

Attention Design: Eight issues to consider

Sharon Wood*, Richard Cox, Peter Cheng

*Representation & Cognition Group, Department of Informatics,
University of Sussex, Falmer, BN1 9QH, UK.*

Abstract

In HCI research there is a body of work concerned with the development of systems capable of reasoning about users' attention and how this might be most effectively guided for specific applications. We present eight issues relevant to this endeavour: What is attention? How can attention be measured? How do graphical displays interact with attention? How do knowledge, performance and attention interact? What is working memory? How does doing two things at a time affect attention? What is the effect of artificial feedback loops on attention? Do attentional processes differ across tasks? For each issue we present design implications for developing attention-aware systems, and present a general discussion focussing on the dynamic nature of attention, tasks (number, nature and variety), level of processing, nature of the display, and validity of measures. In conclusion, we emphasise the need to adopt a dynamic view of attention and suggest that attention is a more complex phenomenon than some designers may have realised; however, embracing the multi-faceted nature of attention provides a range of design opportunities yet to be explored.

Key words: attention, attention aware systems, cognition, visual displays.

1 Introduction

Designing systems that can monitor the degree and focus of a user's attention and then adapt the interface, dynamically modifying it in some way to manage the user's attention is a major challenge for research in HCI and cognitive engineering. This endeavour is attempting to create a new class of interfaces that are unlike existing systems, so the process of design must be more than

* Corresponding author.

Email address: S.Wood@sussex.ac.uk (Sharon Wood).

an incremental extension of existing technologies. From studies of other design oriented disciplines and investigations of the nature of design (Simon, 1981; Goel & Pirroli, 1992) it has been established that successful design in such innovative situations needs strong scientific bases to underpin the enterprise: the exploration of a vast space of imaginable design possibilities must be constrained. Fortunately for the new area of attention design in HCI there are bases (in psychology and cognitive science) upon which to build. The aim of this paper is a preliminary foray into some of that work. The theories and findings that may be relevant are presented as eight issues, in order to stimulate debate rather than attempt to (prematurely) provide a coherent framework.

The structure of this paper is as follows - under eight separate headings we address issues concerned with the nature and measurement of visual attention, the role of external representations, human error, human memory and task variables. For each we also provide some specific design implications. We conclude by considering the implications of research findings in these areas for the design of adaptive interfaces for managing the user's attention.

2 Eight Issues for Attention-Aware Design

2.1 Issue 1: What is attention?

Attention is a process of selection and selective processing, required because the brain has a limited information processing capacity, (for example, Allport, 1993). In this paper we discuss visual attention rather than auditory attention. This distinction is necessary because the nature of attentional processes in these two sensory modalities appears to differ in some important respects. For example, broadly speaking humans demonstrate a greater capacity for subconsciously monitoring (and processing) unattended events in the auditory modality compared to the visual modality (for example, Cocktail party phenomenon, Cherry, 1957). The visual attentional system seems, in contrast, to be prone to *inattentive blindness* (see *Issue 2* below and Pylyshyn, 2003 for a review).

Several metaphors have been used as the basis for theorizing about the nature of visual attention: a spotlight with a moving fixed size diameter focus (Treisman, 1986); attention as a glue that binds together features of things that are being processed (Cowan, 2001); attention as a 'coherence field': a viewer-centred nexus maintained in short term memory, stabilising links between a selected set of low level visual structures (proto-objects) (Rensink, 2000a); attention as a bottleneck through which a limited amount of information can be filtered (Broadbent, 1958); attention as a limited capacity for

information processing (Broadbent, 1971); attention as a multimodal process in which different processing modes (top-down, bottom-up) are implicated at different stages (Pashler, Johnstone & Ruthruff, 2001). The marked differences between some metaphors, and similarity of others, is indicative of the lack of thoroughgoing theoretical consensus in this area.

One body of research suggests that visual attention operates in a manner akin to a camera's zoom lens - ie. a 'lens' of variable focal-length that can alter the scope of visual attention between 'wide-angle' and 'telephoto'. However a fixed amount of processing resources are applied to any particular scene within the range. Thus processing can either be intensive and local - dedicated to a small part of a scene with reduced attention to the rest of the scene, or more evenly distributed across a wider area (Eriksen & St James, 1986; LaBerge, 1983 - see Eysenck & Keane, 2000 for a review).

Moreover, recent research suggests that visual attentional profiles can be shaped within a scene. For example a study by Juola, Bowhuis, Cooper & Warner (1991) required subjects to identify a target letter ('L' or 'R'). The letter was presented in any one of three concentric rings. Participants fixated the centre and were cued about which ring (inner, middle or outer) the stimulus was to appear in. Performance was superior when the letter appeared in the ring that was cued. The spotlight model of visual attention would predict that performance would be best for stimuli presented nearest the centre fixation point, but that was not the finding. Instead it seems that visual attention, in this case, was deployed in an 'O' shaped pattern. In an earlier study, Farah (1989) used a cueing paradigm based on cells within 5x5 grids describing 'H' or 'T' letter shapes. Results suggested that subjects could spread their attention voluntarily over particular shapes (scene sub-regions). Other research suggests that such 'shaped attentional regions' usually correspond to the shapes of known visual objects rather than to randomly shaped areas. For reviews of shaped attentional allocation see Pylyshyn, (2003) and Eysenck & Keane (2000).

Actively ignoring display elements (for example, items in a peripheral part of display, or shapes overlaying a target stimulus) can result in that stimulus taking longer to identify later. For example Rock & Gutman (1981) required participants to attend to simple line drawings of objects such as a house. Participants were instructed to make aesthetic judgments about the figures that were drawn in a particular colour but to ignore other figures in the display. The target figures were overlaid with other, differently coloured, figures (for example, the outline of a tree). In a subsequent recall task participants recalled the attended figures (for example, house, figure) but unattended figures (for example, tree) were not recalled and were not recognised as having previously appeared in the 'aesthetic judgment' experiment. This phenomenon has been termed 'negative priming'.

Most visual attention laboratory research has tended to employ external 2D visual stimuli and has tended to ignore the individual's motivational state and goals. Visual attention can certainly be captured by strong stimuli in the environment (a non-inhibitable orienting response) or it can be directed under voluntary control (we are capable of 'paying attention'). A well-documented example of involuntary attentional capture is the 'weapons focus' effect reported in the forensic psychology literature. The presence of a weapon at a crime scene draws an eyewitness's attention away from other features of the scene such as the culprit's face (see Wells & Olson, 2003, for a recent review).

There seems to be reasonable consensus for a conception of visual attention as a system that is capable of switching between such wide-angle and telephoto modes as a function of task demands. The question of what drives such switching processes is a central concern of this paper.

Overviews of this extensive area can be found in Underwood (1993) and (more recently) in Pashler, *et al* (2001). The latter reviewers conclude that the last 10 years of research suggests that attention is controlled by top-down, cognitively driven processes to a much greater extent than was believed even 10 years ago. There are complex interactions between voluntarily adopted mental 'sets' and the attention-capturing attributes of stimuli. Research is also tending to show that the effects of practice upon attention and performance are less pronounced than was hitherto believed. Designers need to comprehensively target both top-down (cognitively-driven) and bottom-up (stimulus-driven) processes when they design support systems or information display systems.

Issue 1: Implications for design. This research suggests that designers must view attention as a dynamic process when designing displays for attention management. Visual display events that temporally precede the current display state might need to be considered. Why? There are several reasons. First, to use zoom lens metaphor, previous displays may have influenced the user's current attentional 'zoom setting'. Secondly previous display states may have resulted in the user shaping his/her current attentional region to a particular object's form. Thirdly, there may be negative priming effects from previous stimuli that the user has sought to actively ignore.

The findings on attentional shaping and resource allocation also suggest that gaze-based methods of acquiring information about a user's attention (Baudisch DeCarlo, Duchowski & Geisler, 2003) are probably of quite limited value. These methods use eye-tracking technology to monitor the user's gaze - this work is usually based on visual search paradigms. However visual search is one task among many that design-for-attention must address. A user might be engaged in any number of tasks other than visual search - vigilance tasks requiring the detection of an infrequent signal, integrating disparate pieces of information from one or more representations, and so on.

2.2 Issue 2: How can attention be measured?

The user's failure to apprehend pertinent information in a timely manner may in some circumstances have serious implications for the user's task whilst in others may simply slow activity down (Rensink, 2002). Designing systems to compensate and support users operating under the constraints imposed by limitations in attention involves understanding the basis of these limitations and being able to measure attention effectively.

Research into the psychophysical factors affecting attention (for example, Pylyshyn & Storm, 1988; Kahneman, Triesman & Gibbs, 1992; Rensink, 2000b; Hayhoe, 2003; Triesch, Ballard, Hayhoe & Sullivan, 2003) reveals that visual activity predominantly arises as a consequence of actively scanning the field of view in a task directed manner. The phenomenon of *inattentional blindness* (Hayhoe, 2003; Mack and Rock, 1998; Rensink, 2000b) or *inattentional amnesia* (Rensink, 2000b; Wolfe, 1999) demonstrates the selective nature of vision. Even though entities are clearly within view, if they are not central to the task in hand, they frequently remain unseen (Rensink, 2000b). By visually pursuing the selection of information about those entities central to the current task, other entities in the visual scene are actively ignored, no matter how conspicuous they may seem to the non-task-oriented viewer (Simons and Chabris, 1999).

A related phenomenon of *change blindness* further demonstrates aspects of natural vision which result in failure to notice changes to entities in the visual scene when these take place during a saccadic eye movement (Rensink, 2000b; Grimes, 1996 - see Simons, 2000, for a review). It appears that change can only be detected when the changing object is fixated (Rensink, 2000b; Rensink, O'Reagan & Clark, 2000).

This phenomenon has also been demonstrated in a virtual reality setting during activities in which the changed feature is central to the task in hand (Hayhoe, 2003). Participants asked to pick up blocks (pink or blue), and to place these in specific locations according to colour, failed to notice when the selected (virtual) object changed colour between initial selection and final placement. Most often the object was placed in the location appropriate for its colour during initial selection, rather than for the colour to which it had changed. Hayhoe (2003) argues that this demonstrates the 'micro-structure' of vision: that fixation of an object is not sufficient for apprehension of all the visual information associated with it. It would appear that during initial selection of the object, participants pay attention to colour, whilst during subsequent fixations they appear to be concerned with location in guiding the object to its resting place (Ballard, Hayhoe & Pelz, 1995).

In a separate study, Triesch, *et al* (2003) varied the nature of the task in relation to the object feature (this time, size) undergoing change. The number of changes noticed varied significantly between tasks, depending upon whether size of object was relevant to the task and, in particular, whether this attribute was relevant to carrying out the task both before *and* after the change took place. Contrastingly, a detailed analysis of direction of gaze during all activities revealed that patterns of gaze did *not* vary between tasks, despite the evidence that participants varied significantly in terms of the information they visually apprehended. Consequently, it appears the user's awareness of a given situation is highly selective and task-oriented and, in particular, it seems that direction of gaze alone is not a sufficient indicator of the precise nature of the information penetrating the user's awareness.

Issue 2: Implications for design. Limitations on inferring focus of attention from direction of gaze have implications for the development of systems which seek to enhance user-modelling through identifying focus of attention. Approaches to supporting the user in their task through enhanced user-modelling have attempted to contextualise eye-movements/gaze through various additional sources of information. These include users' interactions with software and devices, their prior interests and patterns of activity, and even using information from users' on-line calendars. Current studies seek to investigate the potential these combined sources of information offer for identifying focus of attention more accurately than through gaze alone (for example, Horvitz, Kadie, Paek & Havel, 2003).

Limitations on inferring focus of attention from direction of gaze also have implications for studies where measures of attention form part of the experimental design. Studies whose primary focus is the study of visual attention in relation to eye movement consider it a crucial aspect of experimental design, for the subject to perform activities contingent on attention, such as a discrimination task (for example, Sneider & Deubel, 2002). If direction of gaze is not necessarily synonymous with focus of attention, studies will need to validate focus of attention through further evidence, for example the transfer of information within a given task context, whereby evidence of the deployment of knowledge indicates the actual focus of attention.

2.3 Issue 3: How do graphical displays interact with attention?

The contribution in the early Twentieth Century by Gestalt psychology to our understanding of how visual displays are perceived, what structures we appear to naturally see in images, is well known. The various Gestalt laws of perceptual organization, which concern the visual forms that are particularly salient, can be found in any introductory cognitive psychology text (for ex-

ample, Eysenck & Keane, 2000). For instance, the law of similarity holds that when multiple stimuli are present there is a tendency to see things that are similar as forming a group, other things being equal. There has also been much recent interest in how the perception of graphical objects influences inferences we make. For example, Cleveland & McGill (1985) found that the accuracy of quantitative comparison judgments varied substantially with the particular visual properties being used to represent the quantities. One might assume that findings about the nature of perception and their direct impact on inferences are a good place to begin considering interface design for attention. However, in the field of cognitive science there have been substantial advances in our understanding on the nature of representations and their design to support inference and problem solving, which goes beyond the early perceptually focused work (for example, Glasgow, Narayanan & Chandrasekaran, 1995; Anderson, Cheng, & Haarslev, 2000; Blackwell, Marriot, & Shimojima, 2004).

In their seminal paper on the cognitive benefits of problem solving with diagrammatic versus sentential (textual or propositional) representations, Larkin & Simon (1987) demonstrated that the way an external representation encodes information is critical to how easy it will be for the user to find relevant information and recognize what inferences can be made. In diagrams, information that is needed for each inference step is often located in the same region of the diagram and, as a consequence, the user's attention can be focused on just one part of the diagram during each step of the problem solving process. In contrast, with sentential representations, such information is typically distributed throughout the representation, which requires the user to shift their attention many times to search out that information in different sentences, even within a single step of the problem solution. The *locational indexing* of information in graphical representations means that attentional resources do not need to be spent attempting to match symbolic labels: a necessary demand of sentential representations.

Building on this classic paper, the reasons for the benefits of graphical versus sentential representations have been investigated to some depth in cognitive science. Consider two examples. When presented in a graphical representation, certain information is often readily apparent as particular patterns or shapes. Such features may be emergent properties that seem to pop out of graphical representations, catching the user's attention, in a way that does not happen when exactly the same information is presented in, say, a tabular or list format. For instance, in a scatter plot, the degree and nature of the correlation between the two variables is visually apparent from the overall shape and orientation of the data points. One explanation for this phenomenon is given by Shimojima (1999), who claims that such *free rides* and *derivative meaning* occurs when the constraints on the relation between local and global structure in the representation and in the target domain, are satisfied by the semantic conventions of the representation. In the scatter plot the overall distribution

of data points is a consequence of the rules for drawing such plots but this is compatible with the manner in which relations between pairs of values are considered in assessing correlations. Such matching of constraints is absent in tables of data.

The potential benefits of graphical representations extend beyond the support they give to the search and recognition processes. Cheng (2004) showed that when users attended to the global structure of a diagram they identified symmetries and common configurations within the diagram. This knowledge was then used to adopt more effective problem solving strategies than had been adopted by users who merely focussed on the individual components or regions of a diagram.

Issue 3: Implications for design. Even from this small sampling of findings in cognitive science about the nature of graphical displays, there are clearly some useful insights for the attention designer. First, basic perception may be one aspect to be considered when designing a display that interacts with attention, but the impact of the processes of search, recognition and strategy selection will also have substantial effects. A poor representation can place considerable demands on attention, at worst requiring it to be deployed merely to find relevant information. Secondly, graphical representations are often good because they happen to gather together at the same location information that is needed for each inference. Such location indexing of information reduces the number of shifts of attention at each step in problem solving. Thirdly, information can be made more or less salient by the manner in which it is encoded in a representation. Consideration should be given to how the structure of target relations are to be captured by the relations among the representing elements in a display. Designed well, a display can make relevant information appear to pop out in as highly salient feature. Designed poorly, a display may need special devices (flashing icons, colour changes) to deliberately attract attention to particular information, rather than having that information naturally emerge. Fourthly, graphical representations may be attended to at different levels: locally focussing on particular elements or regions, or more globally as a whole. Different information can be obtained at the different levels and this will affect the user's understanding of the domain and even the particular task strategies they adopt. Hence, the designer should take care to understand what is the appropriate level for different stages in a task and that the mechanisms used to direct attention is aimed at the appropriate level.

2.4 Issue 4: How do knowledge, performance and attention interact?

Reason (1990) proposed a generic error modelling system (GEMS) framework for conceptualising the range of error types that humans exhibit when engaged

in complex tasks such as controlling real-time, dynamic, multivariate systems. Three basic error types are proposed - skill-based slips, rule-based mistakes and knowledge-based mistakes.

The difference between the three levels of processing has been nicely encapsulated in an example provided by Felciano (1995). The routine task of opening a door is usually a *skill-based* activity performed unconsciously without the need to ‘think about’ it. If, however, the door refuses to open when we turn the handle then we switch to *rule-based* processing. We reason that the door might be locked or that maybe we should try pulling the door rather than pushing it. If such rule-based responses fail, we then move to the *knowledge-based* level of processing. We must now call upon previous experience to troubleshoot the problem with the jammed door. We might entertain the possibility that the door is being held shut by someone on the other side, has had its lock changed, or has become physically obstructed. The important point is that the depth of processing changes within the task over its time course.

Rules and knowledge-based processing are both conscious processes. Knowledge-based processing requires a mental model of the problem, and analysis of more abstract relations between structure and function. Experts differ from novices crucially in terms of rule-based and knowledge-based problem solving. Experts tend to have a larger rule-base and they also represent their knowledge at a more abstract level. For example seminal work by Chase & Simon (1973) in the domain of chess showed that chess experts ‘chunk’ board configurations in much larger units than novices. This enables them to recall board states more accurately than novices and to readily distinguish chess piece configurations from real games from random chess piece configurations.

The process of shifting from rule-based to knowledge-based levels of problem solving can be due to the detection of countersigns. These are present when inputs indicate that a more general rule is inapplicable (Reason, 1990). In the door example used above, the fact that the door doesn’t open (despite unlocking it and trying both pushing and pulling) acts as a countersign and prompts a shift from rule-based to knowledge-based levels of processing.

In some contexts (for example, safety critical process control) there are some aspects of expertise that can be disadvantageous. Reason (1990) suggests that skilled individuals are more likely than novices to exhibit ‘strong but wrong’ response errors at the skill-based and rule-based levels of performance. The application of a strong-but-wrong rule was exemplified in the Oyster Creek nuclear power plant accident of 1979. A water level indicator showed a reading that had become decoupled from the actual water level in a tank. Normal water levels were indicated but the actual level in the tank was too low. This situation had never occurred before and took a long time to diagnose - meanwhile inappropriate and catastrophic countermeasures had already been instigated

on the assumption that the indicator reading was accurate.

The attention and change blindness literatures have major implications for several aspects of Reason's model, particularly those concerned with the need to switch from skill-based to rule-based levels when error states in a system are detected. For example, change blindness may result in failure to detect 'countersign' information. Reason (1990) argues that the difficulty of detecting countersigns is further compounded by 'information overload' of the cognitive system by a high volume of information. Change blindness phenomena and top-down/bottom-up interactions in attention provide a useful basis for operationalising Reason's rather vague concept of 'information overload' in his GEMS model.

Issue 4: Implications for Design. Designers should be aware that the ostensibly similar behaviours of a user can, in fact, have very different processing states underlying them at different points in time. These processing states differ in their 'depth' and can range from automatic (skill-based) processing modes that consume few attentional resources through to states that demand much higher degrees of conscious attention (rule-based and knowledge-based responding).

Another implication is that display designs which make countersigns salient are clearly desirable in safety-critical contexts where it is crucial to find out quickly whether an indicated state is due to an instrumentation failure or real system error. An example of good practice in this respect is provided, as is so often the case, by aviation instrumentation design. In many aircraft cockpits it is possible to test that the instruments and warning lights are working via an 'instrument test' button. This provides a means of testing that the display warning light globes, and so on, are working independently of the aircraft control and engine systems.

2.5 Issue 5: What is working memory?

The nature of working memory is a major and active area of research in psychology and cognitive science and should be central to thinking about designing for attention. An excellent overview of what is known, and what is far from certain, about working memory is Miyake & Shah (1999). They posed eight questions about the form, mechanisms, functions and implementation of working memory to leading researchers in the area. The answers concerning the relation of working memory to attention and consciousness are especially pertinent here. Some consider working memory, attention and consciousness as largely synonymous constructs. Others makes distinctions between them or consider them as overlapping or being in subset relations. What is clear is that understanding and accounting for working memory is important.

It is widely accepted that humans exhibit a limited working memory capacity. Miller's (1956) classic paper put the number of chunks that can be stored in working memory at seven plus or minus two chunks. A chunk consists of items of information that are strongly associated with each other and weakly associated with items of information not in the chunk. A chunk may be considered, in general terms, as a concept. Since Miller's paper, the actual capacity of working memory has been challenged, particularly in circumstances that are not ideal. Cowan's (2001) review of findings puts the realistic capacity at four chunks.

There are many explanations for the limited capacity of working memory, including limited supply of activation to spread over concepts, finite processing speed, decay of chunks, interference among concepts and others. Rather than providing details of all of the different theories, what appears to be useful for those interested in designing for attention is a model that operationally reflects what is known about the relation of working memory and attention. A good candidate for this is the one often adopted in classical models of the human cognitive architecture (for example, Stillings, Weisler, Chase, Feinstein, Garfield & Rissland, 1995). Memory is considered to consist of propositions in a network where more or less closely associated propositions are linked. Working memory can be considered to be that part of long-term memory that is currently active and readily available for processing. When particular propositions are being used, for example during problem solving, they are active and their activation spreads to the others that are associated with them. The likelihood that those associated propositions can then be recalled from memory is thus increased. The propositions whose activation rises above threshold can be considered to be in working memory. As problem solving progresses, new information is obtained, which makes other propositions active and available to be processed. The content of working memory is substantially determined by the context of the current task in which the user is engaged: it provides cues, pieces of information, to retrieve related chunks. Only part of the contents of working memory will be subject to conscious attention and engaged with immediate task activity. The activation of chunks in working memory will fade and no longer be available to attention, unless they are reactivated by rehearsal or the activation of related chunks brings them back.

Issue 5: Implications for design. There are various implications for attention designers. The limited capacity of working memory places a constraint on the amount of information to which a user can reasonably be expected to attend. Simply, if system design is not sensitive to this limit then mechanisms used to make the user aware of task relevant information are likely to perform poorly, because they overwhelm the user. If an estimate of the number of chunks required to do a task can be made, this may indicate the cognitive load placed on the user of an interface, and hence may indicate whether attention is substantially deployed and should not be further loaded. This limitation can be

overcome by grouping information and, because of the high association of elements within a chunk, once it is attended to, the constituent elements are likely to be easily retrieved. The system designer should be mindful to both support and exploit the user's natural propensity to chunk information. Information in working memory quickly decays, or gets displaced by other information. So the system designer should consider whether appropriate support should be available to allow the user to easily refresh their working memory with information that is needed but not attended to constantly. A designer may wish the user to attend to certain information and do so by making it explicit by some means. However, an alternative strategy could be employed, knowing that memory is associative and that spreading activation determines, in part, what ends up in working memory. So by making elements related to the target element more active it may be possible to induce the user to recall the target element without explicitly displaying it. This technique might be useful in a situation where the associated elements are already present in a display, and could be highlighted in some way, but where it is undesirable to add yet more items to an already crowded display. A reminder of situational or goal information may be as effective as deliberate attempts to (re)focus the user's attention on task relevant information.

2.6 Issue 6: How does doing two things at once affect attention?

Performance on divided attention tasks often improves with practice. Evidence is provided by divided-attention and dual-task research. In a typical dual-task paradigm, a participant might be required to verbally indicate whether a tone is low or high in pitch while concurrently indicating (via keyboard key selections) which letter appears on a computer screen. The detrimental effects of dual-task interference upon performance diminish with practice but not to the point where performance levels are identical to those observed when the tasks are performed singly (see Pashler et al., 2001).

Factors that determine how well two tasks can be performed together are dissimilarity (for example, walking and talking, driving and listening to music, and so on). It seems that tasks interfere more when they are similar in terms of sensory modality (auditory, visual), when they involve similar information processing stages (perceptual, processing, production) and when they involve similar modes of representation in memory (visual, verbal). The degree to which the response modes are similar is also important - a motor response plus a vocal response is easier to perform than two motor responses for example (see Eysenck & Keane, 2000 for a recent review).

Issue 6: Implications for Design

The implications for design are that tasks should be presented serially rather than concurrently whenever possible. If dual-tasks cannot be avoided then the tasks should be as dissimilar as possible in terms of their characteristics and response requirements in order to minimise interference effects. Some performance improvement with practice can be expected but performance levels on each of the two tasks will not reach the level that would be observed if each task was practiced separately.

2.7 Issue 7: What are the potential effects of introducing artificial feedback in systems designed to monitor the user's attention?

Some 'attention-aware' systems incorporate eye-trackers and modify the display on the basis of gaze direction (for example, Baudisch, *et al*, 2003; Zhai, 2003). In some respects such a configuration is similar to biofeedback, where feedback of information about a body system¹ is provided to a person in the form of an information display with which s/he learns to acquire control over that system. It has been demonstrated that the introduction of an (artificial) feedback loop can, in some circumstances, cause variables that are usually highly correlated to become decoupled. For example Cox & Matyas (1983) conducted a study in which participants underwent a training regime designed to increase the strength of isometric arm extension. In one condition subjects were provided, during training, with feedback of information about muscle motor unit recruitment (EMG) in the triceps muscle. Subjects in another group were provided with feedback from a force transducer which measured the strength of their contraction in terms of downward force at the wrist. A control group received neither kind of feedback. EMG and force measures were recorded on all trials for subjects in all three groups. The results showed that EMG and force measures, which are usually highly correlated, became somewhat decoupled in subjects in the EMG biofeedback condition but not in subjects in the force feedback condition. In other words subjects in the EMG feedback condition learned to control the EMG feedback display by recruiting muscle motor units in ways that did not produce changes in target force.

Issue 7: Implications for design Designers must be careful when designing displays that are 'driven' by users' psychophysiological input. Display events that are associated with one pattern of behaviour at early stages can result from quite different user responses after extended periods due to learned strategies. In the context of a gaze-direction sensing attention-aware system the user might develop an 'unnatural' strategy such as producing a sequence of rapid

¹ Such as autonomic nervous system, for example, heart rate; motor responses, for example, muscle contraction/relaxation; and so on.

saccades in order to refresh parts of the display.

2.8 Issue 8: Do attentional processes differ across tasks (e.g. vigilance vs diagnosis)?

Monitoring the state of variables in a complex system, such as nuclear power plant's control panel for example, entails stimulus-driven (bottom-up) processing to a greater extent than a task like troubleshooting a fault electrical circuit.

In tasks such as *monitoring*, subjects are typically required to keep track of displays and gauges. If they are well-designed, such displays make error states salient to the point that bottom-up, stimulus-driven capture of attention is inevitable. Recent evidence (for example, Pashler *et al*, 2001) suggests that novel objects capture attention to a greater degree than other manipulations such as abrupt changes in luminance².

In contrast, a problem-solving task such as *diagnosing* the cause of a fault entails more user-system interaction. Attention interacts with mode-of-execution. The process is primarily cognitively-driven. In terms of Reason's (1990) GEMS model, the person doing the diagnosis must first notice an anomalous state and would typically attempt a rule-based solution in the first instance. This could be followed by a sequence of cycling to and from a knowledge-based solution level if necessary. Reason (1990) states that a key feature of the GEMS model is that "human beings are strongly biased to search for and find a prepackaged solution at the rule-based level before resorting to the far more effortful knowledge-based level, even where the latter is demanded at the outset." (p. 65). In terms of the 'door' example cited earlier, this might take the form of a person persevering with a 'door is locked' (rule-based) hypothesis in which a person assumes that s/he has the wrong key when, in fact, it is the correct key.

Issue 8: Implications for Design Designers need to conceptualise tasks in terms of their different cognitive demands on the user. Adaptive systems need to differentiate between tasks in order to choose an appropriately 'strong' (attention grabbing) intervention. Adaptive attention-aware systems probably need to intervene more proactively on tasks in which the user is engaged at knowledge-based levels of cognitive processing. Detecting and preventing 'strong but wrong' responses is desirable. In more stimulus-driven situations

² These findings are commensurate with the research cited earlier which shows that recognisable objects can prime attentional 'regions' within a scene and allow the categorical identity and spatial location attentional subsystems to be used concurrently and without extra processing cost.

such as vigilance tasks, the system can be more subtle and employ novel stimuli to capture and direct the user's attention.

3 Discussion

We have posed eight questions highlighting issues which might usefully inform a theoretical framework for designing for attention, and drawn out the implications of each for the development of attention-aware systems. Consideration of these implications has enabled us to identify five themes underlying the issues we have discussed. These themes are depicted in Table 1; the issues which address each theme are also identified.

The first theme concerns the need to understand attention as a dynamic process; the time course of events interact with attentional mechanisms and are therefore pertinent to attention-aware interface design. This time course is affected by cognitive factors such as perceptual set, contextual awareness (preceding and current activities) and task level. The focus of attention is shaped by preceding events and activities engaging the user, including visual aspects of their display, and by the current visual context in which elements compete for the users attention. Context acts to predispose the user to notice some aspects of the visual scene and to actively ignore others. Added to this, limitations on working memory constrain the number of items users can be expected to attend to, so steps should be taken either to avoid overloading the user, or to provide information to prompt the user to remember current tasks that may have slipped working memory.

The second theme concerns aspects relating to the users task(s) and attention. Primarily, the users task influences cognitively based information seeking processes that are part and parcel of attention, affecting what gets noticed and what gets ignored. However, the nature of the tasks being undertaken, their quantity and the users mode of interaction, impact heavily on attentional processes, with implications for the users ability to carry out activities effectively. Understanding the user's task can guide the appropriate level of generality or detail in the graphical forms then used to support apprehension of information.

The third theme concerns levels of processing, in particular, the task driven (top-down) nature of the users attentional processes means these will interact with depth of processing changes within the task over its time course. In particular, knowledge based and rule based reasoning affect the users ability to detect countersigns. Additionally, the nature of the external representation will influence task strategy and recognition, and hence impact the depth of processing.

The fourth theme concerns the nature of the visual display. The tendency of users to ignore crucial information where its role in a problem solving episode is not appreciated requires novel methods to highlight pertinent information that will otherwise be ignored. We can exploit our understanding of cognitive factors to contrive display environments that may make it more likely that relevant elements will be attended. For example, by avoiding cognitive overload and streamlining displays to essential features. Similarly, attention can be aided by appropriate use of representations to support the extrapolation of pertinent information. Furthermore, there are some indications that information display designers should consider using novel and recognisable visual objects for attention capture rather than abstract shapes or information channels such as colour. Designers might also consider using displays for slowly changing data that decouple the display from the information it depicts, by using periodic step changes in the display in place of real-time tracking, as these are likely to be more readily noticed.

The fifth and final theme concerns the validity of methods for determining attention. Gaze is not synonymous with attention; studies of attention require detection through the use of discrimination tasks or application of knowledge indicating apprehension of visual information. Furthermore, gaze applies only to visual search processes, not to monitoring, nor to information integration activities associated with attention. Lastly, there is a serious concern about the effects of biofeedback which may act to change the very nature of attentional behaviour, and the need to anticipate this in designing systems which exploit attentional cues such as gaze.

In essence, the view presented here is that research into the design of attention-aware systems needs to engage with the deeper understanding that exists regarding the nature of attention, as reviewed here. The change blindness and visual attention findings are not yet integrated sufficiently to allow prescriptive design recommendations to be made. To quote Rensink (2002): "If a display is to be designed so that the observer can interact with it optimally, it is essential to understand how attention operates. Unfortunately, there is still much about attention that is not known..." (p. 67). It is also clear that the very nature of attention makes it a difficult phenomenon to study. Eye gaze data must be validated by triangulation with additional behavioural protocol data.

It also seems clear that design decisions, made with the aim of managing the user's attention, should build on existing, well-established principles for the design of effective representations. For example, by exploiting graphical free rides, and through the judicious selection of representational systems.

References

- Allport, A. (1993) Visual attention. In M.I. Posner (Ed.) *Foundations of cognitive science*. Cambridge, MA: MIT Press.
- Anderson, M., Cheng, P.C.-H. & Haarslev, V. (2000) *Theory and Application of Diagrams: First International Conference, Diagrams 2000*. Berlin: Springer.
- Ballard, D., Hayhoe M. & Pelz, J. (1995) Memory representations in natural tasks, *Journal of Cognitive Neuroscience*, 7, 66-80.
- Baudisch, P., DeCarlo, D., Duchowski, A., & Geisler, B. (2003) Focusing on the Essential: Considering Attention in Display Design. *Communications of the ACM*, 46(3), 60-66.
- Blackwell, A., Marriot, K. & Shimojima A. (2004) (Eds.), *Proceedings of the 3rd International Conference on Diagrams, 2004*. Berlin: Springer-Verlag.
- Broadbent, D.E. (1958) *Perception and communication*. London: Pergamon.
- Broadbent, D.E. (1971) *Decision and stress*. London: Academic Press.
- Chase, W.C., & Simon H. (1973). Perception in chess. *Cognitive Psychology*, 4, 55-81
- Cheng, P. C.-H. (2004). Why diagrams are (sometimes) six times easier than words: benefits beyond locational indexing. In A. Blackwell, K. Marriot & A. Shimojima (Eds.), *Diagrammatic Representation and Inference: Third International Conference, Diagrams 2004*, (pp. 242-254). Berlin: Springer-Verlag.
- Cherry, C. (1957) *On human communication*. (2nd ed.) Cambridge, MA: MIT Press.
- Cleveland, W. S. & McGill, R. (1985) Graphical perception and graphical methods for analysing scientific data. *Science*, 229, 828-833.
- Cowan, N. (2001) The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behaviour and Brain Sciences*, 24(1), 87-114.
- Cox. R.J., & Matyas, T.A. (1983) Myoelectric and force feedback in the facilitation of isometric strength training: a controlled comparison. *Psychophysiology*, 20(1), 35-44.
- Eriksen, C. W., & St. James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception and Psychophysics*, 40, 225-240.
- Eysenck, M. & Keane, T. (2000) *Cognitive Psychology: A student's handbook*. 4th Edition. New York: Psychology Press.
- Farah, M.J. (1989) Mechanisms of imagery-perception interaction. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 203-211.
- Felciano, R.M. (1995) Human Error: Designing for Error in Medical Information Systems. Web URL
<http://www.smi.stanford.edu/people/felciano/research/humanerror/>
Accessed 13 Dec 2004.
- Glasgow, J. Narayanan, N. H. & Chandrasekaran, B. (1995) *Diagrammatic Reasoning: Cognitive and Computational Perspectives*. Menlo Park, CA: AAAI Press.

- Goel, V. & Pirolli, P. (1992) The structure of design problem spaces. *Cognitive Science*, 16, 395-429.
- Grimes, J. (1996) On the failure to detect changes in scenes across saccades. In K.A. Atkins, (Ed.) *Perception, Vancouver Studies in Cognitive Science 5*, OUP, Oxford.
- Hayhoe, M.M. (2003) What guides attentional selection in natural environments? *Abstract Proceedings of Fifth Workshop on Active Vision*, University of Sussex, UK.
- Horvitz, E. Kadie, C. Paek T. & Havel, D. (2003) Models of Attention in Computing and Communication: From principles to Applications. *Communications of the ACM*, 46(3), 52-9.
- Juola, J.F., Bowhuis, D.G., Cooper, E.E., and Warner, C.B. (1991): Control of attention around the fovea. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 315-330.
- Kahneman, D. Triesman, A. & Gibbs, B. (1992) The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24, 175-219.
- LaBerge, D. (1983). Spatial extent of attention to letters and words. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 371-379.
- Larkin, J. H. & Simon, H. A. (1987) Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65-99.
- Mack A. & Rock, I. (1998) *Inattentive blindness*. Cambridge, MA: MIT Press.
- Miller, G. A. (1956) The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. *Psychological Review*, 63, 81-97.
- Miyake, A. & Shah, P. (1999) (Eds.) *Models of working memory: Mechanisms of active maintenance and executive control*. New York: Cambridge University Press.
- Pashler, H. Johnstone, J.C. & Ruthruff, E. (2001) Attention and performance. *Annual Review of Psychology*, 52, 629-51.
- Pylyshyn, Z.W. (2003) *Seeing and visualizing: It's not what you think*. Cambridge, MA: MIT Press.
- Pylyshyn, Z.W. & Storm, R.W. (1988) Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, 3, 179-197.
- Reason, J. (1990) *Human error*. Cambridge, UK: Cambridge University Press.
- Rensink, R.A. (2000a) The dynamic representation of scenes. *Visual Cognition* 7, 17-42.
- Rensink, R.A. (2000b). Seeing, sensing and scrutinizing. *Vision Research*, 40, 1469-1487,
- Rensink, R.A. (2002) Internal vs. external information in visual perception. *Proceedings of the 2nd ACM International Symposium on Smart Graphics* (pp, 63-70).
- Rensink, R. A. O'Regan, J.K. & Clark, J.J. (2000) On the failure to detect changes in scenes across brief interruptions, *Visual Cognition*, 7, 127-145.
- Rock, I., & Gutman, D. (1981). The effect of inattention on form perception. *Journal of Experimental Psychology-Human Perception and Performance*, 7(2), 275-285

- Schneider, W.X. & Deubel, H. (2002). Selection-for-perception and selection-for-spatial-motor-action are coupled by visual attention: A review of recent findings and new evidence from stimulus-driven saccade control. In W. Prinz & B. Hommel *Attention and Performance XIX: Common Mechanisms in Perception and Action*. Oxford: Oxford University Press, (pp. 609-627).
- Shimojima, A. (1999). Derivative meaning in graphical representations. Proceedings of 1999 IEEE Symposium on Visual Languages, (pp. 212-219).
- Simon, H. A. (1981) *Sciences of the Artificial*, 2nd ed. Cambridge, MA: MIT Press.
- Simons D.J. (2000) Current approaches to change blindness. *Visual Cognition*, 7, 1-45.
- Simons D.J. & Chabris, C.F. (1999) Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28, 1059-1074.
- Stillings, N. A., Weisler, S. E., Chase, C. H., Feinstein, M. H., Garfield, J. L., Rissland, E. L. (1995) *Cognitive Science: An Introduction* (2nd ed). Cambridge, MA: MIT press.
- Treisman, A. (1986) Features and objects in visual processing. *Scientific American*, 254, 114-124.
- Triesch, J. Ballard, D.H. Hayhoe M.M. & Sullivan, B.T. (2003) What you see is what you need. *Journal of Vision*, 3, 86-94.
- Underwood G. (1993) (Ed). *The Psychology of Attention (Vol 1)*. Aldershot: Elgar.
- Wells, G.L. & Olson, E.A. (2003) Eyewitness testimony. *Annual Review of Psychology* 54, 277-295.
- Wolfe, J.M. (1999) Inattention blindness. In V. Coltheart (Ed.) *Fleeting Memories*, Cambridge, MA: MIT Press.
- Zhai, S. (2003) What's in the Eyes for Attentive Input? *Communications of the ACM*, 46(3), 34-39.

Themes	Issues
Dynamic Nature of Attention	1, 4, 5.
Tasks (number, nature and variety)	2, 6, 8.
Level of Processing	1, 3, 4.
Nature of the Display	1, 2, 3.
Validity	2, 7.

Table 1: Themes underlying issues