

Visualizing Data with Motion

Daniel E. Huber*
Northrop Grumman

Christopher G. Healey†
North Carolina State University

ABSTRACT

This paper describes an experimental study of three perceptual properties of motion: flicker, direction, and velocity. Our goal is to understand how to apply these properties to represent data in a visualization environment. Results from our experiments show that all three properties can encode multiple data values, but that minimum visual differences are needed to ensure rapid and accurate target detection: flicker must be coherent and must have a cycle length of 120 milliseconds or greater, direction must differ by at least 20° , and velocity must differ by at least 0.43° of subtended visual angle. We conclude with an overview of how we are applying our results to real-world data, then discuss future work we plan to pursue.

CR Categories: H.1.2 [Models and Principles]: User/Machine Systems—*Human factors, human information processing*; I.3.3 [Computer Graphics]: Picture/Image Generation—*Display algorithms*; I.3.6 [Computer Graphics]: Methodology and Techniques—*Interaction techniques*

Keywords: direction, flicker, motion, multidimensional, perception, velocity, visualization

1 INTRODUCTION

Visualization converts information into images that can be used to explore, discover, and analyze. The amount of data being collected and recorded has grown rapidly in recent years. Unfortunately, our ability to understand this data has improved more slowly. Building images that allow viewers to analyze their data offers one promising approach to increasing the amount of information we can manage effectively.

Our interests in this paper focus on multidimensional datasets. Such a dataset D stores information about multiple data attributes $A = (A_1, \dots, A_n)$ via numerous sample points or data elements e_i . Specifically, $D = \{e_1, \dots, e_m\}$ and $e_i = (a_{i,1}, \dots, a_{i,n})$, $a_{i,j} \in A_j$. The goal of visualization in this context is to convert the dataset into images that allow viewers to rapidly and accurately complete their analysis tasks.

One way to build *effective* visualizations is to construct data to visual feature mappings based on guidelines from human perception. The low-level human visual system can identify certain visual properties in a scene or an image very rapidly, often in 200 milliseconds (msec) or less [12]. When combined properly, these properties can be used to represent information in ways that allow viewers to perform high-level analysis tasks like locating target elements and spatial boundaries, enumerating, or tracking groups of elements as they move and change over time. The ability to complete these tasks quickly on large, multi-element displays allows analysis to occur in real-time, dramatically increasing the amount of information a viewer can process.

*email: dan.huber@gmail.com

†email: healey@csc.ncsu.edu

Past research has documented how colour and texture can be decomposed into perceptually salient visual features like hue, luminance, and saturation, or size, orientation, density, and regularity of placement [5]. One method for visualizing multidimensional data is to attach a unique visual feature to each data attribute, for example, to use hue to represent the first attribute, orientation to represent the second attribute, and so on. How we map the data to the different visual features depends critically on three constraints: (1) the capabilities of each visual feature; (2) the makeup of the data; and (3) the analysis tasks the viewer needs to perform.

Each visual feature has a certain information carrying capacity, representing its ability to encode data in isolation. Different features can also interact with one another in various ways when they are displayed together in a common image. For example, colour properties like hue and luminance initially dominate texture properties, forcing colour patterns to be seen first by the low-level visual system. These types of issues must be considered when we design a visualization's data-feature mapping.

A third visual feature we would like to use is motion. Experimental results from psychophysics show that properties of motion are detected by the low-level visual system. Motion occurs, both implicitly and explicitly, in numerous visualization techniques. However, the detailed knowledge needed to construct perceptual guidelines on the use of motion in visualization has not been documented. This paper describes a set of perceptual experiments we conducted to investigate the information carrying capacity of three perceptual dimensions of motion: flicker, direction, and velocity.

The remainder of this paper is organized as follows. First, we provide an overview of previous work from psychophysics on motion, together with past research on the use of motion in visualization. Next, we discuss three experiments that studied the basic abilities of flicker, direction of motion, and velocity of motion to encode information. We describe our initial attempts to use our results to visualize real-world data. Finally, we conclude with a summary of our findings and a discussion of areas of future work.

2 MOTION

Like colour and texture, motion is a visual feature that is known to be perceptually salient. Examples of motion in visualization include the animation of particles, dye, or glyphs to represent the direction and magnitude of a vector field (e.g., fluid flow visualization), or the implicit use of motion transients to highlight changes in a dataset across a user-selected data axis (e.g., over time for a temporal dataset, or along the scanning axis for CT or MRI slices).

Although properties of motion are commonly used in visualization, formal studies of the strengths and weaknesses of these properties are less numerous. As with colour and texture, our goal is to identify perceptual dimensions of motion and apply them in an effective manner. Three basic properties described in the perceptual literature were of particular interest to us: flicker, direction of motion, and velocity of motion.

Flicker refers to a repeating on-off pattern applied to an image or an object, and is normally measured as the frequency of repetition F in cycles per second (cps). A common use of flicker research in computer graphics is the critical flicker frequency (CFF), the rate at which images must be redrawn to appear continuous. Below the

CFF, individual clear-and-display operations can be perceived, producing the impression of a “flickering” image. Although $F = 60$ cps is an often-cited rule of thumb for the CFF, this number varies depending on the colour, brightness, or size of the object being displayed, and on its eccentricity, which measures the distance in visual angle from the viewer’s current focal point to the object. Our interest is in flicker frequencies that are perceived as discrete flashes by the viewer, that is, frequencies below the CFF. Experiments by Mowbrey and Gebhard, later discussed by Brown, suggest that frequency must vary from 2 to 5% to produce a distinguishable difference in flicker ($1.02 \leq \Delta F \leq 1.05$) [3, 9]. As objects move away from the focus of attention, however, ΔF increases dramatically, with differences of 100% or more needed for high F at large eccentricities. Separate findings describe how differences in luminance and size can affect ΔF .

Other vision scientists have studied both direction and velocity of motion. Tynan and Sekuler reported that a decrease in a target object’s velocity or an increase in its eccentricity increased the time required to identify it [13], although in all cases viewers responded rapidly. Two hundred to 350 msec were needed to identify targets in the periphery, and 200 to 310 msec for targets in the center of focus. van Doorn and Koenderink confirmed that higher initial velocities produce a faster response to a change in the velocity [14]. They claim this is due to the need for the target to traverse a “critical distance” before it can be detected. Follow-on work by Mateeff et al. [8] showed that for a baseline velocity V_1 and a target velocity $V_2 = 2V_1$, approximately 100 msec is needed to see the velocity change from V_1 to V_2 for slow V_1 (1° per second) and approximately 50 msec for faster V_1 (2° per second or higher). Recent results by Hohsbein and Mateeff suggest that changes in the direction of motion may be detected based on the perceived velocity change they produce [6]. Related work suggests at least 15° of difference is needed to distinguish between static elements with different orientations [16].

Researchers in psychology have used properties of motion to extend a viewer’s ability to perform basic exploration tasks. Nakayama and Silverman showed that coherent motion or stereoscopic depth can be used to perceptually group elements [10], allowing viewers to search each group independently and in parallel. Driver et al. showed that oscillation can also be used to separate elements into independent visual groups, but only if the oscillation pattern is coherent [4]. More sophisticated motion patterns have also been analyzed, although with less success in terms of achieving high-speed search performance. Braddick and Holliday studied both divergence and deformation [2]. Although the basic motion properties being shown can be rapidly identified in isolation, the combinations that produce deformation and divergence were not detected by the low-level visual system.

Properties of motion have been extended to visualization design. Animated motion is used in flow visualization to show the direction and speed of different flow patterns (e.g., in van Wijk [15]). Kerlick proposed the use of animated glyphs to visualize 2D and 3D multidimensional datasets [7]. He designed a set of “boids” to encode attribute values at specific locations in the dataset, for example, a sphere boid to query data values at a user-selected location, or pyramid and dart boids that animate over a vector field to visualize its shape. Bartram et al. studied the use of variations in colour, shape, and motion to “notify” viewers while they were engaged in a separate, attention-demanding task [1]. Results showed that applying motion to a static glyph was significantly easier to recognize, compared to changing the glyph’s colour or shape. This finding held both when the glyph was near the center of focus, and when it was located on the periphery of the viewer’s gaze. Experiments were conducted to measure how distracting each secondary motion cue appeared to a viewer. Flicker was the least distracting, followed by oscillating motion, then divergence, and finally movement over

long distances. The authors concluded by confirming that different motion paths can be used to perceptually group glyphs in a manner similar to the work of Nakayama and Silverman and Driver et al.

3 FLICKER EXPERIMENT

The flicker experiment studied a viewer’s ability to distinguish the presence or absence of a small group of target elements that flicker at a rate different from the background elements. Each experiment trial contained a 20×20 regular grid of yellow squares on a black background that filled a 19-inch screen. For target present trials, a 3×3 group of elements was randomly selected to flicker at a target flicker rate f_t different from the background flicker rate f_b . For target absent trials, all the elements flickered at the same background flicker rate f_b .

Observer accuracy and response time were used to measure viewer performance. The experiment was constructed to test three separate conditions:

1. *Cycle length*: the duration of the target elements’ cycle f_t in milliseconds. This allowed us to test whether different cycle lengths produced different viewer performance.
2. *Cycle difference*: the absolute difference between the target and background elements’ cycle lengths $\Delta F = |f_t - f_b|$. This allowed us to determine how much of a difference in cycle length is needed to generate a perceived difference in flicker rate.
3. *Coherency*: whether elements with the same cycle length flicker in phase or not. When flicker is coherent, elements with the same cycle length turn on and off together. For example, given two elements e_1 and e_2 with cycle lengths of 120 msec, both e_1 and e_2 have an on-off pattern at timesteps $t = 0, 60, 120, 180, \dots$. When flicker is non-coherent, elements with the same cycle length can turn on and off at different times. Given the same elements e_1 and e_2 with cycle lengths of 120 msec, e_1 could have an on-off pattern of $t = 0, 60, 120, 180, \dots$, while e_2 could have a different on-off pattern of $t = 20, 80, 140, 200, \dots$. In other words, even though the length of the on-off cycle is the same 120 msec, the onset of each cycle can occur at different times. This allowed us to test for the effect of coherency on viewer performance.

We tested four target-background cycle differences: $\Delta F = 120, 240, 360,$ and 480 msec. The experiment included four target cycle lengths for each ΔF : $f_t = \Delta F, \Delta F + 120, \Delta F + 240,$ and $\Delta F + 360$ msec. For each f_t two different background cycle lengths $f_b = f_t \pm \Delta F$ were selected, except when $f_t = \Delta F$, since $f_b = f_t - \Delta F = 0$ in this situation. To avoid $f_b = 0$ while still maintaining balance within the experiment, $f_b = f_t + \Delta F$ was used twice when $f_t = \Delta F$. This produced eight (f_t, f_b) pairs for each ΔF . For example, for $\Delta F = 240$, we tested the following target-background pairs:

- $f_t = 240$ msec: (240, 480) and (240, 480)
- $f_t = 360$ msec: (360, 120) and (360, 600)
- $f_t = 480$ msec: (480, 240) and (480, 720)
- $f_t = 600$ msec: (600, 360) and (600, 840)

Target-background pairs for the other three ΔF were generated in an identical fashion. This resulted in 64 different trial types (two coherency types by four ΔF by four f_t by two f_b). Each trial type was shown to a viewer six times during the experiment, for a total of 384 trials. Half the trials for a given trial type were randomly selected to contain a target group of elements flickering at the target flicker rate. The remaining trials contained no target, with all the elements flickering at the background flicker rate.

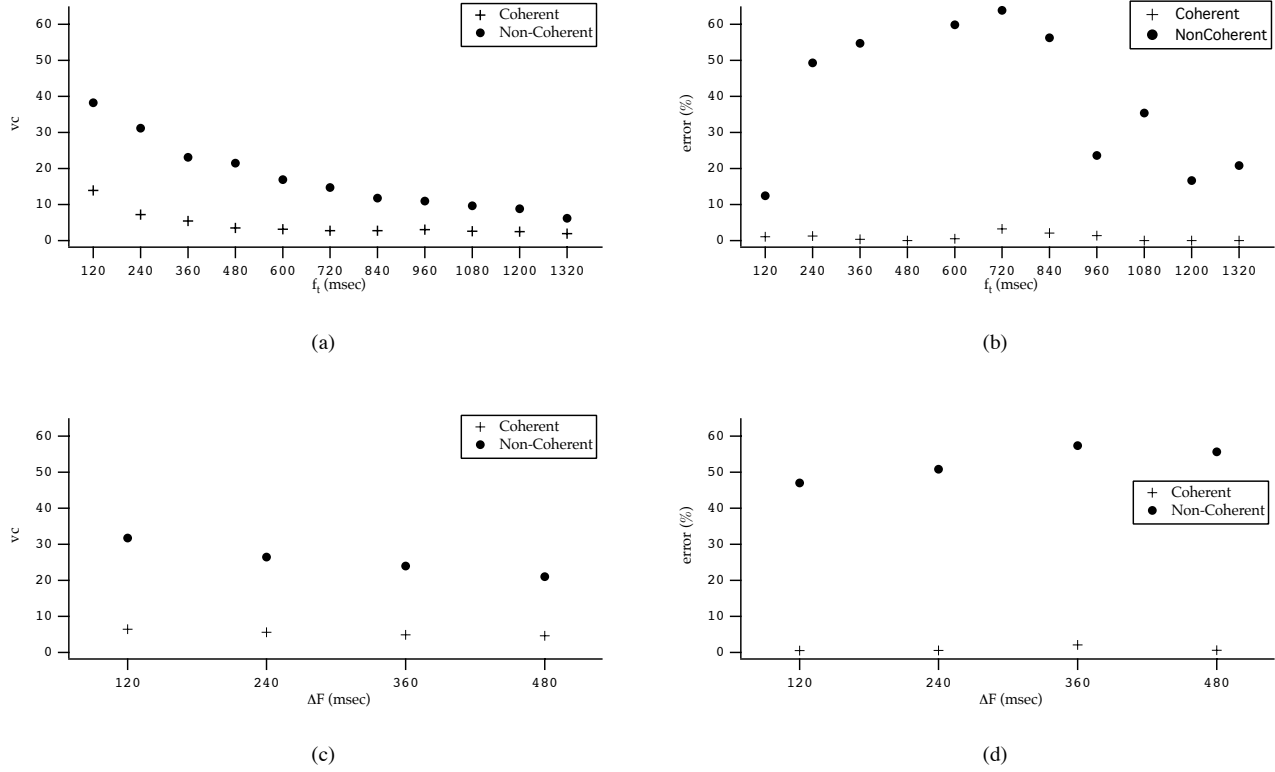


Figure 1: Average viewer performance during the flicker experiments for coherent and non-coherent trials: (a) \overline{vc} versus f_t ; (b) \bar{e} versus f_t ; (c) \overline{vc} versus ΔF during target present trials; (d) \bar{e} versus ΔF during target present trials

3.1 Procedure

Eight members of the Computer Science department (faculty, staff, and graduate students) volunteered to participate during the experiment. Viewers were naive to the purpose of the experiment, although some students had knowledge of related work conducted in our laboratory. All viewers had normal or corrected to normal vision. The experiment was run on a SunBlade 150 connected to a 19-inch Sun Microsystems LCD monitor.

At the beginning of the experiment viewers were instructed to search each display for a small group of target elements flickering at a rate different from the background elements. Viewers were told to hit a key representing “present” if the target group was present, or a key representing “absent” if the target group was absent. Viewers were asked to answer as quickly as possible, while still maintaining a high level of accuracy. Feedback was provided for one second after each viewer response: a plus-sign for a correct answer, or a minus-sign for an incorrect answer. This allows viewers to gauge their performance, and to ensure they are entering their answers properly. The next trial followed immediately after the feedback display. Coherent and non-coherent experiment trials were separated into two parts. Viewers completed each part separately. Half the viewers started with the coherent trials. The other half started with the non-coherent trials. Trials within each part were presented in a random order.

Prior to starting an experiment part, viewers completed 32 practice trials. Sixteen practice trials were randomly selected to contain a target group. Each of the 16 possible $(\Delta F, f_t)$ pairs were shown once within these trials. The remaining 16 trials did not contain a target group. Viewers were allowed to repeat the practice session as often as they needed to become comfortable with the procedure and speed of the experiment. Viewers completed all 384 experiment

trials plus the 64 practice trials (192 experiment trials plus 32 practice trials for both the coherent and non-coherent experiment parts) during a single one hour session. Viewers were offered an opportunity to rest after every 96 experiment trials. A viewer’s present or absent response for each trial, together with the time taken to respond, were saved for later analysis.

3.2 Results

Two performance metrics were recorded during the experiment: mean viewer error rate \bar{e} and mean viewer response time \overline{rt} . Although \bar{e} can be analyzed directly, \overline{rt} may not provide a reliable direct measure for comparison, particularly for coherent trials across different ΔF . In order for a viewer to perceive a difference between the target and the background flicker rates, the target elements (or background elements) must complete at least half a cycle (i.e., either the target or the background must “turn off” while the other “stays on” at least once before a flicker difference is observable). The time required for half a cycle to complete is longer for larger ΔF . Therefore, we would expect \overline{rt} to also be longer for larger ΔF .

To address this issue, we combined ΔF and rt to compute the number of visibility changes vc that occur before a response is made. vc tells us how many times the target (or background) turned on and off before a viewer responded. vc is calculated using the faster of the two cycle lengths, $f_{vc} = \min(f_t, f_b)$. Given rt and f_{vc} , the number of visibility changes vc is:

$$vc = \lfloor 2 \frac{rt}{f_{vc}} \rfloor \quad (1)$$

Trials were divided along the conditions of target present or absent, coherency, target cycle length f_t , and cycle difference ΔF , with

\overline{vc} and \overline{e} for all viewers collapsed and averaged over the resulting categories. t -tests and analysis of variance (ANOVA) with a 95% confidence interval were used to identify statistically significant variations in performance. In summary, we found:

1. Coherent trials had significantly lower \overline{vc} and \overline{e} , compared to non-coherent trials.
2. In most cases, \overline{e} for non-coherent trials approached the chance rate of 50% (i.e., viewer responses were no better than simply guessing whether the target was present or absent).
3. For coherent trials, there was no effect of either ΔF or f_t on \overline{e} .
4. For coherent trials, \overline{vc} was significantly higher for smaller ΔF , particularly at $f_t = 120$ msec.
5. For non-coherent trials, \overline{e} was significantly lower for $f_{vc} = 120$ msec and $\Delta F \geq 360$, and for $f_{vc} \geq 960$ msec.

Both \overline{vc} and \overline{e} varied significantly between coherent and non-coherent trials, with $t(50) = 5.79$, $p < 0.0001$ and $t(50) = 13.56$, $p < 0.0001$, respectively. For coherent trials, $\overline{vc} = 6.8846$, $\overline{e} = 0.009$. For non-coherent trials, $\overline{vc} = 28.9231$, $\overline{e} = 0.465$.

Since \overline{e} approached the chance rate of 50% for non-coherent trials, we restrict our remaining analysis to coherent trials only. Target presence or absence during coherent trials had no effect on either \overline{vc} or \overline{e} , $t(24) = 0.091$, $p = 0.928$ and $t(24) = 0.034$, $p = 0.973$, respectively. For coherent target present trials, target cycle length f_t had no effect on either \overline{vc} or \overline{e} , $F(6, 9) = 1.469$, $p = 0.290$ and $F(6, 9) = 0.528$, $p = 0.775$, respectively. Finally, for coherent target present trials, cycle difference ΔF had a significant effect on \overline{vc} , but not on \overline{e} , $F(3, 12) = 56.391$, $p < 0.0001$ and $F(3, 12) = 0.524$, $p = 0.674$, respectively.

3.3 Interpretation

Our experiment results suggest that a difference in coherent flicker is easy to detect, at the cycle lengths we tested. Non-coherent flicker, however, is very difficult to see. When elements with the same cycle length flicker at different onsets, they normally cannot be grouped together by the visual system. The visual result is similar to a set of blinking Christmas lights: they seem to turn on and off randomly, and the fact that their cycle lengths may (or may not) be identical is no longer apparent. This is similar to Driver et al.'s finding that coherent oscillation is salient, but that non-coherent oscillation is difficult to detect [4].

For coherent trials, \overline{vc} was stable across ΔF for a given f_t . However, \overline{vc} was significantly higher when $f_t = 120$ msec, compared to other f_t values (Fig. 1a). During the experiment viewers needed time to formulate their answer, plus an additional 250 to 500 msec to react and enter their response. If f_t falls below the minimum reaction rate viewers are penalized by some additional visibility changes, because they cannot react quickly enough to avoid the changes that happen as they enter their response. This explains why \overline{vc} was higher for $f_t = 120$ msec. Less extreme examples of this effect can be seen for $f_t = 240$ and $f_t = 360$ msec.

One case where non-coherent flicker differences were easy to detect occurred when $f_{vc} = 120$ msec and $\Delta F \geq 360$ msec (Fig. 1b). Here, either the target elements or the background elements (but not both) flickered very quickly. Visually, this looked like either a rapidly changing region within a relatively stable background, or a rapidly changing background with a relatively stable region inset within it. The slower-changing region needed a cycle length of 360 msec or more to appear stable. When $f_{vc} > 120$ msec, the high-speed region no longer had the rapidly changing, blurred appearance needed to trigger this phenomenon.

Non-coherent error rates were also lower for $f_{vc} \geq 960$ msec. Since the largest f_t tested during target present trials was 840 msec,

any $f_{vc} > 840$ msec implies $f_t = f_b$, a target absent trial. When non-coherent trials alone are tested, target presence or absence is significant, $t(31) = 2.26$, $p = 0.031$, with target absent trials ($\overline{e} = 0.352$) more accurate than target present trials ($\overline{e} = 0.535$). This explains the lower error rates for these non-coherent trials.

4 DIRECTION EXPERIMENT

The direction experiment studied a viewer's ability to distinguish the presence or absence of a small 3×3 group of target elements that move in a direction d_t different from the background elements d_b . The velocity v of both target and background elements was fixed to a constant rate. Trials were otherwise identical to the flicker experiment. Observer accuracy and response time were used to measure viewer performance. The experiment was constructed to test two separate conditions:

1. *Motion direction*: the direction of the target elements' motion d_t measured in degrees rotated counterclockwise from the horizontal axis. This allowed us to test whether different directions produced different viewer performance.
2. *Motion difference*: the absolute difference between the target and background elements' directions of motion $\Delta D = |d_t - d_b|$. This allowed us to determine how much of an angular difference is needed to generate a perceived difference in motion direction.

In order to simulate motion, elements "walked" within constant-size screen regions (or cells) centered about an element's starting position. Neighboring cells adjoined with one another, so there were no gaps between the cells. As an element reached the edge of its cell, it wrapped around to the opposite side. When adjacent elements have the same direction of motion d and velocity v , and the same relative starting position within their cell, the impression of a continuous stream of motion is created. Each element reaches the edge of its cell and wraps around simultaneously, with the space left empty by one element being immediately filled by its adjacent neighbor. Thus, there is no visible separation between neighboring cells. When a target group is present in the display, however, a discontinuity will occur at the boundary between target and background cells, since the elements within the cells have different directions of motion $d_t \neq d_b$. This introduces a potential problem, since viewers may identify a target group by looking for the boundary discontinuity, and not by searching for different directions of motion.

We investigated a number of solutions to this problem. The simplest approach is to do nothing. Unfortunately, this produces a noticeable "popping" effect when elements reach a target-background edge. We tracked the center of an element e_i to determine when it should wrap around. When e_i 's center crosses a target-background edge, the entire element disappears from one side of the cell and reappears on the opposite side. Since there is no corresponding element to fill the space vacated by e_i , a visual discontinuity appears in the display. We next tried gradually increasing e_i 's transparency as it approached a target-background edge, and gradually decreasing its transparency as it moved away from the edge. Although this removed the visual discontinuity, we worried that viewers might identify targets by searching for variations in luminance, rather than differences in direction of motion. Our final solution was to use the stencil buffer to mask the group of cells containing the target elements. As a target element e_i passes over a target-background cell boundary, it is masked so that only the portion of e_i within the target cell region is drawn. The stencil also clips background elements as they pass into the target region. The visual effect is similar to placing an array of background elements on a plane with a hole cut through it, then drawing a second array of target elements underneath the plane. The target elements can be seen through the hole,

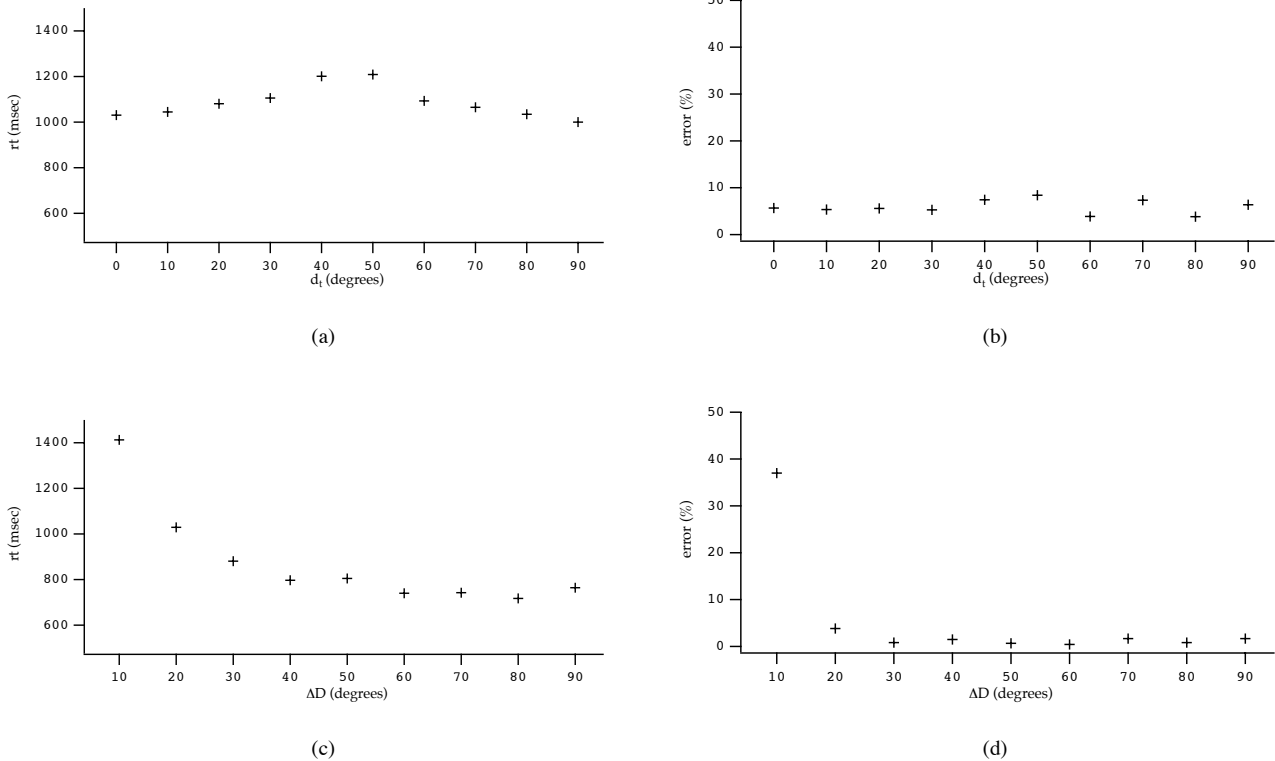


Figure 2: Average viewer performance during the direction experiments: (a) \overline{rt} versus d_t ; (b) \overline{e} versus d_t ; (c) \overline{rt} versus ΔD during target present trials; (d) \overline{e} versus ΔD during target present trials

moving smoothly beneath the background elements. By making elements appear and disappear in this way, no popping or change in luminance was visible. This forced viewers to use changes in the direction of motion to identify the presence or absence of a target group.

We studied ten different motion directions ranging from 0° to 90° rotated counterclockwise from the horizontal axis in 10° steps. Every possible combination of motion directions (d_t, d_b) was tested, producing 90 cases where $d_t \neq d_b$ (target present trials) and 10 cases where $d_t = d_b$ (target absent trials). Motion differences $\Delta D = |d_t - d_b|$ ranged from 10° to 90° during the target present trials. Each target present trial was shown three times, while each target absent trial was shown 27 times. This produced a total of 540 trials during the experiment.

4.1 Procedure

Ten members of the Computer Science department (faculty, staff, and graduate students) volunteered to participate during the experiment. Experiment details were similar to the flicker experiment, except that viewers were instructed to search each display for a small group of target elements moving in a direction different from the background elements.

Prior to starting the experiment, viewers completed 36 practice trials. 18 trials contained a target group (with each of the 9 possible motion differences included twice), and 18 trials did not. Viewers were allowed to repeat the practice session as often as they needed to become comfortable with the procedure and speed of the experiment. Viewers then completed 540 experiment trials presented in a random order during a single one hour session. Present or absent response for each trial, together with the time taken to respond, were

saved for later analysis.

4.2 Results

Two performance metrics were recorded during the experiment: mean viewer response time \overline{rt} and mean viewer error rate \overline{e} . Trials were divided along the conditions of target present or absent, motion direction, and motion difference, with \overline{rt} and \overline{e} for all viewers collapsed and averaged over the resulting categories. t -tests and ANOVAs with a 95% confidence interval were used to identify statistically significant variations in performance. In summary, we found:

1. \overline{rt} was significantly longer for target absent trials, compared to target present trials.
2. Neither \overline{rt} nor \overline{e} varied significantly across motion direction d_t .
3. Both \overline{rt} and \overline{e} varied significantly across motion difference ΔD .

\overline{rt} varied significantly between target absent and target present trials, $t(78) = 17.33$, $p < 0.0001$. The difference in \overline{e} was not significant, however, $t(78) = 1.41$, $p = 0.163$. For target absent trials, $\overline{rt} = 2202$ and $\overline{e} = 0.007$. For target present trials, $\overline{rt} = 914$ and $\overline{e} = 0.065$.

For target present trials, neither \overline{rt} nor \overline{e} varied significantly over target motion direction d_t , $F(9, 60) = 0.176$, $p = 0.996$ and $F(9, 60) = 0.107$, $p > 0.999$, respectively. Both \overline{rt} and \overline{e} did vary significantly over motion difference ΔD , however, $F(8, 61) = 112.9$, $p < 0.0001$ and $F(8, 61) = 95.9$, $p < 0.0001$, respectively.

4.3 Interpretation

Results showed that viewers had longer \overline{rt} for target absent trials, compared to target present trials. This type of finding is not uncommon, and represents a viewer's attempt to confirm his or her answer. Viewers can respond as soon as they detect a target during target present trials. For target absent trials, they will often wait for a short period of time to confirm no target is present before responding. The lack of a significant difference in \overline{e} across trial type shows that viewers did not have more difficulty answering correctly during target absent trials.

There was no significant difference in either \overline{rt} or \overline{e} for different target motion directions d_t . Motion in certain directions (e.g., cardinal directions like 0° or 90°) is no easier to detect. \overline{rt} and \overline{e} did vary significantly based on motion difference $\Delta D = |d_t - d_b|$, however. Figs. 2c and 2d show that \overline{rt} is high for $\Delta D \leq 20^\circ$, and that \overline{e} is high for $\Delta D = 10^\circ$. Both values drop significantly past these points. This finding is supported by previous research which shows that at least 15° of rotation is needed to distinguish differences in glyph orientation in a static image [16]. Our results suggest a similar difference in motion direction is needed to accurately distinguish between target and background elements.

5 VELOCITY EXPERIMENT

The velocity experiment studied a viewer's ability to identify the presence or absence of a small 3×3 group of target elements that move at a speed v_t different from the background elements v_b . The direction of all elements d was identical for a given trial. Trials were otherwise identical to the previous two experiments. Observer accuracy and response time were used to measure viewer performance. The experiment was constructed to test three separate conditions:

1. *Target velocity*: the target elements' velocity v_t . This allowed us to determine if a minimum target velocity was needed to generate perceptually distinguishable motion.
2. *Velocity difference*: the absolute difference between the target and background elements' velocities $\Delta V = |v_t - v_b|$. This allowed us to determine how much of a difference in subtended visual angle is needed to generate a perceived difference in velocity.
3. *Motion direction*: the direction of the elements' motion d measured in degrees rotated counterclockwise from the horizontal axis. This allowed us to test whether different directions produced different viewer performance.

In order to simulate motion, elements "walked" within constant-size screen regions (or cells) centered about an element's starting position. Neighboring cells adjoined with one another, so there were no gaps between the cells. As in the direction experiment, we used the stencil buffer to eliminate visual artifacts at the target-background edge during target present trials.

We tested four different velocities: 10, 18, 26, and 34 pixels per second. Given an average viewing distance of 24 inches, a screen width and height of 14.5×11.5 inches, and a screen resolution of 1280×1024 pixels, these velocities corresponded to subtended visual angles of approximately 0.27° , 0.49° , 0.70° , and 0.92° . Previous research showed that velocities of 1° or greater are easily perceived [8]. We therefore chose values below this ceiling to search for limitations on a viewer's ability to perceive velocity differences.

Every possible combination of velocities (v_t, v_b) was tested, producing 12 cases where $v_t \neq v_b$ (target present trials) and 4 cases where $v_t = v_b$ (target absent trials). Each (v_t, v_b) pair was shown eight different times, with the elements moving in eight different directions d ranging from 0° to 315° counterclockwise about the

horizontal axis in 45° increments. This produced a total of 96 target present trial types, and 32 target absent trial types.

Velocity differences $\Delta V = |v_t - v_b|$ were either 8, 16, or 24 pixels per second for the target present trials. Trials with a given ΔV were shown 96 times, divided equally across each trial type. Each target absent trial type was shown 9 times. This produced 288 target present trials and 288 target absent trials, for a total of 576 trials during the experiment.

5.1 Procedure

Ten members of the Computer Science department (faculty, staff, and graduate students) volunteered to participate during the experiment. Experiment details were similar to the flicker experiment, except that viewers were instructed to search each display for a small group of target elements moving at a velocity different from the background elements.

Prior to starting the experiment, viewers completed 24 practice trials. 12 trials contained a target group (with each of the 12 possible (v_t, v_b) pairs included once), and 12 trials did not. Viewers were allowed to repeat the practice session as often as they needed to become comfortable with the procedure and speed of the experiment. Viewers then completed 576 experiment trials presented in a random order during a single one hour session. Present or absent response for each trial, together with the time taken to respond, were saved for later analysis.

5.2 Results

Two performance metrics were recorded during the experiment: mean viewer response time \overline{rt} and mean viewer error rate \overline{e} . Trials were divided along the conditions of target present or absent, target velocity, velocity difference, and motion direction, with \overline{rt} and \overline{e} for all viewers collapsed and averaged over the resulting categories. t -tests and ANOVAs with a 95% confidence interval were used to identify statistically significant variations in performance. In summary, we found:

1. \overline{rt} was significantly longer for target absent trials, compared to target present trials.
2. \overline{e} was significantly higher for target present trials, compared to target absent trials.
3. Neither \overline{rt} nor \overline{e} varied significantly across target velocity v_t or target motion direction d .
4. Both \overline{rt} and \overline{e} varied significant across motion difference ΔV .

\overline{rt} varied significantly between target absent and target present trials, $t(110) = 14.920$, $p < 0.0001$. The difference in \overline{e} was smaller but also statistically significant, $t(110) = 3.335$, $p = 0.0012$. For target absent trials, $\overline{rt} = 1339$ and $\overline{e} = 0.005$. For target present trials, $\overline{rt} = 901$ and $\overline{e} = 0.043$.

For target present trials, neither \overline{rt} nor \overline{e} varied significantly over either target velocity v_t or target motion direction d_t . For target velocity, $F(7, 72) = 0.3205$, $p = 0.9424$ and $F(7, 72) = 0.6803$, $p = 0.6881$ for \overline{rt} and \overline{e} , respectively. For target motion direction, $F(7, 72) = 1.2632$, $p = 0.2931$ and $F(7, 72) = 1.3971$, $p = 0.2502$. Both \overline{rt} and \overline{e} did vary significantly over velocity difference ΔV , however, $F(2, 77) = 370.68$, $p < 0.0001$ and $F(2, 77) = 52.167$, $p < 0.0001$, respectively.

5.3 Interpretation

As in the direction experiments, \overline{rt} was longer for target absent trials than for target present trials, suggesting that viewers again waited some period of time to confirm no target was present before responding. Interestingly, \overline{e} was slightly higher for target present

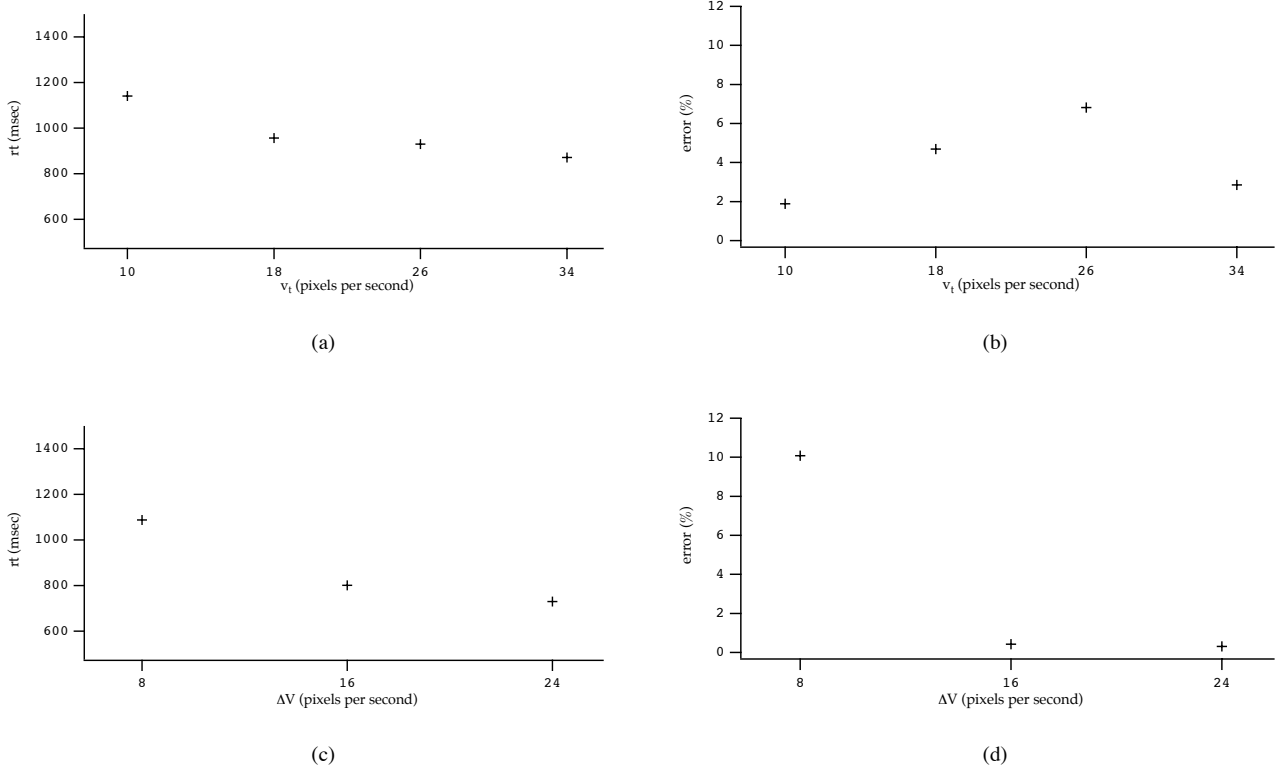


Figure 3: Average viewer performance during the velocity experiments: (a) \overline{rt} versus v_t ; (b) \overline{e} versus v_t ; (c) \overline{rt} versus ΔV during target present trials; (d) \overline{e} versus ΔV during target present trials

trials ($\overline{e} = 4.3\%$) compared to target absent trials ($\overline{e} = 0.5\%$). This error occurred almost entirely during trials with a velocity difference $\Delta V = 8$ pixels per second. Viewers mistakenly responded that no target was present during some of these trials.

There was no significant difference in either \overline{rt} or \overline{e} for different target velocities v_t or target motion directions d . This suggests that even a 10 pixel per second velocity (0.27° of subtended viewing angle) can be rapidly detected as visible motion. As well, motion in certain directions was no harder or easier to detect.

\overline{rt} and \overline{e} did vary significantly based on velocity difference $\Delta V = |v_t - v_b|$. Figs. 3c and 3d show that \overline{rt} and \overline{e} are high for $\Delta V = 8$ pixels per second. Both values fall quickly past this point. Our results suggest that a velocity difference of 0.22° of subtended visual angle is not sufficient to differentiate target and background elements. At 0.43° (16 pixels per second), however, viewers could rapidly and accurately locate the target group (\overline{rt} of 1087 msec versus 801 msec and \overline{e} of 10% versus 0.4% for 8 and 16 pixels per second, respectively).

6 PRACTICAL APPLICATIONS

We were interested in applying our experimental results to real data. One collaboration in our laboratory involves astrophysics researchers who are investigating various aspects of supernovas. Part of this effort includes simulating how a supernova collapses. We were provided with a number of 2D slices through a 4D volume (x, y, z and t) generated by their simulation. We want to visualize this data in a way that is perceptually salient, and that also minimizes the amount of on-screen information used to represent flow directions and velocities. Our current technique carefully seeds a small number of dots throughout the flow field. The dot

positions are selected using an algorithm proposed by psychologists that models how visual the system perceives local orientation around each dot [11]. We start with a sparse regular grid of dots over a target flow field, then iteratively optimize the dot positions to converge to a solution that produces a perceived local orientation that matches the flow direction at each dot's position in the flow field. Level-of-detail hierarchies are built to display both a high-level overview and increasing amounts of detail within each field.

In order to reinforce flow direction, and to visualize the corresponding flow velocities, we animate each dot. We are experimenting with different data representations. In one example, continuous flow directions are binned into 18 categories: $0^\circ, 20^\circ, \dots, 340^\circ$, and flow velocities into six categories: 0, 10, 18, 26, 34, and 42 pixels per second. This guarantees perceptually noticeable differences between each category. We are also visualizing the continuous flow directions and velocities without discretizing. This allows viewers to see the exact direction and speed of the flow, and to distinguish between locations with differences at or above the thresholds of 20° and 0.43° of subtended visual angle. We move each dot along a line through its position, with a direction and velocity based on the underlying flow field values. The length of the line is fixed to keep it within the perceived local neighborhood of each dot. Dots fade in and out at the endpoints of their lines to avoid visual discontinuities (i.e., similar to the use of the stencil buffer during our experiments). The result is a smooth, continuous animation that displays flow direction and velocity throughout the flow field. Fig. 4 shows an example of visualizing data around two vortices and a shear boundary, with dot patterns from a twenty-frame interval composited into a single image to highlight the paths that the dots traverse. The astrophysics researchers are using our technique together with existing tools (e.g., volume visualization, LIC, and animated movie sequences) to study different aspects of their




Figure 4: An example of a Stevens flow visualization, with four static snapshots over a one second interval, and the composite of 20 timesteps over the same one second interval (center image); longer streamlines (e.g., at the top of the composite image) show regions of higher velocity

simulation data. They have expressed an appreciation for the hierarchical abilities of our technique, which allows them to rapidly identify high-level patterns, then conduct more in-depth analyses of the underlying flow details.

We have also applied our motion results to data without an inherent motion context, for example, to visualize temperature and pressure gradients and areas of high precipitation in a meteorological dataset. Although the use of flicker, direction, and velocity work as expected, there were interactions between the properties (e.g., some flicker frequencies interfered with the detection of direction or velocity patterns). The same type of visual interference exists between texture and colour properties [5]. Further experiments are needed to classify these effects in a way that will allow us to minimize or avoid them when designing multidimensional visualizations.

7 CONCLUSIONS

This paper describes an experimental study of the perceptual properties flicker, direction of motion, and velocity of motion. Results suggest that a coherent flicker difference of at least 120 msec, a direction difference of at least 20° , and a velocity difference at least 0.43° of subtended visual angle are needed to distinguish between different values of the three properties.

There are a number of additional issues we plan to address. We are designing follow-on experiments to investigate visual interference, first between individual motion properties, and next between motion, colour, and texture. This knowledge is needed to apply motion more effectively during visualization. Higher order motion properties like acceleration and non-linear motion paths may also hold the ability to encode certain types of data.

REFERENCES

- [1] L. Bartram, C. Ware, and T. Calvert. Moticons: Detection, distraction, and task. *International Journal of Human-Computer Studies*, 58(5):515–545, 2003.
- [2] O. J. Braddick and I. E. Holliday. Serial search for targets defined by divergence or deformation of optic flow. *Perception*, 20:345–354, 1987.
- [3] J. L. Brown. Flicker and intermittent stimulation. In Clarence H. Graham, editor, *Vision and Visual Perception*, pages 251–320. John Wiley & Sons, Inc., New York, New York, 1965.
- [4] J. Driver, P. McLeod, and Z. Dienes. Motion coherence and conjunction search: Implications for guided search theory. *Perception & Psychophysics*, 51(1):79–85, 1992.
- [5] C. G. Healey and J. T. Enns. Large datasets at a glance: Combining textures and colors in scientific visualization. *IEEE Transactions on Visualization and Computer Graphics*, 5(2):145–167, 1999.
- [6] J. Hohnsbein and S. Mateeff. The time it takes to detect changes in speed and direction of visual motion. *Vision Research*, 38(17):2569–2573, 1998.
- [7] G. D. Kerlick. Moving iconic objects in scientific visualization. In *Proceedings Visualization '90*, pages 124–130, San Francisco, California, 1990.
- [8] S. Mateeff, G. Dimitrov, and J. Hohnsbein. Temporal thresholds and reaction time to changes in velocity of visual motion. *Vision Research*, 35(3):355–363, 1995.
- [9] G. H. Mowbray and J. W. Gebhard. Differential sensitivity of the eye to intermittent white light. *Science*, 121:137–175, 1955.
- [10] K. Nakayama and G. H. Silverman. Serial and parallel processing of visual feature conjunctions. *Nature*, 320:264–265, 1986.
- [11] K. A. Stevens. Computation of locally parallel structure. *Biological Cybernetics*, 29:19–28, 1978.
- [12] A. Triesman and S. Gormican. Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95(1):15–48, 1988.
- [13] P. D. Tynan and R. Sekuler. Motion processing in peripheral vision: Reaction time and perceived velocity. *Vision Research*, 22(1):61–68, 1982.
- [14] A. J. van Doorn and J. J. Koenderink. Temporal properties of the visual detectability of moving spatial white noise. *Experimental Brain Research*, 45:179–188, 1982.
- [15] J. J. van Wijk. Image based flow visualization. In John Hughes, editor, *SIGGRAPH 2002 Conference Proceedings*, pages 745–754, San Antonio, Texas, 2002.
- [16] C. Weigle, W. G. Emigh, G. Liu, R. M. Taylor, J. T. Enns, and C. G. Healey. Oriented texture slivers: A technique for local value estimation of multiple scalar fields. In *Proceedings Graphics Interface 2000*, pages 163–170, Montréal, Canada, 2000.