HUMAN-COMPUTER INTERACTION, 2004, Volume 19, pp. 389–422

Unseen and Unaware: Implications of Recent Research on Failures of Visual Awareness for Human–Computer Interface Design

D. Alexander Varakin and **Daniel T. Levin**
Vanderbilt University

Roger Fidler
Kent State University

ABSTRACT

Because computers often rely on visual displays as a way to convey information to a user, recent research suggesting that people have detailed awareness of only a small subset of the visual environment has important implications for human–computer interface design. Equally important to basic limits of awareness is the fact that people often over-predict what they will see and become aware of. Together, basic failures of awareness and people's failure to intuitively understand

D. Alexander Varakin is a graduate student with an interest in visual memory and awareness; he is working toward a Ph.D. in cognitive psychology in the Department of Psychology and Human Development of Peabody College at Vanderbilt University. **Daniel Levin** is a psychologist with an interest in the interface between perception and cognition; he is an associate professor of psychology in the Department of Psychology and Human Development of Peabody College at Vanderbilt University. **Roger Fidler** is an educator and researcher with an interest in information design and digital publishing; he is a professor in the School of Journalism and Mass Communication at Kent State University.

CONTENTS

them may account for situations where computer users fail to obtain critical information from a display even when the designer intended to make the information highly visible and easy to apprehend. To minimize the deleterious effects of failures of awareness, it is important for users and especially designers to be mindful of the circumscribed nature of visual awareness. In this article, we review basic and applied research documenting failures of visual awareness and the related metacognitive failure and then discuss misplaced beliefs that could accentuate both in the context of the human–computer interface.

1. INTRODUCTION

One of the primary ways computers convey information to human users is through the presentation of visual information on a display. Thus, the simple act of seeing is often the user's first step in gaining information from a system. Usually, the appropriate courses of action depend on both the user's task and the state of the computer system. It is therefore important for computer displays to be designed such that users can accurately apprehend information about what a computer is capable of doing at any given time. The consequences of incomplete or inaccurate understanding of a computer display's contents range from inefficiency in an office setting, when employees miss important links on a company's Web site, to major catastrophes, when a pilot crashes a passenger jet due to misconstruing a mode setting. An obvious conclusion is that research on visual awareness has important implications for the human–computer interface. What is not obvious is how visual awareness can fail.

Recently, scientists studying how people acquire and use visual information have documented a series of increasingly surprising lapses of visual awareness. Of course, psychologists have long studied limits in the ability to process information but in the area of visual information processing and perception, a working assumption of much past research has been that humans are relatively cognizant of the visual environment. This assumption, reasonable as it might seem, was based more on intuition than empirical observation and several recent demonstrations have proven it to be fundamentally flawed. Two phenomena in particular have convinced many, if not all, vision scientists of the highly circumscribed nature of visual awareness. One of these phenomena, called *Inattentional Blindness* (IB), occurs when observers fail to notice the presence of unattended stimuli, even when these stimuli are presented within an observer's field of view and occupy the same location in space as attended and consciously perceived stimuli. The second phenomena is *Change Blindness* (CB), which is the difficulty people have detecting changes to visual stimuli that occur across views, sometimes even when the changing object is being attended.

Failures of visual awareness such as CB and IB demonstrate the extent to which people can be unaware of visual information and may help explain some of the difficulties human–computer interface designers have faced in their attempts to create displays from which users can reliably obtain information. For example, modern graphical user interfaces are designed to make information easily accessible. Many tasks, such as searching for a specific link on a Web site or initiating an application require little more from a user than finding and clicking on an appropriate place on the monitor. Nonetheless, such tasks are often completed only after the user has been explicitly told where to click, even in cases where the relevant area of the monitor is otherwise quite salient.

Part of what makes failures of visual awareness so intriguing from a psychological point of view, and so troublesome from an interface-design standpoint, is their counterintuitive nature. People are not only unaware of great amounts of visual information but they are also unaware of the extent to which they may be unaware of visual information. This metacognitive failure has become a topic of research in its own right and might further explain why interfaces designed to make information easily accessible can turn out to make it paradoxically inaccessible.

In this article, we review the psychological literature documenting IB and CB and the related metacognitive failure and then discuss how this research might inform the theory and practice of computer interface design.

2. FAILURES OF VISUAL AWARENESS

2.1. Inattentional Blindness

Psychologists interested in attention have generally argued that focusing attention on one thing reduces the degree to which other, unattended things are processed. Early research focused on auditory tasks in which participants were asked to pay attention to one channel of information (e.g., a speech stream presented via headphones to the participant's left ear; see Cherry, 1953), although ignoring another channel of information (in this case, the speech stream presented to the right ear). Generally, when participants do this, they have considerable difficulty reporting anything about the meaning of the information in the unattended ear. In the earliest demonstration of IB, Neisser and Becklen (1975) decided to set up a visual analog to these auditory attention experiments.

In these original demonstrations, participants viewed two superimposed videotaped events with instructions to monitor one of the events. In one case, Neisser and Becklen (1975) showed participants a video of two people playing a hand-slapping game superimposed over a video of three people passing a basketball back and forth. The participants had to count the number of either hand slaps or the passes. After several trials of monitoring one event, critical trials were introduced on which unexpected events occurred in the unattended video (e.g., the people would stop playing the hand-slapping game momentarily to shake hands). Usually, participants failed to detect the unexpected events, as long as they occurred in the unattended video. In another version of the study, participants were instructed to attend to basketball passes while a woman carrying an umbrella walked through the scene. Again, despite the salient and bizarre nature of the event, many subjects failed to detect it at all.

In Neisser and Becklen's (1975) original demonstrations, the attended and unattended events were both partially transparent due to the superimposition. However, recent work has confirmed that transparency is not a necessary condition for inducing IB. Simons and Chabris (1999) had participants watch a video of two teams (one in white shirts and one in black) chaotically moving about the screen while passing basketballs back and forth. The observers' task was to count the number of passes made by members of one of the two teams. Near the end of the video, a person wearing a gorilla suit walked into the middle of the two teams, turned and faced the camera, beat her chest and then proceeded to walk out of view (see Figure 1). In an experiment using similar videos, Wayand and Levin (2001) had participants perform virtually the same task and, midway through the video, a woman walked into the camera's view and scraped her fingernails down a chalkboard creat-

Figure 1. **Still from video used by Simons and Chabris (1999). Reprinted from "Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events" by Simons & Chabris (1999), *Perception, 28*(9), 1059–1074. Copyright 1999 by Pion Ltd. Reprinted with permission.**

ing the well-known and highly noxious screech noise. In both Simons and Chabris's and Wayand and Levin's experiments, about half of the observers failed to notice the unexpected stimulus. What makes these failures surprising is that the undetected stimulus was present for several seconds, occupied virtually the same space as attended stimuli, and was bizarre or annoying.

Other recent demonstrations of IB have used more tightly controlled stimuli to demonstrate equally compelling visual failures. In Mack and Rock's (1998) inattention paradigm, participants were again instructed to perform an attentionally demanding cover task and were later asked questions about unexpected and unattended stimuli. In this case, participants were asked to judge whether the horizontal or vertical arm of a briefly presented cross was longer (see Figure 2). Participants completed two to three trials of the cover task before the critical inattention trial. On the inattention trial, an additional stimulus appeared on the screen in one of the quadrants of the cross (see Figure 2). Immediately after these stimuli were removed from the screen, participants were asked about various attributes of the unexpected stimulus (e.g., color, shape, etc.). The original purpose of this line of research was to investigate how inattention affected the perception of these various attributes but Mack and Rock stumbled onto something they found to be even more interesting. In a series of experiments using variants of the inattention paradigm, Mack and Rock consistently found that a large proportion of participants

Figure 2. Diagram of the critical inattention trial similar to that used by Mack and Rock (1998). Circle (A) shows the fixation cross, circle (B) shows the stimulus with the unexpected object in the lower left quadrant. In trials preceding the inattention trial, the lower left quadrant was empty. Circle (C) show the poststimulus mask.

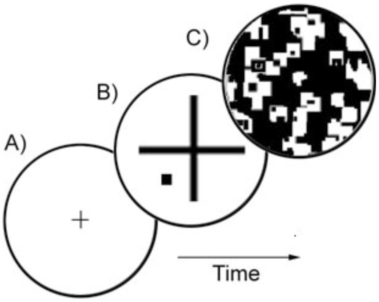

were unaware of anything other than the target cross and could report nothing about the unattended stimulus' attributes. Mack and Rock therefore stopped asking participants about various aspects of the unattended stimuli and instead asked if participants had noticed anything other than the cross at all. Changing the question did not change the basic finding; a large proportion of people remained inattentionally blind.

A number of other recent demonstrations of IB have combined the dynamic stimuli of the Neisser and Becklen (1975) paradigm with the more controlled stimuli characteristic of the Mack and Rock (1998) work. For example, Most et al. (2001) had participants view dynamic displays of four white and four black shapes (T's and L's or circles and squares) independently moving on a computer screen. Periodically, these shapes would bounce off the edge of the display. The participants' task was to count the number of times shapes of a designated color (either black or white) touched the edge of the display and to ignore shapes of the other color. As in Mack and Rock's inattention paradigm, participants performed a few trials before a critical trial on which another unexpected object (a black, white, or red cross) moved across the screen, taking about 5 sec to do so. Immediately after the critical trial, participants were asked if any other objects were present on the screen and a large proportion failed to have any recollection of the unexpected stimulus, so long as it was not the same color as the shapes to which the observer was attending. Even when the unexpected stimulus was completely novel in shape and color (e.g., the attended objects were white circles, ignored objects were black circles and the unexpected object was a red cross) about 30% of observers failed to notice its presence.

So, several demonstrations using static images, dynamic displays, and videotaped events have shown that observers often fail to explicitly detect unexpected stimuli. Although this work has clear precedence in a long history of research exploring attentional limits, the specific failures it documents are, in many important ways, novel. Most important, these experiments demonstrate not only that people fail to comprehend unattended stimuli; they can also fail to register their existence entirely. In the earlier auditory dual-channel experiments that inspired the current work, participants were aware of the *presence* of the unattended message and were even able to report changes to some of that message's basic qualities (e.g., detect a change between a male and female speaker; Cherry, 1953). Research on IB therefore suggests that the cost of attending to some stimulus can be quite high. However, IB, in and of itself, leaves open the possibility that awareness of attended items might be fairly complete. The next section describes research documenting the phenomena of CB, which suggests that observers may not even be completely aware of attended stimuli.

2.2. Change Blindness

Often, when something in the visual world suddenly changes, viewers notice it because the change produces a sudden perceptual transient that calls attention to itself. For example, if a lamp suddenly changed position, or a sofa suddenly changed into a chair, most people's attention would be attracted to a sudden apparent movement, or "pop" in their environment. Although the visual system appears to be hardwired to detect these transients, they are not always present when something changes. When transients do not occur, or when they are somehow obscured, visual changes can be remarkably difficult to detect and the result is the phenomenon of CB (see Rensink, 2002; Simons, 2000; Simons & Levin, 1997, for reviews).

Many early demonstrations of CB relied on the natural disruption of visual input that occurs each time a large eye-movement (a *saccade*) is made. While a person's eyes are fixated, the light reaching the retina creates a pattern stable enough to allow for coherent perception of the world. But several times a second the eyes move and the stable image on the retina is "smeared." For the 30 to 50 msec it takes to complete the eye movement, very little useful visual information is obtained. Thus, if a change in the visual array occurs while the eyes are moving, the smearing on the retina will obscure the transient associated with the change and the change may be missed. For example, McConkie and Zola (1979) had participants read text in alternating cases (e.g., LiKe ThIs) while a computer attached to an eye tracker recorded their eye movements. This eye-tracking system was able to detect the onset of a saccade and to make a change to the display before the saccade ended. Accordingly, each time the reader made a saccade, every single letter in the text changed case (e.g., from LiKe ThIs to lIkE tHiS). Under these conditions, reading times remained unaffected and readers rarely noticed that the display was changing. Grimes (1996) demonstrated that even changes to realistic scenes are often missed when the change is made during an eye movement. Grimes had participants view pictures of real scenes, with instructions to study for a recognition test. As the participants viewed the pictures, large and sometimes ridiculous changes were made to them. Grimes found that as long as the change was completed while the eyes were in motion, the changes were often missed. For example, in one scene, 50% of observers failed to notice that two people switched heads.

Later research demonstrated that saccades are not the only thing capable of masking a change. Almost anything that occludes the change-induced motion signal or makes it less salient can induce CB. For example, changes are often missed when an observer's eyes remain still and the picture on the monitor moves (Blackmore, Brelstaff, Nelson, & Troscianko, 1995). In this case, the movement of the picture on the monitor mimics the retinal displacement caused by an eye movement. However, CB can be induced even without this kind of large-scale image displacement. Blackmore et al. (1995) also discovered that blank screens inserted in between presentations of two different versions of a picture are also effective. For example, in the flicker paradigm (Rensink, O'Regan, & Clark, 1997) participants view a cycle of two versions of a scene presented sequentially at the same location on a monitor with a blank screen inserted in between each presentation to mask the motion signal associated with the change. Even though participants are explicitly told to look for changes, which are often quite large (e.g., the disappearance of a building in a city skyline), it often takes many cycles before the change is noticed.

Eye movements, simulated eye-movements, and blank screens create a global distraction that disrupts all visual input from the changing display. But not even a global mask is necessary to induce CB. The "mud-splash" paradigm (O'Regan, Rensink, & Clark, 1999) is similar to the flicker paradigm, except that several small patches, resembling mud splashing on the windshield of a car, are flashed over the changing image simultaneously with each change in the scene. In the mud-splash paradigm, the motion signal associated with the change is not occluded but is one of several motion signals that occur simultaneously. The result is the same as in the flicker paradigm; it often takes several cycles before the changing region of the display is located.

Like IB, CB is not an artifact that occurs only when people view static displays. Levin and Simons (1997) had people watch short video clips of two women sitting at a table and talking. The clip contained several shots taken from different angles (see Figure 3). Although editing the shots together, they

Figure 3. **Stills from video used by Levin and Simons (1997). In shot (A), the woman in pink is wearing a scarf, which disappears in the second shot (B), and reappears in shots C and D. Reprinted from "Failure to Detect Changes to Attended Objects in Motion Pictures," by Levin and Simons (1997), *Psychonomic Bulletin and Review, 4,* 501–506. Copyright 1997 by the Psychonomic Society. Reprinted with permission.**

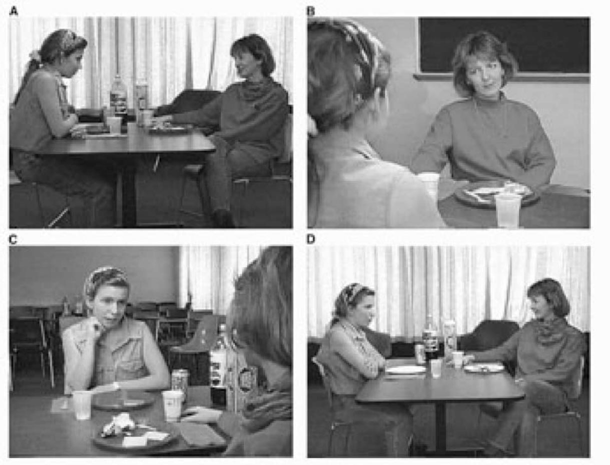

intentionally created sequences where successive views of the same scene contained inconsistent details. For example, one shot showing a woman wearing a scarf cut to another shot of the same woman without the scarf (see Figure 3). The film cut proved to be an effective mask, as every participant who viewed the tape failed to notice the scarf's disappearance. In a second experiment, Levin and Simons showed participants a video in which the sole actor in a two-shot video changed clothes and identities across a film cut. Even though most observers were presumably attending to the changing actor, about 50% failed to notice the substitution when it was unexpected. Simons and Levin (1998) even demonstrated that CB can occur in a real-world social interaction; about 50% of naive participants failed to notice the replacement of one conversation partner with a completely different person when the switch took place behind a passing door (see Figure 4).

One of the most basic and straightforward conclusions that one might draw based on these findings is that our visual system does not automatically store a detailed memory of previously attended items (O'Regan & Noe, 2001; Rensink, 2000, 2002; Simons & Levin, 1997). Accordingly, visual attention is central to our awareness of visual scenes. In the absence of an automatic detail-tracking system, to detect a change in the absence of a transient, an observer *must* be attending to the changing region as the change occurs (Rensink et al., 1997). This view seems to be supported by experiments that find a link

Figure 4. An unsuspecting participant in the real-world person-switch experiment (Simons & Levin, 1998). In box A, the initial experimenter approaches a pedestrian. In box B, the initial experimenter walks away behind a passing door, and another experimenter finishes the interaction (box C). Box D shows the switching experimenters side by side. Reprinted from "Failure to Detect Changes to People in Real-World Interactions," by Simons and Levin (1998), _Psychonomic Bulletin and Review, 5_, 644–649. Copyright 1998 by the Psychonomic Society. Reprinted with permission.

between change detection and attention. For example, experiments using the flicker paradigm find that changes to items that capture attention (e.g., a horizontally oriented item in an array of vertical items) are detected more quickly than changes to items that don't (Scholl, 2000) and changes are usually not detected until the eyes, and therefore attention, fixate a changing region (Hollingworth, Schrock, & Henderson, 2001).

As an explanation for CB, hypothesizing that attention is necessary to be aware of a change is intuitive and at some level probably correct. However, this kind of explanation may pose more questions than it answers, especially when considering the need to find application for this research. In the next section, we describe a few attempts both to gain a more usefully articulated understanding of the causes of failed visual awareness and to understand what, exactly, these failures mean.

2.3. Failures of Visual Awareness Clearly Happen, But What Do They Mean?

The research described earlier documents a series of striking failures but it is critical to ask, "failures of what?" Throughout this article, the phrase *failures* *of visual awareness* has been used deliberately. Possible alternatives are to refer to these phenomena as failures of vision, or failures of visual representation. So, if an observer did not see their conversation partner change into another person right in front of their eyes, it might be tempting to conclude that they have not internally stored any visual information. Instead, their visual systems may have assumed that the visual world generally remains stable, and therefore have entirely foregone the expense of creating and storing visual representations of the world. This might be a relatively radical assumption but it is consistent with IB and CB, and it has good precedence in the psychological literature (Gibson, 1979; O'Regan, 1992).

However, there is little reason to commit to such representational minimalism because observers who demonstrate CB or IB may nonetheless have internalized quite a bit of information about the visual world. There are two fundamental ways in which a rich representation might exist in the face of these failures. First, unreported changes and unattended stimuli may be represented using a memory system that generally does not make contact with awareness. Thus, viewers might represent the information but be unable to verbally report it. Another important alternative in the case of CB is that observers might represent pre- and post-change information in a consciously accessible manner but fail to compare it. This section explores these alternatives and describes circumstances where each seems most relevant.

The idea that people have more information represented than they can report has a history in psychology about as extensive as that belonging to the study of attentional limits. Generally, most cognitive psychologists would agree that there are circumstances in which observers can experience something and have it affect their behavior although being unable to consciously report the experience. For example, in the phenomenon of priming participants are often faster to respond to some stimulus the second time they have seen it despite their inability to recall the initial presentation. Observations such as these have led researchers studying memory to hypothesize the existence of an implicit memory system that operates largely outside of awareness.

Given this history, it is not surprising that researchers exploring failures of visual awareness have tested for these traces of experience. In the case of IB, these experiments do appear to have revealed the effects of stimuli to which participants were blind. Mack and Rock (1998) performed a series of experiments investigating whether undetected visual stimuli were implicitly perceived and identified. In these experiments, participants performed the same task as in the original inattention paradigm. On the critical inattention trial, a word was presented in addition to the cross (e.g., chart). After asking participants if they detected anything other than the target cross, Mack and Rock had participants perform a stem completion task. For the stem completion task, participants were given the first few letters of the unexpected word (e.g., *cha*___) and were asked to complete the stem with the first two words that came to mind (e.g., *chart*, *chafe*, *champ*, *chair*, *change*, etc.). Mack and Rock found that participants who were initially inattentionally blind to the unexpected word were much more likely to complete the stem with the unexpected word than were participants who had not been in the IB experiment. In another experiment, Mack and Rock asked participants to choose a line drawing of an object from a line-up after the inattention trial, without any further explanation. The line-ups contained a picture of the word that was presented on the inattention trial (e.g., if the word was "flake," the line-up had a picture of a snow flake). They found that participants who reported being unaware of the word were more likely to choose the picture of the word than were participants who were given the same lineup but were not in the inattention part of the experiment. These results constitute fairly strong evidence that participants who were unaware of the unattended word not only implicitly perceived it but also implicitly identified it. These implicit effects are compelling because they suggest that even when people are unaware of a certain stimulus, their behavior is still affected by its presence. Thus, it is possible that people may have picked up information that will help them learn about a visual scene, even though they are unaware of it (Chun & Nakayama, 2000).

Although there is some reasonably strong evidence that subjects have implicit memories of unattended stimuli, the evidence for implicit representations of change is less clear and it is important to note at the outset that the former is not evidence for the latter. The claim for implicit representation of change involves both implicit representations of the pre- and post-change objects themselves and also an implicit representation of the difference between them. It is this implicit difference-representation that is controversial (Mitroff, Simons, & Franconeri, 2002), although some evidence suggests it exists (e.g., Fernandez-Duque, Grossi, Thornton, & Neville, 2003; Fernandez-Duque & Thornton, 2000; Hollingworth & Henderson, 2002). The controversy exists because behavioral evidence for implicit change detection can often be explained by explicit mechanisms, such as observers' hesitance to explicitly report a change due to low confidence (Mitroff et al., 2002) and experiments using measures of neuronal activity (e.g., ERP and fMRI) have produced inconsistent support for implicit change detection (e.g., Beck, Rees, Frith, & Lavie, 2001; Fernandez-Duque et al., 2003; Turatto, Angrilli, Mazza, Umilta, & Driver, 2002).

Based on the evidence, it is probably fair to suggest that in the case of IB, there is good evidence to suggest that viewers may have implicit memories of objects they were not aware of, although in the case of CB, the jury is still out. However, even in the case of CB, viewers may have more in their heads than a report of "no-change" might suggest. Change detection requires not only having a sufficiently detailed representation for a pre-change item, but also comparing the pre- and post-change representations across views. Thus, CB might occur not because the visual representation is lacking, but because the comparison process is lacking in some respect (Simons, 2000). In an analog to the implicit research discussed earlier, one straightforward way of testing this hypothesis is to determine whether people who have missed a change can consciously recognize features of the pre-change object. In one experiment (Angelone, Levin, & Simons, 2003), participants viewed videos in which an actor's clothing changed between shots. In this case, subjects were not specifically told to look out for changes, so the experiment tests incidental change detection. Across several experiments participants who missed the change were, indeed, able to recognize how the pre-change actor looked when presented with a lineup. In fact, subjects who missed the change were no worse at recognizing the changed person than subjects who saw the change. Thus, change detection was completely unassociated with an ability to recognize the changing objects. Other recent research suggests that people are capable identifying previously viewed objects on a forced-choice recognition test given several minutes after viewing, even when the alternative choices are from the same category (e.g., the original and foil items are both cups; Hollingworth & Henderson, 2002; Hollingworth, Williams, & Henderson, 2001) and when the recognition test is unexpected (Castelhano & Henderson, in press; Varakin & Levin, 2003). These results suggest that people can indeed retain specific visual information in memory because such information is presumably required to distinguish two objects from the same semantic category. Furthermore, change detection is facilitated when post-change cues allow observers to restrict the comparison process to one region of a scene (Hollingworth, 2003). So, people seem to retain visually specific information in memory despite having difficulty detecting changes.

In contrast to the earlier evidence, there is also evidence suggesting that CB *is* sometimes associated with a broad failure to sufficiently represent information. Levin, Simons, Angelone, and Chabris (2002) ran a series of CB experiments using the real-world person switch described earlier. In this study, participants were approached on the street and were unaware that they were even in an experiment when their conversation partner was switched in front of them. As in previous work, between 28% and 53% of participants failed to detect the change, and this time, participants who missed the change were completely unable to pick either the pre- or post-change experimenter out of the lineup. In contrast, participants who did see the change were able to recognize both experimenters. In this case, then, seeing the change was associated with recognition of the changing objects, suggesting that CB (especially as compared with change detection) is associated with a failure to represent the changing objects sufficiently to succeed at a recognition test.

Clearly, future research will need to determine when CB is a sign of a representational failure and when it results from a more subtle comparison failure. This is important because it will help to explain when and why people create visual representations. However, this research does suggest that people do not automatically create large numbers of visual representations and may only do so when the specific situation calls for it. Accordingly, one might hypothesize that visual representations require some effort to create and therefore are not created unless they are really needed. Note that this hypothesis does not preclude the possibility that certain tasks might lead to the creation of detailed visual representations even if observers are not intentionally trying to remember lots of visual information (e.g., Castelhano & Henderson, in press; Varakin & Levin, 2003).

Other follow-ups inspired by the CB and IB phenomena have attempted to fill in the blank left by the visual attention hypothesis. Recall that the most basic explanation for failures of visual awareness is that viewers do not see (i.e., consciously perceive) things they do not attend to. As suggested earlier, this is a reasonable explanation so far as it goes but to really understand what people see, an explanation is needed of how, specifically, attention is guided in different situations. Again, there is a long history of exploring this question in psychology but only recently has this research explored attentional guidance as it relates to the highly circumscribed awareness that appears to characterize vision in natural settings.

To begin answering this question, a number of researchers have been exploring the structural and contextual factors that guide attention using CB paradigms. For example, Rensink et al. (1997) found that changes to regions of a scene that are rated to be of central interest are detected more quickly than changes to regions that are rated as marginally interesting. Expertise is another factor that may affect how attention is deployed and how objects in a scene are represented. Werner and Thies (2000) found that American football experts detected changes to game-relevant portions of a scene faster than novices, although there was no difference in how quickly non-game-relevant changes were detected (see also Reingold, Charness, Pomplun, & Stampe, 2001, for analogous evidence using chess experts and change detection). Thus, experts may focus their attention more effectively on the most meaningful parts of a scene, although being no better than novices when exploring less meaningful parts of a scene (and related research on radiologists suggests that experts may be even worse than novices in remembering non-central information; Myles-Worsley, Johnston, & Simons, 1988).

Other experiments have shown that change detection is affected by the level of categorization at which an object is learned (Archambault, O'Donnell, & Schyns, 1999). That is, participants who learned to categorize a set of objects at a specific level (e.g., a set of computers as Mary's computer, John's computer, etc.) were subsequently able to detect changes to those objects more quickly than objects learned at a general level (e.g., as a computer). An observer's prior knowledge about the probability of a given change may also affect change detection. Changes to scene properties that have a high probability occurring in the real world (e.g., a lamp shade turns from off to on across views) are detected more often than changes to scene properties that have a low probability occurring (e.g., a lamp turns from blue to green across views; Beck, Angelone, & Levin, in press).

The CB literature suggests that people's awareness of the visual environment is highly incomplete. Whenever the normal sequence of events associated with a change in the environment is disrupted somehow, it is likely that observers will have difficulty detecting the change. Much of the research following up on these initial findings has explored whether these failures really mean that viewers do not represent visual information or that viewers do represent it but fail to be aware of it or compare it across views. Generally, this follow-up research suggests that viewers do, in some cases, represent visual information but may not do so automatically or indiscriminately. One of the natural outgrowths of this research arises from the fact that even if we do represent some information sometimes, the represented world is a very small subset of the visual world. Therefore, it becomes critical to explain why we represent when we do. Researchers have begun to explore this question by understanding how visual attention is guided by knowledge, context, and visual properties.

3. VISUAL COGNITION AND VISUAL METACOGNITION

Sometimes, when cognitive scientists present their research to the public, it is so arcane that the typical listener simply ignores it. Other times, it is relevant to people's experience but is consistent with their intuition, so people insist that they already knew what researchers had taken pains to demonstrate. In stark contrast to these cases, much of the research demonstrating CB and IB confounds these intuitions and can be quite surprising, if not shocking, to other scientists and the lay public alike. Accounts of incredulous participants and audiences abound within the community of researchers who do this work and researchers have recently begun to focus on the counterintuitive nature of these findings as a research topic in itself. After all, if people are so surprised by these findings, it suggests a potentially interesting mismatch between people's naive understanding of vision and the way it actually works. Recently, the well-known philosopher Dan Dennett (2001) commented about the utility of just such a mismatch:

> Surprise is a wonderful dependent variable and should be used more often in experiments; it is easy to measure and is a telling betrayal of the subject's *having*

> *expected something else.* … These behavioral responses are themselves data in good standing, and in need of explanation. (p. 982)

If CB only occurred for subtle, near threshold changes, then the phenomena would not be so interesting. In the experiments described earlier, large changes that seem to draw attention to themselves once an observer knows exactly what the change is and where it will occur are the kinds of changes that are often missed. In fact, in many CB experiments, participants often do not believe that they could have missed a change after discovering, or being told, what the change is. To demonstrate this phenomenon empirically, Scholl, Simons, and Levin (in press) had participants search for changes in the flicker paradigm (see earlier). After participants detected a change, they were asked when in the cycle the change was inserted. Most of the time, participants would insist that the change was inserted right before it was detected, when in fact the change had been there throughout the viewing cycle. This experiment is one of a growing number of studies investigating the discrepancy between the realities of visual awareness and people's beliefs about visual awareness.

It appears, then, that people think they will always see things that they actually will not, an effect known as "change blindness blindness" (CBB). In an initial attempt to verify that people's predictions far outstrip actual performance, participants who had not heard about CB were asked to predict whether they would see the changes that were actually missed by participants in other CB experiments (see Figure 4 and Figure 5 for examples of such changes; Levin, Momen, & Drivdahl, 2000). For example, the participants were told to imagine that they were watching a movie in which two actresses were conversing, and that on one of the film cuts the scarf disappeared. The participants were also shown the stills in Figure 4 to make sure they knew exactly what the objects looked like. Despite this, about 90% of the participants predicted that they would see the scarf change when in fact *none* of the participants in the actual change detection experiment detected that particular change! Change detection ability was overestimated to a similar degree when participants were predicting other people's performance (Levin et al., 2000). Furthermore, this metacognitive error is not a methodological artifact owing to participants' misunderstanding instructions and imagining a situation in which a change would actually be detected due to the presence of a motion transient (Levin, Drivdahl, Momen & Beck, 2002). In these experiments, participants were shown the actual videos used in the Levin and Simons (1997) change detection experiments and imagined scenarios in which delays of up to 1 hr were inserted in between the pre- and post-change views. These delays were intended to make clear that no motion transients would be available.

Despite this, participants still overestimated change detection performance to the same degree.

Although research into the causes of CBB is still in its infancy, experiments are beginning to shed light on the reasons why people overestimate change detection ability. One factor that correlates with CBB is a person's beliefs about how visual attention works (Levin & Beck, in press). For example, in Levin, Drivdahl et al. (2002) participants were asked to provide justification for why they would or would not detect an object change if they weren't attending to it. Based on these open-ended justifications, it seemed as though some participants believed that they could detect changes by simultaneously attending to an entire scene, rather than attending to individual objects within the scene. Later experiments revealed that people who think that their attentional "spotlight" was particularly broad predict that they will see more changes (Levin & Beck, in press). These findings imply that people may believe that attending to a scene allows them to apprehend a large proportion of the scene all at once and to ignore the fact that it is often necessary to systematically sample small portions of the scene to really perceive it. It is interesting to note that research in developmental psychology has arrived at a similar conclusion about children's putatively immature understanding of visual attention. Flavell, Green, and Flavell (1995) argued that children do not understand the correct "spotlight" model of attention whereby attention must be focused on one thing at a time, although other things outside the focus of the spotlight are not seen. Instead, Flavell et al. suggest that children adhere to a "lamp" model of visual attention in which intending to perceive a scene allows one to see it wholesale just as one illuminates an entire room by turning on a lamp. The metacognitive research described earlier therefore suggests that this immature understanding of visual attention might, in fact, characterize adult thinking as well.

4. WHY THIS MATTERS: EVIDENCE FOR FAILURES OF VISUAL AWARENESS IN THE HUMAN–COMPUTER INTERFACE

Typically, computer users are exploring visual information and using it to signal meaningful options and mode or state changes. Therefore, they may be far more aware of visual information than participants in experiments investigating visual awareness. PC software users are clearly aware that the screen contains important task-relevant information and therefore may be more engaged than participants in experiments investigating incidental change detection who passively watch films. Pilots landing an airplane know to check a variety of displays and scan certain parts of the environment for trouble, unlike participants in experiments exploring IB who are unwittingly engaged in a task that is deliberately designed to be distracting. Even participants given the task of detecting changes might be exploring visual information at a very superficial level, unlike the computer user who is exploring visual information that is usually wellstructured and meaningful to the task at hand. A cockpit's instruments have known meanings and are all in well-defined places and the desktop metaphor has been a tremendous boon to even the most naive users because it provides a known visual structure that allows individuals to effectively search the interface.

Combined, these factors might make lapses of visual awareness in the human–computer interface a rarity, especially in situations where participants are knowledgeable professionals or are highly focused on a specific task. Although these intuitions might suggest limits to visual failures in these situations, a variety of anecdotal and experimental evidence confirm that research on IB and CB is, indeed, highly relevant to HCI.

4.1. Anecdotal Evidence

One particularly telling experience occurred during the 1980s when one of the authors (Fidler) was involved in developing a service called PressLink™ that was designed to allow newspaper professionals online access to large image libraries. Today the "You've Got Mail!" voice message is highly recognizable and closely associated with AOL and the e-mail revolution. In the context of failures of visual awareness, it seems reasonable to speculate that AOL's rationale for adding the voice message had to do with a common problem encountered with PressLink's e-mail. In the mid-1980s, most people were still unfamiliar with e-mail. Even the majority of those who had access to PressLink in the early years of the service rarely took advantage of the feature. Consequently, most people would not notice when the arrow appeared over their in box even though it was quite prominent.

The staff at PressLink learned about this underutilization of e-mail the first time the service was taken down for a major upgrade. The staff received numerous calls from angry customers who complained that they had not been informed about the downtime. When the staff told the customers that they had been sent several e-mail messages alerting them to the interruption in service, most insisted that they could not have received the messages because their version of PressLink did not have the e-mail feature. All PressLink subscribers had the same e-mail feature but the only way they could be convinced of this was to get them to reconnect to PressLink. Even when the staff pointed out the In Box with the arrow, some users would argue that it must have just been put there because they were absolutely certain it was not there the last time they used PressLink. This problem kept reoccurring until the simple "You've Got Mail" auditory message was added. From that point on, e-mail became a much more reliable way to get messages to customers, although even then it was not entirely fool proof.

One of the most interesting things about this anecdote is the parallel between it and the earlier research on visual metacognition. In both cases, a limit to visual awareness produced a failure to detect some visual feature. Also, in a striking parallel to Scholl et al.'s (in press) participants who insisted that a visual change that had been present all along had been added just before they saw it, the PressLink users insisted that the company staff had just added their In Box. So, not only do people miss things, they misconstrue how hard it is to see them, thinking that a visual stimulus should be easily detectable when it is not.

Other well-known anecdotes illustrate both limits to visual awareness and the failure to understand the limits in situations where human–computer interfaces fail with devastating consequences. For example, the crash of a French Airbus AT320 near Strasbourg, France in 1992 has been attributed to a failure on the part of the pilots to be aware of a mode-signal on the aircraft's flight control computer (Peterson, 1996). In this case, the aircraft was making a landing approach and mysteriously crashed into a mountain well short of the runway. Later investigation revealed that the aircraft was descending at a far higher rate than that detailed in the flight plan. Alert investigators noticed that if the pilots had programmed the flight control unit (FCU) with the digits specified in the flight plan but had misconstrued its mode setting, the faster (and fatal) rate would have resulted. Apparently, the crew had accidentally selected 3,300 fpm descent rate instead of a 3.3-degree flight path angle. The key to this error is that the FCU signaled its mode visually (with a three-letter abbreviation of the mode next to its setting readout) and that the difference between the two modes was probably well known to the pilots. So, meaningful, task-relevant information was available to pilots in a region immediately adjacent to the readout digits they must have been focusing on when they programmed the computer. Still, they failed to become aware of it and therefore flew their aircraft into a mountain, killing 87 of the 96 people aboard. Disasters such as this provide much of the impetus for research on effective display design.

4.2. Experimental Evidence

Recently, there have been a few experimental investigations concerning failures of visual awareness in the context of the human–computer interface. Benway (1999) observed one such failure that is very similar to IB. She was exploring a Web site designed to provide information to a specific company's employees and discovered a striking flaw in the site's interface. It seems that employees were being encouraged to sign up for a specific training class but to the instructor's dismay few were actually following through. A bit of investigation revealed that many of the employees did, indeed, try to sign up but failed to find the proper page. The interesting thing about this failure was that the page was prominently advertised with a colorful banner that was linked to the sign-up page. The banner however, was not close to other link-rich areas of the display. Many employees had, indeed, found the page containing the banner but still failed to read it so they could follow the link. Benway replicated this failure in the lab and dubbed it "Banner Blindness" because participants appeared to be completely unaware of a prominent signal that was directly related to their current goals. Clearly, the individual who programmed the site intended to make sure that everyone saw the ad, so he or she produced a signal that was supposed to be the most salient thing on the page, but paradoxically, it was instead the least salient.

We have also completed a pilot experiment testing for CB in the interface of e-text reading software. The experiment borrowed portions of a prototype version of software (designed by Fidler) that displayed text and pictures in a magazine-like format. One of the program's key features was a bilevel organization in which a surface level displayed a series of pages including story headlines and two to three sentences describing the contents of the story. Simply clicking on one of the surface-level stories would access the other level that contained the complete stories. Thus, one of the key jobs for the interface was to allow effective navigation by informing users at which level they happened to be. The interface therefore included a visual contrast between the surface and detail levels by changing the format of the page numbers displayed in a 12-point font at the upper right-hand corner of the display. Thus, when one clicked on a story, some different information might appear, and the page numbers would simultaneously change font and format. It is interesting to note that not all of the information changed, often graphs or images were included on both the surface and detail level. From the standpoint of CB, this would be a classic example because the page-number change would be masked both by the brief flicker between the two pages and by the transients caused by changes in the positions of the images and text.

In the pilot experiment, participants were asked to use the interface to find a specific bit of information contained in one of the stories. The existing header change was so small that nearly all of the participants would probably have missed it (see Figure 5). Therefore, the display was altered to make the header change much larger (see Figure 6). All participants were told about the structure of text they were about to look at, with special emphasis on the surface versus detail level distinction. They were also informed that there would be a range of questions about the interface itself subsequent to the completion of their basic task. Participants were then allowed to see the surface page and all located the correct story and clicked

Figure 5. **The original layout and header of the e-text reading software.**

Figure 6. **The altered layout and header used in the change-detection experiment.**

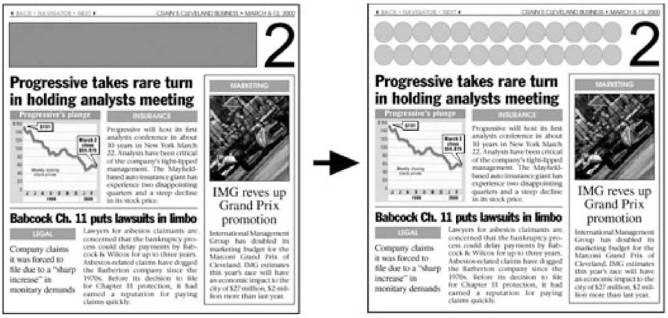

on it. Once they clicked, the page flickered off the screen and reappeared with the same information but a different header. It is important to note that participants did not actually see the detail page because the purpose of the experiment was to investigate if participants would notice anything on the interface change. If the text had actually changed, then questions about change detection question would have been ambiguous. Similar to the Levin and Simons (1997) experiments described earlier, participants were asked a series of questions assessing whether they had seen the change. The most specific asked if they had seen the header change and it included a figure illustrating where the header had been on the page. Despite the enormous size of the changes, only half of participants reported them. These results converge with Benway's results in demonstrating that even in the human–computer interface, perceptually salient information can escape awareness. In both cases however, the participants being tested were not necessarily experienced computer users and, furthermore, were not extensively trained on how to use the interface. In contrast, research from the field of human factors shows that failures of visual awareness are not necessarily a rarity, even for highly trained professionals such as airline pilots.

In the quest to develop ways to more effectively convey flight-related information to pilots, human-factors researchers have produced results that are unmistakably similar to those obtained from the inattention and change detection paradigms described earlier. For example, Haines (1991) asked airline pilots to complete simulated landings using newly developed head-up displays (HUDs). These displays are designed to allow pilots (and drivers; see Tufano, 1997) to maintain focus on the outside world by projecting navigation-relevant information on the windshield of their vehicle. The idea was that this would allow them to spend less time "heads down" looking at their instrument panel so they could more effectively monitor the external world for potential danger. When Haines used such a display in the simulator, the test pilots believed it to be quite useful and commented that they enjoyed the experience. In fact, some HUD designs do facilitate detection of expected events, such as the appearance of a runway upon landing approach (Wickens & Long, 1995). The trouble is that HUDs appear to make the pilots paradoxically less aware of unexpected and potentially dangerous events that can occur outside of an airplane, an effect now known as "cognitive capture" or "cognitive tunneling" (Tufano, 1997; Wickens & Long, 1995). In one case, Haines arranged to have another 747 taxiing onto the runway and found that 2 of 9 pilots would have landed their aircraft right on top of it had the simulation not been terminated. Other research by Wickens and Long (1995) demonstrated that even in cases where pilots using HUD displays noticed an unexpected runway incursion, the time it took them to initiate appropriate emergency maneuvers was significantly longer than pilots using traditional head-down displays.

More recent research has directly examined change detection in the context of a realistic military command and control station and has replicated the phenomenon of CB (DiVita, Obermayer, Nugent, & Linville, in press). In DiVita et al., participants monitored the activity of aircraft in a simulated naval Combat Information Center (CIC) console. Critically, the CIC console used standard military symbology with which the participants were familiar. In a real combat scenario, operators monitoring such displays must have quick access to any updated information so they can make well-informed decisions about what to do next. Although available technology allows such information to be almost instantaneously delivered to a computer display, research on CB suggests that the new information might still escape the user's awareness, especially if the screen blinks or the user's attention is momentarily diverted as information is updated. DiVita and colleagues confirmed this prediction. Not only were their participants unable to detect several task relevant and meaningful changes but they actually detected fewer changes as trials progressed.

Change detection experiments have also been conducted in an aviation cockpit context. Nikolic and Sarter (2001) had experienced pilots (flight instructors) report automated mode transitions while piloting a simulated airplane. The mode transition display of principal interest was located on the primary flight display that also conveyed information about other flight-relevant information. Therefore, pilots necessarily had to look at this display while in manual flight. In addition to responding to the mode transitions, pilots were also asked to respond to the presence of other events that might occur in the flight simulator (e.g., warning lights, traffic, etc.). Despite explicit instructions to respond to a specific mode transition each time it occurred, these pilots missed a significant number of mode transitions during manual and high-tempo periods of flight. Notably, change detection was poorest when the signal was a change from the words *on* to *off,* which is a signal currently used in many cockpits. When the mode transition was signaled by a peripherally salient cue, such as long horizontal light strip located below the primary flight display, mode transition detection was increased. However, pilots did not fare well at responding to the presence of other events, especially when the event occurred simultaneously with a mode transition. Although the magnitude of the changes missed by the participants in these experiments is sometimes less than the changes missed in the CB literature, these data nonetheless confirm that the phenomena of CB and IB do occur in the human–computer interface, even when the participants are familiar with the display.

The anecdotes and research described earlier represent a relatively small literature, in part because research exploring basic failures of visual awareness has only recently been highly active. Therefore, we see a wide range of untapped applications for this research. The next section presents an initial attempt at sketching out a set of relatively broad principles and situations where these failures might be important in user interface design.

5. ILLUSIONS OF VISUAL BANDWIDTH

Many of the problems that stem from visual limits and people's failure to understand them can be expressed as a mistaken belief that visual properties are more efficient at delivering information than they are. That is, people may falsely believe that many more aspects, features, or concepts can be communicated in a visual form than in other symbolic forms. In the human–computer interface context, such beliefs might underlie metacognitive overestimations of viewer's ability to apprehend visually presented information and could accentuate failures of visual awareness. Therefore, in this section three hypotheses are presented concerning "Illusions of Visual Bandwidth" (IVBs). These IVBs are possible manifestations of a belief that turning information into visual properties increases processing capacity more than is actually the case. It is important to note that these illusions do not represent an established theory of visual meta-awareness but rather should be taken as a series of initial hypotheses that might guide future exploration and interface testing.

The three IVBs discussed here reflect three relatively broad ways that one might overestimate visual processing capacity. First, people might overestimate how much visual information can be attended simultaneously. If one thinks of attention as a spotlight (Posner, 1980), then this IVB amounts to overestimating the area illuminated by the spotlight. The second IVB is overestimating the number of locations that a user will typically attend to. More specifically, the second IVB is a misestimation of which regions of a display a user will voluntarily search or of which visual cues will automatically capture a user's attention. And third, people may overestimate the representational consequences for having attended to an object or location. That is, people may assume that attending to an object leads to an exceedingly rich representation of its features or its meaning. The three IVB's are in principle distinct but each one could potentially underlie the general illusion that people are more aware of visual information than is actually the case.

5.1. IVB #1: Overestimate of Breadth

The mistaken assumption underlying the first IVB is that viewers can take in all (or most) of the details of a scene at once. Although phenomenological experience may often suggest to observers that everything in a scene is being made available to awareness simultaneously, research on IB and CB demonstrates that this feeling of global awareness is probably not the product of a detailed and global representation of a scene. As reviewed earlier however, people may believe that the attentional spotlight is broader than it is and this belief predicts the degree to which they over-predict their ability to see visual changes (Levin & Beck, in press). One possible cause of this overestimate is the well-organized nature of many natural scenes. This kind of organization might lead to holistic awareness of the scene and its global meaning (or "gist"; see Simons & Levin, 1997, for review) but not necessarily to a broad awareness of the specific features that make up the scene. If this hypothesis is correct, then coherent visual organization may cause people to mistake holistic awareness of gist for broad awareness of specific features, leading them to believe that looking at one part of a scene is associated with awareness of many nearby things.

In the context of human–computer interaction, consider the fact that the personal computer's rise in popularity was accompanied by a shift from command-based alphanumeric interfaces (e.g., DOS) to interactive Graphical User Interfaces (GUIs; e.g., MS Windows®). Virtually everyone agrees that GUIs are more user-friendly. But is this because a set of visual objects organized into a coherent scene allows users to expand their spotlight of attention and therefore obtain more information from a display? Although these displays might be more aesthetically pleasing than text, it is not clear that adding visual details and organization to a display does anything to improve usability by increasing the width of the attentional spotlight. In fact, the capacity to attend to and remember several objects simultaneously decreases as a function of visual complexity (Alvarez & Cavanagh, 2004). Furthermore, adding more detail to a display places a greater demand on the user to actively ignore task-irrelevant features, which in turn can lead to a decreased awareness of all unattended features (Most et al., 2001; Simons & Chabris, 1999). Thus, it is plausible that adding extra visual features to a display actually leads to a decrease in the amount of information that can be simultaneously attended and could therefore make it harder for users to find specific bits of information.

One place where this IVB might have impact is in the decision to place options in a dialog box or in the regular pull-down menu structure of a program. This decision will likely be impacted by the programmer's hypotheses not only about when the options are useful but also about the degree to which the dialog box represents a spatially circumscribed region that brings the full range of options it contains to the user's awareness. Will a glance at one part of the dialog box bring to awareness options present in other parts of the box? If so, then the dialog box will be useful in presenting options only when they are needed and, if not, then it might make sense to keep commands in the regular menu structure of the program.

5.2. IVB #2: Overestimate of Countenance

The second IVB is the belief that users will attend to a higher proportion of regions in a display than they do. Two assumptions could underlie the second IVB. The first is that observers endogenously shift their attention (i.e., voluntarily) to more regions than is the case and the second is that it is easier to induce an exogenous (i.e., automatic or externally driven) shift of attention than is actually the case. So, people may correctly realize that attention to one thing precludes attending to other things but still overestimate the likelihood that a particular region of a display will be attended through either endogenous or exogenous attentional shifts. Indeed, participants who believe in a more exhaustive attentional search also tend to predict that they would see more changes (Levin, 2001). Although these beliefs are about different mechanisms that can guide attention, both are overestimations of the efficacy of visual information to attract a user's attention and both beliefs underestimate a user's capacity to ignore certain parts of a display.

This IVB could stem from an overgeneralization of the fact that unique visual cues (i.e., colors, shape, locations, movement, etc.) do in fact attract attention in some circumstances. The illusion is that uniqueness attracts attention in all situations. The aforementioned banner-blindness effect (Benway, 1999) is one example of a possible consequence of IVB #2. The designer of the initial banner blindness inducing Web site was obviously not trying to make the information difficult to find. The designer's choice to present the information in a unique way makes sense under the assumption that spatial and featural uniqueness will reliably capture attention. Obviously, the designer was mistaken. Although Benway's results demonstrate that attention may not be deployed to every unique item in a display, she did find that when the target information was located in an expected area of the screen and was highlighted through a distinctive color, users had little trouble locating it. So, in this case uniqueness may have attracted attention. However, Benway stopped short of recommending distinctive color as a general purpose heuristic for drawing user's attention because participants in her study were searching for the information contained in the color-highlighted link.

Other researchers have attempted to address the issue of how to attract a users attention to notify them about dynamic information while they are engaged in primary task that is unrelated to the notification signal. In a study by Bartram, Ware, and Calvert (2003) participants performed a primary task, analogous to the kinds of tasks computer users might typically perform (e.g., sorting numbers in a spread sheet or playing solitaire), along with a secondary task of detecting signals from icons located in a different part of the display. In this situation, participants only made a simple response to the cues but in practice such signals could be associated with more substantive information (e.g., status of a running program, system alerts, messages about various events such as stock market fluctuations). Compared to signals that relied on changes in peripheral icon's color or shape, participants were better able to detect changes that involved motion. Bartam et al. therefore recommended that information be coded in the motion of icons ("moticons"). Nikolic and Sarter (2001) make a similar suggestion for designing effective notification signals in an aviation context. Based on results of their change-detection experiments described earlier, they suggested that peripherally salient changes, such as a strip that suddenly becomes brighter, are more effective signals than other more subtle changes, such as changing the features of an object or word. Even though such signals have the potential to be distracting, these suggestions are good because they represent the kind of transient that will likely capture attention and therefore counteract the mistaken assumptions underlying

IVB #2. However, it is important to bear in mind that saccades will mask about 10% of these transients which may make them less reliable than might be apparent based on their attention-grabbing properties.

Attracting users' attention through either endogenous or exogenous shifts of attention is possible but it is important for designers to realize that users typically have expectations about where in a display to look and may routinely ignore most of what is visible at any given time. Unique visual cues and changing visual properties are not always sufficient to draw attention. Furthermore, even if a user's attention is successfully captured, it does not guarantee that they will apprehend all of the information at the attended location. This consideration leads to the third IVB.

5.3. IVB #3: Overestimate of Depth

The third IVB is perhaps the most general because it reflects the belief that attending to an object leads to more complete and deep coding of the object than is the case. There are at least two ways that one could overestimate the consequences of attending to an object. First, one could assume that all features of an attended object are automatically represented, stored in memory and used to maintain awareness of the environment. Second, one could assume that the act of attending to a visual object leads to a relatively effortless appreciation of the object's deep meaning. These two assumptions overestimate two different consequences of attention to an object but both are overestimations of what an observer will represent and possibly learn as a result of attending to objects and locations in a visual scene.

The first assumption that could underlie the third IVB is that attention to objects leads to awareness and deep representation of its features. Indeed, there is some evidence that people are better at processing and reporting two visual features if they are a part of a single object instead of features of different objects in the same location (e.g., Duncan, 1984). However, attending to an object does not necessarily lead to an awareness of all of its features (Levin & Simons, 1997). Some specific visual information about things attended in the past might be represented in memory (e.g., Castelhano & Henderson, in press), but these memories do not necessarily contribute to an online representation that supports awareness of the environment (Varakin & Levin, 2003). If they did, how would CB in attended objects happen at all? Instead, the selection of features within attended objects is highly dependent on moment-by-moment task demands (Hayhoe, 2000). What this suggests is that even attended pictograms and other visual objects do not necessarily have an advantage over text in terms of allowing a user to maintain awareness of multiple visual cues.

The second assumption that could underlie the third IVB is that attended visual objects are better at communicating meaning to an observer than is actually the case. A wide range of findings do suggest that attending to a picture causes its basic-level category to be activated (see, e.g., Jolicoeur, Gluck, & Kosslyn, 1984; O'Connor & Potter, 2002). That is, looking at a picture probably activates something akin to a one-word label most closely associated with it (e.g., *Dog*, *Chair*, or *Car*; see Rosch, Mervis, Gray, Johnson, & Bayes Braem, 1976). Similarly, attending to a word causes basic semantic (and perhaps grammatical) information to be activated (e.g., as in the well-known Stroop effect; Stroop, 1935). So, it is easy for observers to access the common meaning of both words and objects but it is not clear that attended objects have a necessary advantage over attended words in terms of communicating anything to a user beyond basic meaning.

So why do personal computer interface designers often rely so heavily on pictograms and other visual features in designing interfaces? The toolbar full of pictograms representing various functions is ubiquitous in applications such as database search engines, spreadsheets and word processing packages (e.g., the disk icon commonly represents the function *save*). From a design standpoint, one could argue that pictograms and other visual elements have several advantages over text. One such advantage might be that some (but not all) ideas are difficult to succinctly express in one or a few words. However, the point is often made that it is difficult to succinctly represent new complex and abstract concepts or the purpose of an application using a single visual icon as well (e.g., Böcker, 1996; Goonetilleke, Shih, On, & Fritsch, 2001). So, there are probably cases where designers have opted to use a pictogram, even though a word would likely communicate the intended meaning better than a pictogram and would take approximately the same amount of space in a display. For example, to save a document in many software applications, a user can click on the pictogram that looks like a floppy disk in a toolbar or select the word *save* from a pull-down menu. Since the word *save* and a pictogram of a disk take up approximately the same amount of space (or could be made to), the decision to make the pictogram visible,and the word hidden until a user selects a pull-down menu makes sense only under the assumption that the pictogram confers some advantage in making the option more readily accessible. Of course, a picture of a disk probably does not communicate the function *save* better than the word itself and in fact may be worse because it could just as easily mean *open*. In situations where concepts can be represented as a single or a few words, attended pictograms probably have few advantages over words in terms of helping users remain aware of an interface's functions or communicating these functions to a user.

5.4. The Origins and Implications of IVBs

The three IVBs discussed here might arise from people's experience when they explore the visual environment. Thus, it might not be surprising that a person who is unfamiliar with the research described here holds some of these beliefs. In the real world, when a person is looking for something and finds it, the act often (but not always) seems easy. Accordingly, the transients inherent in studies of CB complicate a search task that is otherwise fairly simple. Therefore, people may sometimes be correctly optimistic in their ability to apprehend visual information. However, the typical human–computer interface is not like the real world, as there are more ways for a computer display to induce the failures described earlier (e.g., overlapping windows, slow Web pages, etc.). In fact, it has been pointed out that some personal computer interface displays, particularly Web pages, are in many ways like displays used in CB experiments (Hudson, 2001). This may explain why some believe that "the user interface took a giant step backwards from the set of applications people were using before [the Internet's] arrival" (Olson & Olson, 2003, p. 500). Because many modern computing systems are designed with the user's abilities in mind (Carroll, 1997), it might seem surprising that displays designed for efficiency and those designed to interfere with efficiency (i.e., IB and CB experiment displays) ended up being so similar. IVBs may help to explain how this happened.

Because IVBs may reflect reasonable beliefs based on everyday visual experience, holding these beliefs would not always prevent a person from successfully navigating in the real world and will not always prevent someone from successfully using a computer. However, each of these IVBs may lead to inefficiency (as in the case of PressLink's unused e-mail) and in some situations disaster (as in the case of the crash of the Airbus). In light of the research described throughout this article, it should come as no surprise that one bit of advice is to avoid the assumption that changing the features of a visual display will be sufficient to ensure that a user gets the message that the visual change is intended to convey.

However, the consequences of failing to detect a change are not always detrimental. For example, the "recycle bin" and "trash" pictograms on the desktops of popular GUIs change when files are dragged into them. The change does not usually affect the way in which a user should interact with the computer, so it makes little difference if the change is or is not detected. Adding new links and animations to a Web site or more icons to the toolbars in a word processor only improves the functionality of the interface if users see that they are there and only reduces the functionality if it is a distraction. The research described in this article suggests that, in many cases, there may be a lot of room for visual interface designers to experiment and be creative without affecting human–computer interactions much one way or the other (see Diaper & Waeland, 2000), as long as the designer does not intend for easily ignored and often missed visual changes to convey something to a user.

6. CONCLUSION

Research exploring many of these dramatic failures of visual awareness is still in many ways just beginning. Therefore, many of the specific recommendations that can be offered are necessarily tentative. The examples provided in the discussion of each IVB should be taken as starting points for applied research on dealing with visual limits and misapprehension of those limits. Fortunately, the methods employed in IB and CB experiments, such as inattention (Mack & Rock, 1998) and flicker (Rensink et al., 1997) paradigms, provide relatively simple to implement procedures for testing how easy it is for observers to notice various kinds of changes, and incidentally have relatively high levels of ecological validity for questions concerning the human–computer interface. Researchers investigating IB were not interested in cognitive capture in an aviation setting or banner blindness in Web pages and researchers interested in CB and CBB were not interested in why PressLink subscribers did not use their e-mail. In a wide range of settings and applications, testing which of an interface's features frequently escape users' notice will constitute a good first step toward designing an interface that is optimized for usability. Similar to how Mack and Rock (1998) shifted their research program from asking questions about how inattention affects conscious perception to whether conscious perception occurs at all without attention, designers can move beyond testing hypotheses about how useable the elements in a visual display are when they are seen and regularly test whether elements in a visual display are seen by users at all.

NOTES

Acknowledgments. We thank James Arrington, James Dallaire, and Joseph Wayand for helpful comments on earlier versions of this article. The comments and suggestions offered by Richard Pew, Joseph DiVita, and an anonymous reviewer were also extremely valuable in helping us to improve the article.

Support. Preparation of this article was supported by NSF Grant SES - 0214969 awarded to DTL.

Authors' Present Addresses. D. Alexander Varakin, Department of Psychology and Human Development, Peabody Box 512, 230 Appleton Place, Nashville, TN 37203. E-mail: alex.varakin@vanderbilt.edu; Daniel T. Levin, Department of Psychology and Human Development, Peabody Box 512, 230 Appleton Place, Nashville, TN 37203. E-mail: daniel.t.levin@vanderbilt.edu; Roger Fidler, Institute for Cyber Information, 203 Moulton Hall, Kent State University, Kent, OH 44242. E-mail: roger@rogerfidler.com

HCI Editorial Record. First manuscript received November 12, 2003. Revision Received June 14, 2004, Accepted by Richard Pew. Final manuscript received July 17, 2004. *–Editor*

REFERENCES

Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science, 15*(2), 106–111.

Angelone, B., Levin, D. T., & Simons, D. J. (2003). The relationship between change detection and recognition of centrally attended objects in motion pictures. *Perception, 32*, 947–962.

Archambault, A., O'Donnell, C., & Schyns, P. G. (1999). Blind to object changes: When learning the same object at difference levels of categorization modifies its perception. *Psychological Science, 10*(3), 249–255.

Bartram, L., Ware, C., & Calvert, T. (2003). Moticons: detection, distraction and task. *International Journal of Human–Computer Studies, 58*, 515–545.

Beck, D. M., Rees, G., Frith, C. D., & Lavie, N. (2001). Neural correlates of change detection and change blindness. *Nature Neuroscience, 4*(6), 645–650.

Beck, M. R., Angelone, B. L., & Levin, D. T. (2004). Knowledge about the probability of change affects change detection performance. *Journal of Experimental Psychology: Human Perception and Performance, 30*(4), 778–791.

Benway, J. P. (1999). *Banner blindness: What searching users notice and do not notice on the World Wide Web.* Unpublished doctoral dissertation, Rice University, Houston, Texas.

Blackmore, S. J., Brelstaff, G., Nelson, K., & Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception, 24,* 1075–1081.

Böcker, M. (1996). A multiple index approach for the evaluation of pictograms and icons. *Computer Standards and Interfaces, 18*, 107–115.

Carroll, J. M. (1997). Human–Computer interaction: Psychology as a science of design. *Annual Review of Psychology, 48*, 61–83.

Castelhano, M. S., & Henderson, J. M. (in press). Incidental visual memory for objects in scenes. *Visual Cognition: Special Issue on Scene Perception.*

Cherry, E. C. (1953). Some experiments on the recognition of speech with one and with two ears. *Journal of the Acoustical Society of America, 25*, 975–979

Chun, M. M., & Nakayama, K. (2000). On the functional role of implicit visual memory for the adaptive deployment of attention across scenes. *Visual Cognition, 7*, 65–82.

Dennett, D. C. (2001). Surprise, surprise. *Behavioral and Brain Sciences, 24*(5), 982.

Diaper, D., & Waeland, P. (2000). World Wide Web working whilst ignoring graphics: Good news for Web page designers. *Interacting with Computers, 13*, 163–181.

DiVita, J., Obermayer, R., Nugent, W., & Linville, J. M. (2004). Verification of the change blindness phenomenon while managing critical events on a combat information display. *Human Factors, 46*(2), 205–218.

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General, 113*(4), 501–517.

Fernandez-Duque, D., Grossi, G., Thornton, I. M., & Neville, H. J. (2003). Representation of change: Separate electrophysiological markers of attention, awareness, and implicit processing. *Journal of Cognitive Neuroscience, 15*(4), 491–507.

Fernandez-Duque, D., & Thornton, I. M. (2000). Change detection without awareness: do explicit reports underestimate the representation of change in the visual system? *Visual Cognition, 7*, 323–344.

Flavell, J. H., Green, F. L., & Flavell, E. R. (1995). Developmental changes in young children's knowledge about the mind. *Developmental Psychology, 5*(1), 1–27.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston, MA: Houghton Mifflin.

Goonetilleke, R. S., Shih, H. M., On, H. K., & Fritsch, J. (2001). Effects of training and representational characteristics in icon design. *International Journal of Human–Computer Studies, 55*, 741–760.

Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In K. Atkins (Ed.), *Vancouver studies in cognitive science: Vol. 2. Perception* (pp. 89–110). New York: Oxford University Press.

Haines, R. F. (1991). A breakdown in simultaneous information processing. In G. Obrecht, & L.W. Stark (Eds.), *Presbyopia research: From molecular biology to visual adaptation* (pp. 171–175). New York: Plenum Press.

Hayhoe, M. (2000). Vision using routines: A functional account of vision. *Visual Cognition, 7*, 43–64

Hollingworth, A. (2003). Failures of retrieval and comparison constrain change detection in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance, 29*, 388–403.

Hollingworth, A., & Henderson, J. M. (2002). Accurate visual memory for previously attended objects in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance, 28*, 113–136.

Hollingworth, A., Schrock, G., & Henderson, J. M. (2001). Change detection in the flicker paradigm: The role of fixation position within a scene. *Memory and Cognition, 29,* 296–304.

Hollingworth, A., Williams, C., & Henderson, J. M. (2001). To see and remember: Visually specific information is retained in memory from previously attended objects in natural scenes. *Psychonomic Bulletin and Review*, 8, 761–768.

Hudson, W. (2001, November/December). Designing for the grand illusion. *SIGCHI Bulletin*, 8.

Jolicoeur, P., Gluck, M. A., & Kosslyn, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology, 16*, 243–275.

Levin, D. T. (2001). Visual metacognitions underlying change blindness blindness and estimates of picture memory [abstract]. *Journal of Vision, 1*(3), 10a. Retrieved on October 18, 2004 from http://journalofvision.org/1/3/10/, doi:10.1167/ 1.3.10

Levin, D. T., & Beck, M. (2004). Thinking about seeing: Spanning the difference between metacognitive failure and success. In D. T. Levin (Ed.), *Visual metacognition: Thinking about seeing* (pp. 121–143). Cambridge, MA: MIT Press.

Levin D. T., Drivdahl, S. B., Momen, N., & Beck, M. (2002). False predictions about the detectability of visual changes: The role of beliefs about attention, memory, and the continuity of attended objects in causing change blindness blindness. *Consciousness and Cognition: An International Journal, 11*(4), 507–527

Levin, D. T., Momen, N., & Drivdahl, S. B. (2000). Change blindness blindness: The metacognitive error of overestimating change-detection ability. *Visual Cognition*, *7*, 397–412.

Levin, D. T., & Simons, D. J. (1997). Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin and Review, 4,* 501–506.

Levin, D. T., Simons, D. J., Angelone, B. L., & Chabris, C. F. (2002). Memory for centrally attended changing objects in an incidental real-world change detection paradigm. *British Journal of Psychology, 93*(3), 289–302.

Mack, A., & Rock, I. (1998). *Inattentional blindness*. Cambridge, MA: MIT Press.

McConkie, G. W., & Zola, D. (1979). Is visual information integrated across successive fixations in reading? *Perception and Psychophysics, 25*(3), 221–224.

Mitroff, S. R., Simons, D. J., & Franconeri, S. L. (2002). The siren song of implicit change detection. *Journal of Experimental Psychology: Human Perception and Performance*, *28*(4), 798–815.

Most, S. B., Simons, D. J., Scholl, B. J., Jimenez, R., Clifford, E., & Chabris, C. F. (2001). How not to be seen: The contribution of similarity and selective ignoring to sustained inattentional blindness. *Psychological Science, 12*(1), 9–17.

Myles-Worsley, M., Johnston, W. A., & Simons, M. A. (1988). The influence of expertise on X-ray image processing. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *14*(3), 553–557.

Neisser, U., & Becklen, R. (1975). Selective looking: Attending to visually specified events. *Cognitive Psychology*, *7*, 480–494.

Nikolic, M. I., & Sarter, N. B. (2001). Peripheral visual feedback: A powerful means of supporting effective attention allocation in event-driven, data-rich environments. *Human Factors, 43*(1), 30–38.

O'Connor, K. J., & Potter, M. C. (2002). Constrained formation of object representations. *Psychological Science, 13*(2), 106–111.

Olson, G. M., & Olson, J. S. (2003). Human–computer interaction: Psychological aspects of the human use of computing. *Annual Review of Psychology*, *54*, 491–516.

O'Regan, J. K. (1992). Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology, 46*, 461–488.

O'Regan, J. K., & Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, *24*(5), 939–1031.

O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change Blindness as a result of "mudsplashes." *Nature, 398*, 34.

Peterson, I. (1996). *Fatal defect: Chasing killer computer bugs*. New York: Vintage.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25.

Reingold, E. M., Charness, N., Pomplun, M., & Stampe, D. M. (2001). Visual span in expert chess players: Evidence from eye movements. *Psychological Science, 12*(1), 48–55.

Rensink, R. A. (2000). The dynamic representation of scenes. *Visual Cognition, 7,* 17–42.

Rensink, R. A. (2002). Change detection. *Annual Review of Psychology, 53,* 245–277.

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8,* 368–373.

Rosch, E., Mervis, C. G., Gray, W. D., Johnson, D. M., & Bayes Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382–439.

Scholl, B. J. (2000). Attenuated change blindness for exogenously attended items in a flicker paradigm. *Visual Cognition, 7,* 377–396.

Scholl, B. J., Simons, D. J., & Levin, D. T. (2004). "Change blindness" blindness: An implicit measure of a metacognitive error. In D. T. Levin (Ed.), *Visual metacognition: Thinking about seeing* (pp. 145–163). Cambridge, MA: MIT Press.

Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition, 7,* 1–15.

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception, 28*(9), 1059–1074.

Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences, 1,* 261–267.

Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people in real-world interactions. *Psychonomic Bulletin and Review, 5,* 644–649.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology,* 18, 636–662.

Tufano, D. R. (1997). Automotive HUDs: The overlooked safety issues. *Human Factors, 39*(2), 303–311.

Turatto, M., Angrilli, A., Mazza, V., Umilta, C., & Driver, J. (2002). Looking without seeing the background change: Electrophysiological correlates of change detection versus change blindness. *Cognition, 84,* B1-B10.

Varakin, D. A., & Levin, D. T. (2003). The relationship between incidental and intentional change detection and long-term recognition [Abstract]. *Journal of Vision, 3*(9), 587a. Retrieved on October 18, 2004 from p://journalofvision.org/3/9/ 587/, doi:10.1167/3.9.587.

Wayand, J., & Levin, D. T (2001). Ignoring a merciless act [Abstract]. *Journal of Vision, 1*(3), 12a. Retrieved on October 18, 2004 from http://journalofvision.org/1/3/12, DOI 10.1167/1.3.12

Werner, S., & Theis, B. (2000). Is "change blindness" attenuated by domain specific expertise? An expert-novice comparison of change detection in football images. *Visual Cognition,* 7, 163–174.

Wickens, C. D., & Long, J. (1995). Object versus space-based models of visual attention: Implications for the design of head-up displays. *Journal of Experimental Psychology: Applied, 1*(3), 179–193.