Multiple Independent Highlighting Techniques

Colin Ware[1] and Nicholas J. Pioch[2]

1. ***Center for Coastal and Ocean Mapping, University of New Hampshire, Durham, NH, USA***
2. ***Systems & Technology Research LLC, Suite 6500, 600 W. Cummings Park, Woburn, MA 01801***

Abstract

Interactive highlighting is a common component of many techniques used in visualization such as brushing and dynamic queries. Sometimes these may be used in combination necessitating that two different highlighting methods be simultaneously applied. The challenge of effective highlighting is to design methods that make a subset of the items on a display stand out clearly without overly interfering with other information on a display. This is especially difficult when more than one subset of displayed symbols must be simultaneously highlighted. Three experiments are reported that investigate four different highlighting methods: 3D vs 2D symbols, encirclement, oscillatory motion and blinking. These are applied to the nodes in node-link diagrams. The first experiment was designed to evaluate the highlighting methods used separately and the results showed all four techniques to be effective. The second experiment evaluated combinations of highlighting methods. E.g. can we easily find a node that is both moving AND 3D in a set of nodes some of which are 3D and some of which are moving. The results showed that combinations including motion were the most effective. The third experiment was designed to determine which highlighting methods, used both separately and in combination supported the rapid counting of small numbers of targets. Again, combinations using motion were the most effective.

Introduction

Interactive visualization frequently requires that symbolic representations of data entities be highlighted. For example, in the technique of brushing [5], selecting a set of data entities in one view, e.g. a scatter plot, causes those same entities to be highlighted in both that view and in another view, e.g. a map. Brushing is one of the oldest and most useful interactive visualization techniques. There are other technique requiring interactive highlighting; for example, the techniques of dynamic queries [1,20] and degree of relevance highlighting [29,30] where items related to a selected item are highlighted in addition to the item itself.

Sometimes it is necessary that several highlighting methods be used in the same interactive visualization and designing multiple highlighting methods that can work well together is challenging for a number of reasons. The fundamental purpose of highlighting is to support efficient visual search for the highlighted entities and so a basic requirement is that any method should result in visually salient targets. But at the same time, highlighting should not interfere with existing coding of information. Color is already extensively used for data visualization to represent attributes of entities represented by symbols so changing the color of an entity to highlight it is often not an option. Similarly, the shape of symbols is often used to represent their type and so changing shape is likely not to be a good choice for highlighting.

The first goal of the present study was to design a set of plausible highlighting methods which would be useful for node-link diagrams, motivated by results in vision science. The second goal was to evaluate the effectiveness of the methods, both separately and in combination. We explore the use of blink highlighting, motion highlighting, surround highlighting and 3D highlighting using standard methodologies from vision research. We were particularly interested in the use of motion and blinking because of results from the vision research literature suggesting that these cues may be especially effective in conjunction search [7]. The type of visualization used in the study is node-link diagram, but we believe that the results should be widely applicable to other types of visualization, such as interactive maps and scatter plots where multi-attribute highlighting is needed.

Visual Search and Pre-attentive Patterns

There is an extensive literature on visual search and the visual properties that make a symbol easy or difficult to find among other symbols. The term pre-attentive is used to describe simple visual patterns that that can be perceived with minimal effort in complex backgrounds [10,11,12,22]. For example, a red dot in a field of black dots. Another more informal term for the pre-attentive property is "pop-out" since these patterns appear to pop out from the background of other patterns. It is worth noting at this point that the term "pre-attentive" has come to be viewed as something of a misnomer [32] since visual search is guided by a pre-configuration of the visual system for an anticipated target. This pre-configuration is a focusing of attention and so it can hardly be said that pre-attentive search comes before the deployment of attention. Nevertheless, we use the term here because of its widespread acceptance.

The basic method for these studies involved displaying a target, in a field of non-targets, called distractors [22]. A display is flashed up and the study participant is asked to hit a 'yes' button if the target is present and a 'no' button if the target is absent. The number of distractors is varied from trial to trial. When the results are plotted as shown in Figure 1 it is found that for certain kinds of target/distractor combinations the curve is flat. This suggests that there is no additional perceptual cost to processing extra distractors and the search is said to be parallel or pre-attentive. When there is a significant slope to the curve this indicates a cost for additional distractors and the slope of the curve is the processing cost per additional item, typically given in milliseconds. A large number of such studies have been run and the following is a brief summary of the major findings.

Simple features support rapid search. The earliest studies show that simple perceptual properties of simple shapes like color, orientation and size are pre-attentive [18,22] (see Figure 2a and b for examples).

Conjunctive search is generally slow. A conjunctive search is a search for a combination of features. So, for example, searching for a red Z in a field of gray Zs and red Os is a conjunctive search (Fig. 2). The conclusion from this is that early-stage feature-based visual processing is responsible for pre-attentive search [22,24].

There is a continuum of speed of search, from 0 milliseconds per item, or even negative slopes in some cases, up to 50 or more

milliseconds per item in the case of conjunctive search. So preattentivity is not an all or none effect as originally thought [32]. Phenomenologically this is also true, some items clearly popout more than others and some require extremely careful scanning.
As a rule of thumb 10 ms per item has been taken as a criterion for rapid pre-attentive visual search [25].

There are search asymmetries. For example, adding a visual feature leads to fast search, whereas removing a visual feature may not [23,24].

Motion can be pre-attentive [17,19]. Moving or oscillating dots pop out from a field of dots that are not moving, or are moving in a different direction.

There are exceptions to the slow conjunctive search finding. A pre-attentive conjunction of motion and shape has been reported [7]. A pre-attentive conjunction of color and 3D has been reported [8].

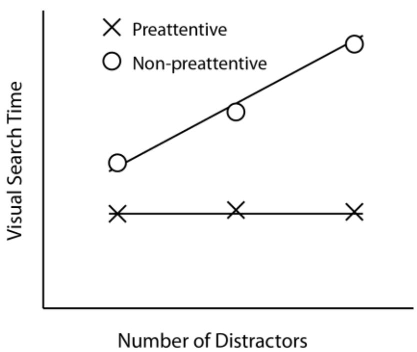

Fig 1. Idealized results for a pre-attentively distinct symbol and a non-pre-attentively distinct symbol. For the non-preattentively distinct symbols the visual search time increases with the number of non targets (called distractors).

The above is by no means a complete list, there are hundreds of studies relating to these phenomena and the reader should consult a review paper such as Huang, and Pashler [13] or Wolfe [32] for more information. Also, some aspects of the above are disputed, for example, although it is indisputable that some kinds of search are fast, and others slow, the idea that this is tied to basic feature processing in early vision is less clear. For example, simple curvature is pre-attentive, but there is no evidence for curvature detectors in early vision.

The concept of pre-attentive processing has direct relevance to design in data visualization [30] because a goal of a well designed visualization is to have certain features be clearly distinguished. A series of studies by Healey et al. has shown the 3D glyph can convey a number of variables simultaneously [10,11,12]. They also showed that such glyphs allow for area estimation based on very short (100 msec) exposures and argued on this basis that they had pre-attentive properties.

A number of papers have argued that motion is an effective cue for highlighting. Motion can be an efficient user interrupting device and the frequency and shape of motion could be used to indicate urgency as well as other attributes [3,27]. Motion cues can support visual grouping of symbols on data displays and for this reason be used as a highlighting method in interactive brushing and filtering tasks [3,4,28,29]. However, none of these studies used the methods of psychophysics to determine if motion was efficient according the methods established by Treisman and others [23,24].

There are two kinds of experimental methods that have been used in determining which shapes and patterns are pre-attentive. The first is the Treisman method described above, involving a variation in the number of distractors and measuring the response time. The second method is to use brief exposure (e.g. 200 ms) to a display containing a target and a set of distractors followed by a mask. In this case the measure is not time to respond, but the number of errors made. The point being that if the display is not processed rapidly the target will not be seen. The mask is a pattern designed to erase any residual activity in the visual system; it is known that an image can persist neurally for a short period what is called iconic store [21]. The short exposure method is based on the fact that participants require several hundred milliseconds to make multiple fixations on a display. This method has more commonly been used by researchers in the field of data visualization [10,14] than the Treisman method. We used the response time method for the first two of the experiments reported here, and the rapid exposure for the third.

A series of studies having particular relevance to the present work is that of Kosara and collaborators [15,16] They developed a method called semantic depth of field, which involved blurring non target nodes. Using the rapid exposures method they found that conjunctions between non-blurred targets and target orientation could be pre-attentively processed. This, for the first time, showed an efficient conjunction search in the context of a visualization task.

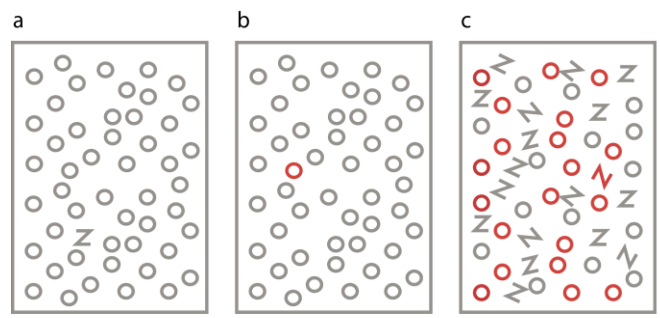

Fig. 2. (a) The Z pops out. (b) The red circle pops out. (c) The red Z does not pop out.

Highlighting Methods

We chose to investigate four highlighting methods: 3D vs 2D, adding a circle surround, motion highlighting where the symbol moved with a circular path, and blink highlighting. We did not include color or shape because for many complex visual analytics designs such as social or abstract semantic graphs [31] these are already being used extensively to represent attributes of symbols. We studied the application of these methods to highlighting nodes in node-link diagrams, one of the most common kind of visualization used in analytics. Some additional considerations are as follows.
1) 3D vs 2D: Three dimensional shaded shapes have been shown to be pre-attentively distinct from 2D shapes [9] although [32] suggests this may be a weak effect. A 3D rendered shape can have the same outline as a 2D shape, thus it can preserve that attribute.
2) Circle highlighting: Additions are asymmetrically pre-attentive [23,24]. An encircled node will stand out more clearly in a set of non-circled nodes than a non-circled node in a set of circled nodes.
3) Motion highlighting: Motion is especially interesting because it may support conjunction searches and therefore dual highlighting. One of the earliest examples of pre-attentive conjunction was Driver et al.'s [7] finding that moving Xs stood out in a field of distractors consisting of static Xs and moving Os.

4) Blink highlighting: Blink highlighting has been shown to be an effective way of directing attention [4,26]. It might also be the case that blink highlighting, like motion would allow for conjunction search since both involve dynamic changes in the image.

Non-highlighted nodes were given fuzzy borders by the simple method of creating a radial transparency profile. This has the advantage of creating a strong visual separation between the 3D highlighted nodes and non-highlighted nodes. It is similar to the semantic depth of field method developed by Kosara et al.

Design Rationales: Experiments 1, 2 and 3

All three experiments used node-link diagrams as stimulus patterns in order to find out if theoretical results obtained with simplified stimuli in the psychology lab would generalize to this common type of visualization. To add to the visual complexity of the displays, nodes were randomly colored, although the colors were irrelevant to the task.

The first experiment was designed to obtain some baseline data on the four highlighting methods and also to determine if they are reasonably independent; in other words, do the highlighting methods interfere with one another? To evaluate independence subsets of the distractors used various combinations of the non-target highlighting methods. For example, if the target was a circled node, a randomly determined subset of the nodes would be moving and an overlapping subset would be rendered in 3D. The subset could include the target node itself.

The second experiment was designed to look at the specific question of whether the different methods used *in combination* would allow for pre-attentive conjunction search; for example, is it possible to do a rapid search for 3D AND motion. In this case subsets of the nodes would be moving and a different subset would be 3D. On half the trials there would be exactly one target that was both 3D and moving and on half there would be none.

For experiments 1 and 2 we used the method developed by Treisman, varying the number of non-targets and measuring reaction time. The task consisted of the presentation of a search target, following which the display was briefly presented. Study participants respond by pressing either a 'yes' or a 'no' button depending on whether they see the target. On different trials the number of non-targets (called distractors) was varied. The time to respond was plotted against the number of distractors and the slope measured.

The third experiment was designed to look at whether the different highlighting methods support rapid quantity estimation (subitizing) both separately and in combination? It used a different method where the patterns were briefly exposed followed by a mask.

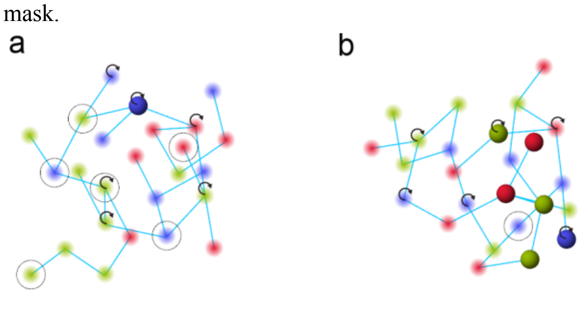

Fig 3. Two of the conditions for Experiment 1. Circular arrows represent moving nodes. (a) The target is the 3D sphere node, distractors include irrelevant circled nodes and moving nodes. (b) The target is the circled node. Distractors include 3D nodes and moving nodes.

Experiment 1: Effectiveness and Independence

The first experiment was designed to determine the extent to which each of the highlighting methods is pre-attentive with a complex set of distractors. Prior research from the perception research literature suggests that they are, but these studies were carried out with simplified displays where the distractors were made up only of of simple shapes or colors. We were interested in complex displays that allowed for multiple simultaneous highlighting methods to be simultaneously used with a visually varied set of disctractors. For example, is 3D highlighting pre-attentive even when a random set of the non target items are moving, circled and randomly colored?

Study Participants

The study participants were 14 undergraduate students at the University of New Hampshire. They were paid for participating. One participant was dropped because the data file was incomplete for an unknown reason.

Conditions

There were six conditions. These were chosen to test the different highlighting methods with different combinations of distractors where the distractors also contained some of the other candidate highlighting methods besides the one being used on a particular trial.

1. Blink target: circle and 3D randomly assigned distractors
2. Moving target: circle and 3D randomly assigned distractors
3. 3D target: circle and motion distractors (3D A)
4. 3D target: circle and blink distractors (3D B)
5. Circled target: 3D and motion distractors (Circle A)
6. Circled target: 3D and blink distractors (Circle B)

Each highlighting condition was run with three different sizes of node-link diagram (13, 25 and 37 nodes). In the case of target present conditions there was always a single target and 12, 24, and 36 non-targets. Two types of task irrelevant highlighting were applied to the nodes. 25% of the nodes were given each type of irrelevant highlighting. These were assigned with independent processes. This meant that both target and distractor nodes could have task irrelevant highlighting.

Temporal Sequence

Prior to each trial, a blank screen first appeared for one second, followed by + fixation point in the center of the screen for 0.5s, followed by a blank for 0.5s. Then the node-link diagram appeared and remained on the screen until the participant responded. Simultaneously the single node illustrating the target type was shown at the left hand side of the screen as a reminder.

Participants pressed the (specially labeled) keyboard m-key (labeled 'yes') if the target was present and the v-key (labeled 'no') if the target was absent. They had their left and right forefingers pre-positioned on these keys. If they made a mistake, the computer emitted a beep. Participants were told to try to keep errors below 5%.

Trial Blocks

The participants began with a training session where the task was explained. They were given each of the conditions for four trials (twice with target present, twice with target absent). During this time the experimenter would correct them if they appeared to misunderstand the task. Following this there were two sets of blocks of trials as follows.

Trials were given in blocks of 16 for a given highlighting method/graph size condition. At the start of a block a single node of the target type was shown at the center of the screen. The conditions were given in random order. This yielded 6x3x16 = 288 trials (6 highlighting methods: 3 diagram sizes: 16 trials) in a block set. On half the trials the target was present and on half it was absent and this was also randomized. Participants could take a break between any block of trials, they used the space bar to advance to the next block. An experimental session consisted of two block sets (576 trials).

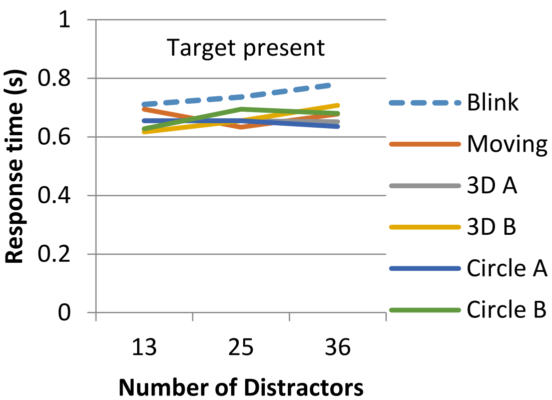

Fig 4a. Experiment 1 results summary – target present.

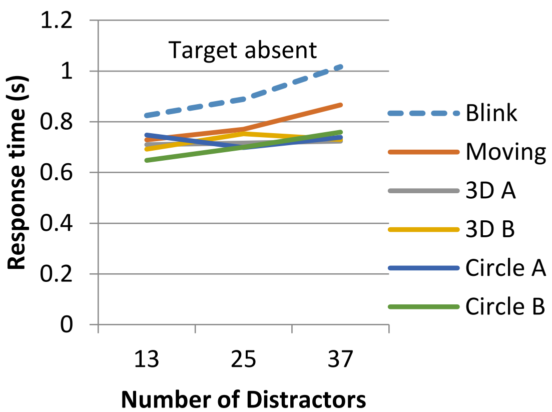

Fig 4b. Experiment 1 results summary – target absent.

The graph display

A graph was algorithmically generated for each trial using simple spring layout. On average each node was connected to 2.5 others. In order to make the different sized graphs display in roughly equivalent areas the mean spring length was varied depending on the graph size. 2.0 for the 13 node graph, 2.4 for the 25 node graph and 3.7 for the 37 node graph. Typically the graph would fit within an area approximately 12 cm in diameter. The nodes were drawn approximately 0.8 cm in diameter. The graph display window was 27cm x19cm and was viewed from approximately 56 cm (at this viewing distance 1 cm = approx. 1 degree of visual angle). The background on which the graph was drawn was white with a luminance of 81 cd/m^2.

Highlighting

For 3D highlighting the nodes were rendered as spheres using OpenGL shading with 10% ambient illumination, 70% diffuse illumination and a 100% specular illumination with an exponent set at 30. Circle highlighting was accomplished with a 1.3 pixel anti-aliased circle 50% larger than the node size. The moving nodes oscillated at 2.0 Hz following a circular path with a diameter that was 67% of the node diameter. Blinking was also sinusoidal at 2.0 Hz and it changed the node contrast relative to white background from 25% to 100%.

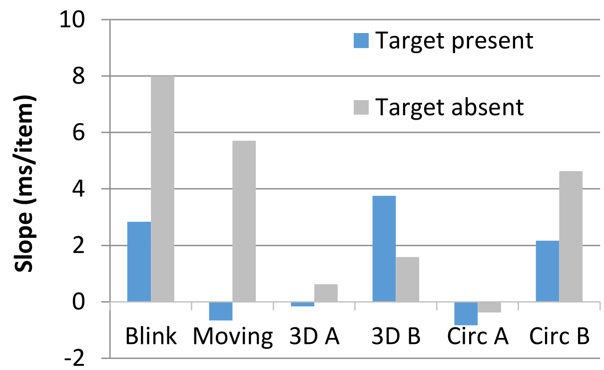

Fig 5. Experiment 1: Visual search times displayed in terms of ms/item.

Results from Experiment 1

The results are summarized in Figures 4(a&b) and 5. A three-way ANOVA was run for highlighting condition, number of distractors and whether target was present or absent. This showed significant main effects for condition ($F[5,60] = 10.745$; $p < 0.001$), number of distractors ($F[2,24] = 10.66$; $p < 0.001$) and whether the target was present or absent ($F[1,12] = 37.4$; $p < 0.001$). Overall the mean response time was 0.670 s for the target present conditions and 0.761 s for the target absent conditions. All two way and three way interactions were also significant. To examine these interactions in more detail we broke the results down into target present and target absent categories and ran separate ANOVAs for each with Tukey HSD test to separate significant effect in the conditions. For the target present conditions, blinking resulted in the slowest responses (0.742 s) from a Tukey HSD test and there were no significant differences between the other conditions (average response time = 0.666 s). For the target absent conditions, a Tukey HSD test showed that blinking again resulted in the slowest responses (0.909 s), motion was the next slowest (0.788 s), and all of the other conditions (3D and circle) were faster and not significantly different (0.718 s).

Figure 5 summarizes the search times per distractor (the slopes of the response time plots). This shows all of the highlighting methods to be pre-attentive following the commonly accepted criterion that slopes of < 10 ms per item should be taken as pre-attentive.

Discussion of Experiment 1 results

The first experiment confirmed our design decisions in the sense that all of the methods were effectively pre-attentive by the usual, if arbitrary criterion of 10 ms/item. This suggests that any of them can be used effectively in conjunction with techniques such as brushing, or degree of relevance highlighting. The blinking method was the slowest but only by < 100 ms. In addition the motion method was slightly slower than the 3D and Circle methods but only for the target absent conditions.

Experiment 2: Are there highlighting combinations supporting pre attentive conjunction search?

The second experiment was designed to see if conjunctions of the highlighting methods could result in efficient searches. For example, can we rapidly find a target that is moving AND three dimensional? This has practical applications in cases where the user is interested in entities contained in the intersection of two highlighted sets. For example, in a date set if symbols representing females are highlighted using one method (e.g. 3D), and symbols representing video game players are highlighted using another method (e.g. motion) can we rapidly see symbols representing female video game players.

Experiment 2 Method

The method was almost identical to that of the first experiment except that in this experiment the visual search target always involved a conjunction of two highlighting methods. The conditions were as follows.

1. Blink 3D target: blinking flat and 3D not blinking distractors
2. Moving 3D target: moving and 3D not moving flat distractors
3. Blink circled target: blinking flat and 3D not blinking distractors
4. Moving circle target: moving and not circled distractors
5. Circle 3D target: circled flat and 3D distractors (illustrated in Figure 6).

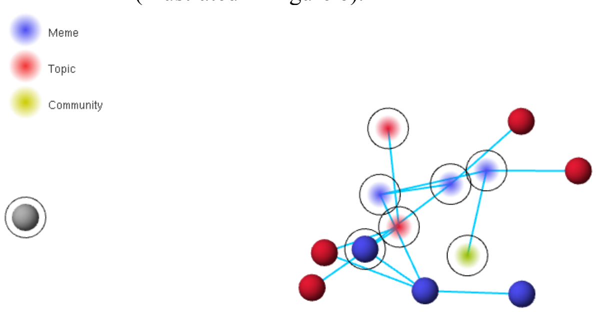

Fig. 6. Example diagram for Experiment 2. The search target is a node that is both 3D and circled.

The Graph Display

Graph layout was the same as for Experiment 1. To generate a conjunction half the nodes (excluding one) were highlighted using highlighting method A, and the other half highlighted using method B. On target present trials, the additional node was highlighted using both methods. On target absent trials it was randomly assigned one of the two highlighting methods.

Participants

There were 16 participants in the study. All were undergraduate students paid to participate. Following a preliminary analysis data two were eliminated because of high error rates (approx. 10%).

Procedure

The procedure was almost identical to that of the first experiment. The participants began with a training session where the task was explained and they were given each of the conditions for four trials. During this time the experimenter would correct them if they appeared to misunderstand the task. Following this there were two blocks of trials as follows. As in Experiment 1, there were three graph sizes with 13, 25 and 37 nodes. On half the trials the target was present and on half it was absent. The trial blocking was the same as for Experiment 1 only there were somewhat fewer trials because there was one less condition.

Results from Experiment 2

The results are summarized in Figures 7 (a&b) and 8. A three-way ANOVA was run for highlighting condition, number of distractors and whether target was present or absent. This showed significant main effects for condition ($F[4,52] = 36.69$; $p < 0.001$), number of distractors ($F[2,26] = 48.8$; $p < 0.001$) and whether the target was present or absent ($F[1,13] = 23.9$; $p < 0.001$). Overall the mean response time was 1.129 s for the target present conditions and 1.67s for the target absent conditions. All two way and three way interactions were also significant. To examine these interactions in more detail we broke the results down into target present and target absent categories and ran separate ANOVAs for each with Tukey HSD test to separate significant effect in the conditions. For the target present conditions, the moving circle and the moving 3D conditions formed fastest group (mean time = 0.94s), followed by 3D blinking (1.05s), followed by 3D circle (1.24 s) followed by blinking circle (1.37s). For the target absent conditions, a similar progression of grouped conditions was observed; the moving circle and the moving 3D conditions formed the fastest group (mean time = 1.35 s), followed by 3D blinking (1.55s),,followed by a group containing the 3D circle and the blinking circle conditions (2.04s).

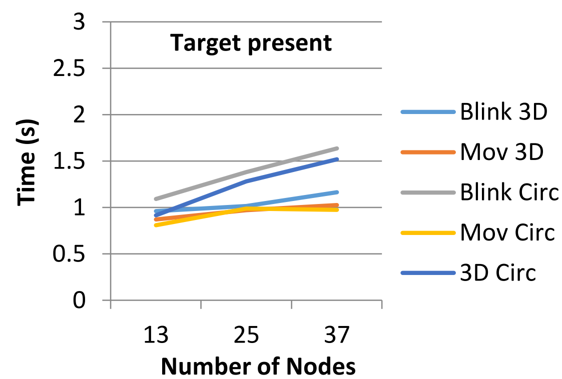

Fig 7a. Experiment 2 results summary – target present.

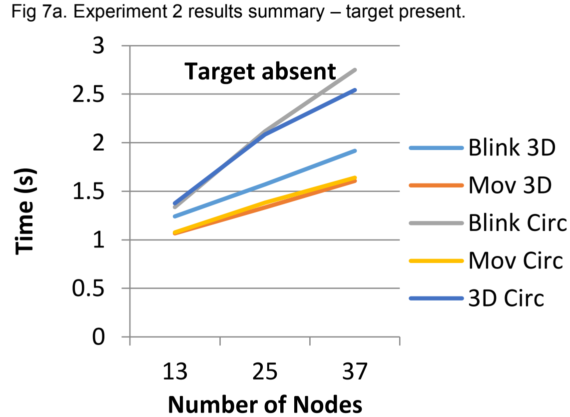

Fig 7b. Experiment 2 results summary – target absent.

Discussion of Experiment 2 Results

The results from the second experiment fall into two discrete groups. With the target present, the blinking 3D, the moving 3D and the moving circle conditions all had low processing times of < 10 ms/item meeting the criterion for pre-attentive search. The blinking circle and the 3D circle had intermediate processing times of > 20 ms/item. In the target absent conditions response times were more than twice as long. This is what is to be expected from the theory of sequential (i.e. non-preattentive) search because for the target present trials on average a target will be found when half the symbols have been processed whereas in the target absent trials detection of a non-target will require that all the symbols be processed. Another reason for slower responses to non-targets is that detecting a non-target has greater uncertainty associated with it.

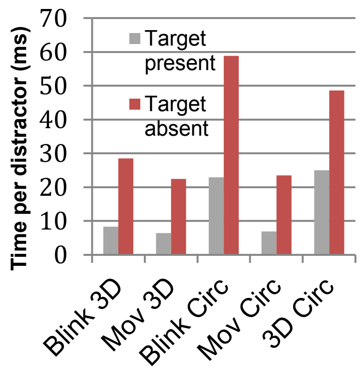

Fig 8. Search time per distractor for target present and target absent conditions (Experiment 2).

Experiment 3: Do the highlighting methods support rapid estimation of quantities

Experiment 3 was designed to evaluate the same four highlighting methods using a different task, namely the rapid estimation of quantity (called subitizing). With many visualization tasks a number of objects will be simultaneously highlighted and rapid quantity estimation is an important requirement. It is well known that humans, and some animals, can "at a glance" estimate quantities typically up to four items [6] and we designed a task based on this ability. We also decided to use an entirely different methodology: brief exposure plus masking [10,14]. This methodology is more suited to the subitizing task because the study participant must enter a number on each trial and this is not compatible with the rapid yes/no responses used in the first two experiments.

Study Participants

The study participants were 15 undergraduate students at the University of New Hampshire. They were paid for participating.

Procedure

Experiment 3 was carried out in two parts. In the first (Expt 3a), the task was to estimate the number of targets highlighted using a single highlighting method. For example, targets might be 3D or in motion. As with Experiment 1 a subset of the nodes were highlighted using one of the other highlighting methods to test for perceptual independence. For example, in the 3D highlighting, there would be between 0 and 4 nodes that are 3D and the subject's task will be to estimate how many after a very brief exposure. Similar to Experiment 1, a subset of the nodes had a second highlighting method applied that the participant had to ignore. The stimuli for experiments 3a and 3b were the same as for experiments 1 and 2, except that the thickness of the line in the circle highlighting method was increased to two pixels. This was done because in experiment 2, circle highlighting produced worse results than 3D highlighting.

In the second part (Expt 3b) the task was to estimate the number of nodes highlighted using a conjunction of two different highlighting methods. For example, approximately half the nodes might be highlighted using 3D and approximately half might be highlighted using motion. The targets would be a subset of 0-4 nodes highlighted using both motion and 3D.

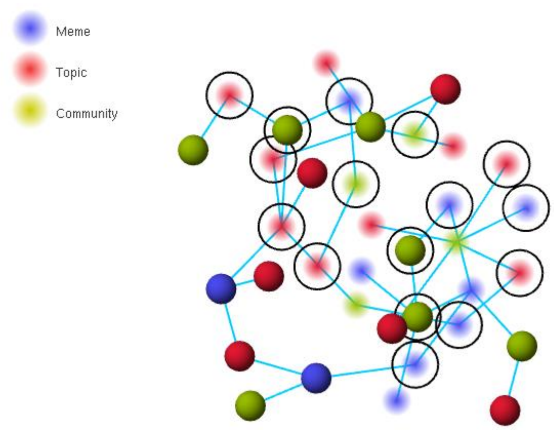

Fig 9. An example diagram from Experiment 3b. There are three nodes that are both 3D and circled.

Table 1: Experiment 3a: Single highlight	
Target highlighting	Distractors highlighting
BLINK	CIRCLE
BLINK	3D
MOTION	CIRCLE
MOTION	3D
3D	MOTION
3D	BLINK
3D	CIRCLE
CIRCLE	MOTION
CIRCLE	BLINK
CIRCLE	3D

Table 2: Experiment 3b: Conjunction highlight	
Highlight 1	Highlight 2
BLINK	CIRCLE
BLINK	3D
MOTION	CIRCLE
MOTION	3D
CIRCLE	3D

Trials

A single trial consisted of a 1.0 sec presentation of a single node representing the target. This was followed by a 1.0 second blank followed by a 1.0 cross hair, followed by a 1.0 second blank followed by the node-link diagram for 250 msec. The node-link diagram was followed by a mask consisting of a grid of squares to prevent retention of the image in iconic memory. This remained on the screen until the participant responded. The participant responded by entering a number in the range 0-4 using the keyboard.

The node-link diagram always had 36 nodes and the layout method was the same as for Experiment 1. On each trial there were between 0 and 4 targets highlighted.

In Experiment 3a, 18 of the distractors were coded with a second highlighting method. As shown in Table 1.

In Experiment 3b the size of the node set highlighted with method A was randomly set to be in the range 14 – 17. The size of the node set highlighted with method B was independently and randomly set to be in the range 14 – 17. The size of the intersection set (A&B) was set to be in the range 0 – 4.

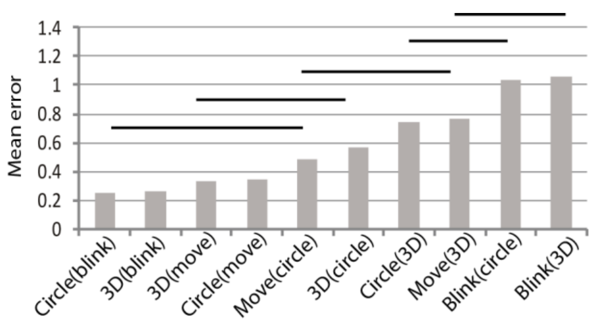

Fig 10. Experiment 3a results summary. Brackets give distractors. The horizontal bars give homogeneous groups according to a Tukey HSD test.

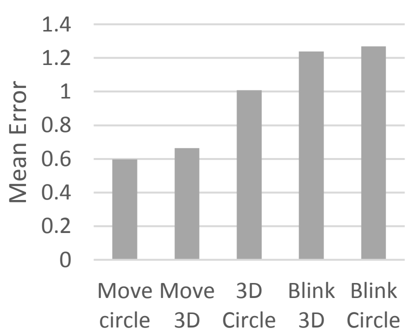

Fig 11. Experiment 3b results summary

Results

Because of the discrete nature of the errors they were averaged within subjects and conditions. A one way ANOVA run on the averaged data from Experiment 3A showed a highly significant main effect of conditions. ($F[9,140] = 20.8$ $p < 0.001$). A Tukey HSD test placed revealed groups represented by the horizontal bars in Figure 10. These show that (for example) the best five conditions were significantly better than the worst four conditions.

The results of Experiment 3b are summarized in Figure 11. There was a highly significant effect of condition on errors ($F[4,70] = 57.6$; $p < 0.001$). A Tukey HSD test placed the results in three groups. The worst group was the combinations with blinking with an average error of 1.25. The best group contained the combinations with motion with a mean error of 0.63. The 3D circle condition was intermediate (mean error = 1.01).

Discussion of Experiment 3

If a participant were simple to guess that there were 2 targets on every trial, the mean error would be 1.2. So, on the trials with the blinking targets, the participants performed approximately at chance in Experiment 3b and only a little better in Experiment 3a. The conjunction test given in Experiment 3b once again supports the finding that motion is effective in supporting conjunction search.

Conclusion

As a guiding principle for design, when two highlighting methods are needed to support different interactions it should be easy to attend to each highlighting method independently because sometimes we will want to attend only to a single subset of symbols, typically one that is newly selected, e.g. from a checklist of attributes. In addition, it should be easy to see symbols highlighted using conjunctions of the two methods.

Experiment 1 showed that all four of the highlighting methods provide a reasonable degree of perceptual independence when there was a single target. They all met the (10 ms/distractor) criteria for pre-attentive search in the target present condition, and all but the blinking method met the criterion for the target absent condition though the blinking method came close. These search speeds were accomplished despite the fact that the entire set of nodes (distractors and the target when present) had two of the alternative highlighting methods randomly applied, so that they had to be ignored. It is also noteworthy that blinking resulted in slower searches overall than the other methods, potentially due to the need to for the blinking to manifest visually. At the 2 Hz blink rate used, one half period would take 250 msec.

Experiment 3a also addressed visual search for targets defined by a single highlighting method, but with a single alternative highlighting method applied that had to be ignored. In this case blinking performed poorly, not much better than chance.

Finding combinations of two highlighting methods that support efficient searches was the primary motivation for this research and we now turn our attention to results that relate to this. Experiment 2 showed that the best combinations involved motion, although the combination of blinking and 3D also was effective. For the target present conditions all gave search speeds of < 10 ms/item. For these combinations, searches in target absent trials were slower, around 20 ms/item. However, when people interactively highlight items it may be a reasonable assumption that detecting the presence and number of targets is more important than detecting a complete absence of targets.

Experiment 3b also showed that the best combinations for dual highlighting involved motion where the task was rapid estimation of quantity. In this experiment blinking resulted in scores at the chance level. Better results came from the combination of 3D and circle, perhaps because of the increased boldness of the circle.

The results from experiments 2 and 3b suggest that using motion in combination with a static highlighting method, such as 3D or a circle surround is an effective choice. However, there was a discrepancy between the two experiments with respect to blinking. We speculate that the reason why blinking performed relatively worse in Experiment 3b than Experiment 2 was due to visual transients inherent in the method used for Experiment 3b.

Short exposures to the node-link diagrams would have produced strong onset and offset transients that may have interfered more with the blinking than with the other methods. Such transients would not normally occur under normal conditions of use and so this result may be not be relevant to actual applications.

This study by no means exhausts the possibilities for combinations of highlighting techniques. For example, stereoscopic depth can be used for highlighting node link diagrams [2] and it is quite likely that the conjunction of this with other cues will be pre-attentive.

To summarize: our main finding is that the combination of motion and other static highlighting methods, either 3D rendering or a surround circle can be very effective and meets a commonly accepted criterion for pre-attentiveness [25]. These combinations are recommended for cases where interaction requires two independent highlighting methods to be applied. Another contribution is the application of the Treisman methodology in information visualization where the number of non-targets is varied and response time is measured. This method has the advantage that it allows for measurements of the speed of perceptual processing for different kinds of symbol features, used both separately and in combination.

Acknowledgments

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. W31P4Q-14-C-0023. The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. (Approved for Public Release, Distribution Unlimited).

References

[1] Ahlberg, C., Williamson, C., & Shneiderman, B. (1992, June). Dynamic queries for information exploration: An implementation and evaluation. In Proceedings of the SIGCHI conference on Human factors in computing systems (pp. 619-626). ACM.

[2] Alper, B., Höllerer, T., Kuchera-Morin, J., & Forbes, A. (2011). Stereoscopic highlighting: 2d graph visualization on stereo displays. Visualization and Computer Graphics, IEEE Transactions on, 17(12), 2325-2333.

[3] Bartram, L., & Ware, C. (2002). Filtering and brushing with motion. Information Visualization, 1(1), 66-79.

[4] Bartram, L., Ware, C., & Calvert, T. (2003). Moticons:: detection, distraction and task. International Journal of Human-Computer Studies, 58(5), 515-545.

[5] Becker, R. A., & Cleveland, W. S. (1987). Brushing scatterplots. Technometrics, 29(2), 127-142.

[6] Beckwith, M., & Restle, F. 1966. Process of enumeration. Psychological Review, 73(5), 437.

[7] Driver, J., McLeod, P., & Dienes, Z. (1992). Motion coherence and conjunction search: Implications for guided search theory. Perception & Psychophysics, 51(1), 79-85.

[8] D'Zmura, M., Lennie, P., & Tiana, C. (1997). Color search and visual field segregation. Perception & psychophysics, 59(3), 381-388.

[9] Enns, J. T. 1990. Three-dimensional features that pop out in visual search. In Visual Search, D. Brogan, Ed. Taylor and Francis, New York, 37– 45.

[10] Healey, C. G., Booth, K. S., & Enns, J. T. (1993, May). Harnessing preattentive processes for multivariate data visualization. In Graphics Interface (pp. 107-107). CANADIAN INFORMATION PROCESSING SOCIETY

[11] Healey, C. G., Booth, K. S., & Enns, J. T. (1996). High-speed visual estimation using preattentive processing. ACM Transactions on Computer-Human Interaction (TOCHI), 3(2), 107-135.

[12] Healey, C. G. and Enns, J. T. (1999) Large datasets at a glance: Combining textures and colors in scientific visualization. IEEE Transactions on Visualization and Computer Graphics 5, 2 145–167.

[13] Huang, L. and Pashler, H. A boolean map theory of visual attention. Psychological Review 114, 3 (2007), 599–631.

[14] Kosara, R., Miksch, S., & Hauser, H. (2001). Semantic depth of field. In IEEE Symposium on Information Visualization, IEEE Computer Society. 97-97.

[15] Kosara, R., Miksch, S., & Hauser, H. (2002). Focus+ context taken literally. Computer Graphics and Applications, IEEE, 22(1), 22-29.

[16] Kosara, R., Miksch, S., Hauser, H., Schrammel, J., Giller, V., & Tscheligi, M. (2002, May). Useful properties of semantic depth of field for better f+ c visualization. In ACM International Conference Proceeding Series (Vol. 22, pp. 205-210).

[17] Nakayama, K. (1985). Biological image motion processing: a review. Vision research, 25(5), 625-660.

[18] Nothdurft, H. C. (1993). The role of features in preattentive vision: Comparison of orientation, motion and color cues. Vision research, 33(14), 1937-1958.

[19] Seymour, K., Clifford, C. W., Logothetis, N. K., & Bartels, A. (2009). The coding of color, motion, and their conjunction in the human visual cortex. Current Biology, 19(3), 177-183.

[20] Shneiderman, B. (1994). Dynamic queries for visual information seeking. Software, IEEE, 11(6), 70-77.

[21] Sperling, G. (1960). The information available in brief visual presentations. Psychological Monographs 74: 1–29

[22] Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. Cognitive psychology, 12(1), 97-136.

[23] Treisman, A., & Southern, J. (1985). Search asymmetry: a diagnostic for preattentive processing of separable features. Journal of Experimental Psychology: General, 114(3), 285.

[24] Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: evidence from search asymmetries. Psychological review, 95(1), 15.

[25] Trick, L.M. and Enns, J.T. (1997) Measuring preattentive processes: when is pop-out not enough. Visual Cognition, 4(2) 163-198.

[26] Waldner, M., Le Muzic, M.P Bernhard, M. Purgathofer, W.;Viola, I.,Attractive Flicker — Guiding Attention in Dynamic Narrative Visualizations," IEEE Transactions onVisualization and Computer Graphics, 20(12), 2456-2465,

[27] Ware, C., Bonner, J., Knight, W and Cater, R. (1992) Moving icons as a human interrupt. International Journal of Human-Computer Interaction. 1(4): 331-357.

[28] Ware, C., & Bobrow, R. (2004). Motion to support rapid interactive queries on node--link diagrams. ACM Transactions on Applied Perception (TAP), 1(1), 3-18.

[29] Ware, C., & Bobrow, R. (2005). Supporting visual queries on medium-sized node–link diagrams. Information Visualization, 4(1), 49-58.

[30] Ware, C. (2013) Information Visualization: Perception for Design. 3rd Edition. Morgan Kaufman.

[31] Ware, C., Wright, W. and Pioch, J. Thinking Design Patterns. In Distributed Multimedia Systems (DMS) Proceedings, Knowledge Systems Institute,150-155.

[32] Wolfe, J. M. (1994). Guided search 2.0 a revised model of visual search. Psychonomic bulletin & review, 1(2), 202-238.

[33] Wolfe, J.M. and Horowitz, T.S. (2004) That attributes guide the deployment of visual attention and how do they do it. Nature Neuroscience 5. 1-7.

Author Biographies

Colin Ware is a Professor of Computer Science at the University of New Hampshire. He as a PhD from the University of Toronto and an MMath from the University of Waterloo.

Nicholas J. Pioch is a Sr. Member of Technical Staff at Systems & Technology Research, in Woburn, MA. He has a B.S and M.Eng. from the Massachusetts Institute of Technology. His current research interests include visual analytics for event discovery and monitoring in social media, and mixed initiative machine learning for situational awareness and decision support in complex domains.