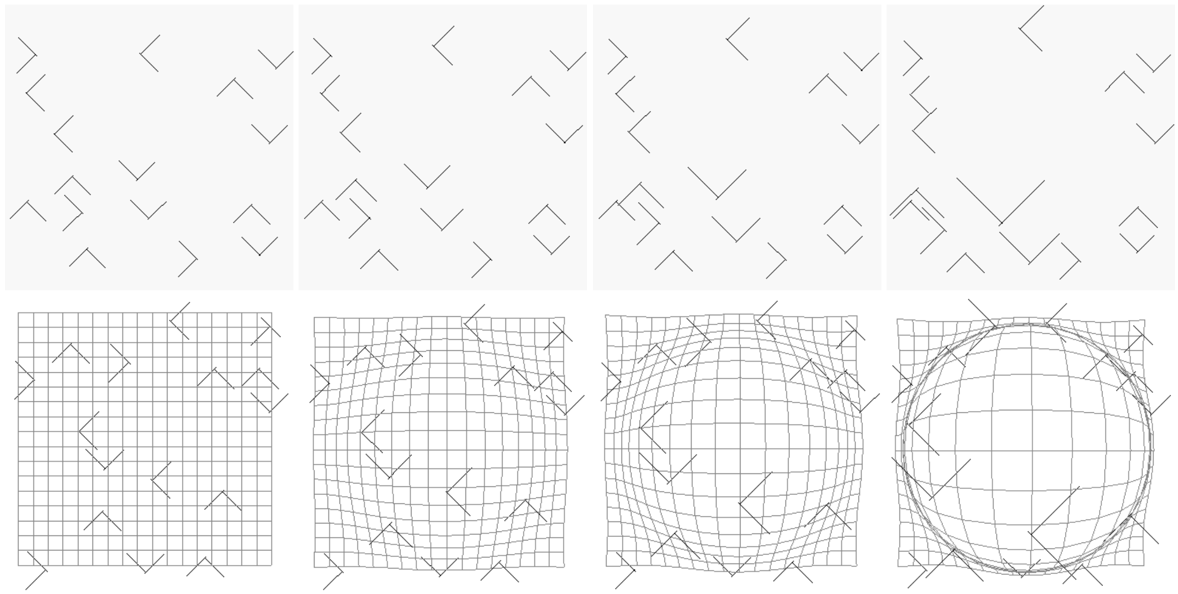

Figure 3: Example randomly-generated images shown to subjects. The top row shows the images without a background grid, and the bottom row shows the grid condition. **Column 1:** Static image. **Column 2:** Magnification level 1, c=0.9. **Column 3:** Magnification level 2, c=1.2. **Column 4:** Magnification level 3, c=1.5.

being kept as high as possible. This task is in many respects similar to what is involved in information visualization, with a common and dominant element being the scanning of a large number of items. As such, the results obtained by using this task are highly relevant to the purposes here.

Performance in visual search is generally measured by examining how RT varies with the number of items present; that is, the set size. RT is usually a linear function of set size, and so the natural measure is the search slope, typically expressed in terms of milliseconds per item (ms/item). If performance is unaffected by a particular transformation, the ratio of slopes for the two conditions will be one. More generally, the extent to which this ratio departs from one will reflect the perceptual cost of the transformation. The *static cost* is measured via the ratio of slopes for the transformed and original static cases; this can be arbitrarily high. The *dynamic cost* is measured via the ratio of the average slope for the two static cases [1] to the slope of the alternating case.

Using this approach on a simple search for a T-shaped item among a set of L- shaped items, it was found that there is no static cost for translations up to 4° of visual angle, rotations up to 90°, and scalings up to a factor of 4 [Rensink 2004]. There is also no dynamic cost for translations up to 4°. However, dynamic cost is incurred only for rotations of 30° or more, and for scalings of 3 or greater. Such invariance indicates a *no-cost zone* of considerable extent in terms of allowable translational, rotational, and scaling transformations. For example, a sudden change in size by a factor of 2 has no effect at all on performance. This clearly indicates the that mechanisms involved differ from those that carry out mental rotation and mental scaling. It also raises the possibility that there exist similar no-cost zones for nonlinear transformations of the kind used in Focus+Context systems.

[1] Since the search rate is expressed in milliseconds per item, and each display is on for an equal length of time, the appropriate average of the two static slopes m_1 and m_2 is the harmonic mean H, where $1/H = 1/2(1/m_1 + 1/m_2)$.

2 Experiment Protocol

The purpose of this experiment is to determine if the shaker paradigm can provide a useful estimate of perceptual cost for nonlinear fisheye transformations. Since this is only an initial investigation of the general feasibility of this approach, only 3 levels of magnifications are examined to test for the existence of a no-cost zone. We also examine whether the existence of a background grid affects performance. Consequently, a total of 6 experiments are needed.

2.1 General Method

Observers were asked to find a T-shaped target amidst a group of L-shape distractor items, as shown in Figure 3, with all items being at various orientations. Targets were present on half the trials, chosen randomly, and absent on the remainder. Observers were asked to determine the presence or absence of the target item as quickly as possible, while keeping errors below 10%. During the experiment, observers were seated at a viewing distance of approximately 55cm from the display, and kept their hands on the keyboard, with their right index fingers on the “p” key and their left index fingers on the “a” key. They were instructed to hit the key labelled “p” if a target was present, and the one labelled “a” if absent. Feedback was provided after each trial.

Each observer ran all three conditions (original static, transformed static, and alternating) of given level of magnification, with 3 blocks of 60 trials in each condition. To minimize any effects of learning, conditions were run such that the alternating case was always run second, with the static cases counterbalanced as to which was run first and which was run last.

2.2 Observers

12 observers were used in each of the six experiments. All were students at the University of British Columbia, with ages ranging from 18 to 35, and with normal or corrected-to-normal vision. Most observers were naïve with respect to visual search methodology.

2.3 Stimuli

Stimuli were arrays of randomly-positioned items as in Figure 3, where an item was either a T or an L. Images contained 16, 24 or 32 items. These numbers were chosen after several test runs to ensure that the task would not be too difficult, and so wear out the observers, or too easy, and thus not provide enough time for several alternations of the images. Density of the items was kept constant by adjusting the area of the image. The three image sizes subtended visual angles of 8.5°, 10.5°, and 12.5°. In the grid conditions, the space between the grid lines subtended a visual angle of 0.5°. Each item subtended a visual angle of 1.0°. Items could be in any one of the four possible orientations: 45°, 135°, 225° or 315°. The position and size of the items were affected by the transformation, but they remained locally rigid and maintained an orientation parallel to the viewing plane.

Static conditions simply had a single image, transformed or original, remain visible until the observer responded. In alternating conditions, the transformed and original images alternated every 480 ms until the observer responded.

The images were generated beforehand with an OpenGL program. They were created such that the eye point was located directly above the image plane, and the height values were transformed by the magnification factor c as in the following equation:

$$z = ce^{-(x^2+y^2)}(c - c(x^2+y^2)) \tag{1}$$

This equation yields the standard center bulge and gentle slope near the edges that are standard features found in all fisheye transformations [Keahey and Robertson 1997; Leung and Apperley 1994]. Three levels of magnification were used, and they corresponded to $c = 0.9$, $c = 1.2$, and $c = 1.5$, where c is the magnification factor. $c = 1.5$ was defined to be the maximum possible magnification, where increasing c any further would push the transformed plane beyond the eye point. For the untransformed case, c was equal to 0. Therefore, the domain $[0, 1.5]$ encapsulated all points from the least to the greatest magnification level. Thus, using an increment of 0.3 and excluding 0, five values of c were possible. However, we omitted the values 0.3 and 0.6 in our analysis because they were found to be very close to the static case during our test runs. Therefore, level 1, 2, and 3 had c values equal to 0.9, 1.2 and 1.5 respectively.

2.4 Analysis

The approach here is based on the time taken to determine if a target is present or absent in an image. To ensure that no speed-accuracy tradeoffs interfered with the analysis, observers were asked at the beginning of each experiment to be as accurate as possible. Observers were removed if they had an error of 35% or greater for any set size. This resulted in 7 of the 72 observers being removed. Accuracy of the remaining observers was good, with an average error of 5.5% across all conditions.

Response times for each observer for a given condition were calculated beginning with the set of average response times (RTs) for target present for each set size of that condition. Error compensation occurred by dividing each average by the accuracy for that condition. Results were qualitatively the same using uncorrected data, except that the level 1 + grid variant had a strongly significant difference between static original and static transformed. The same general pattern, although somewhat noisier, was also found using target-absent responses; consequently, these will not be discussed further here.

Search slopes and baselines were determined by fitting least-squares lines to the three corrected RT averages. To ensure that observers did not complete the search task before at least one alternation of images, observers were checked to see that they had a search speed faster than 15 ms/item. All of the observers met this criterion.

The cost of transformation was analyzed based on the logarithm of the ratio of the slopes for the two conditions under investigation. The ratio of the slopes provides a way to compensate for individual differences in speeds. The logarithmic transformation allows symmetry of these ratios to be re-established. For example, a 2:1 relationship would result in a ratio of 1/2 or 2, which are not equally spaced from 1; the logarithms of these values are equally spaced. More generally, examination of the data shows that the log ratios of the corrected slopes provide the closest fit to a normal distribution; statistical tests based on this are generally more robust. Two-sided z-tests were used to determine if the average log ratio for each set of observers differed significantly from zero; that is, if the average ratio differed significantly from one. Paired t-tests on absolute slopes provide much the same pattern of results as the z-tests. The repeated measures design is maintained because the ratios of each observer's performance are used, and we use slope ratios as the most suitable basic unit of analyis.

3 Results

Baselines did not differ for any of the conditions examined, and will not be discussed further. Absolute slopes for the no-grid conditions are shown in Figure 4 Left. The perceptual costs of the transformation, in terms of slope ratios, are shown in Figure 4 Right. The following table summarizes the statistical significance of the differences for the no-grid conditions.

Level	Static	p	Dynamic	p
1	1.53±.30	<.01	0.98±.11	>.6
2	1.40±.25	<.02	1.12±.09	=.05
3	2.69±.50	<.0001	1.10±.25	>.3

From these figures, two interesting trends appear. First, there is a considerable cost in regards to the static aspect of the transformation, with a reliable reduction in speed even at the lowest magnitude of transformation investigated. This is likely due to items in the transformed images often being extremely close together and so interfering with search, as by overlapping each other. It could also be due to the transformed items simply being more difficult to distinguish, and thus requiring more inspection time per item.

Of more interest is the second trend, which is the minimal effect of the dynamic aspect of the transformation; that is, the effects caused by the actual change itself. Although the difference in speed approaches significance for level 2, this is only a marginal effect, and is at most a 10% slowdown. Thus, even though the two static images differ considerably, there is evidently little cost in sudden switches between them, even at the highest magnitude of the transformation. The dashed horizontal line in the graph shows the worst-

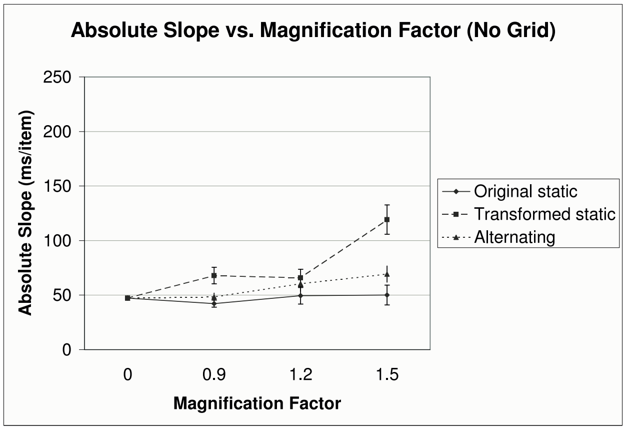

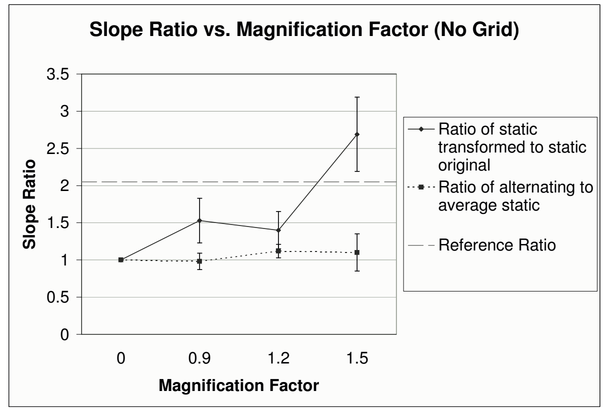

Figure 4: Results for the no-grid condition. **Left:** Absolute slopes. The value of the absolute slopes at $c = 0$ is defined to be the average of the original static slopes at $c = 0.9$, $c = 1.2$ and $c = 1.5$. **Right:** Ratios between slopes show perceptual costs of the transformation: the upper solid line shows the static transformation cost, and the lower dashed line shows the dynamic cost.

case dynamic cost of 2, the point at which the observer could simply use the easiest display if the other one was too difficult. In fact, our observed slope ratios are far less than this. The following table summarizes the statistical significance of the differences for the grid conditions.

Level	Static	p	Dynamic	p
1	$1.27\pm.23$	$=.07$	$0.92\pm.12$	$>.7$
2	$1.84\pm.34$	$<.0001$	$1.13\pm.14$	$>.1$
3	$2.79\pm.37$	$<.0001$	$1.16\pm.14$	$=.08$

The effect of a grid on performance is shown in Figure 5, with the absolute slopes on Figure 5 Left. Figure 5 Right shows this data in terms of the perceptual costs involved.

The same two interesting trends appear again. Although the static cost at the lowest level is only marginally significant, all the other conditions show a strong slowdown in search speed, exactly as in the case of the no-grid conditions. Moreover, a two-way ANOVA of level and grid showed a main effect of level on slope ratios ($F(2,66) = 12.9; p < .0001$), but no main effect of the grid ($F(1,66) = 0.117; p > .7$), nor any significant interaction of the grid with level ($F(2,66) = 1.34; p > .25$), indicating, that the effect of the grid on both original and transformed images is the same. However, a three-way ANOVA of the log slopes shows that the grids increased the value of slopes significantly ($F(1,132) = 32.27; p < .0001$), although there was no interaction with level ($F(2,132) = 2.11; p > .1$) or type of image ($F(1,132) = 0.12; p > .7$), nor were there any significant 3-way interactions ($F(2,132) = 1.08; p > .3$). This indicates that the grid interferes with the task entirely by slowing down search, with the grid perhaps simply acting as a form of visual clutter.

Importantly, the second trend is also found: the minimal dynamic costs are incurred by the transformation, even at the highest magnitudes examined. A two-way ANOVA of level and grid showed no main effect of level on dynamic slopes ($F(2,66) = 1.75$; $p > .15$), no main effect of grid ($F(1,66) < 0.001$; $p > .9$), nor any significant interaction of the grid with them (F(2,66) = 0.14; $p > .85$), indicating that the grid has little effect on dynamic cost.

4 Discussion

The experiments described here provide three interesting results in regards to the use of nonlinear transformations in Focus+Context systems.

First, they show that such transformations can have a considerable static perceptual cost, with performance slowed down by a factor of almost 3 under some conditions. An important question concerns the cause of this slowdown. It may be that it is entirely due to the crowding found in those areas away from the focus, or that it is due to the distortions in the items themselves. If so, it may be possible to redesign systems so as to avoid these pitfalls. The question of how to remove crowding and clutter yet maintain sufficient context is at the heart of Focus+Context techniques. The question of whether items should be themselves be distorted is as yet unresolved. Nonlinear transformations can be shown visually through any combination of global movements of the surrounding space and local changes that affect the objects as well. At one extreme, the size and shape of items would be completely unaffected by the transformation of the space, so only their position would change. This set of experiments investigated a middle ground where a rigid 2D local transformation was applied to each item so that both position and scale changed, but the item remained parallel to the image plane. The physical metaphor for this behavior would be pinning an item to the bulging background surface at a single spot, so that the bulge brings it closer to the viewer. Another approach used in some Focus+Context systems is to apply a nonlinear local distortion transformation to the shape of the individual items, as if they were painted directly onto the bulging surface rather than being pinned at one only point. These issues would benefit from further investigation.

Another interesting result is the existence of a no-cost zone of considerable extent, at least in terms of the dynamic aspects of the transformation. There appears to be no effect of sudden transformations for magnitudes of up to approximately c=1.0 (i.e., about midway between levels 1 and 2). There is also a large low-cost zone beyond that, where performance is only marginally affected. Importantly, this zone covers the entire range of magnitudes examined here, which include those capable of producing large distortions of

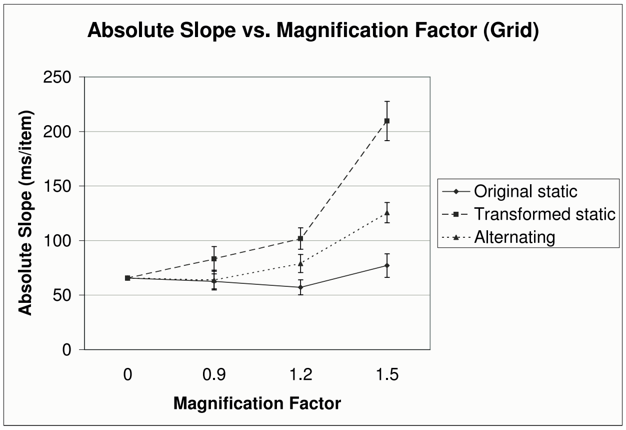

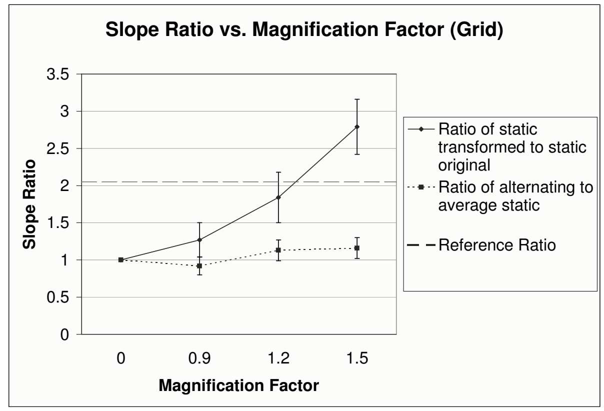

Figure 5: Results for the grid condition, showing no improvement in performance and slowdown in search speed. **Left:** Absolute slopes. The value of the absolute slopes at $c = 0$ is defined to be the average of the original static slopes at $c = 0.9$, $c = 1.2$ and $c = 1.5$. **Right:** Ratios between slopes show perceptual costs of the transformation.

the image. If, as seems likely, the static costs of such transformations could be reduced, this would mean that at least some forms of nonlinear transformation would not carry a high perceptual cost, even at high magnitudes of transformation.

Regardless of whether static costs can be reduced, this finding is also important in that it undermines the common belief that immediate transformations to an image, for example orientation or size, are difficult for observers to track and thus that transitions between views should be smoothly animated [Robertson et al. 1989]. As some of these transformations do not seem to affect performance, it may not be necessary to compensate at all. However, the common belief may be motivated by reasons other than the ability or inability to compensate for dynamic image transformation, so further investigation would be of interest.

The third interesting result is the almost negligible positive effect of the grid on performance. An intuitive prediction might be that the grid would enhance performance in the form of a visual cue which explicitly conveys the distorted space and reduces the cognitive load of the user. It could also provide a consistent frame of reference between the original and transformed images, thus bridging the two. However, our results show little effect on performance apart from a general slowing down of search. There are several possible explanations for this phenomenon. Since the grid is rectangular and the transformation radial, this mismatch may have been enough to suppress the effect. Secondly, the grid lines might have been too thin or too far apart. However, we know that this is not true in our experiment because the grid was so salient that it was was a source of visual clutter. Finally, it may be that grid lines only help when the transformation cannot be coped with easily. Given that the transformations here could be compensated for by the human visual system, the grid lines may have been superfluous, and so act as clutter. In that case, more sophisticated visual layering to make the grid less visually obtrusive than the foreground items might avoid the visual clutter when the grid is unnecessary, yet provide sufficient information when needed. This issue requires further investigation.

Although it is worth keeping in mind that these results are based on visual search tasks, there are good grounds for believing that they can be applied more generally. To begin with, visual search is a common component of many of the visual operations in information visualization [Ware 2000], so that the results obtained here should apply to many, if not most, of its aspects. In addition, visual search appears to draw upon representations that are created relatively early in the visual system, forming a common base to many other visual tasks [Rensink 2000]. If the representations involved with visual search have a particular property, such as being invariant to a particular transformation, it is therefore likely that this property will be inherited by those processes that also rely on these representations. Ultimately, of course, the real test of these results will be whether Focus+Context systems can usefully take advantage of them.

It may be worth pointing out that the shaker paradigm could also be used in a more adventurous way: to determine which transformations can be compensated for by the human observer. It has been found that some transformations, such as size changes, can be compensated for, while others like rotations of 30° or more cannot [Rensink 2004]. Those transformations that can be compensated for might be considered as *natural* ones; indeed, the invariance found here for fisheye transformations suggests that these too are natural, and may be the reason why Focus+Context systems do not confuse viewers. In any event, it is clear that the techniques described here can be applied to other transformations as well. As such, they can provide a ranking of transformations in terms of their perceptual cost, along with a determination of their no-cost and low-cost zones. And, perhaps, one day, these techniques may even provide the means of determining optimal transformations.

5 Conclusions

This work has shown that the shaker paradigm [Rensink 2004] can be used to investigate the effects of nonlinear distortions in Focus+Context systems. In particular, it can provide measures of the perceptual cost of both the static and dynamic aspects of such transformations. The experiments described here show that there can be a significant perceptual static cost to these transformations, likely due to the compression and distortions of the items in the image. They also show that observers are largely unaffected by the dynamic aspect of the transformation, with performance only marginally impaired by sudden switches between original and transformed images. In addition, they have provided a first approximation of the extent of this no-cost zone. Results also indicate no improvement when a rectangular grid that distorts to make the transformation effect more visually obvious in the image. The grid slows down the search task in both the distorted and the undistorted views, apparently acting as visual clutter. The shaker paradigm is a powerful tool to investigate Focus+Context systems, and further studies with this paradigm may provide useful guidance for better, perhaps even optimal, designs of information visualization systems.

6 Acknowledgements

We wish to thank Sandy Chuang and David Loron for their assistance in running the experiments. Thanks to Ciarán Llachlan Leavitt for technical writing contributions. We also wish to thank the Natural Sciences and Engineering Research Council (Canada) and Nissan Motor Co. Ltd. (Japan) for their support of this project.

References

BEDERSON, B. B., CLAMAGE, A., CZERWINSKI, M. P., AND ROBERTSON, G. G. 2004. DateLens: A Fisheye Calendar Interface for PDAs. *Trans. Computer-Human Interaction (TOCHI) 11*, 1 (March), 90–119.

BENNETT, D. J. 2002. Evidence for a Pre-Match 'Mental Translation' on a Form-Matching Task. *Journal of Vision 12*, 50.

BUNDESEN, C., AND LARSON, A. 1975. Visual Transformation of Size. *Journal of Experimental Psychology: Human Perception and Performance 1*, 214–220.

CARPENDALE, M. S. T., AND MONTAGNESE, C. 2001. A Framework for Unifying Presentation Space. In *Proc. Symp. User Interface Software and Technology (UIST)*, ACM, 61–70.

CARPENDALE, M. S. T., COWPERTHWAITE, D. J., AND FRACCHIA, F. D. 1995. 3-Dimensional Pliable Surfaces: for the Effective Presentation of Visual Information. In *Proc. Symp. User Interface Software and Technology (UIST)*, ACM, 217–226.

CUTTING, J. E. 1987. Rigidity in Cinema Seen from the Front Row, Side Aisle. *Journal of Experimental Psychology: Human Perception and Performance 13*, 323–334.

CUTTING, J. E. 1991. *On the Efficacy of Cinema, or What the Visual System did not Evolve to Do*. London: Taylor & Francis, 486–495.

FURNAS, G. W. 1986. Generalized Fisheye Views. In *Proc. SIGCHI*, ACM, 18–23.

GUTWIN, C., AND FEDAK, C. 2004. A Comparison of Fisheye Lenses for Interactive Layout Tasks. In *Proc. Graphics Interface 2004*.

GUTWIN, C., AND FEDAK, C. 2004. Interacting with Big Interfaces on Small Screens: a Comparison of Fisheye, Zoom, and Panning Techniques. In *Proc. Graphics Interface*.

GUTWIN, C., AND SKOPIK, A. 2003. Fisheyes are Good for Large Steering Tasks. In *Proc. SIGCHI*, ACM, 201–208.

GUTWIN, C. 2002. Improving Focus Targeting in Interactive Fisheye Views. In *Proc. SIGCHI*, ACM, 267–274.

HORNBAEK, K., BEDERSON, B., AND PLAISANT, C. 2002. Navigation Patterns and Usability of Zoomable User Interfaces with and without an Overview. *Trans. Computer-Human Interaction (TOCHI) 9*, 4, 362–389.

KEAHEY, T. A., AND ROBERTSON, E. L. 1996. Techniques for Non-Linear Magnification Transformations. In *Proc. Symposium on Information Visualization (InfoVis)*, IEEE, 38–45.

KEAHEY, T. A., AND ROBERTSON, E. L. 1997. Nonlinear Magnification Fields. In *Proc. Symposium on Information Visualization (InfoVis)*, IEEE, 38–45.

KOBSA, A. 2003. User Experiments with Tree Visualization Systems Technical Report. Tech. Rep. UCI-ISR-03-11, University of California Irvine Department of Computer Science.

LAMPING, J., RAO, R., AND PIROLLI, P. 1995. A Focus+Content Technique Based on Hyperbolic Geometry for Viewing Large Hierarchies. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI '95)*, ACM, 401–408.

LEUNG, Y. K., AND APPERLEY, M. D. 1994. A Review and Taxonomy of Distortion-Oriented Presentation Techniques. *Trans. Computer-Human Interaction (TOCHI) 1*, 2 (June), 126–160.

MUNZNER, T., GUIMBRETIERE, F., TASIRAN, S., ZHANG, L., AND ZHOU, Y. 2003. TreeJuxtaposer: Scalable Tree Comparison using Focus+Context with Guaranteed Visibility. In *Proc. SIGGRAPH 2003, published as Transactions on Graphics*, vol. 22, ACM, 453–462.

MUNZNER, T. 1997. H3: Laying Out Large Directed Graphs in 3D Hyperbolic Space. In *Proc. Symposium on Information Visualization (InfoVis)*, IEEE, 2–10.

RAO, R., AND CARD, S. K. 1994. The Table Lens: Merging Graphical and Symbolic Representations in an Interactive Focus+Context Visualization for Tabular Information. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI '94)*, ACM, 318–322.

RENSINK, R. A. 2000. The Dynamic Representation of Scenes. *Visual Cognition 7*, 17–42.

RENSINK, R. A. 2004. The Invariance of Visual Search to Geometric Transformation. *Journal of Vision 4*. [Vision Sciences Society, Sarasota, FL, USA. May 2004.].

RISDEN, K., CZERWINSKI, M. P., MUNZNER, T., AND COOK, D. 2000. An Initial Examination of Ease of Use for 2D and 3D Information Visualizations of Web Content. *International Journal of Human Computer Studies 53*, 5 (November), 695–714.

ROBERTSON, G. G., AND MACKINLAY, J. D. 1993. The Document Lens. In *Proc. Symp. User Interface Software and Technology (UIST)*, ACM, 101–108.

ROBERTSON, G. G., CARD, S. K., AND MACKINLAY, J. D. 1989. The Cognitive Coprocessor Architecture for Interactive User Interfaces. In *Proc. Symp. User Interface Software and Technology (UIST)*, ACM, 10–18.

SARKAR, M., AND BROWN, M. H. 1994. Graphical fisheye views. *Communications of the ACM 37*, 12 (Dec.), 73–84.

SARKAR, M., SNIBBE, S. S., TVERSKY, O. J., AND REISS, S. P. 1993. Stretching the Rubber Sheet: A Metaphor for Viewing Large Layouts on Small Screens. In *Proc. Symp. User Interface Software and Technology (UIST)*, ACM, 81–91.

SCHAFFER, D., ZUO, Z., GREENBERG, S., BARTRAM, L., DILL, J., DUBS, S., AND ROSEMAN, M. 1996. Navigating Hierarchically Clustered Networks through Fisheye and Full-Zoom Methods. *Trans. Computer-Human Interaction (TOCHI) 3*, 2 (June), 162–188.

SHEPARD, R. N., AND METZLER, J. 1971. Mental Rotation of Three-Dimensional Objects. *Science 171*, 701–703.

TARR, M. J., AND PINKER, S. 1989. Mental Rotation and Orientation-Dependence in Shape Recognition. *Cognitive Psychology 21*, 233–282.

TREISMAN, A., AND GORMICAN, S. 1988. Feature Analysis in Early Vision: Evidence from Search Asymmetries. *Psychological Review 95*, 15–48.

TREISMAN, A. 1985. Preattentive Processing in Vision. *Computer Vision, Graphics, and Image Processing 31*, 156–177.

WARE, C. 2000. *Information Visualization: Perception for Design*. Morgan Kaufmann/Academic Press.

ZHANG, J. 1991. The Interaction of Internal and External Representations in a Problem Solving Task. *Proc. Thirteenth Annual Conference of Cognitive Science Society*.