ELSEVIER

Available online at www.sciencedirect.com

SCIENCE DIRECT

Int. J. Human-Computer Studies 58 (2003) 515–545

International Journal of Human-Computer Studies

www.elsevier.com/locate/ijhcs

Moticons: detection, distraction and task

Lyn Bartram[a,*], Colin Ware[b], Tom Calvert[c]

[a] *Colligo Networks, Inc., 2nd Floor, 1071 Mainland Street, Vancouver BC, Canada V6B 5P9*

[b] *Data Visualization Lab, University of New Hampshire, Durham, NH 03824, USA*

[c] *School of Computing Science, Simon Fraser University, Burnaby, BC V5A 1s6, Canada*

Received 1 May 2002; accepted 1 November 2002

Abstract

In this paper, we describe an empirical investigation of the utility of several perceptual properties of motion in information-dense displays applied to notification. Notification relates to awareness and how dynamic information is communicated from the system to the user. Key to a notification technique is how easily the notification is detected and identified. Our initial studies show that icons with simple motions, termed *moticons*, are effective coding techniques for notification and in fact are often better detected and identified than colour and shape codes, especially in the periphery. A subsequent experiment compared the detection and distraction effects of different motion types in several task conditions. Our results reveal how different attributes of motion contribute to detection, identification and distraction and provide initial guidelines on how motion codes can be designed to support awareness in information-rich interfaces while minimizing unwanted side effects of distraction and irritation.

Keywords: Notification user interfaces; Motion; Animated user interfaces; Attention

1. Introduction

Users typically function in multi-task environments in which information is distributed across windows and applications and is not necessarily exclusive to the

*Corresponding author. Tel.: +1-604-685-7962; fax: +1-604-685-7969.

E-mail addresses: lbartram@colligo.com (L. Bartram), colinw@cisunix.unh.edu (C. Ware), tom@sfu.ca (T. Calvert).

1071-5819/03/$ - see front matter
doi:10.1016/S1071-5819(03)00021-1

task at hand. For example, a financial analyst may be monitoring stock market quotes while reviewing a client's portfolio and evaluating performance patterns over time. Simultaneously, she may be cued as messages to do with office administration arrive from her colleagues. Alternately, a telecommunications manager may be planning scenarios for equalizing phone traffic across variably loaded channels, while sporadic alarms indicate overloads on current routes. In both cases, these users are being made aware of dynamic information outside the specific scope of the data they are using for their current tasks. In some cases, certain types of dynamic information are contained in a dedicated display which the user must constantly monitor, such as a stock ticker or message flag. These displays are typically located on the periphery of a screen (Czerwinski et al., 2000; Maglio and Campbell, 2000; McCrickard, 2000). In other cases the changing information can be located anywhere in the visual field, such as mode information or element state directly tied to data objects in the displays (Sarter and Woods, 1995; Mitchell and Sundstrom, 1997).

Monitoring dynamic information can be a cognitively strenuous task which requires the user to examine the currently displayed information and decide whether it has changed, so it is preferable to explicitly alert the user to a change by a *signal*. Signals are graphical events which indicate to the user that something has happened in some area in the display. Replacing an "empty mailbox" icon with a "full mailbox" to show e-mail status or animating the transformation of old text into new in a peripheral display (McCrickard, 2000) are examples of signals. Signals are incorporated into peripheral awareness tools (Maglio and Campbell, 2000; McCrickard, 2000), messaging (Parsowith et al., 1998; Cutrell et al., 2000), state changes, system events or alarms (Adams et al., 1995). They can also be used as navigation markers or guides to dynamically emphasize relevant points in a display.

Current information visualization interfaces rely heavily on graphical coding devices (also termed *display dimensions*) such as shape, colour, size, texture, orientation and position (Ware, 2000). These schemes can be very effective in enabling information analysis because they are *mentally economical* (Woods, 1991; Healey et al., 1995): rapidly and efficiently processed by the preattentive visual system rather than attentive effort. However, only a small amount of information can be encoded in each visual dimension. For example, a typical recommendation is that no more than eight colours be used to define information categories (Shneiderman, 1986; Gilmore et al., 1989). For this reason there is a shortage of perceptually efficient codes than can be used in information-rich user interfaces.

One promising way of visually coding information is to use simple motion. Motion has a unique ability to attract attention over a large visual field and offers a rich graphical vocabulary. Its use has only recently become feasible due to the advent of fast graphics processors and supporting software technologies. However, compared with the use of colour coding, which is supported by a large literature of design guidelines based on decades of experimental studies, there has been little research relating to the effective design of motion codes. Such work is urgently needed because available technologies such as Javascript and image animation have led to a riot of moving and jiggling icons that compete for our attention. The notification studies described in this paper investigated the effects of moving icons, which we call *moticons*, as alerting mechanisms in situations where the user is engaged in a primary task and needs to be made aware when an event occurs outside the task area.

2. Background

Research into human perception tells us a great deal that is relevant to the use of icon motion as a coding mechanism. Of special interest is the fact that motion triggers a kind of orienting response attracting a user's attention, even when it appears in the periphery of the visual field (Faraday and Sutcliffe, 1997).

Motion compares very favourably to colour and shape if we are concerned with designing icons to attract a user's attention at the edge of a computer screen. The human visual system is very non-uniform with respect to our ability to resolve detailed information. For example, we can only resolve about one-tenth of the detail about 10° of visual angle to the side of the point of fixation (Smith and Atchison, 1997). Thus, icons that rely on detailed shape to convey their meaning are only effective if directly fixated. Our ability to discriminate colour information is also very non-uniform across the visual field. In fact, in the periphery we are almost colour-blind (Wyszecki and Stiles, 1982). One of the great potential advantages of motion as an attention-getting device is that our ability to perceive motion falls off much less towards the periphery of the visual field. Peterson and Dugas (1972) confirmed this in an applied setting that showed static targets to be virtually invisible in the far field, whereas moving targets were easily detected.

Our ability to see things at the edge of a computer screen may vary with our level of attention to the task we are performing. A "searchlight metaphor" has been used to model how attention falls off in the visual periphery as a function of the cognitive load or the stress level of an operator (Wickens, 1992). A phenomenon known as tunnel vision occurs under conditions of very high stress, but even under relatively low stress conditions the focus of attention narrows considerably (Williams, 1985). Focusing attention in a visually "noisy" field requires the user to both maintain control of where she is attending and awareness of potentially interesting areas as conditions change. To theoretically explain this, Wolfe and Gancarz (1996) suggest that the visual system has a fixation map in which regions of excitation specify the locations of potentially important visual information. These areas become inhibited after saccades are directed to them.

Woods defines several criteria for signals he terms *cognitive tools* to support control and direction of attention (Woods, 1995): *accessibility* (i.e. the user should be capable of picking them up without losing track of current activities); *partial information* (the signal should carry enough partial information for the user to pick up whether to shift attention to the signalled area) and *mental economy* (the representation should be processed without cognitive effort).

Because other information can remain on the screen while a user attends to a moving signal, it may provide the required accessibility. Because motion can be registered in the periphery, it may be ideal for conveying partial information, and because motion is pre-attentive it may have the required characteristic of mental economy. However, all of these qualities require experimental verification in task-related studies.

Motion may have the additional advantage of adding an extra coding dimension to information visualization. Bartram and Ware (2002) found that motion codes seemed independent of colour and shape codes, suggesting that motion has the additional advantage of adding extra dimensionality to the display.

The human visual system is very good not only at perceiving but also at tracking and predicting movement. We can track up to five moving objects in parallel (Pylyshyn et al., 1994) without effortful context-switching. Hillstrom and Yantis (1994) suggest that it may not be motion per se that attracts attention, but rather the appearance of a new object in the visual field. These findings suggest that introducing extensive motion into user interfaces may be problematic. When a new object gains the attention of the tracking system, another object will typically be lost. This can lead to problems occurring with distracting irrelevant items. In particular, the moving banner animations that grace many web pages may be particularly effective in distracting us.

2.1. Previous studies with moving icons

Blinking can be considered to be an elementary form of motion and much use has been made of blinking in user interfaces to attract and direct visual attention. In many systems, it is the primary visual cue for alarm conditions. However, anecdotal evidence indicates that people find blinking excessively annoying and visually ineffective when too many items are flashing (who has not cursed the WWW HTML blink function?) In large-scale systems where alarms tend to propagate rapidly, over-flashing not only reduces effective alarm information but also renders the displays visually disturbing, distracting users from effectively perceiving the needed information from other representations (Gilmore et al., 1989; Sarter and Woods, 1995; Woods, 1995). Ware investigated the use of a simple moving icon as a "human interrupt" signal in the interests of seeing if this would evoke the same direct pull of attention as blinking or flashing without causing the associated irritation (Ware et al., 1992). Participants performed a primary task and were told to respond by hitting a key when they noticed movement of one of two small icons on either side of the top of the display. The icon was a small bar which grew and shrank vertically in an smooth, oscillatory fashion. Amplitude, side and velocity of the movement were varied. There was no effect for amplitude or side, but increases in velocity led to an increase in the number of quick responses and a decrease in the number of long ones. The good average response times indicated that participants had no trouble noticing the interruption without any reported irritation factor. Even the slowest times were acceptable and were reported to be less irritating.

One way of using simple animation is in illustrating a simple procedure. Baecker and Small animated icons to identify and explain their function (Baecker and Small, 1990). The advantages of animation were particularly noticeable when the small size of icons meant a low resolution of information (i.e. intricate depiction was impossible). Ambiguity was reduced and users remembered the function of the particular icon better.

McCrickard and Stasko investigated how animated information could be used to maintain peripheral awareness (McCrickard et al., 1999). While participants browsed through on-line text, additional information would appear in a secondary window in one of three ways: the words would fade slowly in, would scroll across the window, or would suddenly appear (“blast in”). None of these cues was found sufficiently distracting to impede the primary task, but as might be expected from the theory of perceptual onset (Hillstrom and Yantis, 1994) the blast was the most effective in getting attention.

The authors have found that simple moving icons are better detected than colour and shape cues, especially in the periphery (Bartram et al., 2001). However, certain types of motion appear to be more distracting and irritating than others. This paper examines these issues in more depth.

3. Motivation

The experiments in the detection and distraction studies reported here were designed to address a number of questions that relate to the use of motion signals as a cognitive tool (Woods, 1995) for managing attention. A dual task design was used to simulate situations common in current desktop environments where the user is engaged on a primary task that takes most of her attention. We are interested in the kind of situation awareness where a change in an icon is used to signal some event, such as the arrival of a message, new users in a conference, or system events like a printer jam. We divided our studies into two areas of investigation:

1. *Determine the suitability of motion cues for attracting and directing attention*: Colour, label and form are often employed to show dynamic state changes and alerts: mailbox flags go up to indicate new messages, for example, or valve icons show red or green to indicate open or closed. Such cues are often poorly seen in peripheral conditions (Sekuler et al., 1981). We were interested in comparing motion to colour and shape as a signal for directing attention to particular points in a display. We were also interested in comparing different properties of simple motion as a notification mechanism. Were there effects on detection and indentification accuracy or response time? In Experiments 1 and 2, we examined the detection and identification performance of several simple moticons compared to traditional colour and shape cues.
2. *Examine how motion contributes to distraction*: The power of movement to attract attention can be distracting and disruptive. The recent explosion of scrolling, popping and jiggling icons competing for our attention on WWW pages is a well-known example. Yet humans are able to process simultaneous and unrelated motions seamlessly in our natural environment. Researchers investigating peripheral awareness tools disagree about the distraction aspect of animation

(Maglio and Campbell, 2000; McCrickard, 2000), but their investigations were concerned with the monitoring of animated text rather than with the distraction component of the movement itself. Our intuition was that motion-induced distraction may be affected by the type of motion and vary with task load. Some interactions between motions and task effects have been observed in uncontrolled studies (McCrickard, 2000), but more rigorous evaluation was required. In Experiment 3, we measured detection and distraction effects of different moticon types in several task conditions.

The motion attributes investigated in this research all applied to periodic motions. Linear direction and velocity (frequency) were most often reported in the perceptual research; see (Bartram, 2001) for a review. Empirical studies also showed phase to be noteworthy (Limoges and Ware, 1989). Interpretative studies and disciplines reveal that type (the outline traced by the motion), amplitude and temporal continuity ("smoothness"), all contribute to different impressions (Bartram, 2001).

4. Experiments 1 and 2: signal detection and identification

There were two experiments in the detection study, investigating how moving icons compare to both colour and shape in attracting a user's attention. A large screen display was used to address the issue of how far in the periphery motion can be effective compared to colour and shape. Large displays are becoming increasingly common as the focus for work-group activities; we believe that most of the results also apply to desktop displays, especially in multiple-monitor configurations. Such peripheral issues arise in interfaces any time users need to see dynamic information outside the foveal area. This happens frequently in several circumstances: multiple monitor configurations, large screens, or small peripheral awareness displays (see McCrickard (2000) for a review). Indeed, even standard desktop 21″ displays where information is located in tool- and menubars along the edges of the screen can result in poor information detection and discrimination due to peripheral effects (Maglio and Campbell, 2000). Finally, in any environment in which users are not guaranteed to be looking straight ahead at a display all the time peripheral vision is a factor.

Experiment 1 compared how well simple motion cues were detected compared to colour and shape cues, respectively, and how these varied with distance from the locus of attention (the primary task window). Both detection rates and detection times were measured. Experiment 2 took a more ecological approach: it combined colour, shape and motion cues and investigated how well they performed in a more complex screen in both detection and identification: i.e. how accurately the participant could identify which object in the visual field had changed.

4.1. Experiment 1: detection

This experiment had the purpose of evaluating the effectiveness of different kinds of motion, relative to colour change, or shape change, in signalling users across a wide visual field populated with icons. A wide field of view was used to investigate both near-field and far-field conditions. The cues were colour change, shape change and two linear motions (up and down) of the same frequency but different amplitudes. The first part (Experiment 1A) compared colour cues to motion cues. The second part (Experiment 1B) compared shape cues to the same motion cues. We measured the error rate and response time for detection (noticing that something had changed). Participants performed both 1A and 1B in a single session; ordering was counter-balanced. We had two main hypotheses for this experiment. First, we expected that as distance from the centre of attention increased, colour and shape detection rates would drop off (cues would be missed) and response times would increase. Second, we anticipated that motion detection rates and response times would be significantly better than those of colour and shape and would not be adversely affected by distance from the centre of attention. Additionally, we thought that motion amplitude might have an effect: smaller motions would lead to lower detection rates and longer detection times: this last expectation was based only on intuition. Most of the method was the same for 1A and 1B. We describe this common method first.

4.1.1. Method

4.1.1.1. Screen layout. A wide rear-projected screen was used. The dimensions of the projector screen were 116 cm × 87 cm. The resolution of the screen was 800 × 600 pixels which translated to approximately 145 mm/pixel. The participant was seated in front of the screen at 95 cm as illustrated in Fig. 1. The primary task area of approximately 11.6 cm was placed at either the left or the right of the stimulus window such that it was vertically centred and had a horizontal margin on each side of roughly 2.9 cm. Fig. 2 shows one screen layout with the primary task located on

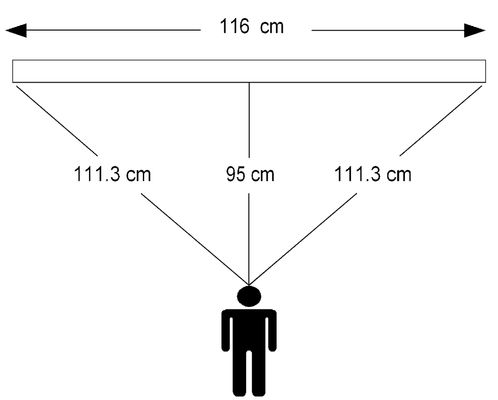

Fig. 1. Participant position and screen width.

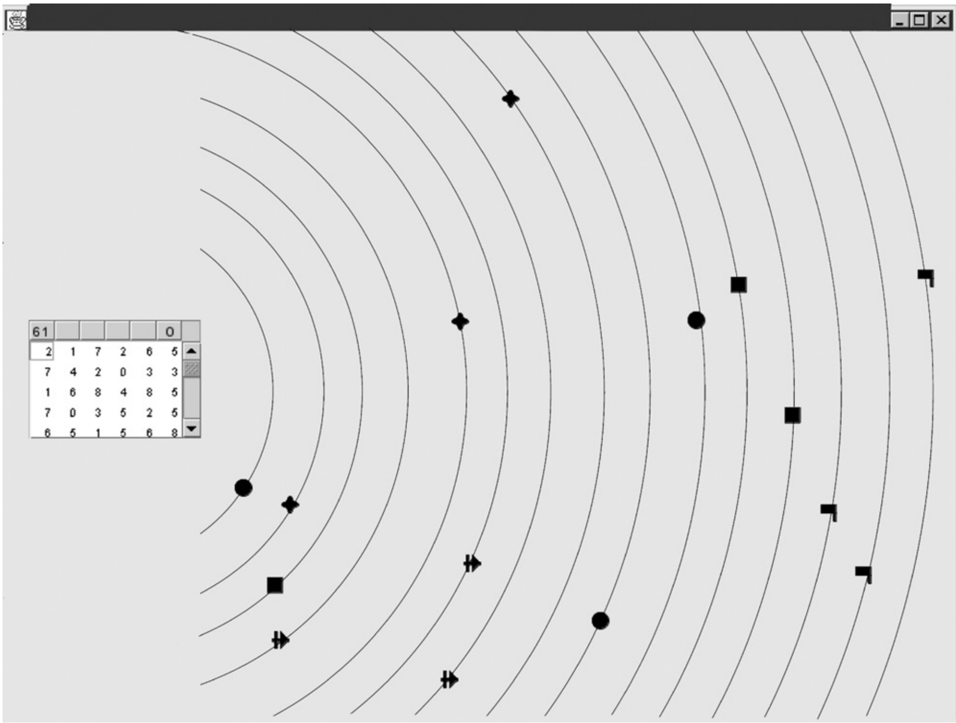

Fig. 2. One trial of Experiment 1A. The arcs were not displayed.

the left. The range of stimulus positions subtended a field of view from approximately 7° to 52° of visual angle from the centre of the primary task window.

The screen was located in a graduate research lab and the overhead fluorescent lights were dimmed for the experiment. No other special considerations were applied, so that participants were often doing the experiment while other work was going on the lab with normal (moderately quiet) noise levels.

Fig. 2. illustrates the screen layout.

4.1.1.2. Primary task. Participants were asked to carry out a simple editing task in a window which was located either to the left or the right on a larger window (see Fig. 2). The small window contained a scrollable table of numbers from 0 to 9 and the participant was instructed to find all the 0's in the table and replace them with 1's. A static counter in the upper left-hand corner showed the total number of 0's in the table; a running counter in the upper right-hand corner indicated the number of 0's currently found and replaced by the participant. Participants could use the arrow keys and/or the mouse and scrollbar to navigate through the table.

4.1.1.3. Secondary task. The larger window contained 15 icons, of which one might change according to one of the cues (motion, colour or shape). In a fifth of the trials nothing would happen. Upon detecting a change participants were instructed to hit the CONTROL key, thereby ending the trial. Each of the 15 icons was randomly positioned on an arc at fixed radial distance from the centre of the task window, as illustrated in Fig. 2. Icons were positioned in near and far conditions such that the NEAR targets were positioned on arcs 1–5 and the FAR targets were positioned on arcs 11–15. The NEAR targets subtended a field from approximately 7° to 20° of visual angle, while the FAR targets subtended a field from approximately 40° to 52° of visual angle. The target was randomly determined for each trial from the respective set of five icons. None of the icons on the arc 6 through 10 was ever a target. Each icon was bounded by a rectangle of $2\,\text{cm} \times 2\,\text{cm}$. This corresponded to roughly 1° of visual angle at the viewing distances we used.

Cue onset occurred between 5 and 20 s after the trial started. Cue onset was randomly selected from this 15 s range for each trial. Each cue lasted 5 s. If no detection was indicated, the trial timed out after 30 s.

4.1.1.4. Experiment participants. Twelve SFU students in either Computing Science or Psychology participated in the study. There were six males and six females. None were colour-blind, although five wore glasses. Participants were paid.

4.1.1.5. Experiment conditions. The 4 (cue type) $\times$ 2 (distance) factor design produced eight cue conditions. An additional two no cue (NO_CUE) control conditions were added resulting in 10 conditions overall. Thus, in 20% of the trials nothing happened and the trial timed out. A trial block contained two repetitions of each condition for a total of 20 trials. The ordering of conditions was randomized within a block independently for each participant. Trials were combined into four blocks of 20 trials each. Blocks were repeated 4 times totalling 80 trials per participant.

The task position was counter-balanced for left and right positions and changed with each block. An equal number of participants started with the task in the left position and in the right position. Each block was participant-initiated, so that participants could pause between blocks if desired.

All participants were given a training block of 10 trials with all cues before each study so that the participant was comfortable with the primary task and had seen all cues before the experiment.

4.1.1.6. Motion cues. There were two motion conditions. In the high amplitude condition (HIGH_AMP) the icon moved smoothly up and down along a path its own height (2 cm) with sinusoidal motion, a distance which covered approximately 1° of visual angle. In the low amplitude condition (LOW_AMP), the path length was the half the height of the icon (1 cm) or approximately half a degree of visual angle. Frequency was roughly 3 Hz.

4.2. Experiment 1A: colour vs. motion

In Experiment 1A two colour signals were evaluated together with the two motion signals. Six icon shapes were used and all were initially coloured black. Icon shapes

Fig. 3. Icon shapes used in Experiment 1A.

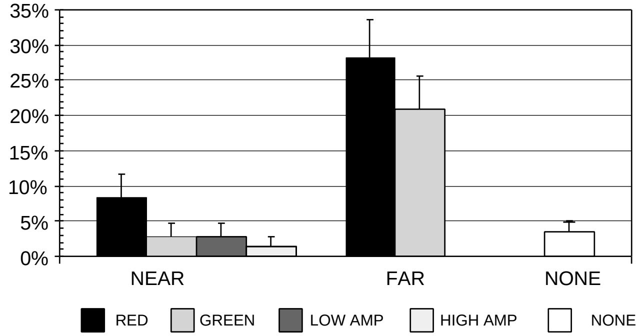

Fig. 4. Mean detection error rates: colour vs. motion (Experiment 1A).

were randomly assigned on each trial from a set of six canonical shapes shown in Fig. 3. The colour signals were a colour change to RED (RGB 255,0,0) or GREEN (RGB 0,255,0).

4.2.1. Results

Detection times were faster with the moving icons and error rates were lower. Detection error was measured as failing to detect a cue (in the four cue conditions) or falsely indicating detection in the NO_CUE condition. A 5 (cue) $\times$ 2 (distance) analysis of variance (ANOVA) showed both cue ($F(4,48) = 14.098$, $p < 0.0001$) and location ($F(1,12) = 14.846$, $p < 0.0002$) effects to be highly significant for detection accuracy (errors are shown in Fig. 4). As we expected, there was a significant interaction between cue and location: $F(4,48) = 7.304$, $p < 0.0001$. The moving icons were equally well detected in both the near and far visual fields. There was nearly a 0% error rate for motion in both the near and far conditions. However, as predicted, colour detection fell off in the periphery for both red and green icons, while motion detection remained constant. The average error rate for colour was 5.5% in the near condition and 24% in the far condition.

The difference between colour and motion is also evident in the detection times shown in Fig. 5. A similar 4×2 ANOVA[1] showed significant effects for cue and location: cue $F(3,36) = 40.278$, $p < 0.0001$; and location $F(1,12) = 7.135$, $p < 0.01$. All of the detection times for motion were around 1 s, but the detection times for colour averaged 2.3 s in the near condition and 4.6 s in the far condition. Again, the interaction between cue and location was significant: $F(4,48) = 2.695$, $p < 0.05$. There

[1] Obviously, we did not measure timing effects for the NO_CUE condition.

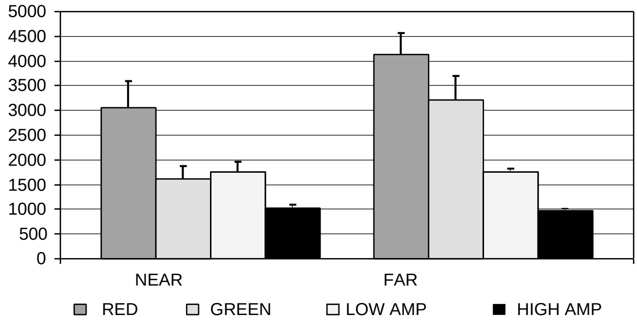

Fig. 5. Mean detection times: colour vs. motion (Experiment 1A).

Fig. 6. Icon shapes used in Experiment 1B.

was no effect of motion amplitude on detection rate or time, contrary to our expectation.

4.3. *Experiment 1B: shape vs. motion*

In Experiment 1B, two shape-change signals were compared to the two motion signals from Experiment 1A. In this experiment, all the icons had the same circular shape but colour varied. Icon colour was randomly assigned from a set of six RGB colours: red (255,0,0), green (0,255,0), blue (0,0,255), cyan (0,255,255), yellow (255,255,0), white (255,255,255) or black (0,0,0). The two shape cues were an X and an upright flag shape (Fig. 6).

4.3.1. *Results: Experiment 1B*

The pattern of result for shape is very similar to that obtained for colour and is summarized in Figs. 7 and 8 As predicted, the moving icons had higher detection rates and overall faster detection times than the shape cues. A multiway 5 (cue) $\times$ 2 (location) ANOVA showed that cue and location have highly significant effects on detection: cue $F(4,48) = 13.19$, $p < 0.0001$; and location $F(1,12) = 15.92$, $p < 0.0001$. While motion has a constant detection rate in the near and far locations, the detection rate for the two shape cues falls off in the periphery. A multiway ANOVA (cue $\times$ location) shows this interaction to be statistically significant ($F(4,48) = 7.37$, $p < 0.0001$).

Detection times were only measured for correct responses. Again, cue and location had strong effects on detection times, from a 4 (cue) $\times$ 2 (location) ANOVA: cue $F(4,48) = 19.55$, $p < 0.0001$; and location $F(1,12) = 31.33$, $p < 0.0001$. The interaction

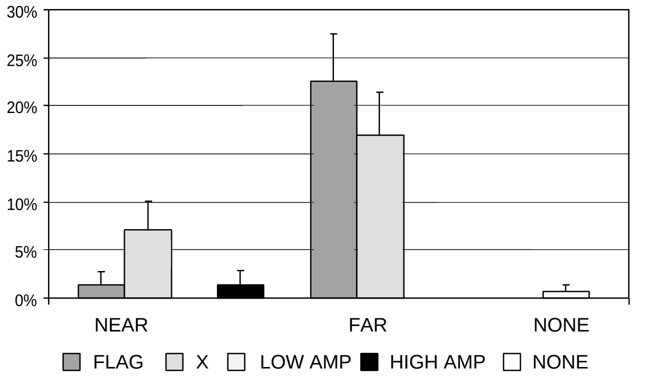

Fig. 7. Mean detection error rates: shape vs. motion (Experiment 1B).

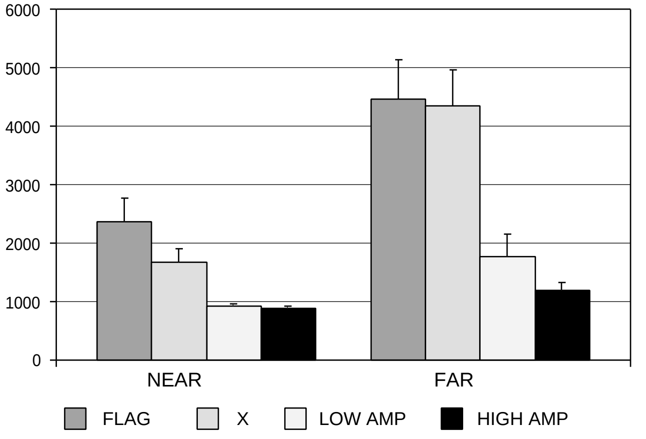

Fig. 8. Mean detection times: shape vs. motion (Experiment 1B).

between cue and location is also statistically significant ($F(4,48) = 4.66$, $p < 0.05$). The detection time results (Fig. 8) showed a large near–far difference for shape but not for the high-amplitude motion condition, increasing from 2.0 to 4.4 s, while times for that motion type remained constant at about 1.0 s. Interestingly, there is a location effect in this experiment for the low-amplitude motion: detection time almost doubles in the far condition, from 0.92 to 1.78 s (which is still only about 40% of the shape detection times), while there was no such effect in Experiment 1A.

4.4. Discussion: Experiment 1

The results of Experiment 1 show that even relatively slow motion is more effective than either colour or shape change in attracting user attention, especially in the periphery of the visual field. Our expectation that motion amplitude would affect detection accuracy was not supported: there was no significant difference in either experiment between the two motions. However, detection times were influenced by motion amplitude, but only in Experiment 1B in the far condition. As both experiments (1A and 1B) were performed by a participant in a single session, and as experiment ordering was counter-balanced, this result is puzzling. Two possible explanations relate to speed and distance. The LOW_AMP motion moves more slowly than the HIGH_AMP motion, as it covers half the distance in the same amount of frames. Ware et al. have shown that faster movements provoke quicker responses (Ware et al., 1992). Another reason may be that there is a distance threshold an object has to exceed before it is perceived as moving. It is conceivable that these properties have greater influence over the detection of motion in the periphery than in the central area of vision. Nonetheless, this effect was still substantially below the detection times of the static cues. This indicates that very small linear oscillations of limited duration are appropriate notification mechanisms in a dense display environment.

4.5. Experiment 2: cue identification in a heterogeneous environment

Experiment 1 demonstrated the efficacy of oscillating motion as a signal in simple displays with only a few variables (motion and shape or colour). In the real world, however, displays are crowded with multiple colours and shapes. A signal must not only alert the user to some event; it must also be immediately identifiable, so that the user knows which visual element on the screen changed and thus which event the signal represents. We adapted the previous experiment method to a more ecologically valid approach to simulate more realistic environments, in which colour and shape are combined, and in which change requires the user to not only recognize that something has happened but also to find the relevant object.

4.5.1. Method

Experiment 2 combined colour, shape and motion cues to investigate how well they performed in a more complex screen in both detection and identification: i.e. how accurately the participant could find which icon in the visual field had changed. In addition, four distinct motion types were investigated, of different shapes and smoothness (continuity). Accuracy and response time in both detection and identification were measured. Specifically, it extended the approach used in Experiment 1 in the following ways.

- In Experiment 1A, the icons varied in shape but were all initially of the same colour. In Experiment 1B, the icons varied in colour but were all initially the same shape. Experiment 2 icons varied in both colour and shape.
- In Experiment 1, the only secondary task was detection: all that the participant had to do was indicate he or she had seen the cue. Experiment 2 required the participant to both indicate detection and to identify (by pointing at it with the mouse) the icon that had changed.

- Experiment 1 had only two dependent variables: detection error and detection response time. Experiment 2 measured four dependent variables: detection error, detection time, identification error and identification time.

We expected motion cues to be better detected than colour and shape cues as in Experiment 1. In addition, we thought motion continuity, or "smoothness", would affect detection.

We also anticipated that identification would follow the same detection patterns as Experiment 1: identification errors and times would increase with distance for colour and shape cues, while motion identification would not be affected by distance.

4.5.1.1. Screen layout. The same screen layout and task window as in Experiment 1 were used. The screen contained 15 icons to which the shapes of Experiment 1A (Fig. 3) and the colours of Experiment 1B were randomly assigned. Icon size, placement and NEAR/FAR assignment were handled as in the previous experiment. Green (0,255,0) and the X shape were reserved for cues.

4.5.1.2. Experiment tasks: primary and secondary. As in Experiment 1, the primary task consisted of editing a table of numbers to replace 0's with 1's. The secondary task consisted of a detection and an identification stage. In the first stage, the participant performed the editing task and was instructed to hit the CONTROL key upon detecting any change. In 25% of the trials, no change occurred. Cue onset occurred at some time between 5 and 20 s after the trial started and lasted for 5 s. Onset was randomly selected from this 15 s envelope for each trial. Indicating detection ended this stage of the trial. If no cue was detected, this stage timed out after 30 s. Thus if there was no cue and this stage of the trial timed out, detection was logged as TRUE. The timing envelope ensured a buffer of at least 5 s between when the cue stopped and when the participant was prompted to identify it.

When the cue was detected (the end of the detection stage), two things occurred simultaneously. The cue target changed back to its original appearance, meaning it either stopped moving or changed colour or shape back to its previous state. This meant that in the case of the colour and shape cues, there was actually a "motion-like" event (a single-frame change) immediately after detection was indicated when the object state changed back, whereas in the motion cues the object immediately returned to its starting position. As well, a prompt appeared over the task table instructing the participant to identify the changed stimulus using the mouse or to indicate "no change" with CONTROL. The identification stage was not time-constrained. Completing the identification task ended the trial.

4.5.1.3. Cues. Six signals were used as cues: a colour change, a shape change and four motion types. An icon could change to GREEN or change shape to an X. (The GREEN and X cues were chosen as they were the best "performers" in their categories). Alternately, it could move according to one of two shapes and one of two smoothness conditions, as follows:

- SMOOTH LINEAR: Similar to Experiment 1, the icon moved smoothly up and down a path its own height (14 pixels) with sinusoidal motion (henceforth, *smooth* will refer to sinusoidal motion).
- JUMPY LINEAR: The icon moved up the same linear path smoothly and then "jumped" in one frame back to the starting point.
- SMOOTH ZOOM: The icon smoothly grew and shrank between 100% and 200% of its original size, centered on its origin.
- JUMPY ZOOM: The icon smoothly grew to 200% of its original size and then "jumped" in one frame back to its starting size. This gave the impression of "popping".

4.5.2. Experiment conditions

Adding two NO_CUE conditions to these six cues gave an 8 (cues) × 2 (location) design, resulting in 16 conditions of which each participant performed five repetitions, totalling 80 trials per participant. Condition ordering was independently randomized for each participant. over the entire 80 trials. Trials were grouped into blocks of 20. As in the previous experiment, task position was counter-balanced for left and right positions and changed with each block. An equal number of participants started with the task in the left and right positions. Each block was participant-initiated. All participants were given a training block with 16 trials covering all cue types before each study and could practise with it until the participant was comfortable with the primary task and familiar with all cues before the experiment.

Twelve SFU students in either Computing Science, Psychology or Kinesiology participated in the study. None were colour-blind although four wore glasses. Participants were paid.

4.6. Results and Discussion: Experiment 2

Detection error rates and times are shown in Figs. 9 and 10, respectively. A 7 (cue type) × 2 (location) ANOVA was performed on both detection errors and times. Similar to Experiment 1, detection error rates were lower with the moving icons. The moticons significantly outperformed the static cues in both the NEAR and FAR fields ($F(6,72) = 72.76$, $p < 0.0001$) in both detection accuracy and response time. This was to be expected given the findings of the first experiment and that the colour and shape coding dimensions were already "crowded". However, as assigning colour and shape to signals is a common approach in information-rich displays, it is noteworthy how poorly they perform. Location also had a significant effect: ($F(1,12) = 22.67$, $p < 0.0001$). However, there was again a significant interaction between cue and location ($F(7,84) = 12.84$, $p < 0.0001$). Location had no effect on the detection accuracy for any of the moving cues, while detection of the static cues fell off dramatically in the periphery.

Cue had a main effect on detection time ($F(5,60) = 29.52$, $p < 0.0001$). There was no main effect for location, but there was a significant interaction with cue: $F(5,60) = 4.28$, $p < 0.001$. This was due to the effect of location on response time

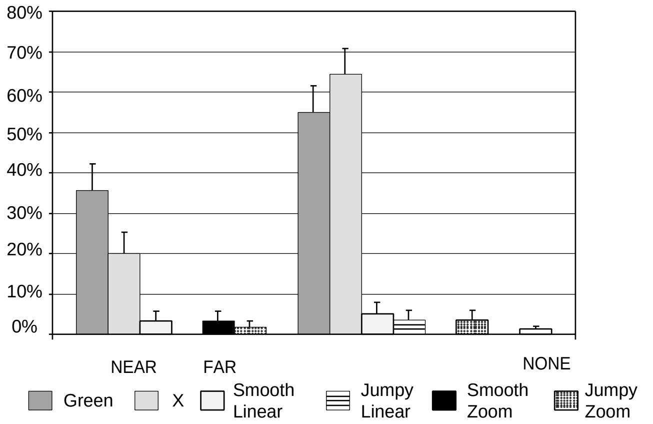

Fig. 9. Experiment 2 mean detection error rates.

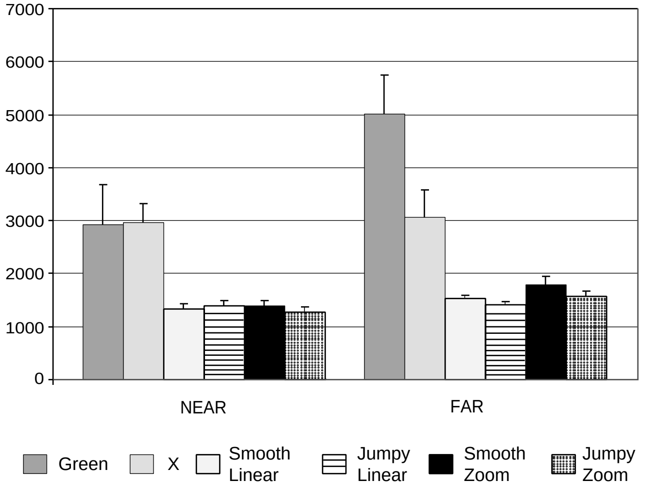

Fig. 10. Experiment 2 mean detection times.

for the colour cue (as in Experiment 1A). The moticons were more quickly detected than the static cues in the NEAR field and more quickly detected than the colour cue in the FAR. Interestingly, while the shape cue and colour cue had similar error rates in the far condition, it took participants approximately the same time to detect the shape cues in both near and far conditions. Contrary to our expectation, the smoothness of the motion had no effect on the detection results. All of the motion cues performed equally well with no significant differences. Detection times were only measured for accurate detections.

Fig. 11 shows the mean identification error rates and Fig. 12 the mean identification times. A 6 (cue) $\times$ 2 (location) ANOVA showed cue to have a significant effect on identification accuracy ($F(5,60) = 8.85$, $p < 0.0001$), while

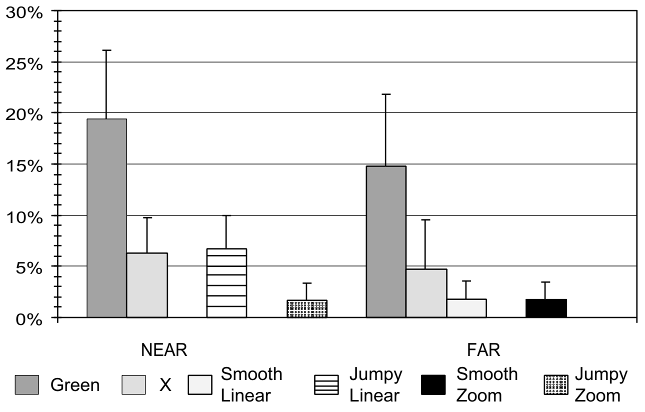

Fig. 11. Experiment 2 mean identification error rates.

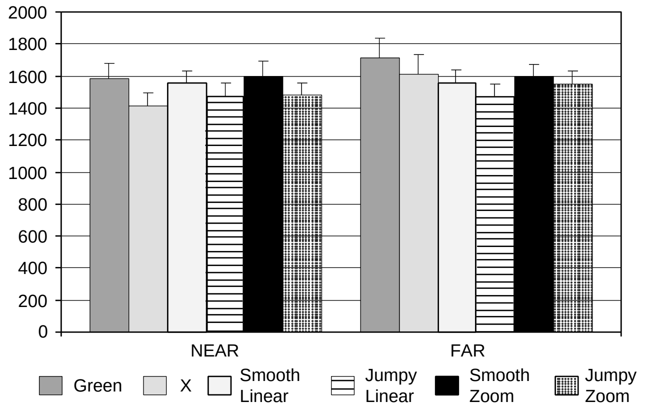

Fig. 12. Experiment 2 mean identification times.

location had no effect and there was no interaction. Accuracy in identifying the colour cues was markedly poor (almost 20% in the near condition and 15% in the far condition). The identification times (Fig. 12) show no significant effects and were all around 1.4 s, contrary to our expectations. It must be noted that identification errors and times were only measured across "detected" responses.

4.6.1. Discussion

While the detection results for colour vs. motion reflected the results of the previous experiment, shape detection did not. In this experiment there was no difference in shape response time due to location: participants took longer than in the previous experiment in the NEAR condition to detect the change and a little less time in the FAR. The longer time in the NEAR condition can be explained by the density of the visual environment, but the drop about 1 s in mean detection time for the FAR condition is puzzling. It is possible that participants could have relied on a polling strategy rather than attentional pull, but in that case we would assume a similar pattern for all cues, including colour.

The identification results suggest several preliminary conclusions. One is that colour cues, which are the most common signals, are substantially unreliable in a visually busy environment. Even when participants detected them, they were unable to accurately identify which stimulus had changed more than 15% of the time. Another is that, excluding colour, it appears that once a signal has been detected, it can be quite dependably found regardless of location in a wide visual area. (We remind the reader that the sample size of identified cues varies, as we only measured detected cues.) Why would this not be the case for the colour cue? Perhaps colour change, unlike motion or shape change, does not evoke onset and does not ‘‘register’’ with the tracking mechanism as effciently as motion and shape. Shape change might, after all, be considered a single-frame animation or a very short motion. More research is indicated.

5. Distraction and detection in different task conditions

The previous results proved the strength of motion signals as notification mechanisms in a particular task environment, but further questions arise concerning the effect on different types of tasks. While we believe that motion has great promise to support Woods' cognitive tools for managing attention (Woods, 1995), the issue of distraction must be addressed. Motion effects are increasingly used in interfaces for a variety of purposes and anecdotal evidence accrues that such effects are often inappropriately distracting and irritating to the user. In one study confirming the common complaints of distraction and task interference from such moving elements, Maglio et al. studied both continuous and discrete, or sporadic, scrolling banners to assess how well the subject could execute a primary editing task while remaining aware of the information in the banner as a secondary monitoring task (Maglio and Campbell, 2000). They found that continuous scrolling was substantially more distracting (impeded the primary task) with no increased benefit in awareness of the information contained therein.

Our intuition was that, while many types of motion are efficient in attracting and directing attention, certain types of motion are inherently more distracting than others. In particular, the moving banner animations on many web pages may be particularly distracting. A related intuition was that this distraction factor would depend to a certain extent upon how engaged the user was in the primary task, following work by Czerwinski on how different types of search tasks are affected by interruption (Czerwinski et al., 2000). These latter researchers found that interruptions are more disruptive to faster, stimulus-driven tasks than to effortful, cognitively taxing tasks.

Researchers have examined the cognitive cost of interruptions to a primary task (Gillie and Broadbent, 1989; Czerwinski et al., 2000; Cutrell et al., 2001) and the interference interaction between the primary task and a secondary task monitoring information conveyed by motion (Maglio and Campbell, 2000; McCrickard, 2000). Our concerns were slightly different and oriented towards the "user experience". Does introducing motion into the user interface inappropriately distract and annoy the user? We were interested in two questions.

1. Are certain motion types more distracting than others?
2. Does the level of distraction differ with task engagement?

Note that we describe "engagement" rather than attention. We borrow a concept from educators here. According to Crook, "task engagement" determines how far cognitive skills get mobilized and deployed (Crook, 2000). Using this concept, we think of more "engaging" tasks as ones in which the user is more committed to actively participating, and less prone to boredom and distraction.

6. Experiment 3: distraction, detection and task

This study was designed to evaluate the detectability, distraction and irritation properties of different types of moticons in desktop environments under different task conditions. We chose three different tasks intended to require different kinds of user response and varying degrees of engagement. This experiment compared how well participants detected various types of motion cues and how distracting and irritating they considered them in these various tasks.

The participant was engaged in a primary task in a window in the middle of the screen with icons distributed around the periphery of the monitor. During the task, different moving cues of different speeds occurred in random locations. The participant was instructed to indicate when the cue was noticed and later, in a post-task stage, to rate each of the cues for how distracting and irritating they were. Distraction was defined to the user as how *attentionally demanding* the cue was: that is, how strongly the participant remained aware of the movement once it had been detected. Participants were asked to rate the cue for these criteria with respect to type of task in the trial: i.e. if the participant had been playing Tetris in the trial, she was asked to rate any of the cues seen with respect to how much attention they had taken away from Tetris. Irritation was explicitly undefined—we wanted to see if participants differentiated between the two concepts. We did not instrument task performance.

Detectability was measured by detection rate and time. Distraction and irritation were rated on separate 5-point Likert scales. We considered four motion types in two major categories: *anchored* and *travelling*. Anchored motions such as those used in the detection experiments are characterized by small trajectories around ("anchored on") the icon origin. Travelling motions involve larger trajectories in which the object leaves its original position and "travels" through several degrees of visual angle. Our intuition is that traveling motions demand more attention because there is a cognitive act of *tracking* involved in addition to detection. Thus the animated icons (moticons) commonly used as place markers on a WWW page would be considered anchored, while the scrolling tickertape banners would be considered travelling. We had the following hypotheses. First, we expected that detection errors, times and distraction ratings would vary with task engagement (the most engaging task would show the lowest distraction rating and highest detection times). This was based on the naive assumption that the highly engaged user would be the least prone to distraction. Second, we anticipated that detection and distraction would be strongly affected by motion type. In particular, we expected that travelling motions would be rated as more distracting than anchored motions, based on previous research with scrolling.

6.1. Method

A dual-task design was used with different primary tasks and a secondary task of detecting moving icons. Participants rated distraction and irritation aspects of the motion types in a post-task rating phase.

6.1.1. Screen layout

The primary task was centred in a single full-screen window which was framed by a border of 32 icons. Icon shapes and colours were the same as those described in the detection experiments. Shape and colour were randomly set each trial but position was constant. Each icon was bounded by a rectangle of 16×16 pixels. A standard 21″ monitor was used. Fig. 13 shows one screen layout with the Solitaire task.

6.1.2. Primary tasks

Three primary tasks were used: browsing and studying on-line text (TEXT), playing a variant of Solitaire called FreeCell (SOLITAIRE) and playing Tetris (TETRIS). In the

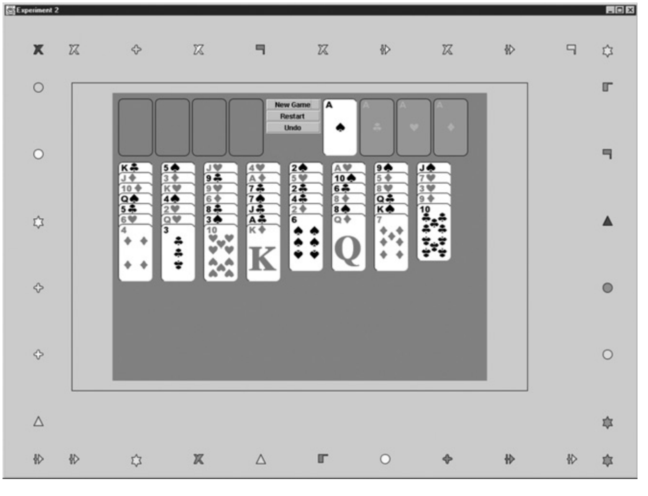

Fig. 13. Experiment screen with Solitaire task.

text task, participants were instructed to imagine that they had a class the next day in the particular subject and that they had to know this material before they went to class. Four different texts were selected to be outside the normal topic scope of the participant population. They dealt with the following areas: climate change due to methane gas production from farming in New Zealand (environment/agriculture); the history of the search for the 10th planet (astronomy); drug addiction treatment and management in Australia (social work); and the Roman–Carthaginian wars (history). Solitaire is a game in which the user can be quite involved but events are completely user-driven; diverting attention pauses the game. Tetris, on the other hand, is played against the machine in a stimulus–response type of interaction in which the user must intervene continuously; the game does not stop when the user does. We hypothesized that Tetris would be the most engaging task and that the text task would be the least engaging.

6.1.3. Secondary tasks: detection and rating

In the first phase of the trial, participants were instructed to simply indicate detection of a moving icon by hitting the CONTROL key. The moticon did not stop moving upon the detection event. In the rating phase of the trial, participants replayed a shorter version of each cue in random order and rated each on a 5-point Likert scale for each of the following four criteria: how distracting the motion was (DISTRACTION); how distracting it would it be if it persisted throughout a similar task (LONG DISTRACTION); how irritating it was (IRRITATION); and how irritating it would be if persisted (LONG IRRITATION). Cue ordering was random and timing was participant-initiated: the next cue was only played when the participant had answered all the questions and requested the next cue.

6.1.4. Motion cues

Four motion types (*shapes*) were used.

- LINEAR, in which the icon moved smoothly up in a sinusoidal motion and then "jumped" back to the origin (similar to the JUMPY LINEAR cue of Experiment 3).
- ZOOM, in which the icon zoomed smoothly from starting size to twice the starting size and then jumped back to the starting size.
- BLINK, a standard on–off cycle.
- TRAVEL, in which the icon moved across the screen at a constant rate either from right to left (if the icon were in the top or bottom location) or from top to bottom (if the icon were in a left or right position), leaving the screen at one side and wrapping around to the other.

Each motion type was implemented in two frequencies: SLOW (30 frames/s, roughly 1 Hz) and FAST (15 frames/s, roughly 2 Hz), resulting in eight distinct motion cues. Amplitude for the anchored motions was the same as Experiment 1, namely the stimulus size of 16 pixels.

6.1.5. Experiment design

The experiment was divided into groups of blocks. A participant performed one task in each block. Blocks were organized into groups of three with each block containing one instance of each task. Each participant performed 12 blocks in a session: four groups of three blocks each. Task ordering was counter-balanced and each participant had one block group starting with each task type and one more randomly selected block group.

Each block comprised eight trials, one per motion type (or *cue*). Thus each participant saw each cue in a given task condition 4 times.

Fourteen students from Psychology, Engineering, Computer Science and Kinesiology at SFU served as participants. There were an equal number of males and females. None were colour-blind. Participants were paid.

6.1.6. Block and trial description

Each block consisted of two phases: the trial phase and the rating phase. In the trial phase the participant engaged in one of three primary tasks for a period of 4.5 min. During this phase there were eight trials, one for each motion (4 cue types $\times$ 2 frequencies). Cue duration was 10 s. Trial timing was evenly distributed across eight 30-s "envelopes" with a randomized onset in each envelope from 5 to 15 s, such that the effect was random but no cues occurred less than 10 s or more than 30 s apart. Since there is evidence that interruptions are more distracting in the beginning stages of a task than when the task is well underway (Czerwinski et al., 2000), we hoped this even distribution would mitigate a task timing effect.

Cue–icon assignment was also random but evenly distributed such that there were always two cues from each of the top, left, bottom and right sides. No icon was used twice in a trial. Participants were instructed to concentrate on the primary task and to indicate detection by hitting the CONTROL key when they noticed any one of the eight cues. Cue detection did *not* stop the cue nor end the task, but it did end the trial. Trials were transparent to the participant, since he or she interacted continuously with the task until the entire block of trials had run.

When the trial stage of the block ended (by timing out), the participant was presented with a rating screen in which a single icon unique to the rating phase appeared in the upper left corner and a rating panel appeared in the centre of the window. Each motion cue was then briefly replayed and rated by the participant.

6.2. Experiment 3: results and discussion

Figs. 14–16 summarize the detection results of this experiment. As predicted, there was a significant task effect on detection errors (Fig. 14), although the differences between the two game tasks were smaller than anticipated. Generally, cues were most accurately detected in the text browsing task, which was hypothesized to be the least engaging of the tasks and therefore the most prone to distractors. The two games had similar results with Tetris showing a slightly higher rate of detection errors overall except for the `FAST BLINK` and `SLOW ZOOM` cues. The `SLOW BLINK` was the least detected in all tasks with the odd exception of the `FAST BLINK` in Solitaire.

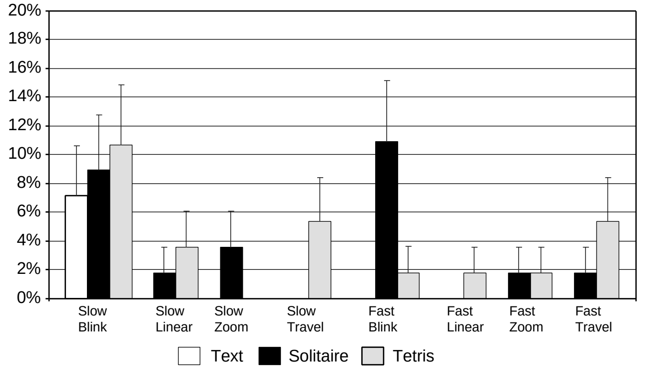

Fig. 14. Cue detection error rates by task.

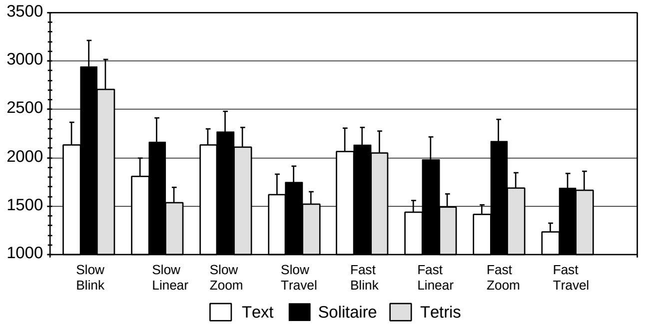

Fig. 15. Cue detection times by task.

	Fast Blink	Fast Linear	Fast ZOOM	Fast Travel	Slow Blink	Slow Linear	Slow ZOOM	Slow Travel
Fast Blink								
Fast Linear	■				✦✢■		■	
Fast ZOOM	■				✢■		■	
Fast Travel	■				✦✢■		■	
Slow Blink								
Slow Linear					✢			
Slow ZOOM								
SlowTravel					✦✢			
✦: Solitaire ✢: Tetris ■: Text								

Fig. 16. Pairwise comparisons of detection times: a task symbol in a cell indicates that the row cue had a faster (i.e. less) detection time than the column cue.

An analysis of variance showed significant effects on detection for cue type ($F(7,84) = 3.84$, $p < 0.001$) and task ($F(2,24) = 4.75$, $p < 0.0001$). There were no significant interactions. There was, however, a significant participant–task interaction for detection ($F(2,26) = 2.82$, $p < 0.0001$). Two participants detected fewer cues in Tetris than in Solitaire, where other participants saw fewer in Solitaire than in Tetris. This may indicate the differing levels of engagement in the two games for different people. Detection for all participants was more accurate in the text task.

With one notable exception (the FAST BLINK in Solitaire) these results match our expectations that task engagement and cue type together affect detection. The slow blink cue is the least well detected. Frequency alone has no effect on error rate. Instead the type of motion overweighs frequency (FAST LINEAR, for example, is not better than SLOW ZOOM). However, the BLINK, LINEAR and TRAVEL cues were less well detected in Tetris than the ZOOM motions, suggesting some interference with the motions already present in the task. The lower rate of detection for the FAST BLINK Solitaire is mysterious.

Figs. 15 and 16 show the detection time results. Overall, linearly moving cues (LINEAR and TRAVEL) were detected most efficiently. Analyses of variance in detection time show task ($F(2,22) = 4.75$, $p < 0.0001$) and cue ($F(7,84) = 3.84$, $p < 0.001$) effects to be highly significant, but this is mainly due to the slow detection of the SLOW BLINK cue, and does not clearly support our hypothesis that travelling motions would be detected more quickly than anchored motions.

We also expected the fast cues would be more quickly detected than the slower cues, but pairwise comparisons of significant differences in cue detection times by task from a post hoc Tukey analysis (Fig. 16) failed to support this except for the slow blink cue, which had a much greater failure-to-detect rate. Task plays a large part in determining how relatively efficiently cues are detected: in particular, the differences between many cues are less pronounced in the two games tasks than in the text condition. Detection was fastest in the text task and slowest in Solitaire.

6.2.1. *Distraction ratings*

Figs. 17–19 summarize the distraction results. There were only minor differences between the four subjective ratings (DISTRACTION, LONG DISTRACTION, IRRITATION and LONG IRRITATION): therefore we only report distraction. Generally, participants rated the cues as increasing somewhat in distraction and irritation for longer persistence, but this was not a statistically significant difference. An analysis of variance on the Likert ratings for each of the three variables showed significant effects on distraction ratings for task ($F(2,24) = 14.62$, $p < 0.0001$); motion type ($F(7,84) = 86.89$, $p < 0.0001$); and frequency ($F(1,12) = 40.18$, $p < 0.0001$).

As evidenced in Fig. 17, motion type, rather than motion frequency, clearly is the dominant factor: it has the most pronounced effect on distraction rating. Blinking is rated as the least distracting, followed by linear, zooming and finally travelling trajectories. Although some participants reported that their judgement of what was distracting changed as the experiment session progressed, there was no indication that when a trial occurred had any effect on the results. While task also had a significant effect, the magnitude of the difference was smaller than anticipated. As expected, cues were generally seen as more distracting in the text task than in the two game tasks. Frequency also influenced distraction rating as shown in Figs. 17 and 18.

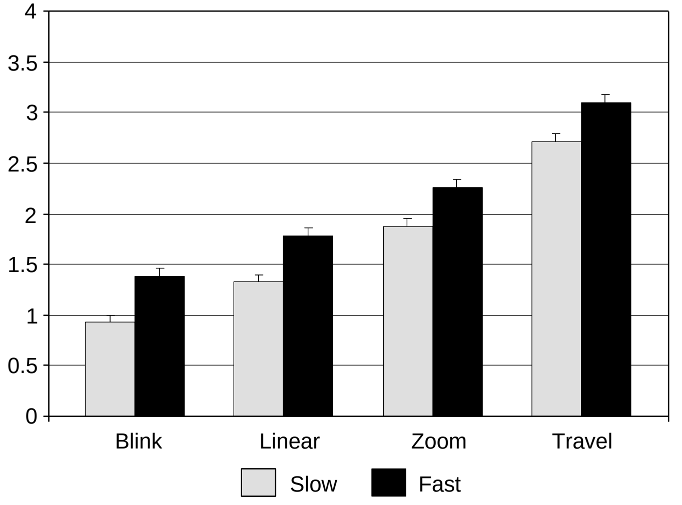

Fig. 17. Distraction rating by type and frequency.

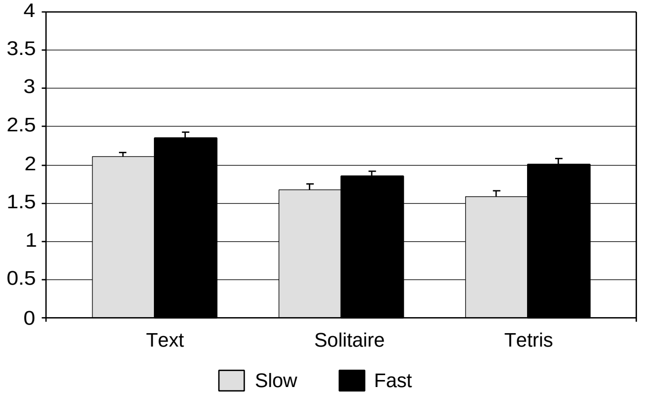

Fig. 18. Distraction rating by task and frequency.

Again, however, while the difference was constant, the magnitude was relatively small. Fast linear motion in tetris is rated as more distracting than in Solitaire.

There was a participant–task interaction for distraction rating ($F(22,242) = 5.3984$, $p < 0.0001$). We believe this was largely due to differences in performance and experience with the two games. However, there was also a participant–cue interaction ($F(2,26) = 2.904$, $p < 0.0001$), which we believe is due to the fact that the 5-point scale is a relative measure. In this case it is more revealing to consider significant differences in cue distraction ratings from a Tukey pairwise comparison (Fig. 19). From this it is immediately apparent that the two travelling cues are by far the most distracting, confirming our hypothesis. The ZOOM cues, especially the fast,

	Fast Blink	Fast Linear	Fast Zoom	Fast Travel	Slow Blink	Slow Linear	Slow Zoom	Slow Travel
Fast Blink		■	✦✢■	✦✢■				✦✢■
Fast Linear			✦	✦✢■				✦✢■
Fast Zoom				✦✢■				✦
Fast Travel								
Slow Blink		✦✢■	✦✢■	✦✢■			✦✢■	✦✢■
Slow Linear			✦✢■	✦✢■			✦	✦✢■
Slow Zoom				✦✢■				✦✢■
Slow Travel								
✦: Solitaire ✢: Tetris ■: Text								

Fig. 19. Pairwise comparisons of distraction ratings: when a task symbol appears in a cell, the row cue is less distracting than the column cue.

were considered the next most distracting. On the other hand, the SLOW BLINK and SLOW LINEAR cues were rated as the least distracting.

6.2.2. Discussion

With one notable exception (the FAST BLINK in Solitaire), these results match our expectations that task engagement and cue type together affect detection and suggest that a measure of detection may be used with similar cues to gauge a user's involvement in a particular task. The two games had similar results with Tetris showing a slightly higher rate of detection errors overall except for the FAST BLINK and SLOW ZOOM cue suggesting that participants were devoting more resources to Solitaire, which ran counter to our expectations of task engagement.

Although we had expected that Tetris would be the task most engaging of attention and our subjects anecdotally said it was, the objective measure did not support this. Indeed the measured response times were longest for Solitaire. A possible explanation for this might lie in the interaction between primary task type and the nature of the distraction. Gillie and Broadbent (1989) provided evidence that the most important factor in determining what makes a distraction disrupt performance is task similarity. Using this we can construct the following argument to explain the unexpected result. Firstly, because distractions came with predictable overall frequency (even though the actual time of occurrence was random) it is possible that the subjects might have adopted a scanning strategy, using eye movements to the side to check for periodic changes. Secondly, visual search theory suggests that near field regions would have been more efficiently and frequently monitored than the far-field regions (Wolfe, 2000). Finally, such a visual search strategy requires spatial memory for icon locations and because Solitaire is also highly demanding on spatial layout memory this might explain the unexpected result. Playing Solitaire may have disrupted the visual scanning strategy more than the other tasks. However, this explanation is highly speculative.

Our expectation that motion type would overwhelm at least the two speeds we tested was strongly borne out. Travelling motions are subjectively judged as substantively more distracting than other kinds of motion and this is of particular concern given their prevalence in current interfaces.

There are obvious questions and issues around interactions between the motions in the primary tasks and the cue motions, notably with the fast blink in Solitaire and linear motion in Tetris, suggesting that there may be issues of perceptual interference between cue motion cues and game motion. This demands further study.

7. Conclusions and discussion

These findings have a number of implications for the use of moticons as signals in human–computer interfaces. They suggest that motion has several advantages as a notification mechanism. It is significantly better than the traditional static codes of colour and shape in designing icons used to attract a user's attention, especially in the periphery. Our results showed that the percentage of undetected targets increased dramatically from 6% to 25% with the peripheral colour targets, whereas the failure-to-detect with motion was less than 2% in both the near and far field. Identification of the moving target was also substantially better with the moving cues than with the static cues across the entire visual field. These results once more emphasize the effect that location has on static and moving cues. Motion seems to be equally efficiently "remembered" in both near and far conditions, while colour and shape are less well noticed and less accurately tracked. The poor performance of the colour and shape cues even in the very sparse experimental displays is particularly noteworthy, as these cues are the most commonly used signals in current computing environments. In complex visualization and control systems where notification is of great importance, the poor detectability and identifiability of these cues argues for much more perceptually distinct signals. Even small, slow linear oscillations of the types reported here are excellent candidates for such notifiers. They have the advantage of being computationally cheap and consume little spatial and temporal resources. Perhaps the greatest advantage of using motion-based signals, however, is that they do not seem to interfere with existing colour and form coding, allowing extra information to be communicated through an object without changing its original codes which represent other variables.

The high rate of detection even in the more engaging tasks suggests that motion is effective over a wide range of locations, types and amplitudes. Even the least efficiently detected cue, the slow blink, had a worst-case mean response time of less than 3 s and was detected 89% of the time within the 10-s window, indicating good accessibility even when the primary task is demanding. As one reviewer stressed, however, participants were expecting cues, and thus the detectability of these cues may vary (although we believe the relative performance will remain the same).

However, we did find differences in the ease with which people could be distracted from the primary task, although the results were not exactly what we anticipated. Although we had expected that Tetris would be the task most engaging of attention and our participants anecdotally said that it was, the objective measures did not support this. Indeed the measured response times were longest for Solitaire, suggesting that this demanded the most resources. However, it must be emphasized that we did not measure primary task performance but instead signal recognition. In any case, our results show that even with highly demanding tasks motion can be readily used as an alert. They further suggest that such motion cues may be used to gauge task engagement.

Finally, while all the tested motions are effective as signalling mechanisms, some are clearly more distracting to the user. Travelling motions which involve both detection and tracking are substantially more distracting than anchored motions. The zoom motions are also (although less) distracting, probably because they elicit sudden perceptual onset (Hillstrom and Yantis, 1994). These findings confirm our experience that animated banners and popping images are not comfortable visual elements on a screen where one is trying to work but are effective if one, in fact, wants to dominate the user's attention. Overall the slow linear motion would appear to be a good compromise. It was rated among the least irritating and distracting, but it elicited good response times and detection rates.

We note that it is not confirmed whether these effects are due to the low-level event of a user's attention being grabbed by onset or motion, or the result of a slightly higher-level visual search strategy in which the user polled the display every few seconds. Users are very good at learning the distribution around cues in experimental conditions and adjusting polling strategies accordingly.[2] Research leads us to believe in the former. There is substantial evidence to indicate that motion does indeed grab attention at a low level (Pylyshyn et al., 1994; Hillstrom and Yantis, 1994) and across the entire visual field (Sekuler et al., 1981). However, at a functional level, we are less concerned with whether detection was due to polling or low-level attentional grab, since we expect the user deployed the same strategies regardless of cue type. We believe that these operational differences between the various cue types are still striking enough to warrant considering moticons as efficient alerting mechanisms, although future research should indeed look at studies with sparser and more managed cue distribution to ascertain whether these results would hold up in vigilance applications where cues arrive unpredictably and the cost of polling is unacceptable.

7.1. Guidelines for motion-based techniques

A review of the research and these results suggests the following preliminary recommendations and guidelines for adding motion-based signalling and awareness integration techniques to an interface.

G1: *Motion does not seem to interfere with existing colour and form coding*, allowing extra information to be communicated through an object without changing its original codes which represent other variables.

G2: *Small periodic motions are generally better signals than colour or shape cues across the entire visual field*. Such motions are more effectively detected and

[2] Thank you to the reviewer who pointed this out.

more accurately identified, especially in the periphery, and in displays where colour and shape are already used for other coding. Neither detection nor identification are affected by location, where colour and shape are poorly seen and identified outside the central area of view.

G3: *Motion amplitudes can be small*—approximately 1° of visual angle is highly detectable even in pronounced peripheral vision conditions. While amplitude does not affect accuracy in detection, smaller amplitudes than 1° may have slower response times.

G4: *Even relatively slow frequencies are effective*. The motions tested had periodic frequencies between 1 and 3 Hz. There is some research to suggest that frequency affects the perception of urgency (Ware et al., 1992), but there was no perceptible effect in the range of frequencies we tested.

G5: *Motion continuity*, or smoothness, appears to have little effect on detection and identification of signals.

G6: *Motion cue detection time*, *and to a lesser degree accuracy*, *is affected by the level of task engagement*. This effect does not suggest motion signals are inappropriate in highly engaging tasks, as the detection results of all motion types tested were good even in the most attentionally demanding tasks. However, it does suggest that signals can be "tuned" according to task if immediate response time is required.

G7: *Motion can contribute strongly to distraction and irritation*. In particular, travelling motions are significantly more distracting and irritating than anchored motions. Popping motions (where the object zooms in and out along the depth axis) were also considered distracting.

G8: *The slow linear oscillation is a good overall signal*. It is accurately detected, elicits good response times, and is not considered intrusive or distracting.

G9: *Motion signals are easily computed*. Refresh rates between 20 and 30 frames/s and durations of a few seconds are sufficient to elicit the perception of a single continuous motion event.

8. Future work

These results are obviously preliminary and introduce many issues and questions for future work. Of particular interest to us are the issues around the utility of motion signals and awareness visualization techniques in complex display situations: crowded visual displays, large and small screens, and heterogenous display configurations (using various devices together in non-desktop situations). Future work in motion signals will need to address two main areas: perceptual interfence with other motions in the visual field, and the actual cognitive cost of a motion-based interruption (as opposed to the user experience). Our currrent work is investigating the first of these, looking at the motion effect of "grouping", where similar motions in an interface are perceived to form a single entity (Bartram and Ware, 2002).

Acknowledgements

This research was carried out when the first author was at the School of Computing Science at Simon Fraser University and was partly supported by the Canadian Telelearning Network Centre of Excellence and by the Natural Sciences and Engineering Research Council of Canada.

References

Adams, M., Tenney, Y., Pew, R., 1995. Situation awareness and the cognitive management of complex systems. Human Factors 37 (1), 85–104.

Baecker, R., Small, I., 1990. Animation at the interface. In: Laurel, B. (Ed.), The Art of Human–Computer Interface Design. Addison-Wesley, Reading, MA, pp. 251–267 (Taxonomy of animation uses in user interface).

Bartram, L., 2001. Perceptual and interpretative properties of motion for information visualization. Ph.D. Thesis, School of Computing Science, Simon Fraser University.

Bartram, L., Ware, C., 2002. Filtering and brushing with motion. Journal of Information Visualization 1, 66–79.

Bartram, L., Ware, C., Calvert, T., 2001. Moving icons: detection, distraction and task. In: Hirose, M. (Ed.), Proceedings of the IFIP International Conference on Human–Computer Interaction (INTERACT) '01, IOS Press, Tokyo, pp. 157–165.

Crook, C., 2000. Motivation and the ecology of collaborative learning. In: Joiner, R., Littleton, K., Faulkner, D., Miell, D. (Eds.), Rethinking Collaborative Learning. Free Association Press, London, pp. 161–178.

Cutrell, E., Czerwinski, M., Horvitz, E., 2000. Instant messaging and interruption: influence of task type on performance. In: Paris, C., Ozkan, N., Howard, S., Lu, S. (Eds.), OZCHI 2000 Conference Proceedings, pp. 356–361.

Cutrell, E., Czerwinski, M., Horvitz, E., 2001. Notification, disruption and memory: effects of messaging interruptions on memory and performance. In: Hirose, M. (Ed.), Proceedings of the IFIP International Conference on Human–Computer Interaction (INTERACT) '01. IOS Press, Tokyo, pp. 263–269.

Czerwinski, M., Cutrell, E., Horvitz, E., 2000. Instant messaging: effects of relevance and time. In: Turner, S., Turner, P. (Eds.), People and Computers XIV: Proceedings of HCI 2000, Vol. 2, British Computer Society, pp. 71–76.

Faraday, P., Sutcliffe, A., 1997. Designing effective multimedia presentations. In: Proceedings of ACM CHI '97. ACM, New York, pp. 272–279.

Gillie, T., Broadbent, D., 1989. What makes interruptions disruptive? A study of length, similarity and complexity. Psychological Research 50, 243–250.

Gilmore, W., Gertman, D., Blackman, H., 1989. User–Computer Interfaces in Process Control: A Human Factors Engineering Handbook. Academic Press, San Diego, CA.

Healey, C., Booth, K., Enns, J., 1995. Visualizing real-time multivariate data using preattentive processing. ACM Transactions on Modeling and Computer Simulation 5 (3), 190–221.

Hillstrom, A., Yantis, S., 1994. Visual attention and motion capture. Perception and Psychophysics 55 (4), 399–411.

Limoges, S., Ware, C., 1989. Displaying correlation using position, motion, point size or point colour. In: Proceedings of Graphics Interface '89. Number TR94-089, pp. 262–265.

Maglio, P., Campbell, C., 2000. Tradeoffs in displaying peripheral information. In: Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2000). ACM, New York.

McCrickard, D.S., 2000. Maintaining information awareness in a dynamic environment: assessing animation as a communication mechanism. Ph.D. Thesis, Georgia Institute of Technology.

McCrickard, D.S., Stasko, J.T., Zhao, Q.A., 1999. Exploring animation as a presentation technique for dynamic information sources. Technical Report 99-47, Georgia Tech GVU Center.

Mitchell, C., Sundstrom, G., 1997. Human interaction with complex systems: design issues and research approaches. IEEE Transactions on Systems, Man and Cybernetics 27 (3), 265–273.
Parsowith, S., Fitzpatrick, G., Segall, B., Kaplan, S., 1998. Tickertape: notification and communication in a single line. In: Asia Pacific Computer Human Interaction 1998 (APCHI '98), Kangawa, Japan.
Peterson, H., Dugas, D., 1972. The relative importance of contrast and motion in visual perception. Human Factors 14, 207–216.
Pylyshyn, Z., Burkell, J., Fisher, B., Sears, C., Schmidt, W., Trick, L., 1994. Multiple parallel access in visual attention. Canadian Journal of Experimental Psychology 48, 260–283.
Sarter, N., Woods, D., 1995. How in the world did we ever get into that mode? Mode error and awareness in supervisory control. Human Factors 37 (1), 5–19.
Sekuler, R., Tynan, D., Kennedy, R.S., 1981. A sourcebook of temporal factors affecting information transfer from visual displays. Technical Report JUN-TR540, U.S. Army Research Institute for the Behavioral and Social Sciences, June 1981.
Shneiderman, B., 1986. Designing the User Interface: Strategies for Effective Human–Computer Interaction. Addison-Wesley, Reading, MA.
Smith, G., Atchison, D., 1997. The Eye and Visual Optical Instruments. Cambridge University Press, Cambridge, UK.
Ware, C., 2000. Information Visualization: Perception for Design. Morgan Kaufmann, New York.
Ware, C., Bonner, J., Knight, W., Cater, R., 1992. Moving icons as a human interrupt. International Journal of Human–Computer Interaction 4 (4), 341–348.
Wickens, C., 1992. Engineering Psychology and Human Performance, 2nd Edition. Harper Collins, New York.
Williams, L., 1985. Tunnel vision induced by a foveal load manipulation. Human Factors 27 (2), 221–227.
Wolfe, J.M., 2000. Visual attention. In: deValois, K.K. (Ed.), Seeing, 2nd Edition. Academic Press, New York, pp. 335–386.
Wolfe, J.M., Gancarz, G., 1996. Guided Search 3.0. In: Lakshminarayanan, V. (Ed.), Basic and Clinical Applications of Vision Science. Kluwer Academic Publishers, Dordrecht, pp. 189–192.
Woods, D., 1991. The cognitive engineering of problem representations. In: Weir, G., Alty, J. (Eds.), Human–Computer Interaction and Complex Systems. Academic Press, New York, pp. 169–188.
Woods, D., 1995. The alarm problem and directed attention in dynamic fault management. Ergonomics.
Wyszecki, G., Stiles, W., 1982. Colour Science: Concepts and Methods, 2nd Edition. Wiley Interscience, New York.