Filtering and Integrating Visual Information with Motion

Lyn Bartram
School of Computing Science
Simon Fraser University
lyn@cs.sfu.ca

Colin Ware
Data Visualization Lab
University of New Hampshire
colinw@cisunix.unh.edu

Tom Calvert
Technical University of British Columbia
calvert@techbc.ca

Abstract

Visualizing information in user interfaces to complex, large-scale systems is difficult due to visual fragmentation caused by an enormous amount of inter-related data distributed across multiple views. New display dimensions are required to help the user perceptually integrate and filter such spatially distributed and heterogeneous information. Motion holds promise in this regard as a perceptually efficient display dimension. It has long been known to have a strong grouping effect which suggest that it has potential for filtering and brushing techniques. However, there is little known about which properties of motion are most effective. We review the prior literature relating to the use of motion for display and discuss the requirements for how motion can be usefully applied to these problems, especially for visualizations incorporating multiple groups of data objects. We compared three shapes of motions in pairwise combinations: linear, circular and expansion/contraction. Combinations of linear directions were also compared to evaluate how great angular separation needs to be to enforce perceptual distinction. Our results showed that shape differentiation is more effective than directional differences (except for 90°). Of the three shapes studied, circular demands the most attention. Angular separation must be 90° to be equally effective. These results suggest that motion can be usefully applied to both filtering and brushing. They also provide the beginnings of a vocabulary of simple motions that can be applied to information visualization.

Introduction

Improvements in computing power, display resolution and graphics and simulation algorithms have nurtured the development of increasingly powerful data visualization applications on standard desktop machines. These systems can represent and update large volumes of data at interactive rates, allowing the user to have multiple views of the same information and to dynamically adjust parameters and evaluate results. With this power and flexibility has come corresponding cognitive complexity. Information and scientific visualization systems are increasingly characterised by complex displays of multivariate data and multiple windows or views with subsets or alternate representations of the same data.

Increasingly, users must spend extra time and cognitive effort to dynamically configure displays so that only the information, and all the information, they require is easily visually available. The appropriate interactive techniques depend on the information extraction task: *filtering* or *linking*. These are often complementary. Filtering is the process of taking away (ignoring) irrelevant data objects and attributes ("nothing else matters"). Linking involves a search to connect (associate) disparate visual elements ("we are now related").

Multivariate data coding combines variable dimensions mapped to graphical attributes such as colour, position, shape, and size, producing dense representations in which patterns and emergent properties reveal characteristics of the data set not evident in lower-dimensional views. Multiple window approaches provide detail in smaller views while maintaining context in the larger display: the user can explore aspects of the data without losing the overall "picture". Multiple views can also be used to simultaneously show different attributes of the same data object: for example users of spreadsheets commonly present data both as a table and as a graph or chart.

However, these approaches carry their own cognitive costs. Multivariate data representations are often visually impenetrable because they are so complex, and because the various graphical codes may perceptually interfere with each other (Healey 1999). Multiple views, on the other hand, allow the user to extract interesting aspects of the data for inspection, but suffer greatly from visual fragmentation. Extra visual elements must be introduced into the display to link the spatially disparate points to avoid a cognitively effortful visual search. These in turn introduce more visual complexity.

Filtering techniques such as *dynamic queries*, introduced by (Ahlberg et al. 1992), limit visible data to subsets of interest. Sliders control the data detail in the display, allowing the user to quickly and dynamically remove and add attributes and data objects from the display. While this is a much more flexible approach than the standard "Show/Hide" options in many systems, it does have a drawback. One has to "unquery" to return to the previous data view. This is a disadvantage if one was trying only to get a quick picture of how certain data in a view are distributed, since the context of the larger view is lost. In some cases a less interventionist method would be preferable, emphasising the data of interest while visually relegating the context to the attentional background.

The common visualization technique of *brushing* (Becker and Cleveland 1987) highlights subsets of the data interactively, most usually with colour. This supports visual linking of heterogeneous objects and thus addresses the visual fragmentation problem. Unlike filtering approaches which remove data from the display, brushing brings forward additional information by superimposing a refined data view (the linked elements) onto an existing structure. However, brushing requires its own coding dimension: that is, the graphical attribute used to highlight the salient points must be reserved, be it colour, shape or texture. These graphic attributes are typically already used in a multivariate visualization. Colour in particular is excellent for categorization, and less effective at showing other attributes of relations. Moreover, as evidenced by our earlier experiments with detection (Bartram et al. 2001), the eye is insensitive to hue changes on the periphery, suggesting that in larger display configurations elements in the "brush set" may not be noticed.

What is needed are perceptually efficient techniques for filtering and brushing which avoid the above drawbacks. Healey has shown that multivariate data visualization benefits greatly from preattentive visual processing (Healey et al. 1995). Perceptually efficient codes and representations can harness the sophisticated low-level properties of the human visual information processing system; conversely, when they are poorly applied, they can cause visual interference and result in displays which are dense, impenetrable and effortful to comprehend. Previous research suggests motion has great potential as a filtering and brushing technique. It is perceptually efficient across the entire visual field, and it has a strong "grouping" property: things which move together in a similar fashion elicit the perception that they are a group. There are, however, many unanswered questions about what it means to move "in a similar way". In Wickens' terms, we are concerned with the perceptual proximity of motion attributes (Wickens and Carswell 1995). What exactly are the determinants of similarity? This is important for two reasons:

1. we need to know how much leeway we have in coding motion to reliably elicit this perception of a group when we want to cause the effect; and

2. inversely, we want to be careful in environments with multiple unrelated moving objects that similarity is not inadvertently evoked resulting in an erroneous judgment.

The second question has particular relevance to the design of modern graphical interfaces and Web pages in which animated objects are increasingly common. It is hardly the intention of the designers to visually associate unrelated animated elements, especially in the case of advertising where the companies are entirely distinct. Nonetheless, careless and haphazard combination of moving icons on these displays and pages can cause false perceptual grouping, in fact negating the original goal of emphasising one particular element.

This paper describes an experiment which investigated the effectiveness of different simple motions in aiding a visual search task. Filtering and brushing, and the problems they address, are related to the fundamental problem of visual search. A brief description of visual search and relevant research results from perceptual psychology follows to motivate the experimental context.

Visual Search

Visual search involves two types of perceptual processing: early, low-level preattentive search, which processes basic aspects of the scene in parallel, and the subsequent higher-level, serial process of *focused attention* which assembles the simple features into meaningful objects. Preattentive search is aided or impeded by the kinds of perceptual features present in the overall display (the visual field). Features are defined as a limited set of basic visual properties which are processed pre-attentively without the need for focused attention (Treisman 1985, Healey 1999). Treisman stated that these features "pop out" in isolation and are processed in parallel. We do not have to serially examine each object to see if it is red in a field of coloured objects, for example; rather, all the red ones are seen together (Treisman 1985). Perceptual features are either simple elements extracted from the actual visual objects or emergent features arising from a combination of simple features. A triangle, for example, has three simple oriented line features and an emergent closure feature. Hue, form, orientation and motion are examples of perceptual features. Feature asymmetry is common: it occurs when a feature A may be found preattentively within distractors of feature B, but the inverse is not true. Research has also shown that some features are more compelling, or *visually important*, than others (Callaghan 1990). Visually important features will interfere with lower priority features. For example, Healey found that in some circumstances hue can overwhelm shape in a boundary detection task (Healey 1999). There is evidence that motion can serve as a kind of feature mask and help the viewer ignore the non-salient (non-target) information (Driver et al. 1992). This suggests that it may be useful as a visual filtering tool, especially when other display dimensions are fully saturated.

In order to study the value of motion as a support for visual filtering, it is necessary to have an idea of the important parameters. For practical reasons, motions that are anchored and oscillate about a point are necessary since screen location must be preserved. For such anchored motions we can consider varying the relative frequency, shape and amplitude of oscillation. Amplitude of motion is a design decision; for reasons of economy the minimal amplitude that can achieve the desired result is to be preferred. We appear to be relatively insensitive to different motion frequencies (Ware and Limoges 1994) which leaves us with motion shape. Simple motions, such as linear oscillation about a point, are processed pre-attentively (Nakayama and Silverman 1986). However some more complex motions are not pre-attentively distinct. For example, (Braddick and Holliday 1987) investigated compound motions (divergence and deformation) comprising several simple motion features. Divergence and deformation were composed of four motion patterns: (a) compression, (b) expansion, (c) horizontal-vertical deformation, and (d) vertical-horizontal deformation. Divergence trials had the target expand and the distractors compress or the target compress and the distractors expand. Deformation had the rectangles variably changing from short and fat to tall and thin. None of these complex compound motions were pre-attentively distinct which implies that for the visual filtering task simple motions are to be preferred, possibly with a fixed frequency.

Conjunctive search

The pop-out effect does not generally occur over a conjoin of features. A search is conjunctive if the target has no unique feature: for example, a red triangle in a field of green triangles and red squares is made up of the conjunction of the red feature and the triangle feature. Numerous studies have shown that in almost all cases we cannot do conjunctive search preattentively: that is, it requires focused attention to look for a conjunction of perceptual features. Once the preattentive system has processed the features, focused attention takes over to serially assemble, identify and categorise them into objects. Thus the preattentive system in effect is a guide for the higher-order process of focused attention (Hunn 2000). In serial search, the time increases with the number of items to be searched.

There are, however, some studies which show preattentive search effects in conjunctive search can occurs under certain special circumstances. Of particular interest is the finding of McLeod, Driver and Crisp who reported that the detection of a moving X among static X's and moving O's can be done in parallel (McLeod et al. 1988). They hypothesised that the visual system is tuned to simultaneously process motion direction and geometric shape orientation (an X is differently oriented than an O). Subsequently Driver reported preattentive results in a conjunctive search task with motion orientation and geometric shape (Driver et al. 1992). These researchers were interested in effects of *guided search.* Is performance improved when targets are perceptually emphasised or excited (*excitatory guidance*) or non-targets are inhibited (*inhibitory guidance*)? This study investigated the conjoin of motion translation and shape over multiple movements and found a strong effect of motion coherence. Subjects were asked to find an X moving vertically in a field of vertically moving Os and horizontally moving Xs. Coherent motion was shown to be a “guiding” feature. When both sets moved coherently (that is, all vertical objects oscillated in phase and all horizontal objects oscillated in phase), results were preattentive as both excitation and inhibition were active. Search results were serial when one group moved incoherently; inhibitory guidance (non-targets coherent) was slightly better than excitatory (targets coherent). Search time increased most dramatically, however, when both groups moved out of phase. These results suggest motion coherence along a direction can perceptually group elements but that out of phase motions weaken the effect. On the other hand, the Braddick results show that linear motions 100% out of phase (opposite linear directions) are preattentively grouped. Perhaps phase is a "weak" perceptual feature in the discrimination of separate motions, where direction is a strong one.

Differentiation and Similarity

More recent research by Treisman and others has blurred the distinction between parallel and serial processing. Treisman proposes these as two ends of a range in which the degree of difference between the target and distractors determines the search time (Treisman and Gormican 1991). Studies by Quinlan and Humphreys (Quinlan and Humphreys 1987) showed that search time depended on two factors: how many distinct features were involved, and how much the target differed from the distractors. Duncan and Humphreys (Duncan and Humphreys 1989) looked more closely at the effects of feature difference. They found that ease of search is influenced by two similarity dimensions: how much the target differs from the distractors (*T-N similarity*); and how much the distractors differ from each other (*N-N similarity*). As long as the target is sufficiently unlike the non-targets (T-N similarity is low), search is efficient. If the non-targets are homogeneous, and thus perceptually grouped (N-N similarity is high), then T-N similarity is less important: search is still efficient until T-N similarity increases to a point where target and non-targets are indistinguishable from each other. Performance decreases as T-N similarity increases and N-N similarity decreases. Non-target differences are most detrimental when T-N similarity is high. Conjunctive search represents a special problem since target and non-targets share at least one feature (relevant attribute).

Motivation: Guidelines for Designing Motion-Assisted Search

What do these results imply for information visualization? The dense displays common to this domain are already crowded with multiple shape, colour and texture codes complicating efficient visual search in the following ways. First, multivariate coding means search is usually conjunctive. Filtering and linking tasks in this environment can be viewed as first steps, selecting a group of related items by a relevant attribute to be subsequently searched for a secondary attribute. Even if the feature coding the relevant attribute is preattentively accessible, there may be masking and interference effects from other codes. Second, the non-targets are varied and heterogeneous: in Duncan's terms (Duncan 1989), finding a target in this visual environment is rendered difficult by N-N dissimilarity (there are many different types of non-targets). And third, of specific relevance to the use of motion, there are likely to be other moving objects in the display, such as those used for peripheral awareness and signals, or for time-based representations (as in animated trend graphs of changing values over time). These conditions suggest three essential criteria as a starting point for motion-based filtering and brushing.

1. The feature used to code the relevant (grouping) attribute should pop-out as strongly as possible, should not be masked by other codes and motions in the display and should not in turn hamper subsequent scanning of the selected subset.

2. High N-N differences require a marked T-N difference, so the grouping feature should preattentively and substantially separate the targets from the non-targets.

3. Given the presence of other motions in the display, target and non-target motions must be clearly distinguishable. Target motions should be coherent, especially in the case where other (non-target) motions are also moving in phase.

Our goal is to establish a set of empirically validated and practical guidelines for motion-assisted search in dense visual displays. "Practical" refers to our concern that the experiments be ecologically valid: i.e., that the test environment reflects "real-world" conditions of feature density. Motion direction has been shown to be a strong perceptual feature, thus causing the objects sharing that motion to pop out of the display in parallel. The fact that motion can be conjunctive with other features is an asset for most applications. The results from (Driver et al. 1992) indicate that motion can serve as a feature mask and help the viewer ignore the non-salient (non-target) information. This property lends itself strongly to both filtering and brushing; if it is proven robust in more applied environments, it can be used to address both issues simultaneously. If motion can mask visual elements that are not of interest, then it can support filtering without explicitly removing those elements from the display. Similarly, if it can cause the elements of interest to pop out in parallel, it can elicit the grouping and linking effect which is the goal of brushing.

Aside from orthogonal motion, however, little work has been done to guide us in motion differentiation. The Braddick experiments with conjunctive motions emphasise the need for a clearer understanding of how motion features and complex motions interact and affect discrimination (the task of distinguishing separate motions). Motion techniques for brushing and filtering will only be useful in densely coded visualisation representations if they are easily distinguishable from other motions in the display. "Easily" may not mean at the parallel end of Treisman's feature differentiation spectrum (Treisman and Gormican 1991), but in a practical sense it should mean with reliable accuracy within a short period of time (a few seconds). In other words, when there are multiple moving objects in a display, they should be designed to efficiently elicit the perception of a group when they are related and to enforce the perception of separateness when they are distinct.

The experiment described in this paper was based on a conjunctive search task, since as we have discussed motion may have special benefits in this area. Given two groups of moving icons (the *target group* and the *distractor group*) in a larger field of static icons, the subject's task was to determine whether an icon in the target group had a particular geometric shape. Target and distractor groups had different motions: both motion cues lasted only a short time (2 seconds). While at least half of the distractor icons were of the desired *target shape*, only one of the target icons would be, and that only approximately half the time. Moreover, there were always icons of that shape in the third, static, group.

This experimental approach maps well to a real-world task of searching members of a group for some secondary attribute and addresses two questions:

1. is the grouping effect robust: i.e., is motion grouping a visually important feature which pops out in visual search? and, if so,

2. when do two different motions "blur" together (i.e., when is the target group indistinguishable from the distractor group)?.

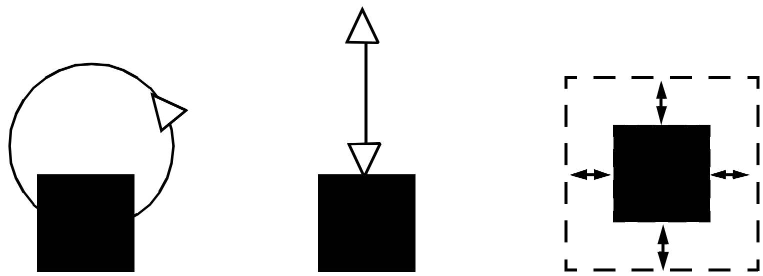

Figure 1. Simple motion shapes used in the discrimination experiment

Motion Properties for Investigation

There are many questions to address in determining effective motion cues for filtering and brushing. The most immediate relate to type, duration and interference. Type refers to which motions (which features) best elicit the popout effect. Duration is concerned with how long the motion cue has to persist to “set” the effect. Interference occurs when one cue either inhibits or enhances visual search for another cue. We were interested if motion overwhelms other features such as form (geometric shape), and how different motions may interact with each other.

Small anchored motions were used, as even small motions are highly detectable (Bartram et al. 2001) and show direction along a small distance (Driver et al. 1992). Both target and distractor groups moved coherently. Thus two perceptual steps were involved: identification of the target motion and search for the target shape. Given the short duration of the motion cues and the abundance of distractor shapes, correct response would depend on whether the motion feature “popped out” and caused the members of the target group to be identified in parallel.

Since motion shape had proven to be a dominant attribute in our previous studies (Bartram et al. 2001), we considered it as a primary candidate for effective discrimination. Three periodic motion shapes were selected for the experiments (Figure 1):

- LINEAR motions traveled smoothly back and forth along a straight path 20 pixels in length. Much of the perceptual research on motion has concentrated on simple linear motions, and they are common in representations of dynamic information, such as trend graphs or oscillating bar graphs.
- CIRCULAR motions traced a circular path of 20 pixels in diameter and were chosen because although the circle is a common simple shape there are few instances of circular motion either in visualization displays or in perceptual studies.
- Expansion/Contraction (henceforth EXPAND) motions had the icon smoothly growing and shrinking between 100% and 200% of its original size centred around its origin. Such motions are increasingly common in displays which incorporate animated elements for attentional pull (e.g. advertising). This shape was similar to the divergence cues reported in (Braddick and Holliday 1987), and combined expansion and compression motions.

In addition, as motion direction has been proven to be a preattentive feature(Nakayama and Silverman 1986, Driver et al. 1992), we needed to investigate if directional differences affect whether motions are seen as a group or as distinct, and if so, at what degree of difference discrimination occurs. Direction was therefore considered as a separate attribute (for the LINEAR and CIRCULAR paths).

Hypotheses

We had five hypotheses.

H1: Motion shape, rather than direction, is the most effective discrimination attribute. Although none of the perceptual studies concentrated on comparisons of path and shape differentiation, the importance placed upon trajectory and path in the expressive and communicative disciplines suggest that shape is deeply important at a fundamental level.

H2: The greater the angular difference, the more discriminable the motions and the more accurate the results. From the Driver study (Driver et al. 1992), however, we anticipated that out-of-phase motions (180 degrees difference and the opposing circular motions) would in fact have less accurate target detection.

H3: In the case of angular differences less than 90 degrees, quadrant orientation will have an effect: when the motions sit on different sides of one of the two major axes errors will be less. This condition, which we term *quadrant change*, is illustrated in Figure 2. The hypothesis was based on an informal pilot study in which subjects reported using vertical and horizontal directions as "guides".

H4: Some motion shapes will command more attention than others. We were also interested in possible asymmetries. A particular motion might be readily discriminated (a positive quality) but also interfere with the detection of other motions (a negative quality).

H5:. We expect a correlation between response time and the number of elements in the moving groups because they will perceptually pop out. Sequential search increases linearly with the number of items to be searched. Thus if the search set is established by motion, search time on it should be affected by element count in the set of moving targets.

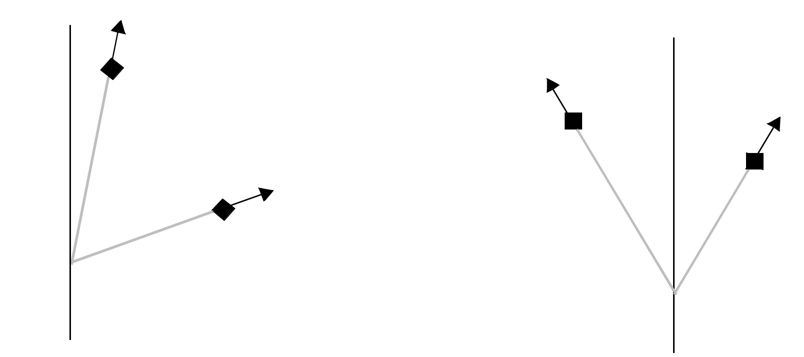

(a) No quadrant change (b) quadrant change

Figure 2. Quadrant change occurs when the angular orientations are in different absolute quadrants

Method

The experiment used a fixation (*target*) icon and a separate trial set of icons in which the search took place. The fixation icon was one of three shapes (circle, X or square) and moved according to a particular motion pattern (shape and angular orientation). The trial set consisted of a table of 35 icons of which some moved according to the *target motion* (the same as the fixation icon), some moved according to another *distractor motion* pattern, and the remaining were static. The subject's task was to decide whether the target motion group contained an icon of the *target shape*. So, for example, if the fixation was a circle moving in a linear vertical motion, the subject's task would be to decide if the icons moving in the same motion pattern contained a circle.

Both the target and distractor groups had an equal number of elements, which was randomly set each trial between 4 and 8. Approximately half of the distractor icons were of the target shape and a random number between 5 and 9 of the static icons were also of the target shape. The target group had an icon of the target shape approximately 50% of the time, randomly assigned.

A two-stage design was used with a fixation phase and a trial phase. In the fixation phase the subject was presented with a single white fixation icon in the middle of a grey window. After 2 seconds, the target icon was removed and the trial set presented. Cue duration was 2 seconds at a frequency of 20 frames/sec. (roughly 1.6 HZ). The motion cue started immediately after the trial phase was displayed. The subject answered YES or NO by hitting assigned keys. These keys were set at the beginning of the experiment according to subject preference. Note that response time was not limited: once the cue finished, the screen remained until the subject made a YES/NO response.

Figure 3 shows one example grouping. In this hypothetical case, the target icon was indicated by an X moving in a circular pattern in the fixation phase. The target group (the circled icons) move in the same circular pattern. The distractors (outlined by rectangles) move in a vertical linear pattern. Note that there are 2 Xs in the distractor group and 4 among the still icons.

Screen Layouts

The trial window was 900x675 pixels in size and was centred in a 1280x1024 screen. An icon in the trial set could be one of the following geometric shapes: circle, square, X, text label ("Hi!"), star, double ellipse, triangle, arrow or flag. Icon colour could be red (RGB 255,0,0), green (0,255,0), blue(0,0,255), cyan (0,255,255) or yellow (255,0,255). Figure illustrates these shapes. None of the trial icons was white; only the fixation icon. This avoided colour interference. Shape and colour were randomly set each trial. Each icon was bounded by a rectangle of 20x20 pixels in an "envelope" or cell of 120x120 pixels, laid out in a regular grid. A standard 21" monitor was used.

Figure 3 shows the screen layout.

Pairwise Motion Combinations

The three motion shapes provided six pairwise cues (the first name is the target motion, the second the distractor): CIRCULAR-LINEAR, CIRCULAR-EXPAND, LINEAR-CIRCULAR, LINEAR-EXPAND, EXPAND-CIRCULAR and EXPAND-LINEAR.

In addition, we considered direction in the CIRCULAR and LINEAR cases. There were two circular directions: clockwise (CW) and counter-clockwise (CCW). We defined direction in the linear

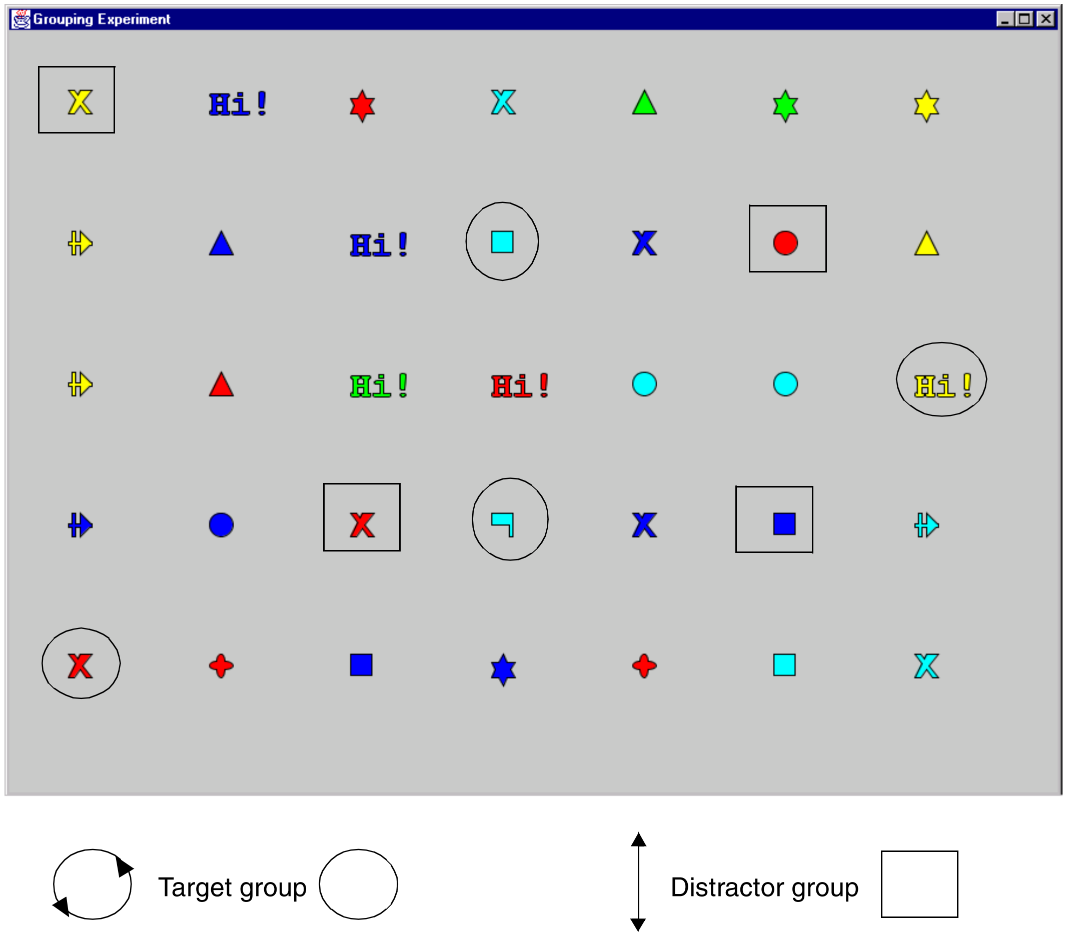

Figure 3. Sample target and distractor groups when the target icon is an X moving in a circular pattern. Distractors oscillate vertically.

case as *angular orientation.* Since we are interested in degree of difference, we constructed 9 LINEAR-LINEAR combinations based on how many degrees of angular orientation separated target and distractor: 2, 4, 8, 16, 32, 45, 90, 135, or 180. These directional combinations provided the remaining 11 trial pairwise combinations: CW-CCW, CCW-CW, LINEAR-2, LINEAR-4, LINEAR-8, LINEAR-16, LINEAR-32, LINEAR-45, LINEAR-90, LINEAR-135 and LINEAR-180. Linear motions in the different-shape conditions and linear target motions had their angular orientation randomly assigned from the set {15, 30, 60, 90, 105, 120, 135, 150, 165, 180}. Distractor linear cues in the direction conditions simply added the angular difference specified in the condition to the angular orientation of the target motion.

Experiment Design

Each condition was a pairwise combination of motions. A block contained a sequence of trials of a single condition. For the first two subjects in each experiment, a block contained 5 trials; later subjects had 7 trials in a block. There were 17 conditions and each subject had two blocks of each condition, resulting in 34 blocks. Block ordering was random. Thus the first 2 subjects in each experiment saw each cue 10 times; later subjects saw each cue 14 times.

Subjects were required to press a YES or NO key as soon as they decided whether the target shape was present or absent. The response was logged as either correct or incorrect; errors were further logged as False Positive (answering "yes" if the target was absent) or False Negative (answering "no" if the target was present). The time from cue onset to response was also recorded.

Before the experiment subjects were trained with a special block of 34 trials, equalling 2 repetitions of each motion pair. Training was not limited: subjects could practise until they felt comfortable with the task. Twenty-four students from Psychology, Computer Science and Kinesiology at SFU served as subjects: 12 for each experiment. There were an equal number of males and females. None were colour-blind. Subjects were paid to participate.

Results

We were interested in both accuracy and response time.

A multivariate repeated measures analysis of variance revealed that motion type had a significant effect on accuracy (error rates). Figure 4 shows the mean error rates for all pairs of motions. The pairwise combinations differed significantly $F(16,176) = 21.5286$, $p<0.0001$. Generally, the different-shape combinations had lower error rates than the directional differences, with the exception of 90° separation.

The number of elements in the moving groups also significantly affected error rates. However, the size of effect from element count was much less substantial than that of motion type. Subject was insignificant in all except both circular direction cues (CW-CCW and CCW-CW).

Our shape hypothesis (H1)) is validated: shape differences are clearly more effective than

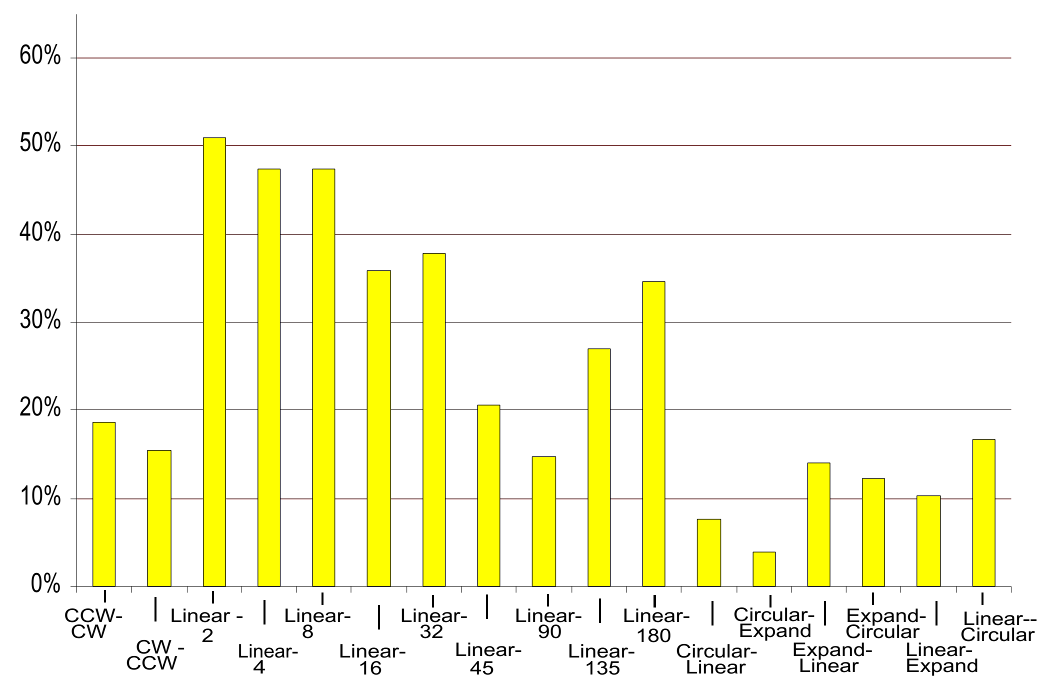

Figure 4. Mean error rate for all conditions. 50% is chance performance

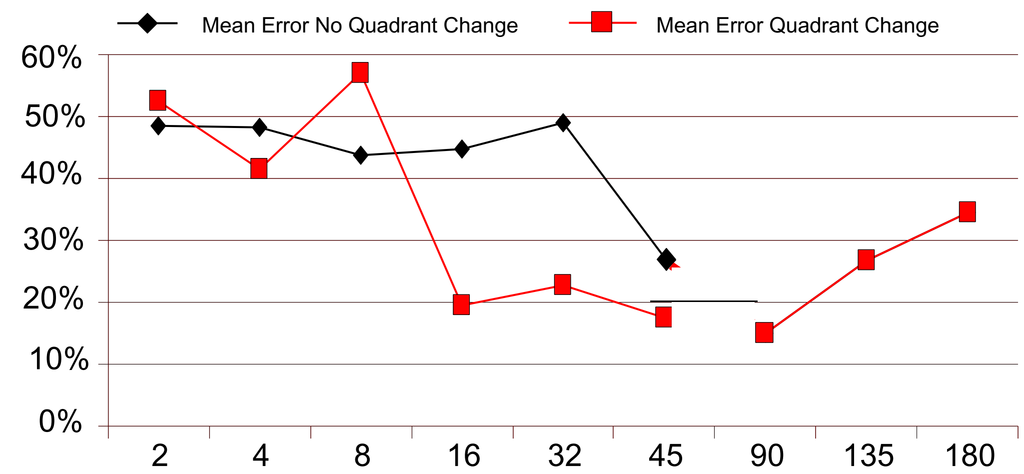

Figure 5. Error rate by angular separation including quadrant change effects

the directional distances, with the notable exception of LINEAR-90. Even the least well detected, LINEAR-CIRCULAR, has an error rate of <17%.

Figure 5 graphs error rate by angular separation. Contrary to our expectation (H2), angular separation did not correlate closely with correct target detection: instead, the most effective angular distances were 90° and 45° (LINEAR-90 and LINEAR-45 respectively). However, surprisingly, the error rates of the distances larger than 90 climb substantially.

Accuracy was much higher for linear motions in separate quadrants for angular separations less than 90° ($F(1,12) = 6.876$, $p<0.0088$) with the major effect occurring in the greater angular separations (LINEAR-16, LINEAR-32 and LINEAR-45). This confirms our hypothesis (H3) that absolute orientation plays a significant role in the discrimination of such motions.

Our hypothesis that the error rates of the 6 shape combinations would be asymmetric was verified. The data show main effects on errors for both target (F(2,24 =6.1991, $p<0.002$) and distractor motion shapes ($F(2,23)=4.4338$, $p=0.012$). The circular target motion (CIRCULAR-LINEAR, CIRCULAR-EXPAND) had less errors than the EXPAND (EXPAND-LINEAR, EXPAND-CIRCULAR) and linear (LINEAR-CIRCULAR, LINEAR-EXPAND) targets. This result points to the nature of the task for which we defined error: that is, to identify and search the target group, rather than to explicitly discriminate members of both target and distractor groups. Figure 6 shows the mean error rates for target and distractor motions. It clearly indicates the relative visual importance of the various motions. The CIRCULAR motion is the most dominant while the EXPAND is the least.

As we expected, response time increased linearly with the number of icons per group at approximately 80 ms./icon and this effect was also significant: $F(3,36)=6.278$, $p<0.0003$. These data suggest that sequential search occurred on the moving group rather than on icon shape: subjects preattentively perceived the motion feature and then searched that subset for the geometric shape. There were no cue-element count interactions.

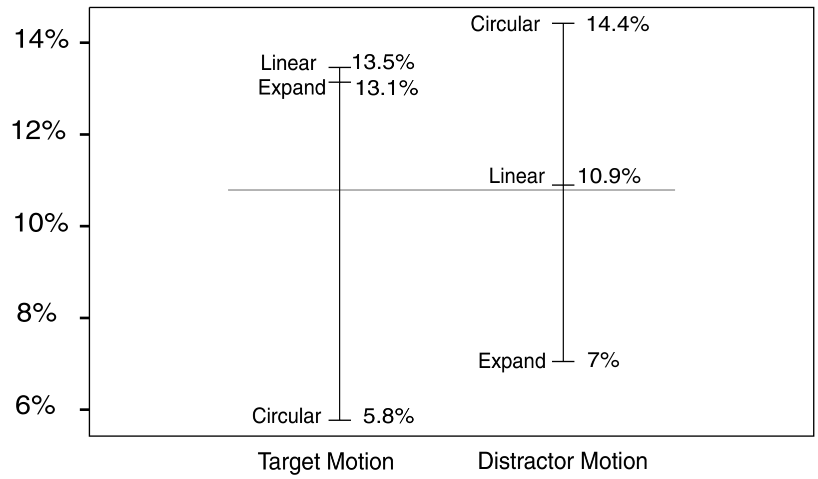

Figure 6. Error rates of target and distractor motion shapes

Discussion

These results show that small, brief and graphically simple motions are sufficient to effectively "pop out" objects in a crowded display which are dissimilar in all except their motion patterns. While there may be debate on whether these are preattentive effects in the pure sense, there is no doubt that they are perceptually efficient in the pragmatic sense. Durations of only 2 seconds were sufficient to visually fix the groups such that they could be searched for a secondary attribute, and the perception of the group persisted after the motion had stopped. In all cases, subjects unanimously reported that once the motion started the static icons effectively "fell out"[1] from the elements to be searched, automatically restricting the search to the moving groups.

Conversely, unrelated moving objects with close timing and similar paths will be erroneously visually associated. It is easy to see that trajectories which have similar shapes and are not directionally distinct will elicit the sense that the elements belong to the same group.

One interesting aspect of such motion-induced grouping techniques is that they combine to support both filtering and linking integration tasks. By relegating the static objects to the attentional background, they satisfy the goal of filtering, which is to remove spurious and therefore distracting information from the view. By eliciting the perception that the scattered objects are part of the same group, they solve the visual fragmentation induced by disparities in position and appearance. Because the motion is superimposed on the existing display, there is no reconfiguration required to restore the display to its previous state before the filtering or linking step. This gracefully supports highly interactive display reconfiguration with a minimum number of steps: users do not have to "unquery" or reset the display to remove residual effects of extracting the pertinent information. It is also useful when a persistent display of the derived information is not required, as in the case when the user wants to search the subset for some secondary attribute rather than retain the subset for further exploratory analysis.

1. This was a direct quote from 2 subjects; all subjects, however, reported a similar effect.

Our results show that motion coding effectively supports visual integration and provide some initial guidelines for practical application in filtering and brushing

G1: **Similar motions are effective in perceptually grouping dissimilar, spatially scattered objects**. Motion can be effectively used to group disparate visual elements for filtering and brushing. In the appropriate conditions, motion clearly pops out and establishes a perceptual grouping between otherwise dissimilar visual objects. This can be reliably elicited even with small and reasonably short motion cues.

G2: **Efficient motion grouping requires coherence between moving elements**. Coherence means common frequency and phase. From the basic psychological research, motion coherence is clearly a basic requirement to perceptually group separate items. The experiment software ensured coherence by redrawing each object in the same frame or step. In more complex applications with distributed redrawing control (such as loosely-integrated program nodules or distinct window operations), some centralised mechanism for coordinating timing will be needed. The error tolerance in coherence remains to be measured: how tightly does frame-to-frame coherence need to be synchronised? Motion timing, on the other hand, does not appear to need excessive control. In our experimental environment we specified a relatively slow motion speed (1 cycle over 20 frames at a frame rate of 30 frames/sec., or approximately 1.6 HZ). System logging showed a moderately variable redraw rate of between 25 and 30 frames a second. Even with this variation in refresh rates the effects were robust, suggesting that redraw rates and standard graphics controls within the operating envelopes of desktop systems are sufficient to support motion-based search techniques.

G3: **Motion shape is an excellent discriminating feature.** In environments with few other motions any of the three shape cues is a reasonable technique. The strong performance of the circular motion shape over the linear and zoom patterns makes it an especially good candidate for filtering. Environmentally, linear motions are common in other animation uses, such as trend graphs and oscillating bars, and expanding motions are also frequently used to attract attention. The dominance of the circular pattern and its infrequency in other representations suggest it as a good brushing cue.

G4: **Direction (angular orientation) can be used to differentiate linear motions under appropriate conditions.** Our results suggest that linear motions are most easily discriminated when they are in different quadrants and separated by at least 16°. Motions separated by 180° are, however, not easily discriminated.

Despite the Driver (Driver et al. 1992) and Braddick (Braddick and Holliday 1987) findings that linear translation in opposite directions was preattentively perceived, our results showed discriminating opposite movements had some problems. In both the CIRCULAR-CIRCULAR conditions and the LINEAR-180 condition, icons moved in opposing directions, 180° out-of-phase. Accuracy was quite good in the circular motion conditions, although subjects complained that these were the most difficult and unpleasant cues to discriminate. However, accuracy in the LINEAR-180 case was worse than in the different-shape conditions or in the larger linear-direction separations (including quadrant change). A possible reason for the discrepancy between our results and those of (Driver et al. 1992) is that they studied much simpler displays and used only two variables.

Clearly phase plays a strong part in motion grouping, since incoherence vastly complicates search (Driver et al. 1992). Other applied studies have reported the effectiveness of mapping data values directly onto phase (Ware and Limoges 1994, Gronquist et al. 1996). Out-of-phase data points were immediately obvious within the moving set without being seen as unrelated. Our results indicate that mapping different groups to opposing directions may be ambiguous and potentially lead to misassociation. Instead, motion phase may be more useful as a secondary attribute to motion shape for coding anomalies in group associations, such as a value that is out of range in a particular subset of interest.

These results are preliminary and much work remains to be done. Nonetheless, they suggest that motion offers a rich graphical vocabulary for visual association and discrimination., and that motion-based techniques for information integration are both viable and promising.

References

AHLBERG, C., WILLIAMSON, C., AND SCHNEIDERMAN, B. 1992. Dynamic queries for information exploration. In *Proceedings of CHI' 92*. ACM, New York, 619–626.

BARTRAM, L., WARE, C., AND CALVERT, T. 2001. Moving icons: Detection and distraction. In *Proceedings of INTERACT 2001*. to appear.

BECKER, R. A. AND CLEVELAND, W. 1987. Brushing scatterplots. *Technometrics 29,* 2, 127–142.

BRADDICK, O. AND HOLLIDAY, I. 1987. Serial search for targets defined by divergence or defomation of optic flow. *Perception 20*, 345–354.

CALLAGHAN, T. 1990. Interference and dominance in texture segregation. In *Visual Search*, D. Brogan, Ed. Taylor & Francis, New York, 81–87.

DRIVER, J., MACLEOD, P., AND DIENES, Z. 1992. Motion coherence and conjunction search: Implications for guided search theory. *Perception & Psychophysics 51,* 1, 79–85.

DUNCAN, J. 1989. Boundary conditions on parallel search in human vision. *Perception 18*, 457–469.

DUNCAN, J. AND HUMPHREYS, G. 1989. Visual search and stimulus similarity. *Psychological Review 96*, 433–458.

GRONQUIST, J., STEHARES, W., ALVARADO, F., AND LASSETER, R. 1996. Animated vectors for visualization of power system phenomena. *IEEE Transactions on Power Systems 11,* 1 (February), 267–273.

HEALEY, C. 1999. *Preattentive Processing: Using the Low-Level Human Visual System*. ACM SIGGRAPH.

HEALEY, C., BOOTH, K., AND ENNS, J. 1995. Visualizing real-time multivariate data using preattentive processing. *ACM Transactions on Modeling and Computer Simulation 5,* 3, 190–221.

HUNN, K. 2000. Preattentive visual information processing. http://home.eunet.no/ khunn/papers/ 2030.html. Research proposal.

MCLEOD, P., DRIVER, J., AND CRISP, J. 1988. Visual search for a conjunction of movement and form is parallel. *Nature 332*, 154–155.

NAKAYAMA, K. AND SILVERMAN, G. 1986. Serial and parallel processing of visual feature conjunctions. *Nature 320*, 264–265.

QUINLAN, P. AND HUMPHREYS, G. 1987. Visual search for targets defined by combinations of color, shape and size: An examination of task constraints on feature and conjunction searches. *Perception & Psychophysics 41,* 5, 455–472.

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

TREISMAN, A. AND GORMICAN, S. 1991. Search, similarity and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception and Performance 17,* 3, 652–676.

WARE, C. AND LIMOGES, S. 1994. Perceiving data displayed through oscillatory motion. Tech. Rep. TR94-089, Faculty of Computer Science, University of New Brunswick.

WICKENS, C. AND CARSWELL, C. 1995. The proximity compatibility principle: its psychological foundation and relevance to display design. *Human Factors 37*, 473–494.