

Perception and Attention for Visualization

By

STEVE HAROZ

B.S. (Clemson University) 2004

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Computer Science

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

David Whitney, Chair

Michael Neff

Ken Joy

Committee in Charge

2013

CONTENTS

Publication List	v
Peer Reviewed Papers	v
Published Abstracts	vi
Abstract	vii
1. Introduction.....	1
2. Natural Visualizations	4
Abstract	4
Introduction.....	4
Background.....	5
Spatial Frequency	5
Fourier Transforms	5
Natural and Unnatural Images	6
Size Distribution	8
Extending Existing Theories	11
Images without occlusion.....	11
Average deviation	14
Measuring Natural Visualizations	15
Testing competition results.....	15
Testing user performance	18
Improving visualizations	19
Utility or art	20

Limitations and future work	20
Conclusions.....	22
3. Temporal Thresholds for Feature Detection in Flow Visualization	23
Abstract	23
Introduction.....	23
Motion Perception	25
Optic Flow in Visualization	27
The Experiment	28
Stimuli.....	29
Task	31
Results	31
Discussion	33
Application.....	34
Conclusions and Future Work	34
4. How Capacity Limits of Attention Influence Information Visualization Effectiveness	36
Abstract	36
Introduction.....	37
Experiment Stimuli	39
Visual Search: Find a Known Target	41
Methods	41
Results.....	42
Visual Search: Find the Oddball.....	43

Methods	43
Results	43
Subitizing: Limited capacity	47
Methods	48
Results	49
Information Visualization Guidelines	50
Grouping greatly helps for some but not all tasks	50
If you cannot group, change the task.....	51
When there are many categories: Less is more	51
Assigning a visual feature to a data dimension	52
Evaluation.....	52
Future Work	53
Conclusion	53
5. Conclusion	56
6. References	57

PUBLICATION LIST

PEER REVIEWED PAPERS

How Capacity Limits of Attention Influence Information Visualization Effectiveness

[Best Paper Award]

[Steve Haroz](#) and David Whitney

InfoVis 2012

Reference repulsion in the categorical perception of biological motion

Timothy D. Sweeny, [Steve Haroz](#), and David Whitney

Vision Research 2012

Perceiving group behavior: Sensitive ensemble coding mechanisms for biological motion in human crowds

Timothy D. Sweeny, [Steve Haroz](#), and David Whitney

Journal of Experimental Psychology: Human Perception and Performance 2012

Seeing the Difference between Cosmological Simulations

[Steve Haroz](#) and Katrin Heitmann

Computer Graphics and Applications 2008

Multiple Uncertainties in Time-Variant Particle Data

[Steve Haroz](#), Kwan-Liu Ma, and Katrin Heitmann

Pacific Visualization 2008

Layout of Multiple Views for Volume Visualization: A User Study

Daniel Lewis, [Steve Haroz](#), and Kwan-Liu Ma

Proceedings of Visual Computing 2006

Natural Visualizations

[Steve Haroz](#) and Kwan-Liu Ma

Proceedings of EuroGraphics Visualization Symposium 2006

PUBLISHED ABSTRACTS

Global - Not Local - Variance Impacts Search

[Best Poster Award]

Steve Haroz and David Whitney

Journal of Vision 2013

Just walk away: Reference repulsion in the perception of crowd behavior

Timothy D. Sweeny, Steve Haroz, and David Whitney

Vision Sciences Society 2012

Seeing the direction of a crowd: Ensemble coding of biological motion

Timothy D. Sweeny, Steve Haroz, and David Whitney

Vision Sciences Society 2011

Temporal thresholds for feature detection in flow visualization

Steve Haroz and David Whitney

Applied Perception to Graphics and Visualization 2010

Free Your Data! Cenimation: Visualization for Constrained Displays

[Creative Winner in InfoVis contest]

Sophie Engle, James Shearer, Michael Ogawa, Steve Haroz, Kwan-Liu Ma

InfoVis 2006

ABSTRACT

This work examines how a better understanding of visual perception and attention can impact visualization design.

In a collection of studies, I explore how different levels of the visual system can measurably affect a variety of visualization metrics. The results show that expert preference, user performance, and even computational performance are all related to how our eyes and brain work to see the world.

1. INTRODUCTION

Visualization is ultimately a translation of binary data to awareness and action using an image as a medium. While the process has many individual components, the true determinant of visualization utility lies in the success of the overall process. Much – arguably most – research in visualization seeks to advance and optimize the computational and rendering techniques necessary to make data into an image. My aim, however, is to improve our understanding of what happens once the information leaves the monitor and enters the eye. This research seeks to inform visualization design using a better understanding of the human visual system (HVS). Put simply, how can understanding the HVS improve preference and performance of visualizations?

The process of perceiving visual information includes many stages of encoding. Lower levels of vision in the retina, lateral geniculate nucleus, and striate cortex encode simple features like edge scale and orientation [Hubel & Wiesel, 1959, 1962]. The receptive field – the region of the visual field where a neuron is sensitive – of neurons implicated in these simple features is typically very small, encompassing a fraction of a degree of visual angle. This information feeds into regions that encode progressively more complex visual features with progressively larger receptive fields. Located in the extrastriate cortex, temporal lobe, and parietal lobe, these ‘middle’ levels of vision encode more complex features such as motion fields, faces, bodies, and places [Chalupa & Werner, 2004]. Attentional selection and working memory, implicated throughout the entire brain [Gazzaniga, Ivry, & Mangun, 2009], incorporate information from the lower and middle levels of vision to filter and analyze the encoded information.

These different regions of the HVS do not strictly feed forward to areas that encode more complex visual features. In fact, the lateral geniculate nucleus only receives ten percent of its input from the retina [Sillito, Cudeiro, & Jones, 2006]. The rest comes from other areas of the brain, demonstrating the significant amount of feedback that occurs. With heavy interconnection throughout the brain, it perhaps comes as no surprise that all components of the HVS impact our ability to perceive information in an image. This work will present three projects that show how visual features ranging from low level scale to higher level motion to scene-wide attention can all inform effective visualization design.

The goal in “Natural Visualizations” [Haroz & Ma, 2006] was to use scale, a feature coded early in the HVS, to see if we could emulate judgments of visualization quality. Because the definition of a “good visualization” is not clear, expert preferences – in the form of the InfoVis contest – were substituted. Inspired by the finding that images with ‘natural’ statistics are more efficiently encoded by the HVS [Olshausen & Field, 1996], we showed that visualizations with more natural image statistics tended to perform better in the contest. Some aspect of the judges’ preference was correlated to how efficiently the HVS encoded the image. Furthermore, we showed that systematically controlling the scale and luminance of objects in a scene could yield these natural statistics. The result of the study was a simple procedure for systematically scaling the objects in a visualization to produce a more efficiently encoded – and perhaps, more preferred and performant – visualization.

Looking beyond lower level perception of scale, we found that taking advantage of higher level areas of visual encoding – namely, motion and optic flow – allows us to improve the computational performance of some visualizations. In “Temporal thresholds for feature detection in flow visualization” [Haroz & Whitney, 2010], we explored how some particle tracers were wasting effort by maintaining orders of magnitude more glyphs for longer than the visual system could actively track them [Pylyshyn & Storm, 1988]. Even if hundreds or thousands of particles are traced and rendered, the visual system is encoding the information as a combination of simple motions [Duffy & Wurtz, 1991] rather than a collection of individual glyphs. Therefore, any computational effort to maintain glyph trajectory continuity may be unnecessary. Past works [Park, Budge, Linsen, Hamann, & Joy, 2005; Helgeland & Elboth, 2006] had described techniques for using short glyph lifetimes in particle tracers to avoid costly overhead of glyph density maintenance, but no empirical work was done to test how brief the glyph lifetimes could be. To answer that question, we ran an experiment in which subjects adjusted glyph lifetime to a threshold in which they could just barely find a feature (such as a vortex) in a flow field. The results show that subjects can detect and locate a feature even when glyphs last for a mere fraction of a second. Such brief glyph lifetimes negate the necessity for computationally expensive algorithms used to maintain glyph density or counteract error buildup over time (e.g. Runge Kutta) in real-time displays. Details, such as glyph continuity, that cannot be perceived don’t need to be calculated or rendered.

The limited ability to attend to multiple individual objects raises the question of how attentional limitations impact user performance for visualizations. In “How Capacity Limits of Attention Influence Information Visualization Effectiveness” and “Global – Not Local – Variance Impacts Search” [Haroz & Whitney, 2012, 2013], we ran a collection of experiments to see how attentional limits impact user performance for a collection of visualization scenarios. Limitations of attentional capacity force people to perform attentionally demanding tasks in a very inefficient serial manner and limit the amount of simultaneously operable information. In other words, having a serial bottleneck can severely slow an otherwise *embarrassingly parallel* process. Our results showed that while some scenarios were unaffected by attentional bottlenecks, others yielded very poor user performance as complexity increased. We concluded with a collection of guidelines for how user task, layout, and visual mapping can all impact attentional demand and, therefore, user performance.

Information about the HVS from low level scale analysis to higher level motion encoding to attentional selection can all give us a better understanding of what happens after a visualization leaves the monitor. As these and many other works show, taking advantage of that information can help make visualizations that are more preferred, computationally faster, and do a better job of getting data to the user.

2. NATURAL VISUALIZATIONS

Published in EuroVis 2006 by Steve Haroz and Kwan-Liu Ma

ABSTRACT

This chapter demonstrates the prevalence of a shared characteristic between visualizations and images of nature. We have analyzed visualization competitions and user studies of visualizations and found that the more preferred, better performing visualizations exhibit more natural statistics. Due to our brain being wired to perceive natural images [Simoncelli & Olshausen, 2001], testing a visualization for statistics similar to those of natural images may help show how well our brain is capable of absorbing the data. In turn, a metric that finds a visualization's similarity to a natural image may help determine the effectiveness of that visualization. We have found that the results of comparing the sizes and distribution of the objects in a visualization with those of natural standards strongly correlate to one's preference of that visualization.

INTRODUCTION

Many have tried to define or explain what properties make for an effective visualization [Tufte, 1990; Sears, 1995; Shneiderman, 1996]. Elements such as aspect ratio, density, color usage, and typeface variation have been examined separately and in combination with varying degrees of success, but no one has given a concrete reason as to why these aspects of a visualization are important [Sears, 1995; Brath, 1997]. Even when measuring proximity and clustering, little evidence is given to explain why these features are helpful [Brath, 1997]. Layout is also a commonly scrutinized aspect of visualizations, but the comparison of layout techniques frequently requires exhaustive user studies involving eye-tracking [Sears, 1995], which is expensive and time-consuming. Though these user studies can give concrete proof that one implementation better conveys information than does another, they are incapable of explaining the underlying reasons behind why some visualizations provide more insight than do others. We therefore aim to explain a more fundamental property that correlates strongly with preferred visualizations.

Visualization is ultimately a field of translation. The goal of every visualization is to convert raw binary data from a machine-readable encoding to a neural encoding understandable to the mind. In order to be successfully converted, this data must pass through the visual system. Therefore, tuning visualizations to the visual system should help in their effective translation or perception. In this chapter, we demonstrate that theories known to neuroscientists who study vision can be effectively utilized to find patterns in the effectiveness and preference of visualizations. As a first step, we replicate computational neuroscience experiments and extend them to show that their image analysis techniques can be applied to images with properties similar to those of visualizations. After inspecting measurements of these images, we apply those measurements to actual visualizations to find a correlative pattern. We then examine how this pattern's categorization of visualizations is similar to that of the InfoVis contest results [Fekete, Grinstein, & Plaisant, 2004; Grinstein, Cvek, Derthick, & Trutschl, 2005] and a timed user study. These techniques could potentially contribute to the advancement and expansion of our understanding of visualization perception and may help influence future visualization developments and applications rooted in neuroscience foundations.

BACKGROUND

For over two decades now, neuroscientists have studied images of nature to better understand how our visual system perceives them [Field, 1987, 1993]. They have found that these natural images contain distinctive statistical regularities that random images lack and that our brains are wired to perceive this natural stimulus [Olshausen & Field, 1996; Simoncelli & Olshausen, 2001]. One of the fundamental distinctions of natural images is their sizing and spacing which can be analyzed by observing the images' spatial frequencies [Field, 1987].

SPATIAL FREQUENCY

Spatial frequencies are similar to sound frequencies. Sound frequencies are a measurement of compression varied over time, whereas spatial frequencies are a measurement of intensity varied over distance [Field, 1987]. Since spatial frequencies can only measure a single intensity value, brightness is commonly used.

FOURIER TRANSFORMS

One way to measure the spatial frequencies of a function is by using Fourier transforms. Essentially, sine and cosine waves of different amplitude and frequency are added together to form the intended function. These sine and cosine functions make up a Fourier series. For a two dimensional image, the Fourier transforms are performed over each line in the horizontal axis then over each line in the vertical axis or vice versa. In turn, to find the two-dimensional Fourier transform of an n -by- n image, one must find $2n$ one-dimensional Fourier transforms.

NATURAL AND UNNATURAL IMAGES

A natural image is any picture of nature [Field, 1987]. Pictures of a forest scene, a mountain, or a dog would be considered natural images. Three examples can be seen in Figure 2-1. This class of images constitutes an infinitely small fraction of all possible images [Reinhard, Shirley, & Troscianko, 2001], yet our visual system is precisely tuned to perceive them rather than some larger range of image types. To measure the spatial frequency distribution of these images, one begins by computing the Fourier transform. The rotational average of the two dimensional result yields a more manageable, one dimensional series also known as a power spectrum [Weisstein, n.d.]. When the amplitude of this spectrum is plotted on a log-log scale as a function of frequency, the spatial frequency distribution can be visualized.

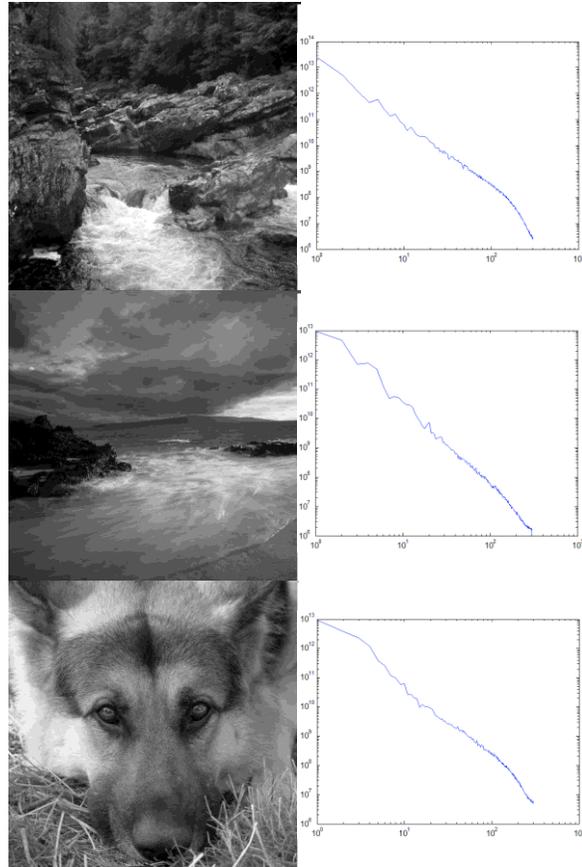


Figure 2-1. Left: These are examples of natural images. Right: These plots correspond to the power spectra of the images on the left. The x-axis is the frequency on a log scale, and the y-axis is the amplitude which is also on a log scale.

On the right of Figure 2-1, plots of the power spectra from the natural images are shown. These plots have nearly straight lines with slopes of approximately -2, which corresponds to an f^{-2} trend. The consistency between the plots is not trivial, as these images appear quite dissimilar. Unnatural images have very different power spectra. Figure 2-2 contains three unnatural images, and on the right their corresponding spatial frequency plots are shown. The distinctness of natural images becomes more evident in these plots, as the unnatural images do not show the f^{-2} trend. Many other papers [Field, 1987] have discussed exhaustive studies on large numbers of natural and unnatural images, and all have repeatedly found the same results. The f^{-2} trend sets apart natural images.

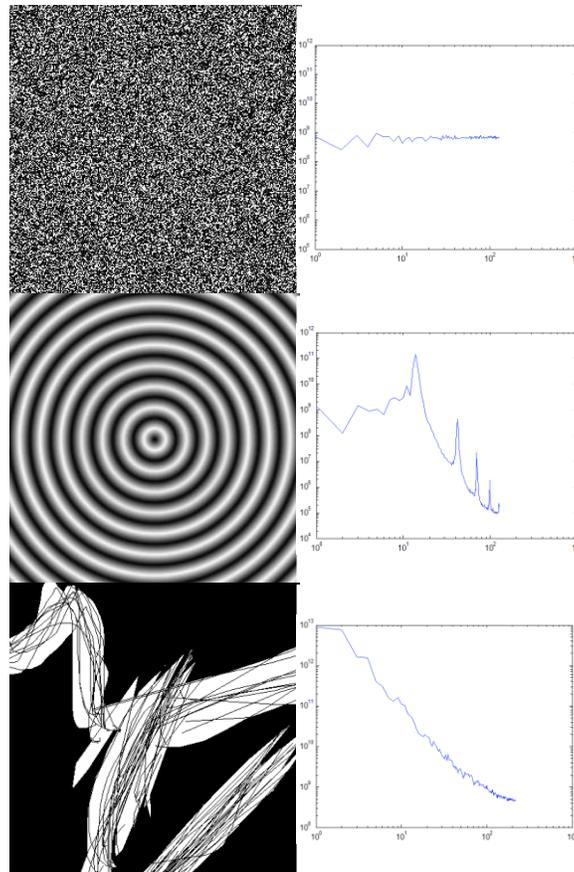


Figure 2-2. Left: These are examples of unnatural images. The first is purely random noise. The second is a radial gradient repeated ten times. The third is just a scribbling. Right: These plots correspond to the power spectra of the images on the left. The x-axis is the frequency on a log scale, and the y-axis is the amplitude which is also on a log scale.

Not all natural images will have a spatial frequency distribution of exactly f^{-2} . Image scaling as well as window size and shape can be used to explain why certain images deviate from the trend, yet whole groups of images such as sky scenes or images at a large scale have also been found to have unusual power spectra. Langer showed that these anomalous groups make up only a small subset of all natural images and tend to cancel out each other when large numbers of images are measured [Langer, 2000]. He also noted that these atypical natural images do not contain structure that is rich or interesting. For visualizations, images without interesting structure would be incapable of providing useful insight, and the corresponding visualization or its scale would be undesirable.

SIZE DISTRIBUTION

Daniel Ruderman [1997] investigated the cause of natural image characteristics being independent of calibration and visual environment. Many assume that natural image traits “result from edges, each with a power spectrum of $1/k^2$ ” [Ruderman, 1997]. Ruderman disproves this belief using contradiction. Instead, he shows that statistically independent ‘objects’ are the cause. These objects are placed randomly on top of each other and have a size distribution that follows a power function. The occlusion resulting from this collage of specifically sized objects produces the observed f^{-2} power spectrum.

Ruderman demonstrated this premise by producing images made up of a collage of squares. These squares are positioned randomly and given a random grayscale value that follows a Gaussian distribution. The size distribution of the squares is given by a power function or an exponential function. The images generated using the power function demonstrate a more natural power spectrum than those generated by the exponential function.

We have replicated Ruderman’s natural image generation test and extended it to include linear and constant distributions along with the power and exponential distributions demonstrated in his experiment. Samples of the generated images can be seen in Figure 2-3. Their corresponding power spectra can be seen in Figure 2-4, and a linear fit was applied to the plots to obtain the slopes of the trends in Figure 2-5. The trends show that the power function size distribution has a slope (-2.5) that is closer to -2 than those of the others (-2.6). In other words, the power function more closely emulates natural images.

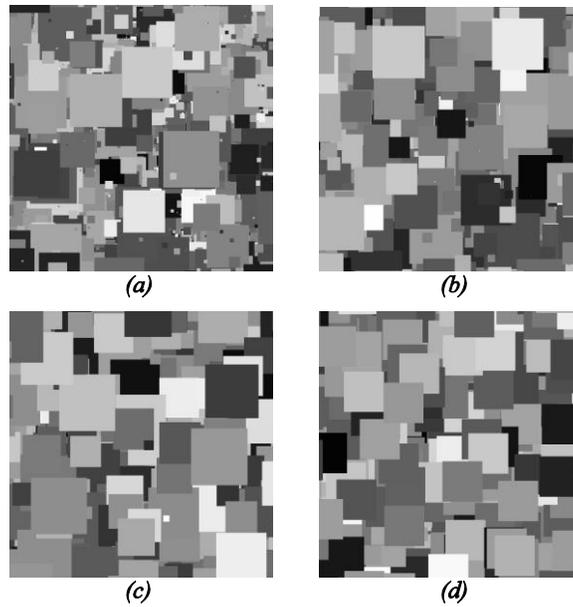


Figure 2-3. These images were generated by randomly placing squares with random Gaussian-fit grayscale values. The size distributions are power (a), exponential (b), linear (c), and constant (d).

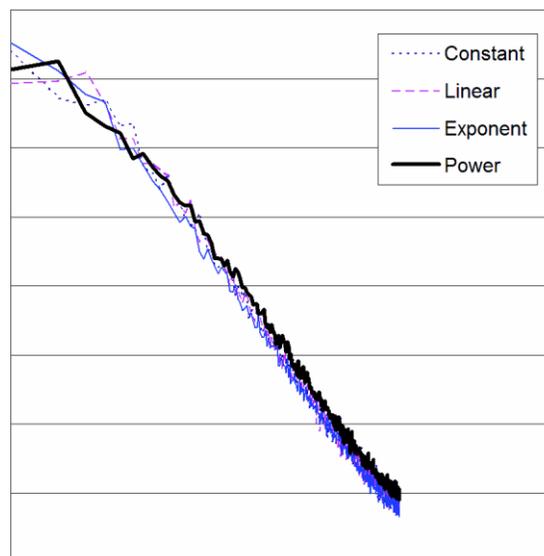


Figure 2-4. The power spectra of images generated by occluding squares. Notice that the power function stays slightly above the others.

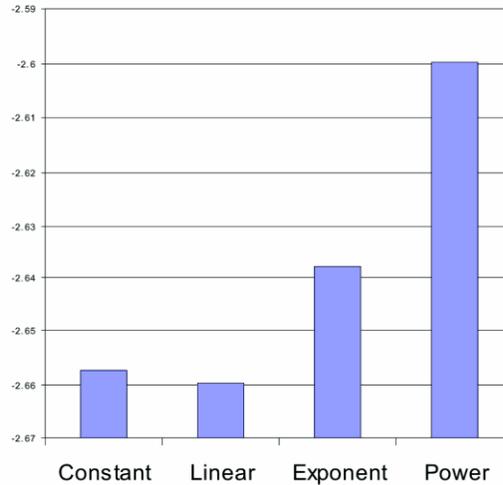


Figure 2-5. The slopes of the linear trends fit to the power spectra in Figure 2-4.

EXTENDING EXISTING THEORIES

In visualization, efforts are generally made to avoid occlusion. According to Ruderman, visualizations would therefore be incapable of having the characteristics of natural images. The implication is that we would have difficulty perceiving visualizations without occlusion because their power spectra would be different than that of our visual system. Accordingly, determining if the same size distribution rules apply to images without occlusion is crucial to knowing if these visualizations can exhibit natural traits.

IMAGES WITHOUT OCCLUSION

We proceeded to revise our image generation program to prevent occlusion. Doing so turned out to be more difficult than creating the original program, as placement cannot simply be random. The squares need to be placed in such a way that no square occludes any other. The resulting image should form an artificial visualization with measurable characteristics. For a given size image, n rows were created, where n is dependent on the size distribution formula. This sizing applies to area, not width, making n reliant on an unexpectedly complex formula:

$$h \geq \sum_{i=0}^n \sqrt{f\left(f^{-1}(\min) + \frac{i \cdot (f^{-1}(\max) - f^{-1}(\min))}{n}\right)}$$

In this equation, f is the size distribution function, h is the height of the image, min is the area of the smallest square, and max is the area of the largest square. The images of non-overlapping squares can be seen in Figure 2-6. The power spectra and their trends (as seen in Figure 2-7 and Figure 2-8) are similar to those of the images with occlusion. The power function's slope is closest to -2 followed by the exponential, linear, and constant functions respectively. Clearly, these images show that images without overlapping objects can have natural characteristics.

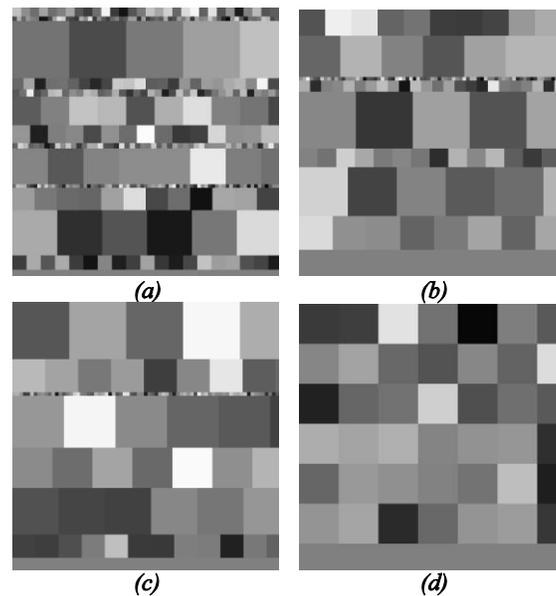


Figure 2-6. These treemap [JS91] resembling images were generated by creating rows of non-overlapping squares with random greyscale values. The size distributions are power (a), exponential (b), linear(c), and constant (d).

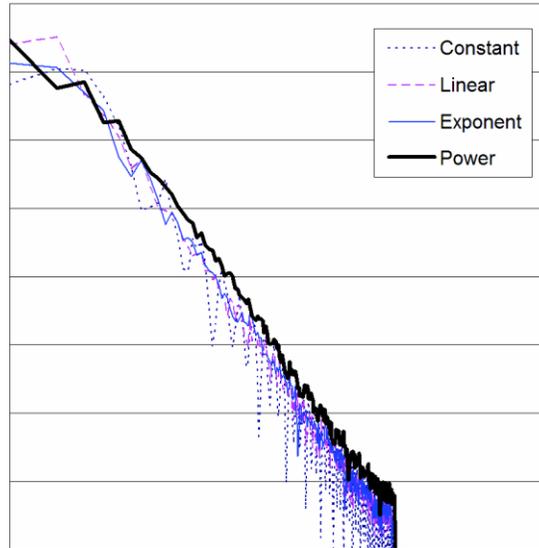


Figure 2-7. The power spectra of images generated by rows of non-overlapping squares

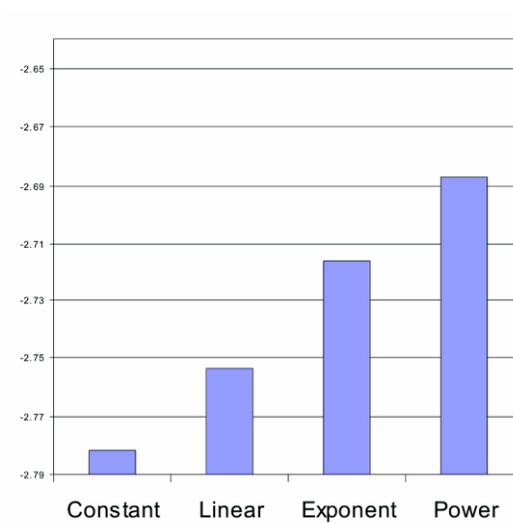


Figure 2-8. The slope of the linear trends fit to the power spectra of the non-overlapping squares.

The trends of these images have slopes that are all within a small range of less than 5%. Such a small range leaves too much room for these findings to result from a mere statistical anomaly. Many seemingly insignificant factors may have been the cause of one type of scaling seeming more natural than did another. We therefore addressed several potentially confounding factors in an attempt to deviate from the original outcome.

- Unusual run: This particular run of the image generating program could have resulted in a fluke, so the program was run several times.
- Row order: Each run had random row placement.

- Orientation: The rotational average of the two-dimensional Fourier transform treats all orientations equally.
- Extra space: The size of the image was set to precisely fit one scaling function for each run, thereby eliminating the black bar for that particular image.
- Image size: Image heights varying from 500 to 1200 pixels were used.
- Shape: Circles and randomized shapes were also used.

Despite the variation of all these factors, the power spectrum trends were not significantly affected. The order of the slopes always remained the same - power, exponential, linear, and then constant.

All of these slopes are well within the range of what could be considered a natural image [Ruderman, 1997].

However, due to their nearly identical values, a means of differentiation besides slope needs to be used in tandem.

AVERAGE DEVIATION

Natural images have power spectra that adhere very closely to a straight line, while the power spectra of unnatural images are more likely to have many spikes and steep dips. To determine how closely a spectrum follows its trend, its average deviation was calculated. The average deviation is found by averaging the absolute value of each point's difference between the actual power spectrum and its corresponding trend. This technique will help determine a plot's linearity. The average deviations for two natural images as well as those for the non-overlapping squares can be seen in Figure 2-9. The natural images are similar to the power and exponential distributions not the linear or constant distributions.

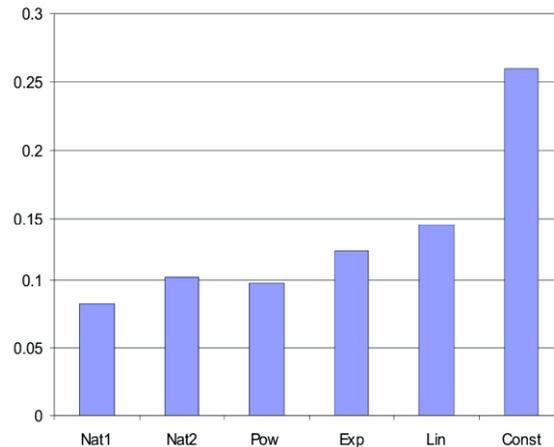


Figure 2-9. The average deviation from the linear-fit trends of the power spectra from two natural images as well as the images from Figure 2-6.

MEASURING NATURAL VISUALIZATIONS

Neuroscientists have shown that our brains most efficiently encode natural images [Field, 1987, 1993; Olshausen & Field, 1996; Reinhard et al., 2001; Simoncelli & Olshausen, 2001; Karklin & Lewicki, 2003]. The distribution of photo receptors in the retina follows an f^{-2} pattern as does the distribution of ganglion cell receptors immediately behind the retina. This pattern is persistent throughout the visual cortex. Therefore, visualizations that are most like natural images should be the most cognitional. In order to determine the extent of a visualization's natural characteristics, we propose measuring the slope of its power spectrum on a log-log scale as well as finding the deviation from a linear-fit trend. A visualization with natural characteristics should have a slope near -2 with a minimal average deviation. The slope is the most important factor because the deviation of a spectrum with a slope far from -2 is unimportant. We have shown that this metric produces predictable, reproducible results for artificial visualizations, so the following examples will demonstrate that natural characteristics correlate closely with the preference and performance of actual visualizations.

TESTING COMPETITION RESULTS

For one test, we looked at the InfoVis 2004 competition results [Fekete et al., 2004]. These visualizations all used the same dataset, which makes the comparison fairly objective. We analyzed an image for each of the first and second place winners (samples of which can be seen in Figure 10), and the results were better than even we

expected. An image of each of the visualizations was taken from the Information Visualization Benchmarks Repository (<http://www.cs.umd.edu/hcil/InfovisRepository/contest-2004/>). The images were then converted to grayscale and truncated to be square in size, as shown in Figure 2-10. Efforts were made to only truncate blank space around the sides. We then performed a spectral analysis of each of the images.

To compare the results, we found the distance of each slope from -2 . Incredibly, all of the first place winners had slopes within $.4$ of -2 , while the second place winners were mostly outside of that range.

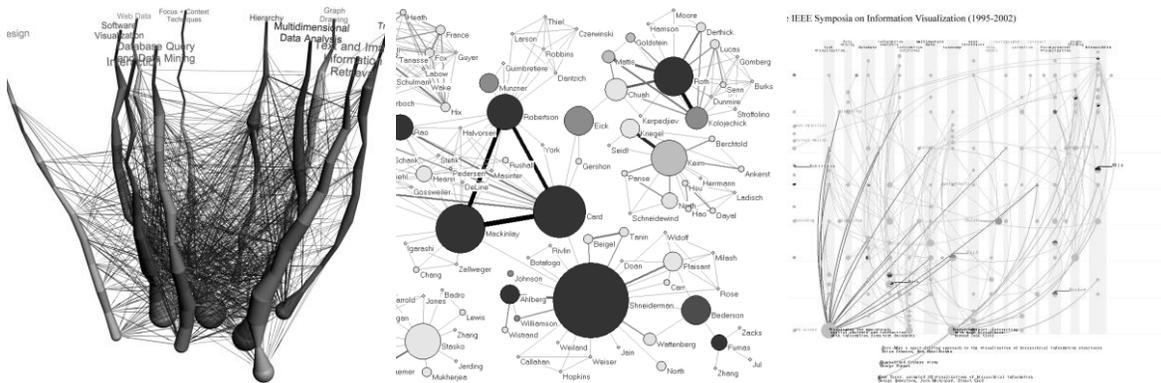


Figure 2-10. Here are select thumbnails from the InfoVis 2004 contest. The left two visualizations received 1st place prizes. The right visualization received a 2nd place prize.

Figure 2-11 shows that the image statistics were strongly correlated with the contest outcome $t=2.2, p<0.05$ (12 entries). The judges must have an unconscious preference for visualizations that are similar to natural images, as their evaluation accurately reflected the visualizations' natural measurements. The spatial frequencies have actually quantified the judges' preferences.

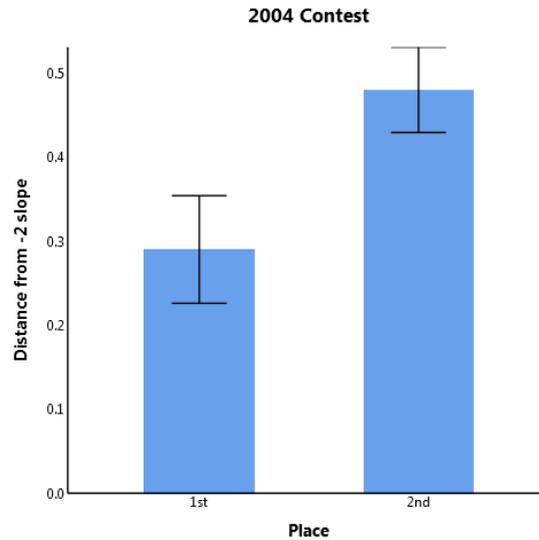


Figure 2-11. This graph shows our analysis results from the 2004 InfoVis contest. We measured the distance of the linear-fit trend from -2. We then took the average of those who came in first place and those who came in second place. The error bars show standard error.

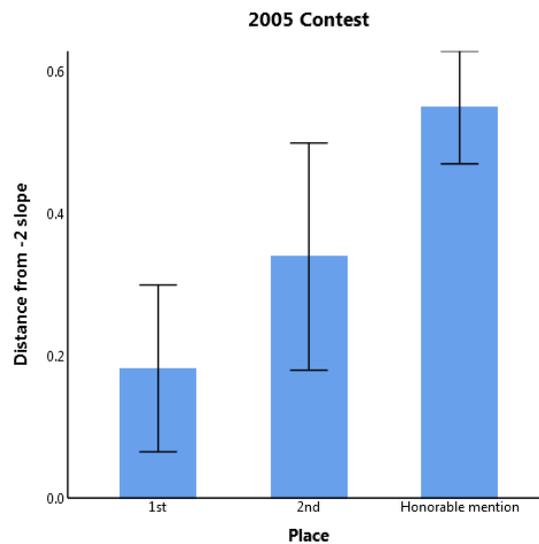


Figure 2-12. In the results from the 2005 InfoVis contest, a pattern between the first place, second place, and honorable mention averages is clearly prevalent. The error bars show the standard error. All of the entrants' images can be found at <http://ivpr.cs.uml.edu/infovis05>

To show that this was not a fluke, we tested the results of the next year's competition [Grinstein et al., 2005] as well, and although 2005 had only ten entries, the results were even clearer ($t=2.6, p<0.05$) due to the inclusion of an 'honorable mention' category (see Figure 2-12). Over both years, the effect of the image statistics on contest outcome was highly significant ($t=3.2, p<0.005$).

TESTING USER PERFORMANCE

Beyond simply predicting competition results, we can also show that a visualization's natural characteristics correlate with users' performance. To test the extent of this property, we analyzed the results of a user study of hierarchical visualizations.

The experiment's purpose was to time a user's ability to find structural similarities and differences within hierarchical data. The experiment looked at three interfaces that implement different hierarchical visualizations. Each user was assigned one interface which they could use to answer to six questions. The experimenters then recorded the time taken to answer each question. An important aspect of this experiment is that, as with the InfoVis competitions, everyone in the experiment used the same data. This aspect helps reduce unforeseen influences on the results.

After examining the average time taken by the users of each interface for each question, the experimenters found that Windows Explorer had the worst times, and the treemap and RINGS interfaces generally had similar times. We then used the information to correlate the time taken with the power spectra of screenshots. In this case, we not only looked for the relative order of the power spectra's slopes, but we also looked at the ratio of those slopes compared to the ratio of the times.

The results of the initial images analysis were predictable; Explorer has a power spectrum that is far from the natural standard, whereas the other two have more natural traits. We then plotted the correlation between the average response time and the distance of the power spectrum slope from -2 (Figure 2-13). Half of the questions had correlations with absolute values of around 90%. After examining the differences between the strongly and weakly correlated questions, we found that the strongly correlated questions required users to look for data that they were unlikely to have seen in a previous question. In other words, most of the weakly correlated questions were follow-up questions. For example, question one asked to find similar folders, and question two asked to find "very" similar folders. A probable cause of the distinction in correlation is that the influence of naturalness on an interface is strongest while the user is unfamiliar and still learning; whereas it is less influential once the user has

memorized some of the information. Our general finding is that a user's ability to extract new information has a strong correlation with the power spectrum of the interface used.

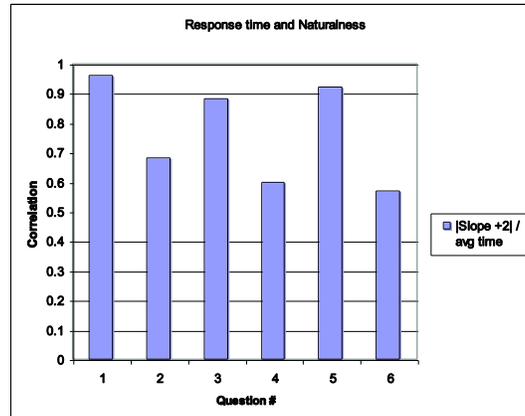


Figure 2-13. The absolute value of the correlation between the user response time for each interface and the respective naturalness of a screenshot of that image.

IMPROVING VISUALIZATIONS

We have demonstrated a technique for analyzing images to retrieve information about their natural characteristics. The resemblance of a visualization to a natural image has been shown to strongly relate with both competition results and user performance. A problem with using image analysis, however, is that it only gives the extent of naturalness, not the cause. A question thereby arises of whether this knowledge can be used as a guideline for the design of visualizations.

A likely culprit for an unnatural visualization is the underlying data. Fitting a visualization's data to a power function can make resulting visualizations more natural, and treemaps can be used to test that theory. Treemaps [Bederson, Shneiderman, & Wattenberg, 2002] are a variation of the unoccluding images in Figure 2-6. These treemap [JS91] resembling images were generated by creating rows of non-overlapping squares with random greyscale values. The size distributions are power (a), exponential (b), linear(c), and constant (d). When applying the color and size distribution methods from that section to treemaps, similar results are produced. Constant, linear, exponential, and power distributions produce power spectra near f^{-2} , and their average deviations decrease respectively. These images and their results can be seen in Figure 2-14. These are treemaps generated

using power (left), linear (middle), and constant (right) size distributions. Their corresponding power spectra are next to them. Notice the low average deviation for the power function. The implications of these findings mean that the size distribution of a visualization’s data can help determine the visualization’s natural qualities.

Consequentially, distorting data to fit a power distribution may improve the resulting visualization.

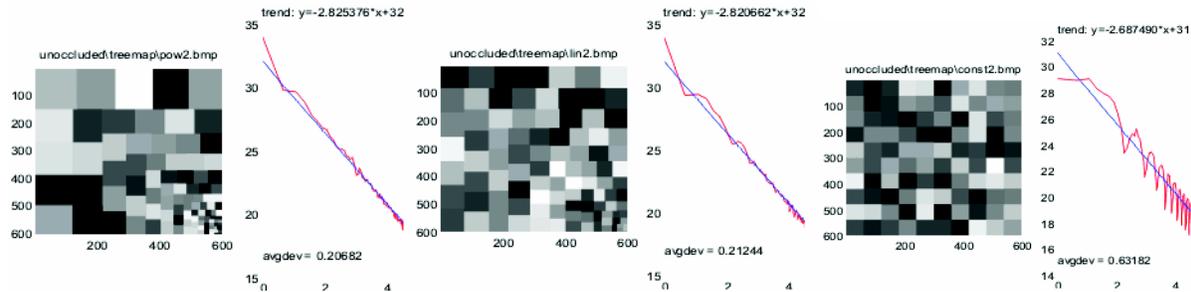


Figure 2-14. These are treemaps generated using power (left), linear (middle), and constant (right) size distributions. Their corresponding power spectra are next to them. Notice the low average deviation for the power function.

UTILITY OR ART

Questioning what is being measured by a visualization’s closeness to a natural image is essential. We have not established any causality between naturalness and visualization quality. They may both be caused by some other factor. A likely candidate for influencing naturalness and our preference of one visualization over another is aesthetic appeal. Art is usually found to have statistics similar to those of a natural image [Schroeder & Wiesenfeld, 1991; Reinhard et al., 2001]. This observation is made evident by our fascination with fractals [Peitgen & Saupe, 1988] which have repeating shapes and frequently have a power spectrum just below f^{-2} [Schroeder & Wiesenfeld, 1991]. A possible implication is that visualizations are preferred due to their appearance rather than their ability to enhance cognition. On the other hand, these qualities may not necessarily be mutually exclusive, as the distinction between art and utility may not necessarily exist. Obviously, this question leaves much room for research.

LIMITATIONS AND FUTURE WORK

One should note that this measurement technique does have limitations. The metric only gauges the sizing and spacing in a visualization. It does not directly evaluate other aspects such as color or font variety. Moreover, it does not even provide appreciably helpful feedback as to the cause of a visualization’s unnatural evaluation. The only

appreciable feedback for improvement given by the power spectrum is a rough estimation of the underrepresented and overrepresented frequencies. To broaden the encompassment of the measurements, future work could be done to observe the power spectrum of each of the colors in a visualization or to study the influence of text sizing and distribution.

Another limitation is the inability to use power spectra as an all-encompassing, exclusive means of visualization measurement. A correlation is not causality. Although natural characteristics show a non-random prevalence in preferred visualizations, some visualizations with very unnatural appearances are useful for certain tasks. A spreadsheet, for example, is probably the most effective visualization if the row and column of the desired information is known. Nevertheless, the spreadsheet has a highly unnatural power spectrum due to its regularity. The natural properties needed for an effective visualization might be determined by the task being performed. An interesting future study could compare the naturalness and performance of visualizations with the type of task being performed. If large numbers of similar visualizations were collected, one could also study how the Fourier transform is affected by the type of visualization and its contents. Fourier transforms have been shown to be capable of categorizing natural images and determining their contents [Torralba & Oliva, 2003]. This approach might be also applicable to visualizations for the purpose of content recognition or automated feature extraction. We should also note that spatial frequency is not the only property of a natural image, and testing for other natural properties might provide more insight.

Further research of this metric's implications may help in understanding how insight is obtained from data. Since this measurement shows how well our visual system is adapted to perceiving an image, we may have a better ability to focus and extract information when that information is presented naturally. Past research has studied perception of natural stimulus and found that it has a very sparse representation in the brain [Olshausen & Field, 1996]. If the interface for data is efficiently processed by the brain, more attention may be available to focus on the data that the interface is trying to convey. By naturalizing a visualization, we may be streamlining the process of perception. With more research, a more concrete neurological foundation for establishing visualization utility might result.

CONCLUSIONS

We have shown a strong correlation between the natural characteristics of a visualization and its preference and performance. This underlying principle has been built up from testing abstract nonrepresentational images and has been combined with work already done in the fields of neuroscience and computer vision. When applied to visualization, a metric based on natural image statistics has been shown to be consistent with assumed preference and competition results regardless of the use of occlusion. This connection has been demonstrated both theoretically and practically, and it shows that we can take advantage of our brain's enhanced receptivity to natural images. Hopefully, we can use this pattern for preferred visualizations to design future visualizations that, by their very nature, can better appeal to the human visual system.

3. TEMPORAL THRESHOLDS FOR FEATURE DETECTION IN FLOW VISUALIZATION

Abstract published in Applied Perception to Graphics and Visualization 2010 by Steve Haroz and David Whitney

Full paper in preparation

ABSTRACT

Optic flow – the coherent motion of a large region of the visual field – is important to our perception of and interaction with the world [Gibson, 1986]. In spite of seeing motion fields that are frequently large and highly varied, the visual system has the incredible ability to rapidly encode this higher order percept. In our study, we seek to determine the visual system’s ability to detect and locate deviations or features within an otherwise coherent flow field. This capability would help determine the suitability of optic flow as an exploitable tool for visualization. Using random-dot moving stimuli, our study investigated the effect of dot lifetime on subjects’ ability to detect a variety of features within a flow field. We found that a dot lifetime of a small fraction of a second can yield a coherently animated display that allows the user to find a feature in a flow field. Furthermore, features with gradually defined rather than sharp edges could be detected with dot motion being the only visible property. Even dot density was maintained due to the extremely short dot lifetime. The results are not only helpful to our understanding of optic flow perception but also provide some insight on the use of optic flow for visualization.

INTRODUCTION

Optic flow is the coherent motion of a region in the visual field, and the visual system is astoundingly effective at perceiving this large-scale and complex percept [Williams & Sekuler, 1984; Gibson, 1986; Duffy & Wurtz, 1991]. Whenever you rotate your head or move through a scene, nearly every part of your visual field moves at a different speed and velocity. For example, as you walk through a forest, the trees in the horizon appear to slowly expand, and the trees on your sides appear to rapidly move to the extremities of your vision visual field and disappear. In spite of such a myriad of local motions, all of this information is concisely summarized as moving forward. Such summary encoding of a large variety of elements makes this percept a potentially useful tool to

exploit for visualization. For that reason, our goal in this study is to test the visual system's ability to use optic flow for a basic visualization task, feature detection.

We know that optic flow can be used to catch our attention, as we are usually interested in the movement of specific objects within a flow field. Whether seeing an animal running up a tree in the forest or a bird flying in front of our car, we are typically most interested in the object whose motion is different from the surrounding field. However, mere directional or velocity differences are not necessarily sufficient to segment the object from its surroundings because every part of the field may be moving differently. Somehow, the visual system detects that the object's motion contrasts with how the rest of the motion field varies. This heightened interest in the outlying object is true for visualization as well. A tornado embedded in a broader weather system or the spike in density and diversity of a city in a population map are examples of such scenarios.

Although we know much about optic flow and the perception of it, search for motion is less understood. Some have examined the ability to search for one type of motion among others [O. Braddick, 1974; Royden, Wolfe, & Klempe, 2001], but less is known about how effectively embedded (rather than adjacent) flow fields are detected and localized. Huber and Healey [2005] recently studied the detection of a patch of linear motion among a background moving in a different direction. However, we want to determine if more complex forms of optic flow such as rotation or expansion can be effectively differentiated especially with a gradual envelope. Furthermore, we want to check for any differences in detection of various types of optic flow (e.g. does the attention-grabbing nature of a looming object over a lateral moving object [Franconeri & Simons, 2003] extend to flow fields).

A common method of generating optic flow is limited lifetime Random Dot Motion (RDM). Similar to a particle tracer, this display places dots at random locations and moves them over time. After a certain number of frames, a dot is moved to a new random location. The independent yet formulaic motion of these dots allows one to create almost any type of motion field. Our goal for this study is to find the lowest temporal threshold of the dot lifetime that allows a person to detect a feature, which we define as a region that is different from its surroundings. We also sought to detect a feature without clear and sharply defined borders. Looking at many heatmaps or vector fields will quickly reveal that not all data sets have distinct and disjoint features. Though one region may be

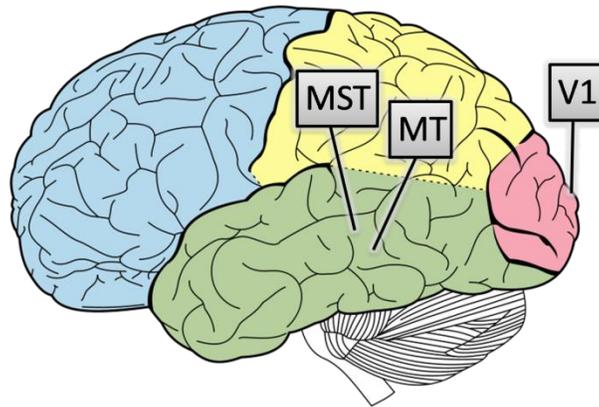


Figure 3-1: This figure shows the Primary Visual Cortex (V1) in the back of the brain as well as the Middle Temporal (MT) and Middle Superior Temporal (MST) regions. MT and MST are known to be important in processing motion. (Brain image vectorized from Grey's Anatomy for Wikipedia by user Mysid).

different from its surroundings, its properties may be defined by a gradual border rather than a high frequency edge. While such a boundary is common in flow field datasets, it is not often present in real-world moving objects, and its study in the field of vision has therefore been limited. Determining the ability of optic flow perception to find features with smooth boundaries will help us better understand what types of data can be visualized by optic flow.

MOTION PERCEPTION

When light hits our eyes, it is sensed via photoreceptors in the retina. After passing through a several layers of neurons, retinal ganglion cells send the signal through the optic chiasm where the signals from both eyes are separated into streams corresponding to the left and right halves of the regions of the visual field. This information, after being relayed in the lateral geniculate nucleus, travels to the back of the brain – called the occipital lobe or V1 (Figure 3-1). Put simply, the neurons in this region respond based on the orientation and scale of contrast gratings (or edges) [Hubel & Wiesel, 1962; De Valois & De Valois, 1988].

One area of the brain that receives input from V1 is the Middle Temporal (MT) lobe. Here, motion in local regions initially sensed in V1 is combined over large parts of the visual field to compute motion of large regions, objects, or scenes [Reichardt, Poggio, & Hausen, 1983; Newsome & Pare, 1988; Britten & Newsome, 1996]. An MT cell detects

that motion occurs in its receptive field and is selective for direction and speed [Albright & Stoner, 1995; Marcar, Xiao, Raiguel, Maes, & Orban, 1995].

Next to MT is the Middle Superior Temporal (MST) area. Cells in this region group the output of MT into even larger receptive fields with more coherent patterns of motion [Nakayama, 1985; Saito et al., 1986; Tanaka et al., 1986; Duffy & Wurtz, 1991]. The three MST neurons are sensitive to linear motion, rotation, and expansion/contraction (Figure 3-2). Some of the MST neurons also encode combinations of these patterns (e.g. rotation and contraction make an inward spiral).

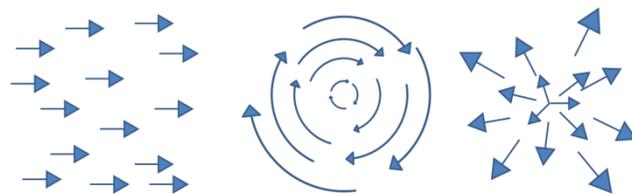


Figure 3-2: The three types of coherent motion encoded by MST are linear motion, rotation, and expansion/contraction. Neurons in MST also encode combinations of these basic motion types.

One type of stimulus used in examining motion perception is moving random dot patterns. Similar to a particle tracer, this display places dots at random locations and moves them over time. In its simplest form, the dots only move for one frame and thereby create apparent motion [Lappin & Bell, 1972]. Early experiments with such stimuli revealed that beyond a maximum displacement (D_{max}), the percept would fail [O. Braddick, 1974], and subsequent studies found D_{max} to vary with distance from the center of the visual field [Nakayama, 1985].

To improve their resilience to noise, the components of the visual system integrate input over time [Barlow, 1958]. For example, MT and MST are capable of filtering or averaging out noise in moving patterns. Burr and Santoro [2001] found that the visual system integrates local motion up to 300 ms, and larger scale higher-order motion can be integrated up to 3,000 ms. Higher order motion is therefore fairly resilient to noise if the motion is sustained.

Depending on the extent of the noise, the visual system can compensate for spatially and temporally varied motion by stochastically sampling [Williams & Sekuler, 1984]. Some channels even average wildly different types of motion into a coherent averaged flow.

OPTIC FLOW IN VISUALIZATION

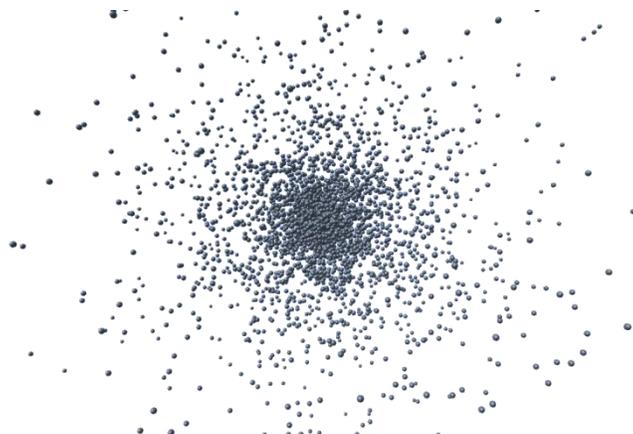
Though many of the psychophysical characteristics of optic flow are well studied, part of our intent is to examine aspects of this percept that could be useful for visualization. Fortunately, optic flow's underlying component of motion and its applicability to visualization have been examined in recent years.

Huber and Healey [2005] conducted three experiments on perceptual properties of motion that could be applied to visualization. Their discrimination tests consisted of finding a patch of 3 x 3 elements that were somehow different from a background field of elements; these differences included flicker frequency, direction of linear motion, and velocity. The experiments respectively found thresholds of 120 ms, 20°, and 0.43° of visual angle and showed that these three features of motion could encode data.

Kinetic Visualization [Lum, Stompel, & Ma, 2002] applied motion as an enhancement technique to show structural detail. In the paper, the authors looked at the difficulty of perceiving slightly curved objects under suboptimal lighting conditions. By moving many small glyphs along a surface, they made use of structure from motion to enhance the perceived curvature of a surface.

Though motion may have many potential uses in visualization, one technique that seems particularly well suited to

Figure 3-3: This image demonstrates the effect of a flow field on glyph density. Here, a sink is pulling all the glyphs in towards the center, and finding details along the sparsely represented outer edges would be difficult. Though the opposite problem of glyphs pushing towards the outer edge is more easily solvable (by simple bounds testing), knowing which problem will arise requires advanced knowledge of the data.



taking advantage of new information about motion and optic flow is particle tracing. Moving particles through a

flow field, or particle tracing, is one of the early and fundamental visualization techniques. It was demonstrated by Lane [1994], Meiselbach & Bruckschen [1996], and others as an effective technique for visualizing vector fields, especially if the field is time-variant. A common limitation with these animated vector field visualizations is that the glyphs or path lines which follow and portray the flow can become cluttered as a result of poor seeding. In regions where the flow coalesces (a sink), the glyphs can become too crowded, while regions that move outward (a source) may not be covered if a seed point does not fall on them (see Figure 3-3). Densely packed regions can easily cause glyphs to occlude one another or block geometry that may be present for reference – such as a beach front in a hurricane visualization. Variations in density can also impact the perceived direction of motion [Williams & Sekuler, 1984], which would obviously be hazardous for the accuracy of a visualization. Park et al. [2005] and Helgeland & Elboth [2006] demonstrated the helpfulness of small limited-lifetime glyphs (in the form of path lines) in alleviating many of these problems, as random glyph placement combined with a limited lifetime achieves a complete and uniform distribution. Their technique has the benefit of avoiding the problems of occlusion that can occur when using large numbers of typical glyphs in 3D. Park et al. also showed that the limited lifetime can almost completely eliminate the problems of path lines degrading in accuracy over time. The technique allows animation as well as color and glyph size to be used for displaying other variables of information. Appreciating the flexibility of using limited-lifetime glyphs, our study seeks to explore the limits of how briefly the glyphs can exist while still effectively conveying the underlying information.

THE EXPERIMENT

Our visual search experiment was conducted on 6 subjects – 4 women and 2 men. Five were graduate students in computer science or psychology, and one was a member of the university staff. All had normal or corrected-to-normal vision.

We used a 21 inch CRT - 19.8 inch (50.3 cm) viewable – running at 1024x768 at 64 frames per second. LCDs were avoided due to the very noticeable ghosting effect which left dots on the screen for additional frames and resulted in glass patterns [Glass, 1969] (or Stevens Dot Pattern [Tateosian, Dennis, & Healey, 2006]). Each subject's head

was placed in a chinrest to keep the view distance at a constant 24 inches (61 cm). Consequently, the visual angle of the screen was 26°, and the view angle of the stimuli was 25°.

The display (see Figure 3-4) consisted of a square random dot motion display with a gray background and approximately 1000 black dots. In visualization and particle tracing terms, these dots were the glyphs which followed a flow dataset. Next to this flow visualization was a single large slider which controlled the average lifespan of the dots. Under the slider was a “next” button to proceed advance to the subsequent trial and a progress bar to let the subject know how many trials have been completed.

We wanted the visualization to make exclusive use of motion, so we could conduct an isolated test of that feature. For this reason we made the visualization as simple as possible by having black dots on a gray background. The dots were seeded randomly with an even distribution, so a static screenshot could not provide any useful information.

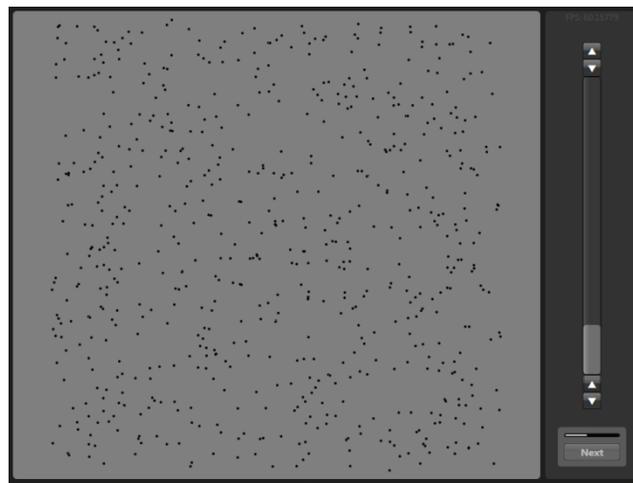


Figure 3-4: This screenshot of the experiment application shows the random dot stimuli on the left and the controls on the right. The slider allowed the user to control the threshold. The next button would appear once a quadrant was selected, and the white and black bar above it showed the progress.

STIMULI

The dataset of each trial was a 2D flow field with two components. The main component of the flow field – the background – followed one of these patterns:

- Left

- Down
- Clockwise
- Expansion
- Contraction

A small region – that we call “the feature” – had a different flow compared with the rest of the field. It followed one of the five flow types but not necessarily the same flow as the rest of the field. It was centered in one of four quadrants and had a Gaussian envelope (Figure 3-5). The transition between the two flow fields –whether via smooth interpolation or sharp borders – created very unusual patterns which made interpreting the feature difficult. To circumvent this problem and provide the feature with a smooth transition, we made use of transparent motion. At any given location, the percent of glyphs which followed the feature vs. the main flow corresponded to the Gaussian function’s value. The Gaussian envelope therefore defined the probability that a particular dot followed the trajectory of the background motion or the feature motion (see Figure 3-5). This effect would be similar to the spinning of a hurricane tapering off as the distance from its eye increases.

A median speed of about $4^\circ/s$ was the same for all types of motion. Rotation was the only motion that had a variable speed, although expansion and contraction obviously had variable directions. The flow of the feature was

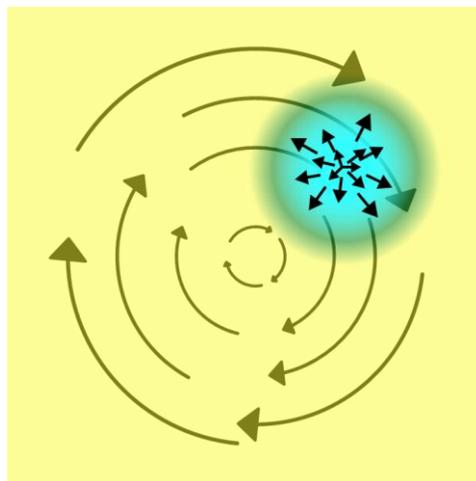


Figure 3-5: A sample trial with rotation as the main flow and expansion as the feature is illustrated here. The main flow field can be seen in yellow, while the small feature is blue. In the partial overlap, the proportion of glyphs following the flow is determined by a Gaussian probability.

exactly the same as an equivalent flow in the background field regardless of size. In other words, the velocity one degree from the center of the background flow was the same as the velocity one degree from the center of the feature. We did, however, make an exception to that rule when the feature and main flow moved linearly in the same direction. For those cases, we increased the speed of feature by 50%, so it could be differentiated.

TASK

We enumerated every combination of main flow (5), feature flow (5), feature location (4 quadrants), and initial dot lifetime (2) for a total of 200 trials. We also incorporated 20 trials without any feature to keep the subjects alert. All 220 trials were ordered randomly for each subject.

For each trial, the subject was presented with a new flow field, and the dot lifetime was reset. We used a method of limits (ascending and descending adjustment), which allowed the subject to control the stimuli and minimized error of habituation. The dot lifetime of each trial was initialized to either 1 frame (no motion) or 30 frames (making the feature easily detectable). The subject then adjusted the dot lifetime to find the lowest value for which the feature was barely visible. To make sure that they were, in fact, finding the feature, they then selected the quadrant where the feature was located and proceeded to the next trial. A progress bar gave the subject an idea of the percent of completed trials.

RESULTS

The results are summarized in Figure 3-6 and Figure 3-7 according to feature type, background type, and location. The lifetimes in the figures are in milliseconds and are averaged across all subjects. We grouped the results into four categories.

		Results in ms					
		Feature					
Background Motion		X	Y	⊙	Exp.	Cont.	
	X	468	121	50	51	57	56
	Y	465	50	130	60	63	64
	⊙	467	59	63	63	63	65
	Exp.	452	60	66	60	65	61
	Cont.	458	61	63	55	64	65

Figure 3-6: This table summarizes the results of the experiment. Each cell contains the mean selected lifespan (ms) for the corresponding type of background and feature motion. The types of motion are left (X), down (Y), clockwise rotation, expansion, and contraction. The leftmost column corresponds to trials where no feature was present. Each cell is colored according to the average time.

(1) The features with more complex types of motion – rotation, expansion, and contraction – make up the majority of the trials. These results can be seen in the three rightmost columns of the results figure. Analyzing the results without averaging between subjects, we found the upper bound of the 95% confidence interval for these trials to be 73 ms. The background motion did have a significant impact ($F(4,20) = 3.843, p < 0.05$), but no other main effects or interactions were significant.

(2) In the most difficult trials, the only difference between the feature and the background was the feature’s 50% speed increase. For these trials, the feature and background had linear motion in the same direction (X-X or Y-Y). The thresholds are noticeably higher for these trials, and the corresponding red cells stand out in the results. This outcome is not surprising as the extent of the directional and speed difference is known to impact the ability to segment.

(3) In some of the other trials with a linear motion feature, a directional difference existed, but it was subtle. For the trials with clockwise background rotation, the right side of the stimuli moved down, so detecting a feature on

the right side with downwards motion was more difficult. The more subtle distinction resulted in slightly higher times and can be seen in Figure 3-7. The cells of the X and Y columns consistently lack symmetry according to location.

(4) The rest of trials with linear features had larger directional differences and resulted in the lowest thresholds. Perpendicular linear motions (X-Y and Y-X) were particularly easy to detect.

DISCUSSION

For category (1), the low and fairly consistent results show promise for the use of optic flow in visualization. However, this study is not an automatic green light for the use of optic flow. For a visual property to be an asset to

		Results in ms											
		Feature											
Background Motion		X		Y		U		Exp.		Cont.			
		X		469	469	112	148	50	44	51	50	55	54
	469		466	124	103	51	56	49	52	56	63	52	57
Y		455	469	48	51	102	124	59	66	72	57	58	68
		469	469	51	49	160	135	55	61	59	65	65	65
U		469	458	48	47	44	72	65	70	64	63	61	70
		467	469	63	79	56	79	57	57	70	57	66	63
Exp.		469	469	72	51	52	51	63	54	55	65	53	63
		461	313	64	52	74	87	51	70	73	70	59	68
Cont.		466	455	44	78	76	69	53	53	66	68	69	61
		344	469	49	73	56	53	56	57	61	60	63	66

Figure 3-7: This table summarizes the results of the experiment according to location. Within each cell, the four values are the times for each feature location. The leftmost feature column represents the trials that had no feature (the locations are arbitrary, and the red text means the subjects picked incorrect locations).

visualization, not only must we be able to detect a variation of that property, but we must also be able to do so with minimal attention. For a visualization to be effective, differentiation and identification of its components must be instant and effortless. The study demonstrated the ability to segment a feature defined exclusively by optic flow, but we did not investigate the attention needed to do so. A study that focuses on attention for feature detection using optic flow would be the next step.

Categories (2), (3), and (4) confirm that directional and speed differences matter for linear motion; detecting a motion boundary defined by orthogonal motion is easier than detecting a speed difference alone. Like every other perceptual property, contrast is key for detecting directional differences. The low spatial frequency edges also confirm what others are beginning to find – that the visual system can detect various spatial frequencies of motion defined edges [Durant & Zanker, 2009].

APPLICATION

Some visualizations (e.g. particle tracers of time variant flows) already make some use of fields of motion, and they could benefit from this empirical study of dot lifetime. Shortening lifetimes to mere fractions of a second could help reduce the dependence on resource intensive algorithms that control glyph density, and this study can hopefully provide more information about optimal thresholds of glyph lifetime. As a consequence, particle tracers may be able to better convey information while reducing resource use. As Figure 3-8 shows, the longer lasting glyphs create distribution problems that leave important details of the field unrepresented in the visualization. One solution to this problem has been to employ the use of computationally intensive seeding or density control algorithm, which could radically reduce the performance of the application. Demonstrating that features could be detected for sub-second glyph lifespans allowed us to create a flow visualization with a consistently even glyph distribution. The outcome is a particle tracer that visualizes all regions of the field evenly with minimal computational requirements.

CONCLUSIONS AND FUTURE WORK

We have shown that a nonlinear optic flow feature can be detected within various other types of optic flow for dot lifetimes of 73 ms (except for aligned linear). This result is interesting for two reasons. First, answering the core question of our study, a dot lifetime of less than one tenth of a second can effectively convey a feature defined exclusively by non-linear optic flow. Though more information is needed, this study suggests that using optic flow for visualization does indeed have potential.

The second interesting aspect of the result is that the dot lifetimes are actually fairly high. We know that two frames of apparent motion are sufficient to perceive optic flow [Burr & Santoro, 2001]. Yet for some reason, the lifespan needed to detect a feature within the background flow field was more than simply apparent motion. If apparent motion were sufficient for feature detection, we should have seen lifetime thresholds of around two frames (31 ms). Instead, even the high contrast (orthogonal) fields almost all needed lifetimes of more than three frames (47 ms). The intriguing possibility is that a capacity limited process, perhaps related to visual attention [O. J. Braddick & Holliday, 1991; Royden et al., 2001], is required to detect features embedded in flow fields, and that having more than one feature could impose a severe tax on the system. The question of our effectiveness at feature detection for complex flow fields leaves room for future study and would have intriguing value to both visual perception and visualization.

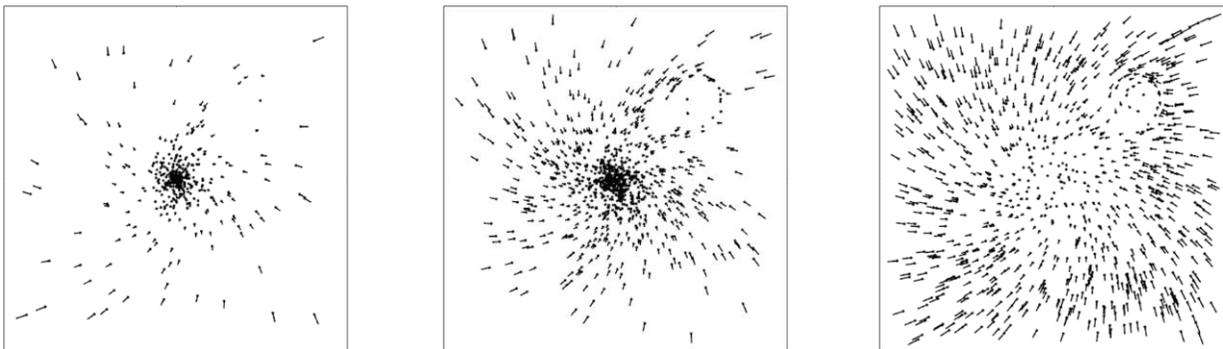


Figure 3-8: These images all show the same flow field visualized using glyphs with different life spans. Small tails were added to the dots to account for the static nature of the images. The field spirals towards a center 'sink' and has a turbulence vortex in the upper right. (Left) With life spans of 3 seconds, these glyphs are drawn toward the sink where they become stuck. As a consequence, the outer edges have very sparse coverage, and all detailed information (such as the turbulence) is not visible. (Middle) The 1 second glyphs are slightly more spread out from the center and have some minimal coverage of the turbulence. Nevertheless, the majority of the glyphs are clustered towards the center. (Right) The 60 ms glyphs have a very even distribution across the field. The outer edges as well as the vortex are clearly visualized, as the glyphs are not trapped in the center.

4. HOW CAPACITY LIMITS OF ATTENTION INFLUENCE INFORMATION VISUALIZATION EFFECTIVENESS

Published in InfoVis 2012 by Steve Haroz and David Whitney

Poster presented at Vision Sciences Society 2013

Best Paper Award at InfoVis 2012

Best Student Poster at VSS 2013

ABSTRACT

In this chapter, we explore how the capacity limits of attention influence the effectiveness of information visualizations. We conducted a series of experiments to test how visual feature type (color vs. motion), layout, and variety of visual elements impacted user performance. The experiments tested users' abilities to (1) determine if a specified target is on the screen, (2) detect an odd-ball, deviant target, different from the other visible objects, and (3) gain a qualitative overview by judging the number of unique categories on the screen. Our results show that the severe capacity limits of attention strongly modulate the effectiveness of information visualizations, particularly the ability to detect unexpected information. Keeping in mind these capacity limits, we conclude with a set of design guidelines which depend on a visualization's intended use.

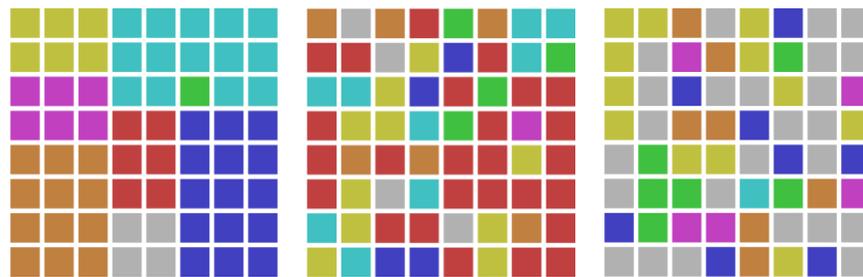


Figure 4-1. These images each have one colored square that is unique within that image. How long does it take you to find each? How many color categories are there in each panel? Why does grouping make both tasks substantially easier?

INTRODUCTION

An information visualization designer aims to present the maximum amount of data without overwhelming the user with complexity and information-overload. The components arranged to form a GUI or visualization are visual features – the properties of any image that the brain is capable of encoding and integrating into a coherent percept [Treisman & Gelade, 1980]. Examples that apply to visualization include position, color, size, orientation, texture [Bertin, 1983], and motion [DiBiase, MacEachren, Krygier, & Reeves, 1992; Bartram, Ware, & Calvert, 2003; Huber & Healey, 2005]. The designer’s role is to effectively associate these visual features with corresponding dimensions or categories in the underlying data [Cleveland & McGill, 1985].

Unfortunately, the speed and capacity of human attention for these visual features are severely limited [Horowitz & Wolfe, 1998; Wolfe & Horowitz, 2004], and these limits may influence the effectiveness of information visualizations. Exceeding the limits of visual attention markedly impairs both the accuracy and timing of one’s response to a visual scene (Figure 4-1). This consequence may seem intuitive (e.g., the benefit of grouping in Figure 4-1 might seem obvious), however visualizations often violate or ignore this intuition in part because it has not been formalized or empirically tested in visualization-relevant tasks. Characterizing and measuring these limitations in the context of data visualization is therefore fundamental and necessary to achieve the goal of conveying information via the human visual system.

Using a visualization should provide information faster or more broadly compared with serially inspecting the raw data in the form of a table or database [Fekete, Wijk, Stasko, & North, 2008]. A perceptual hindrance that restricts a user to serially inspecting each visualized element rather than enabling a rapid summary perception of the whole scene would make the visualization hardly better than a simple table. It is therefore critical to understand how the capacity limits of attention impact various visualization tasks.

To study the effect of these limits, we tested the effect of two types of arrangements or layouts on user performance in visual search and subitizing (or rapid counting [Kaufman & Lord, 1949]) experiments. We also examined how performance in these experiments is influenced by user goals and the variety of a visualization’s

visual features. To measure maximum performance, the experiments all had subjects fully attending to one task as opposed to dividing attention with a peripheral display [Somervell, Mccrickard, North, & Shukla, 2002].

We ran three experiments with the similar stimuli that used either colored or moving features. Each experiment's task, however, corresponded to a different, commonly performed visualization task:

- Detect a unique target with a known appearance (e.g., find the red object)
- Detect a unique target with an unknown appearance (e.g., find a unique or oddball target)
- Determine and compare the number of visual categories (e.g. determine extent of heterogeneity or consistency)

The latter two tasks are of particular interest, as performing them by browsing a table or running a database query is difficult. For example, finding the hour of the day that you receive the most or fewest emails would be nontrivial without some sort of graphical display. Such tasks which assume no prior knowledge of where to look or what to look for are where information visualization excels.

One of our primary goals was to quantify the influence of a grouped arrangement compared with a random arrangement (e.g., Figure 4-1) on visual search and subitizing tasks. We hypothesized that the effect would vary by task and would modulate the impact of capacity limits, so each experiment separately tested both layouts. To show that these effects were consistent across visual feature, we tested both color and motion. Several studies have investigated the impact of variance and grouping on visual search. Duncan and Humphreys showed that visual search of a known target (or one of two known targets) can be impacted by variance among even a small number of objects [Duncan & Humphreys, 1989]. Treisman found that increasing the attentional demand by searching for a conjunction of visual features increases the adverse impact of more groups [Treisman, 1982]. In the context of HCI and visualization, the spatial and sizing consequences of grouping have been proposed as models for design and evaluation [Haroz & Ma, 2006; Rosenholtz, Twarog, Schinkel-Bielefeld, & Wattenberg, 2009]. Furthermore, Tatu et al [Tatu, Bak, Bertini, Keim, & Schneidewind, 2010] examined how grouping in scatter plots influences preference (grouped displays are preferred), but preference and performance (an operational measure

of visualization effectiveness) are not necessarily equivalent. We aim to extend these investigations of search, attention, and capacity and present our findings in the context of information visualization.

For each experiment and each block, we tested how the variety of a feature (color or motion) affected performance. This property corresponds to the number of categories or discrete steps of a nominal axis that are displayed in a visualization. While showing the largest amount of detail or dynamic range for each data axis is ideal, we tested whether adding more detail can actually reduce someone's ability to utilize any of the information.

EXPERIMENT STIMULI

All three experiments used the same stimuli. They consisted of an 8 x 8 grid of 64 squares centered on a gray background (similar to Figure 4-1). The number of squares remained constant for all of the experiments to emulate simple visualizations with the same number of data elements but differing arrangements (correlation with the spatial axes or grouping algorithm) and amounts of variety (categorical detail).

For the color trials, each square had one of the eight highly discriminable colors in Figure 4-2. The large number of easily discriminable colors is what makes color such a frequent and effective tool for visualization [Ware & Beatty, 1988; Healey, 1996]. For the motion trials, each square acted as an aperture onto a monochromatic tiled texture (Figure 4-3) which moved using one of the motions in Figure 4-2. The number of motions was limited to six to ensure easy discrimination of each type of motion. The texture was the same for all squares but varied between trials. All of the textures had the same number of black and white pixels, so overall brightness did not vary.

As the top of Figure 4-2 shows, the stimuli had a limited number of colors or motions. The subset of colors or motions was randomized for each trial. We refer to the number of colors or motions in each stimulus (not including a search target) as the amount of 'variety'. Furthermore, the experiments were divided into grouped layouts and random layouts (bottom of Figure 4-2). For the grouped layouts, all squares of a particular color or motion were grouped into a single cluster. The sizes of the groups varied within a given display, so the size of an individual group provided no information about the number of groups. The random layouts assigned a color or

motion to each square using a 1/n random distribution. For both layouts each color or motion was present on at least four squares.

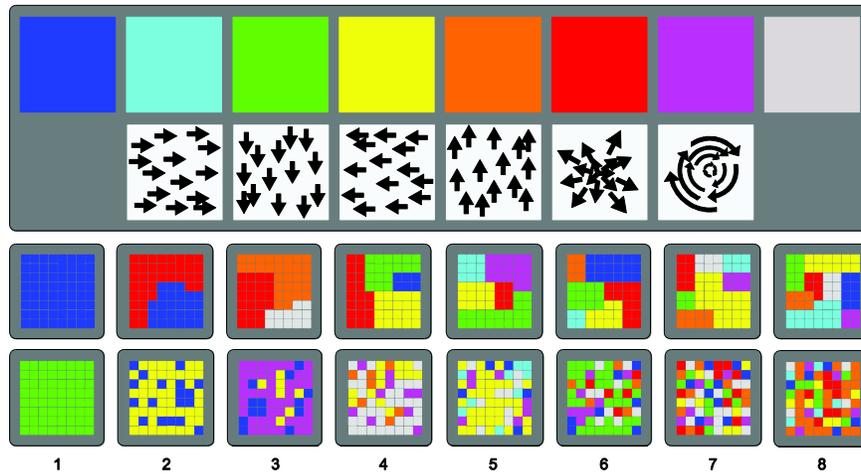


Figure 4-2. (Top) All of the colors and motion types used in the stimuli. The saturation is enhanced in the figures, as the actual colors were more isoluminant. (Bottom) A comparison of the grouped and random layouts for differing variety amounts. Each variety amount (except 1) had 16 possible grouped layouts (manually created by the authors). Each variety had a different number of squares in any given stimulus.

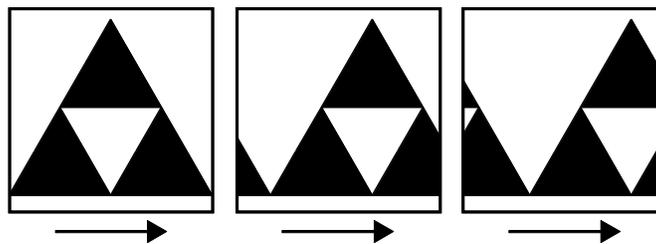


Figure 4-3. A demonstration of the textured motion used in the experiments. The texture moves, but the square's position never changes. All motions cycled at 1 Hz.

All of the experiments were divided into four blocks:

- Color – Grouped
- Color – Random
- Motion – Grouped
- Motion – Random

The order of these blocks was balanced between subjects such that each block went first and last at least once.

The experiment application was written in XAML and C# for the Windows Presentation Framework. The motion component was an HLSL shader that independently applied a transform in texture-space to each textured square. The experiments were conducted on a 24 inch iMac (using the built-in LCD monitor) running Windows 7 via Boot Camp.

The subjects performed the experiment in a dark room to avoid external visual distractions. They maintained a distance of 22.4 inches (57 cm) from the screen via a chinrest. The stimulus was 12.7 inches (32 cm) tall and wide. It consequently encompassed 32.3° of visual angle, while the squares were each 3.7° due to the small gaps between them.

For all the results with response time (RT) analyses, trials with an absent target (50%) or that had an incorrect response were not used.

VISUAL SEARCH: FIND A KNOWN TARGET

The first experiment was the simplest and served primarily as a control. The aim was to simulate the task of finding a known target. Realistic examples include finding a Green Party region on an election map, finding a large file in a directory treemap, or finding rain in a weather visualization. In these scenarios, the user knows the appearance of the target (perhaps via a legend) and is searching for its presence.

METHODS

Five subjects participated in this experiment. Two were female. All were either graduate students in psychology or computer science or trained university staff.

As Figure 4-4 shows, each trial began by displaying a square with the target visual feature in the center of the screen. A progress bar for the elapsed number of trials was also displayed. A one second pause (a gray screen) was then displayed to prevent any aftereffects or apparent motion. Then the stimulus was presented. Using the keyboard, the subject responded whether the target was *present* or *absent*. Answers and response times were recorded.

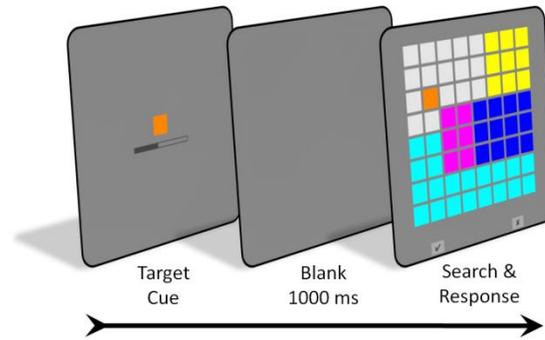


Figure 4-4. The procedure for the visual search experiments. The target and a progress bar were shown to the subject (the oddball experiment only had a progress bar). A blank gray screen was presented to avoid any aftereffects or apparent motion. Finally the stimulus was displayed while awaiting the subject's response.

For each block, each variety count (number of colors/motions in the stimulus) had 40 trials. In half of the trials, a random square was replaced by the target, whereas the other trials had no target. Each of the experiment's color blocks had 280 (7 * 40) trials, and each motion block had 200 (5 * 40) trials for a total of 960 trials per subject (4,800 total trials). The trials were randomly ordered within each block, which began with five practice trials.

RESULTS

The results in Figure 4-5 show several trends (ANOVA results included):

Accuracy was over 95% for all blocks and did not diminish with higher variety counts.

Visual feature: Color had a trend for lower RTs compared with motion for grouped – $F(1,4) = 921, p < 0.0001$ – and random – $F(1,4) = 41, p < 0.01$. This effect was not unexpected and reaffirms that visual feature can impact user performance.

Layout: The grouped layout performed slightly better than the random layout for color – $F(1,4) = 65, p < 0.001$ – and motion – $F(1,4) = 19, p < 0.05$. Size pop-out (when a distinct object perceptually stands out from its surroundings [Treisman & Gelade, 1980]) was possibly the reason because the grouped layout had no other groups as small as the target (1 square). For the grouped layout, detecting if a square was unique amongst its immediate neighbors could confirm it as a target without needing to check the entire screen for confirmation.

Variety: For color and grouped motion, the variety count had almost no impact on performance. The random motion block's RTs appeared to increase until 3 varieties and then plateau thereafter like the other blocks.

Subject: Though between-subject variability was significant, all subjects had the same trend of results (Figure 4-5 error bars).

A purely horizontal slope implies that the visual system can detect the presence of a target in parallel irrespective of the number of varieties. Duncan and Humphreys' study [Duncan & Humphreys, 1989] found a very small performance cost of a few milliseconds per additional category. Our results suggest that this cost remains minimal for a large number of objects.

VISUAL SEARCH: FIND THE ODDBALL

This variant of the visual search experiment served to test users' ability to detect something odd or out of place. It is analogous to finding a unique element in a visualization without knowing to look for it. Examples include noticing that one file's icon is different compared with its neighbors in a file manager or detecting an oddly behaving port in a network security visualization. This type of scenario is where visualizations can be most useful because if the user knows what to search for, a simple database query may yield faster and more accurate results. Detecting a unique element—an oddball, however, is nontrivial to query.

METHODS

Five subjects participated in this experiment. Four were female. All were either graduate or post-graduate students in psychology or trained university staff.

The experiment was identical to the known-target experiment shown in Figure 4-4, but the target was not shown (only the progress bar was visible). Subjects were told simply to detect if a unique square was present. As in the known-target experiment, each subject performed 960 trials (4,800 total trials).

RESULTS

The results in Figure 4-6 show several trends (ANOVA results included):

Accuracy: In contrast to the known-target search experiment, accuracy in finding the oddball was significantly affected by the amount of variety. Although the grouped layouts showed little performance degradation with increased variation, the random layouts had a marked decline as variation increased.

Visual feature: Color again had consistently lower RTs than motion for grouped – $F(1,4) = 120, p < 0.0001$ – and random – $F(1,4) = 101, p < 0.0001$. Nevertheless, they both showed similar trends despite the differing scales.

Layout: The grouped layout performed significantly better in RT and accuracy for both color – $F(1,6) = 109, p < 0.0001$ – and motion – $F(1,4) = 117, p < 0.0001$. In the grouped blocks, subjects were capable of detecting the target without influence from the amount of variety. Increasing the variety in the random layouts markedly impaired performance.

Variety: For the grouped blocks, the variety amount did not affect RT or accuracy. Conversely, the random blocks showed a clear dependence on the number of visual features used – $F(1,4) = 22, p < 0.05$

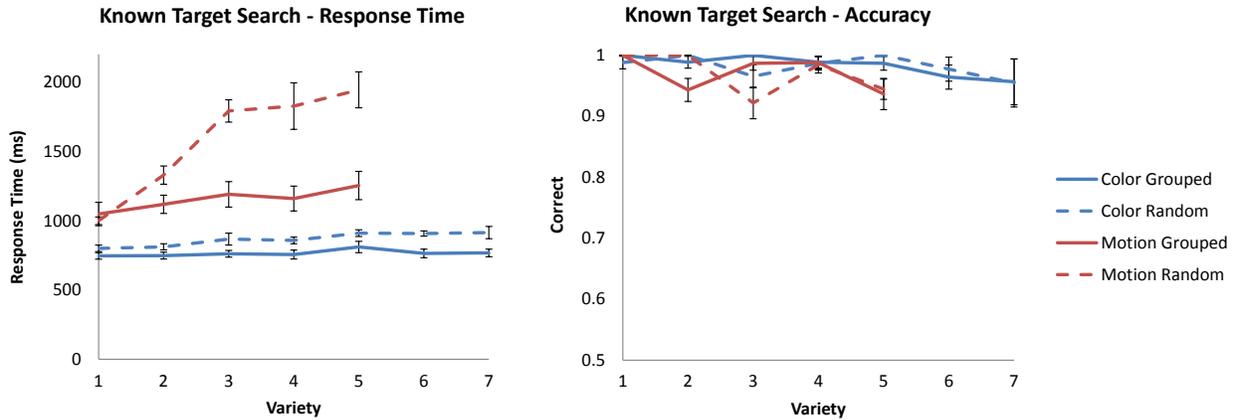


Figure 4-5. Experiment 1 results: visual search for a known target. (Left) The RTs of the visual search experiment as a function of the variety of features (e.g., the number of colors or motion directions visible). Only correct responses for present targets were used in the calculation. (Right) The accuracy was consistently above 95%. Error bars show the inter-subject standard error.

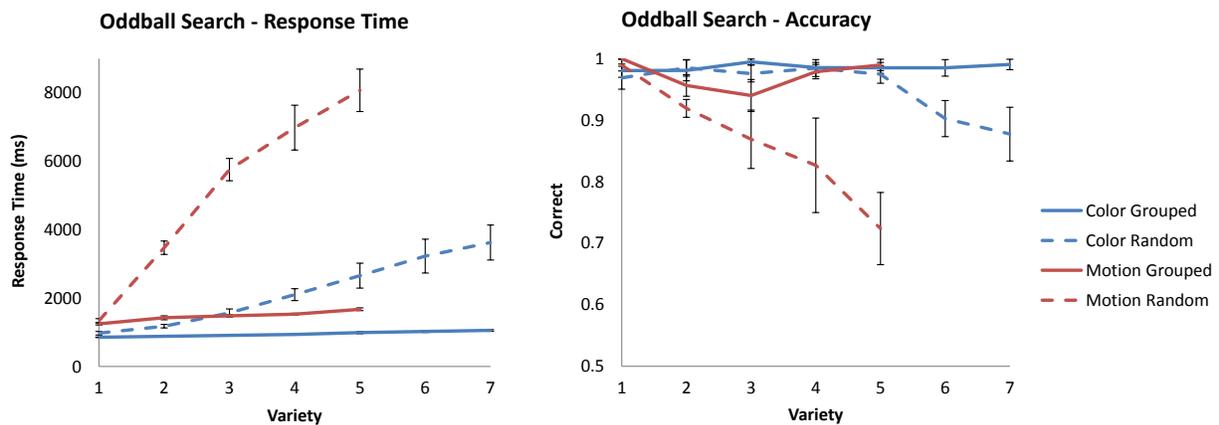


Figure 4-6. Experiment 2 results: visual search for an oddball target. (Left) RTs plotted as a function of the variety of features. Only correct responses for present targets were used in the calculation. Note that the variety of features has little influence on the RTs for grouped layouts, while the RTs for random layouts clearly grow with more variety. (Right) Accuracy for the grouped layouts was consistently over 90%, whereas the accuracy for random layouts dropped significantly as variety increased. Error bars show the inter-subject standard error.

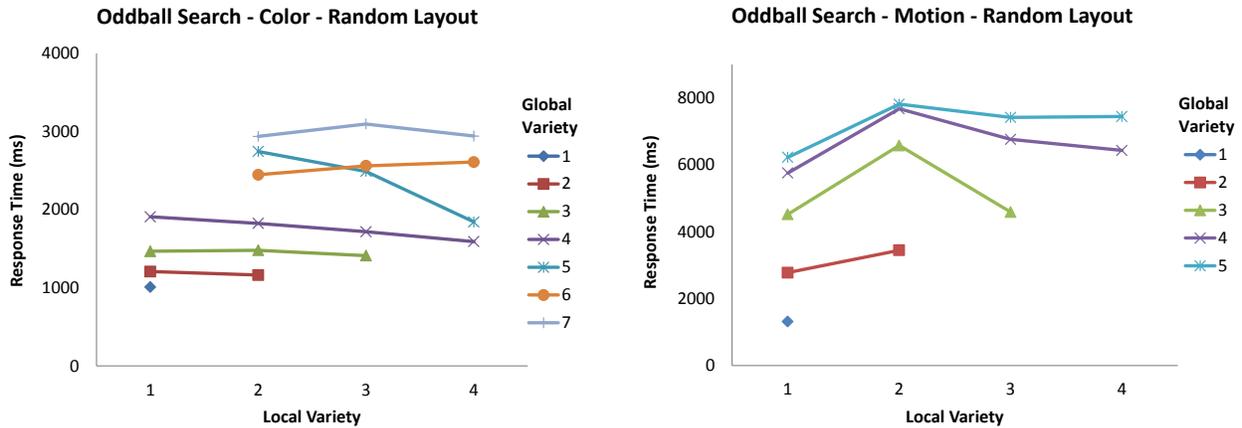


Figure 4-7. Local vs. global variety. These plots show the effect of local variety on performance under different amounts of global variety. Local variety is the number of colors or motions immediately adjacent to the target, whereas global variety is the total number of colors or motions on the screen. The horizontal slopes show the lack of impact of local variety on user performance.

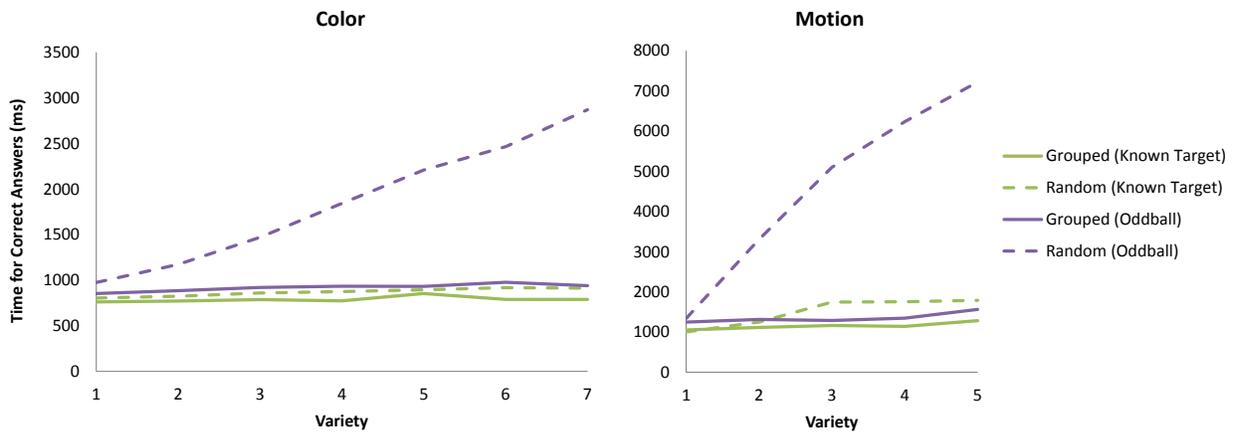


Figure 4-8. The figure shows the response times for the visual search experiments with known and oddball targets and groups them by visual feature. In spite of their different scales, the same trend is clear – variety of visual features has little influence on RT when subjects know the target in advance or have a grouped layout. That is, targets are relatively easy to find when they are predefined or when the distractors are grouped. Finding an unknown target in a random layout, on the other hand, becomes significantly more difficult with increased variety. From this information, we can infer that tasks requiring users to find oddball targets among ungrouped layouts (a likely occurrence with scatter plots, treemaps, or connectivity graphs) would make for an ineffective visualization design.

Subject: Though the between-subject variability was significant, all subjects had the same trend of results (Figure 4-6 error bars).

Local vs. global variety: We sought to examine whether the performance difference was caused by variety in the whole scene or immediately surrounding the target. To calculate local variety, we counted the number of colors or motions in the eight squares surrounding the target. Figure 4-7 shows the RTs for the random layouts plotted by local and global variety. The grouped layouts showed no effect from either type of variety.

The strong distinction between Experiments 1 and 2, searching for a known target versus an oddball target, is apparent in Figure 4-8. Corroborating the results of Nothdurft [1993], color and motion show the same pattern. The variety of visual features does not strongly impact performance when the search target is known, but has a dramatic effect when the search target is not known (is an oddball). Interestingly, grouping the features mostly compensates for this difference, creating a large performance gap between grouped and random layouts for the oddball search. The implication is that subjects are switching from a parallel process in the known-target and grouped odd-ball tasks to a serial process dependent on the number of features. A possible explanation is that the subjects were performing a series of Boolean map filters [Huang & Pashler, 2007] to detect global uniqueness only for the random oddball search. How much a target *popped-out* from its neighbors (due to lack of surrounding variety) had no impact on performance. This process would be repeated until the search yields no similar squares. These Boolean map filters would be applied sequentially for each variety until a unique target was found.

SUBITIZING: LIMITED CAPACITY

The goal of this experiment was to understand how visual complexity affects a user's ability to grasp aspects of the global structure of a display. For example, users may need to extract information about the gist of a visualization (was the market, on average, up or down in Figure 4-11), or they may need to estimate the variety of types of information (how many unique file types are in a tree map? Or, how much variance is there in a map of the market—was the market mixed today, or was it homogeneously down?). The task in our experiment was to subitize (or rapidly count [Kaufman & Lord, 1949]) the relative number of categories in a pair of stimuli.

Understanding the overall structure and number of categories in a dataset is yet another task which would be difficult without a visual aid (e.g., finding how many file types are in a directory is trivial with a nested tree map visualization). Yet even these aids are not always particularly helpful, as many have experienced struggling to understand the overall message of a display with a confusing layout or too many visual categories. This struggle highlights the fact that it is unknown how overall qualitative comprehension is affected by layout or visual feature variety. This experiment sought to provide some insight on the matter. Practical examples are available, however. The NewsMap (<http://newsmap.jp>) in Figure 4-12 shows an example of similar information being either easy or difficult to discern depending on the layout. When the grouping deteriorates, even a basic gist (or ensemble percept [Ariely, 2001; Chong & Treisman, 2003]) of the most significant news category becomes difficult. Likewise in the file treemap in Figure 4-13, using colors to categorize too many types of files makes the information seem too overwhelming to even examine a subset. By reducing or more compactly binning the amount of information, the user may actually have a more efficient or accurate perceptual representation of the display.

METHODS

Five subjects participated in this experiment. Three were female. All were either graduate or post-graduate students in psychology or computer science or trained university staff.

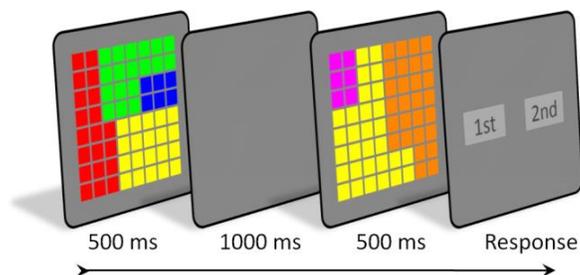


Figure 4-9. The trial time course for the subitizing experiment. The first stimulus was presented for 500 ms followed by a 1000 millisecond gray screen and then the second stimulus for 500 ms. The subjects then responded in a two-interval-forced-choice which of the two intervals had more variety (e.g., more colors). Since the timing was dictated by the procedure, only accuracy was recorded rather than RTs.

Figure 4-9 shows the procedure for the subitizing task. We presented two different stimuli sequentially for 500 ms with a one second blank gray screen in between. The amount of variety in the two displays differed by one, and the order of smaller and larger displays was randomized. The probability of the first or second interval having more

variety was even. The experiment had a two-interval forced choice (2IFC) design, so the subjects simply responded via the keyboard as to which interval had more variety. Due to the timed nature of the procedure, only the accuracy was recorded. Each subject performed 880 trials (4,400 total trials).

RESULTS

Figure 4-10 shows the results (ANOVA results included below):

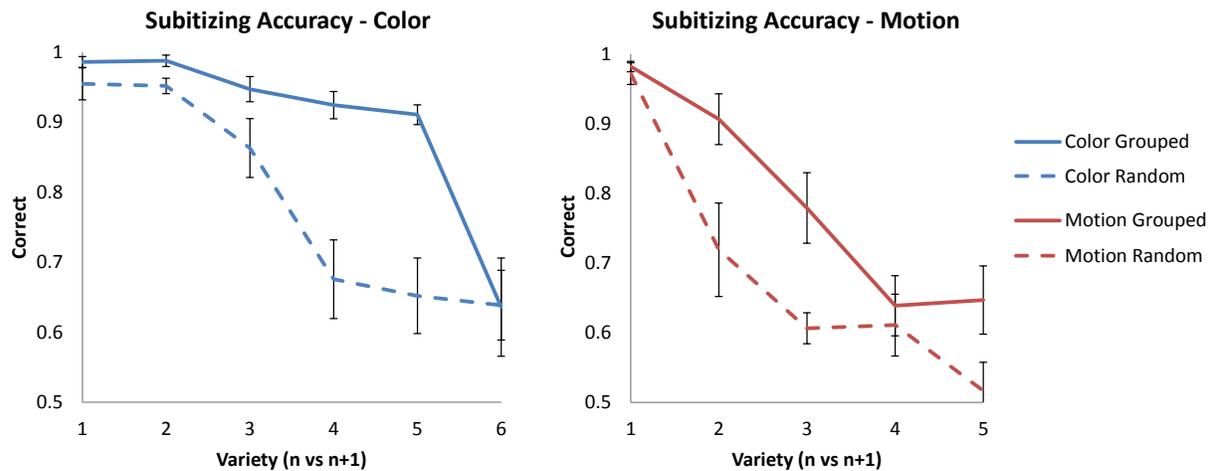


Figure 4-10. The accuracy results of the subitizing experiment follow a sigmoidal shape. The variety amount on the horizontal axis is the smaller of the two stimuli (1 means a comparison between 1 and 2 or 2 and 1). All layouts have nearly 100% accuracy for a comparison with one variety, and accuracy falls towards chance (50%) as variety increases. Grouping the features (color or motion) has a dramatic improvement on performance. Error bars show the inter-subject standard error.

Accuracy and Variety: In contrast to the other experiments, none of the blocks could maintain a high accuracy as variety increased. All accuracies were nearly 100% for one variety but followed a marked, sigmoidal fall-off as variety increased. As additional detail was added, the subjects' ability to understand the structure was severely impaired – $F(1,4) = 96, p < 0.005$.

Visual feature: Using 80% accuracy as a benchmark, the capacity for color variety across subjects was significantly higher at nearly double that of motion for grouped ($F[1,4] = 44, p < 0.0001$) and random layout ($F[1,4] = 18, p < 0.001$). This result is not surprising. Although the attentional capacity of motion with respect variety and layout has little research, color is known to have a high capacity [Luck & Vogel, 1997; Alvarez & Cavanagh, 2004].

Layout: Again using 80% accuracy as a benchmark, the grouped layout had roughly 60% higher capacity for color ($F[1,5] = 44, p < 0.0001$) and motion ($F[1,3] = 18, p < 0.001$). A possible explanation is that visual working memory operates over or is modulated by grouped features. When the layout is complex, more cognitive resources are required thereby leaving less available for visual feature information. Irrespective of the psychological explanation, the data shows that if understanding the overall structure of a visualization is important, a high-capacity visual feature and strong element clustering are critical.

Subject: Though the between-subject variability was significant, all subjects had the same trend of results (Figure 4-10 error bars).

Edwards and Greenwood [Greenwood & Edwards, 2009] performed a similar 2IFC experiment of subitizing motion presented via random dot stimuli (rather than moving texture patches). Their experiment, however, used the same number of elements for each variety on the display. Consequently, the quantity of any one feature type informed as to the number of variants in the stimulus. We purposefully varied the quantity in this and the other experiments. Our subitizing motion results corroborated theirs, demonstrating that the capacity limits are based on the global or scene-wide information rather than the properties of a single element.

While our experiment focused on the accuracy of a quick glimpse, Watson et al [Watson, Maylor, & Bruce, 2005] performed similar subitizing experiments examining how reaction time varied when unconstrained. Combined, our accuracy results and their response time results show an overall degradation of performance without grouping.

INFORMATION VISUALIZATION GUIDELINES

There are several particular results from our experiments that can inform visualization design:

GROUPING GREATLY HELPS FOR SOME BUT NOT ALL TASKS

When searching for an *unknown* oddball target—a target that may or may not, ostensibly, pop-out—the correlation between the spatial and visual feature axes is pivotal. High correlation speeds search; a cluttered

randomly organized arrangement impairs search for the target. The degree to which the targets pop-out in Figure 4-1 and Figure 4-12 is modulated by the extent to which all the other colors are *grouped*.

Though the uses of information visualization can vary widely, we recommend prioritizing the tasks for which visualizations greatly outperform raw data (tables and databases). Extracting summary or contextual information, such as how much variety or how many categories of information are present is difficult with a spreadsheet, and doing so with a visualization relies on attention and visual memory, necessitating spatial grouping by visual feature. By incorporating visual feature into a visualization's layout algorithm, observers could better judge overall qualitative characteristics of the display, such as whether the market had an overall gain or loss or how the success was distributed across industries (Figure 4-11), or which industry had the most variance in performance.

This principle is a matter of prioritization, as optimizations for different visual features will likely conflict with each other. A color-focused layout would likely yield extreme aspect ratios which could hinder performance on size oriented tasks [Kong, Heer, & Agrawala, 2010]. A visualization designer must therefore prioritize the visual feature that benefits the more important tasks.

IF YOU CANNOT GROUP, CHANGE THE TASK

Knowing a target's appearance in advance dramatically improves accuracy and speed of target detection. Knowing to look for a red or rightward moving object drastically improves user performance, even within a highly variable display, suggesting that guided search is an important factor [Wolfe, 1994; Wolfe & Horowitz, 2004]. The intuitive implication is that a legend, key, or any mechanism that helps a user know the visual appearance of a sought after data element greatly simplifies locating it.

WHEN THERE ARE MANY CATEGORIES: LESS IS MORE

A compromise must be made between the number of nominal categories and the perceptual complexity of a visualization. Design options are still available if a high number of nominal categories is needed. Though not always feasible, the spatial axes can be assigned to data dimensions that correlate with the visual feature's dimension [Tatu et al., 2010]. As a result, the feature categories would become more spatially grouped and user performance

would be less limited by the capacity. Because correlation is not always known, a perhaps more deterministic alternative is to limit the visualization to only show a couple of categories at a time (Figure 4-13). The user would be required to interact to see different categories, but the limited information on the display could be analyzed within the limits of attention, enabling more efficient comprehension of the visualization.

ASSIGNING A VISUAL FEATURE TO A DATA DIMENSION

We chose color and motion because we expected their performance and capacity to be fairly different (high for color and low for motion), and our result confirmed that hypothesis. Features should be chosen with care; features like color that have higher capacity should be reserved for data with high dynamic range or many categories.

EVALUATION

Visual search for a known target is not always a sufficient test of visualization effectiveness. Though it may help assess how quickly certain kinds of information are conveyed in a visualization, it provides little or no information about the influence of capacity and other limits of attention. Evaluations of visualizations should make certain to test user performance for more attentionally demanding user goals. While we used oddball and subitizing tasks to examine attention capacity, other attentional limitations (e.g., the spatial and temporal resolution of attention [Intriligator & Cavanagh, 2001; Holcombe, 2009; Franconeri, Alvarez, & Cavanagh, 2013]) should be investigated depending on the users' goals. For example, if comparison between two datasets is the aim of a visualization, the impact of visual short term memory (e.g., change blindness, the inability to detect differences between spatially or temporally adjacent scenes [Rensink, O'Regan, & Clark, 1997; Simons & Rensink, 2005]) should be mitigated [Nowell, Hetzler, & Tanasse, 2001; Haroz & Heitmann, 2008]. This effect can be operationalized using visual short-term memory tasks including change detection [Luck & Vogel, 1997; Alvarez & Cavanagh, 2004]. A limit on our ability to attend to multiple locations is another example of a potential hurdle for visualizations. A visualization that incorporates animations or moving objects should take into account the limits of human multiple object tracking [Pylyshyn & Storm, 1988]. In general, visualization designers need to test visualizations not only for simple pre-attentive tasks but also for tasks which are limited by attention.

FUTURE WORK

We showed the performance impact of grouping on a single, highly discriminable categorical dimension. However, for a scalar, rather than categorical, dimension, the discriminability of the features can be an additional limiting factor. Further, data sets often represent multiple dimensions using more than one feature at a time (e.g., color and size). Examining and understanding how such multiplexing may limit capacity and search in the context of common visualization tasks could help the community understand how to more effectively visualize multiple dimensions.

CONCLUSION

We have three main conclusions: (1) Grouping is far more beneficial for oddball search compared with known-target search. (2) Accessing overall information (like heterogeneity or number of categories) is better for grouped displays. (3) For difficult tasks, aim to reduce variety in the entire view rather than optimizing small regions.

The implication is that the strict limits of attention have profound effects on the ability of observers to extract information from displays. Even performance with a visual feature like color, commonly thought of as 'pre-attentive', can be adversely effected by tasks or arrangements that put a heavy demand on attention and capacity. Together, there is an interaction between perceptual and cognitive limits and task demands. Accounting for these interactions can help design more efficient and effective visualizations.

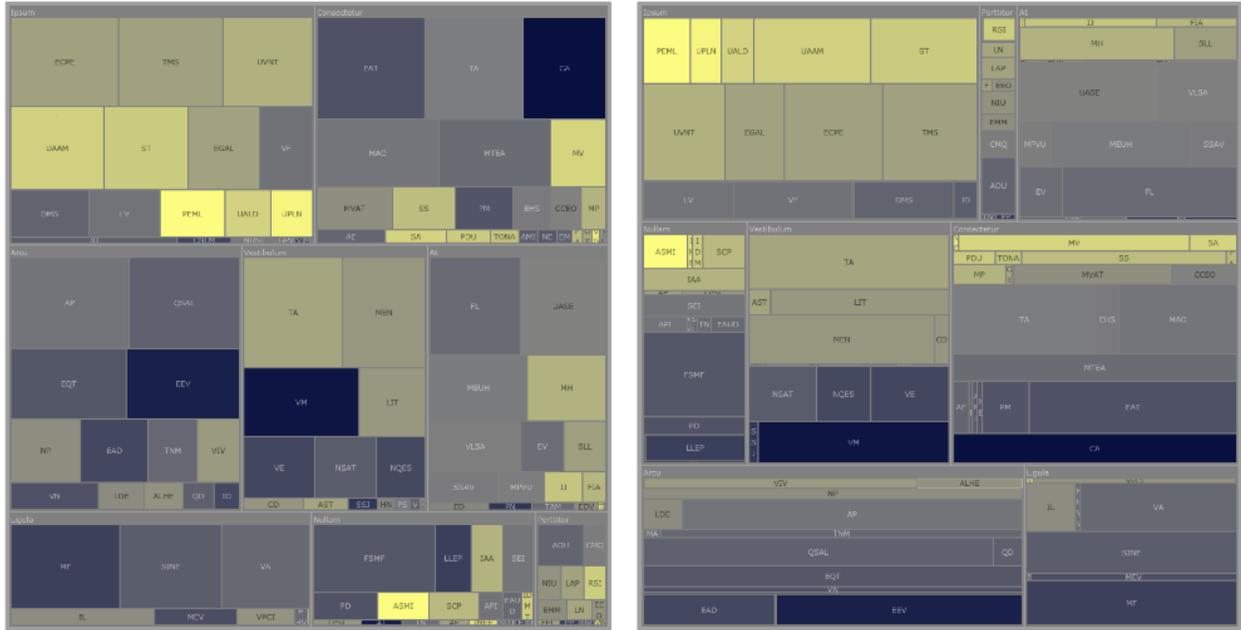


Figure 4-11. In visualizing a stock market using a sorted squarified treemap (left), the impact of prioritizing aspect ratio is a scattering of the colored rectangles. By prioritizing color and size (right), better visual grouping occurs, and more trends can be found by users. In these two stock market visualizations, company capital corresponds to size, color is a function of intraday percent change, and the stocks are segmented by industry. This is a similar approach used by smartmoney.com. Can you tell which industry performed best? Or worst? (Hint: the two visualization use the same data)

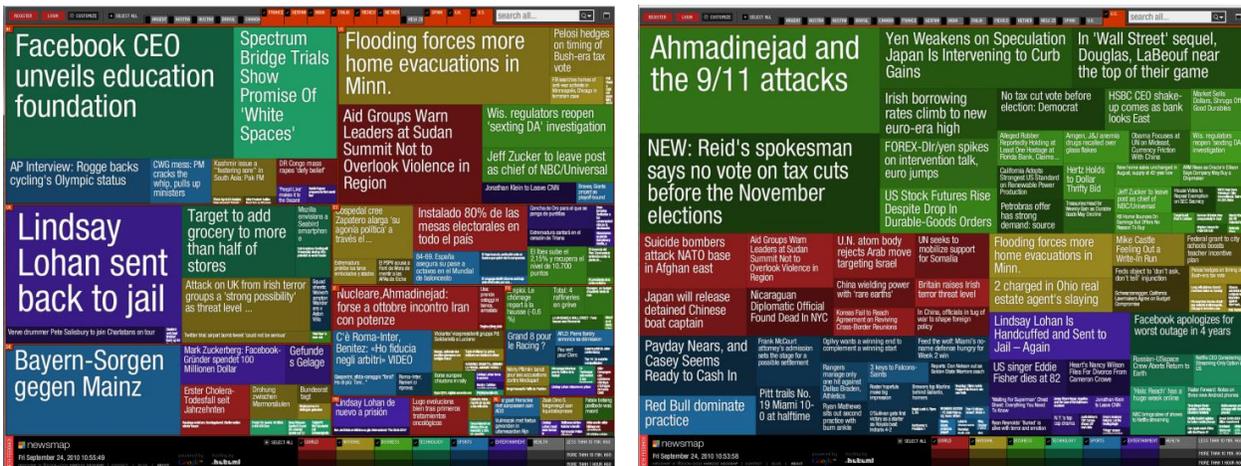


Figure 4-12. The left image with news from multiple countries is technically presenting more information to the user. Yet the complexity of the layout severely degrades basic information (which news category is most prevalent?). On the right the grouping that occurs when only one country is selected simplifies any analysis of the proportions between news categories. Though they both have the same number of categories and (due to screen real estate) are only capable of utilizing the same fixed area for stories, the right image appears more streamlined and useful.

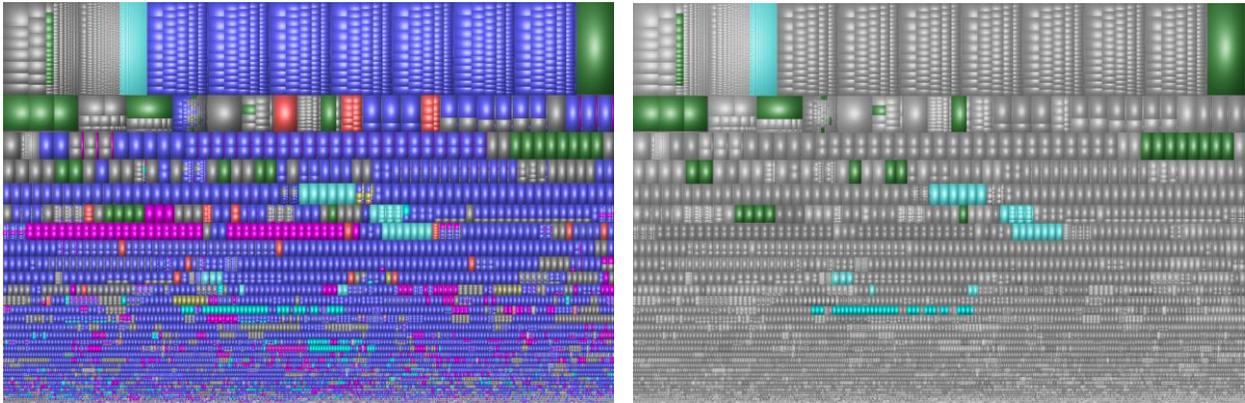


Figure 4-13. (left) This image of WinDirStat shows a complex arrangement of many types of files. Trying to grasp the number of categories present or their relative organization relative to each other is difficult unless attention is focused on a single variant of a feature (attend to green) or the filtration is performed for the user (right).

5. CONCLUSION

Individually, each of these projects acts as an example of the efficacy of exploring the HVS to derive guidelines for visualization. Collectively, these works show that the HVS can impact visualization design in many ways.

At the low level of scale, Natural Visualizations found that scale distribution is correlated with expert preference. However, the impact of low-level vision on visualization and extends beyond preference to outright misinterpretation. For example, Mach Bands – an effect implicated as early as the retina in the misperception of luminance differences near a boundary [Mach, 1865] – can add illusory shadows or highlights to the perception of X-Ray images and heatmaps. Another illusion, the tilt contrast effect, can shift the perceived orientation of a line [Gibson, 1937], which could alter the perceived direction of glyphs in a static vector field visualization. So understanding low-level vision can help assess not just quality but also perceptual accuracy of a visualization.

Moving towards perceptual representations of more complex features, understanding motion perception – as discussed in Temporal Thresholds for Feature Detection in Flow Visualization – can lead to rendering shortcuts that improve the computational performance of visualizations. These shortcuts can also extend to complex types of motion such as character animation. Our rapid perception of groups of people as a collective ensemble (rather than as individuals) [Sweeny, Haroz, & Whitney, 2011, 2013] could allow for animation rendering shortcuts in individuals characters. Meanwhile understanding the misperception of a character's heading when walking towards the viewer [Sweeny, Haroz, & Whitney, 2012a, 2012b] could prevent unintentional misperceptions.

The impact on visualization also extends to the scene-wide but capacity limited attention, which can explain limitations in the user performance [Haroz & Whitney, 2012]. Attentional limits can also result in people missing important information [Simons & Chabris, 1999], which the accuracy drops in our work showed [Haroz & Whitney, 2012]. This inattentional blindness is especially prevalent when spatially or temporally near an element of interest [Mack & Rock, 1998; Kristjánsson & Nakayama, 2002], which speaks to the criticalness of visualization layout.

Overall, the human visual system – from small edge encoding to more complex motion encoding to attentional selection – can inform every facet of visualization design to improve preference and performance.

6. REFERENCES

- Albright, T. D., & Stoner, G. R. (1995). Visual Motion Perception. *Proceedings of the National Academy of Sciences*, 92(7), 2433–2440.
- Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*, 15(2), 106–111.
- Ariely, D. (2001). Seeing Sets: Representation by Statistical Properties. *Psychological Science*, 12(2), 157–162.
- Barlow, H. B. (1958). Temporal and spatial summation in human vision at different background intensities. *The Journal of physiology*, 141(2), 337–50.
- Bartram, L., Ware, C., & Calvert, T. (2003). Moticons: detection, distraction and task. *International Journal of Human-Computer Studies*, 58(5), 515–545.
- Bederson, B. B., Shneiderman, B., & Wattenberg, M. (2002). Ordered and quantum treemaps: Making effective use of 2D space to display hierarchies. *ACM Transactions on Graphics*, 21(4), 833–854.
- Bertin, J. (1983). *Semiology of graphics : diagrams, networks, maps*. University of Wisconsin Press.
- Braddick, O. (1974). A short-range process in apparent motion. *Vision Research*, 14(7), 519–527.
- Braddick, O. J., & Holliday, I. E. (1991). Serial search for targets defined by divergence or deformation of optic flow. *Perception*, 20(3), 345–354.
- Brath, R. (1997). Metrics for effective information visualization. In *Proceedings of VIZ '97: Visualization Conference, Information Visualization Symposium and Parallel Rendering Symposium* (pp. 108–111). IEEE Comput. Soc.
- Britten, K., & Newsome, W. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual ...*, 13(01), 87.
- Burr, D. C., & Santoro, L. (2001). Temporal integration of optic flow, measured by contrast and coherence thresholds. *Vision Research*, 41(15), 1891–1899.
- Chalupa, L. M., & Werner, J. S. (2004). *The Visual Neurosciences. Perception* (pp. 908–923).
- Chong, S. C., & Treisman, A. (2003). Representation of statistical properties. *Vision Research*, 43(4), 393–404.
- Cleveland, W. S., & McGill, R. (1985). Graphical perception and graphical methods for analyzing scientific data. *Science (New York, N.Y.)*, 229(4716), 828–33.
- De Valois, R. L., & De Valois, K. K. (1988). *Spatial Vision*. Oxford University Press.
- DiBiase, D., MacEachren, A. M., Krygier, J. B., & Reeves, C. (1992). Animation and the role of map design in scientific visualization. *Cartography and Geographic Information Systems*, 19(4), 201–214, 265–266.
- Duffy, C. J., & Wurtz, R. H. (1991). Sensitivity of MST neurons to optic flow stimuli. I. A continuum of response selectivity to large-field stimuli. *Journal of Neurophysiology*, 65(6), 1329.

- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*(3), 433–458.
- Durant, S., & Zanker, J. (2009). Characterizing motion contour detection mechanisms and equivalent mechanisms in the luminance domain. *Journal of vision*, *9*(1).
- Fekete, J.-D., Grinstein, G., & Plaisant, C. (2004). IEEE InfoVis 2004 Contest, the history of InfoVis.
- Fekete, J.-D., Wijk, J. J. Van, Stasko, J. T., & North, C. (2008). The Value of Information Visualization. (A. Kerren, J. T. Stasko, J.-D. Fekete, & C. North, Eds.) *Information Visualization*, *4950*, 1–18.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A Optics and image science*, *4*(12), 2379–2394.
- Field, D. J. (1993). Scale-invariance and self-similar “wavelet” transforms: an analysis of natural scenes and mammalian visual systems. In M. Farge, J. C. R. Hunt, & J. C. Vascillicos (Eds.), *Wavelets Fractals and Fourier Transforms* (pp. 151–193). Clarendon Press.
- Franconeri, S. L., Alvarez, G. A., & Cavanagh, P. (2013). Flexible cognitive resources: competitive content maps for attention and memory. *Trends in cognitive sciences*, *17*(3), 134–41.
- Franconeri, S. L., & Simons, D. J. (2003). Moving and looming stimuli capture attention. *Perception & psychophysics*, *65*(7), 999–1010.
- Gazzaniga, M. S., Ivry, R. B., & Mangun, G. R. (2009). *Cognitive Neuroscience: The Biology of the Mind. Corpus* (Vol. 2, p. 666).
- Gibson, J. J. (1937). Adaptation, after-effect, and contrast in the perception of tilted lines. II. Simultaneous contrast and the areal restriction of the after-effect. *Journal of Experimental Psychology*, *20*(6), 553.
- Gibson, J. J. (1986). The ecological approach to visual perception.
- Glass, L. (1969). Moiré Effect from Random Dots. *Nature*, *223*(5206), 578–580.
- Greenwood, J. A., & Edwards, M. (2009). The detection of multiple global directions: capacity limits with spatially segregated and transparent-motion signals. *Journal of vision*, *9*(1), 40.1–15.
- Grinstein, G., Cvek, U., Derthick, M., & Trutschl, M. (2005). IEEE InfoVis 2005 contest, Technology Data in the US.
- Haroz, S., & Heitmann, K. (2008). Seeing the difference between cosmological simulations. *IEEE computer graphics and applications*, *28*(5), 37–45.
- Haroz, S., & Ma, K.-L. (2006). Natural Visualizations. In B. S. Santos, T. Ertl, & K. Joy (Eds.), *Eurovis 2006* (pp. 43–50). Lisbon, Portugal: Eurographics Association.
- Haroz, S., & Whitney, D. (2010). Temporal thresholds for feature detection in flow visualization. In *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization - APGV '10* (p. 163). ACM Press.
- Haroz, S., & Whitney, D. (2012). How Capacity Limits of Attention Influence Information Visualization Effectiveness. *IEEE Transactions on Visualization and Computer Graphics*, *18*(12), 2402–2410.

- Haroz, S., & Whitney, D. (2013). Global - Not Local - Variance Impacts Search. *Journal of Vision*, 13(9), 693–693.
- Healey, C. G. (1996). Choosing effective colours for data visualization. In *Proceedings of Seventh Annual IEEE Visualization '96* (pp. 263–270.). ACM.
- Helgeland, A., & Elboth, T. (2006). High-Quality and Interactive Animations of 3D Time-Varying Vector Fields. *IEEE Transactions on Visualization and Computer Graphics*, 12(6), 1535–1546.
- Holcombe, A. O. (2009). Seeing slow and seeing fast: two limits on perception. *Trends in cognitive sciences*, 13(5), 216–21.
- Horowitz, T. S., & Wolfe, J. M. (1998). Visual search has no memory. *Nature*, 394(6693), 575–577.
- Huang, L., & Pashler, H. (2007). A Boolean map theory of visual attention. *Psychological Review*, 114(3), 599–631.
- Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. (D. E. Dulany, R. L. Devalois, D. C. Beardslee, & M. Winterbottom, Eds.) *The Journal of Physiology*, 148(3), 574–591.
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol*, 160(1), 106–154.
- Huber, D. E., & Healey, C. G. (2005). Visualizing Data with Motion. In *IEEE Visualization 2005 - (VIS'05)* (pp. 67–67). IEEE.
- Intriligator, J., & Cavanagh, J. P. (2001). The Spatial Resolution of Visual Attention. *Cognitive Psychology*, 43(3), 171–216.
- Karklin, Y., & Lewicki, M. S. (2003). Learning higher-order structures in natural images. *Network (Bristol, England)*, 14(3), 483–99.
- Kaufman, E. L., & Lord, M. W. (1949). The discrimination of visual number. *The American journal of psychology*, 62(4), 498–525.
- Kong, N., Heer, J., & Agrawala, M. (2010). Perceptual guidelines for creating rectangular treemaps. *IEEE transactions on visualization and computer graphics*, 16(6), 990–8.
- Kristjánsson, Á., & Nakayama, K. (2002). The attentional blink in space and time. *Vision Research*, 42(17), 2039–2050.
- Lane, D. A. (1994). UFAT-a particle tracer for time-dependent flow fields. In *Proceedings Visualization '94* (pp. 257–264). IEEE Comput. Soc. Press.
- Langer, M. S. (2000). Large-scale failures of f(-alpha) scaling in natural image spectra. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 17(1), 28–33.
- Lappin, J. S., & Bell, H. H. (1972). Perceptual differentiation of sequential visual patterns. *Perception & Psychophysics*, 12(2), 129–134.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279–81.

- Lum, E. B., Stoppel, A., & Ma, K.-L. (2002). Kinetic visualization: a technique for illustrating 3D shape and structure. In *IEEE Visualization, 2002. VIS 2002.* (pp. 435–442). IEEE.
- Mach, E. (1865). Über die Wirkung der räumlichen Vertheilung des Lichreizes auf di Netzhaut. In F. Ratliff (Trans.), *Sitzungsberichte der mathematisch-naturwissenschaftlichen Classe der kaiserlichen Akademie der Wissenschaften* (Vol. 54).
- Mack, A., & Rock, I. (1998). *Inattentional blindness*. MIT Press Cambridge.
- Marcar, V. L., Xiao, D. K., Raiguel, S. E., Maes, H., & Orban, G. A. (1995). Processing of kinetically defined boundaries in the cortical motion area MT of the macaque monkey. *Journal of neurophysiology*, 74(3), 1258–70.
- Meiselbach, C., & Bruckschen, R. (1996). Interactive Visualization: On the way to a virtual wind tunnel. In *Flow Simulation with High-Performance Computers II* (pp. 203–208).
- Nakayama, K. (1985). Biological image motion processing: a review. *Vision research*, 25(5), 625–60.
- Newsome, W., & Pare, E. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J. Neurosci.*, 8(6), 2201–2211.
- Nothdurft, H. C. (1993). The role of features in preattentive vision: comparison of orientation, motion and color cues. *Vision Research*, 33(14), 1937–1958.
- Nowell, L., Hetzler, E., & Tanasse, T. (2001). *Change blindness in information visualization: a case study*. *IEEE Symposium on Information Visualization, 2001. INFOVIS 2001.* (pp. 15–22). IEEE.
- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583), 607–9.
- Park, S., Budge, B., Linsen, L., Hamann, B., & Joy, K. I. (2005). Dense geometric flow visualization. In *EuroVis 2005* (pp. 21–28). Eurographics Association.
- Peitgen, H.-O., & Saupe, D. (1988). *The Science of Fractal Images*. Springer-Verlag New York, Inc.
- Pylyshyn, Z., & Storm, R. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial vision*, 3(3), 179–197.
- Reichardt, W., Poggio, T., & Hausen, K. (1983). Figure-ground discrimination by relative movement in the visual system of the fly. *Biological Cybernetics*, 46(S1), 1–30.
- Reinhard, E., Shirley, P., & Troscianko, T. (2001). Natural image statistics for computer graphics. *Univ. Utah Tech Report UUCS-01*
- Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To See or not to See: The Need for Attention to Perceive Changes in Scenes. *Psychological Science*, 8(5), 368–373.
- Rosenholtz, R., Twarog, N. R., Schinkel-Bielefeld, N., & Wattenberg, M. (2009). An intuitive model of perceptual grouping for HCI design. In *Proceedings of the 27th international conference on Human factors in computing systems - CHI '09* (p. 1331). New York, New York, USA: ACM Press.

- Royden, C. S., Wolfe, J. M., & Klempen, N. (2001). Visual search asymmetries in motion and optic flow fields. *Perception & psychophysics*, *63*(3), 436–44.
- Ruderman, D. L. (1997). Origins of scaling in natural images. *Vision Research*, *37*(23), 3385–3398.
- Saito, H., Yukie, M., Tanaka, K., Hikosaka, K., Fukada, Y., & Iwai, E. (1986). Integration of direction signals of image motion in the superior temporal sulcus of the macaque monkey. *J. Neurosci.*, *6*(1), 145–157.
- Schroeder, M., & Wiesenfeld, K. (1991). Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise. *Physics Today*, *44*(11), 91.
- Sears, A. (1995). AIDE: a step toward metric-based interface development tools. In *Proceedings of the 8th annual ACM symposium on User interface and software technology - UIST '95* (pp. 101–110). New York, New York, USA: ACM Press.
- Shneiderman, B. (1996). The eyes have it: a task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages* (Vol. 0, pp. 336–343). IEEE Comput. Soc. Press.
- Sillito, A. M., Cudeiro, J., & Jones, H. E. (2006). Always returning: feedback and sensory processing in visual cortex and thalamus. *Trends in neurosciences*, *29*(6), 307–16.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual review of neuroscience*, *24*, 1193–216.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: sustained inattention blindness for dynamic events. *Perception*, *28*, 1059–1074.
- Simons, D. J., & Rensink, R. A. (2005). Change blindness: past, present, and future. *Trends in cognitive sciences*, *9*(1), 16–20.
- Somervell, J., Mccrickard, D. S., North, C., & Shukla, M. (2002). An Evaluation of Information Visualization in Attention-Limited Environments. In D. Ebert, P. Brunet, & I. Navazo (Eds.), *VISSYM 2002* (pp. 211–216). Eurographics Association.
- Sweeny, T. D., Haroz, S., & Whitney, D. (2011). Seeing the direction of a crowd: Ensemble coding of biological motion. *Journal of Vision*, *11*(11), 742–742.
- Sweeny, T. D., Haroz, S., & Whitney, D. (2012a). Just walk away: Reference repulsion in the perception of crowd behavior. *Journal of Vision*.
- Sweeny, T. D., Haroz, S., & Whitney, D. (2012b). Reference repulsion in the categorical perception of biological motion. *Vision research*, *64*, 26–34.
- Sweeny, T. D., Haroz, S., & Whitney, D. (2013). Perceiving group behavior: sensitive ensemble coding mechanisms for biological motion of human crowds. *Journal of experimental psychology. Human perception and performance*, *39*(2), 329–37.
- Tanaka, K., Hikosaka, K., Saito, H., Yukie, M., Fukada, Y., & Iwai, E. (1986). Analysis of local and wide-field movements in the superior temporal visual areas of the macaque monkey. *J. Neurosci.*, *6*(1), 134–144.

- Tateosian, L. G., Dennis, B. M., & Healey, C. G. (2006). Stevens dot patterns for 2D flow visualization. In *Proceedings of the 3rd symposium on Applied perception in graphics and visualization - APGV '06* (p. 93). New York, New York, USA: ACM Press.
- Tatu, A., Bak, P., Bertini, E., Keim, D., & Schneidewind, J. (2010). Visual quality metrics and human perception. In *Proceedings of the International Conference on Advanced Visual Interfaces - AVI '10* (p. 49). ACM Press.
- Torralba, A., & Oliva, A. (2003). Statistics of natural image categories. *Network (Bristol, England)*, *14*(3), 391–412.
- Treisman, A. (1982). Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology Human Perception and Performance*, *8*(2), 194–214.
- Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive psychology*, *12*(1), 97–136.
- Tufte, E. R. (1990). *Envisioning Information*. (G. Press, Ed.) *Applied Statistics* (Vol. 41, p. 126). Graphics Press.
- Ware, C., & Beatty, J. C. (1988). Using color dimensions to display data dimensions. *Human Factors*, *30*(2), 127–142.
- Watson, D. G., Maylor, E. A., & Bruce, L. A. M. (2005). The efficiency of feature-based subitization and counting. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(6), 1449–1462.
- Weisstein, E. W. (n.d.). Power Spectrum. *From MathWorld--A Wolfram Web Resource*. <http://mathworld.wolfram.com/PowerSpectrum.html>.
- Williams, D. W., & Sekuler, R. (1984). Coherent global motion percepts from stochastic local motions. *Vision Research*, *24*(1), 55–62.
- Wolfe, J. M. (1994). Guided Search 2.0 A revised model of visual search. *Psychonomic Bulletin & Review*, *1*(2), 202–238.
- Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature reviews. Neuroscience*, *5*(6), 495–501.