

Social Science Information

<http://ssi.sagepub.com/>

Attentive user interfaces: the surveillance and sousveillance of gaze-aware objects

Roel Vertegaal and Jeffrey S. Shell
Social Science Information 2008 47: 275
DOI: 10.1177/0539018408092574

The online version of this article can be found at:
<http://ssi.sagepub.com/content/47/3/275>

Published by:



<http://www.sagepublications.com>

On behalf of:



[Maison des Sciences de l'Homme](#)

Additional services and information for *Social Science Information* can be found at:

Email Alerts: <http://ssi.sagepub.com/cgi/alerts>

Subscriptions: <http://ssi.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://ssi.sagepub.com/content/47/3/275.refs.html>

>> [Version of Record](#) - Sep 5, 2008

[What is This?](#)

Roel Vertegaal & Jeffrey S. Shell

Attentive user interfaces: the surveillance and sousveillance of gaze-aware objects¹

Abstract. *Attentive user interfaces are user interfaces that aim to support users' attentional capacities. By sensing users' attention for objects and people in their everyday environment and by treating user attention as a limited resource, these interfaces avoid today's ubiquitous patterns of interruption. Focusing upon attention as a central interaction channel allows development of more sociable methods of communication and repair with ubiquitous devices. Our methods are analogous to human turn-taking in group communication. Turn-taking improves the user's ability to conduct foreground processing of conversations. Attentive user interfaces bridge the gap between foreground and periphery of user activity in a similar fashion, allowing users to move smoothly in between. The authors present a framework for augmenting user attention through attentive user interfaces. We propose 5 key properties of attentive systems: to (1) sense attention, (2) reason about attention, (3) regulate interactions, (4) communicate attention and (5) augment attention. We conclude with a discussion of privacy considerations of attentive user interfaces.*

Key words. *Attentive user interfaces – Context-aware computing – Eye tracking – Notification – Ubiquitous computing*

Résumé. *Les interfaces utilisateur attentives sont des interfaces utilisateur qui ont pour but de soutenir l'attention de l'utilisateur. En sentant l'attention portée par l'utilisateur aux objets et aux personnes dans l'environnement quotidien et en tenant compte de ce que l'attention de l'utilisateur est une ressource limitée, ces interfaces pallient aux schémas habituels d'interruption pervasive. En faisant de l'attention le canal central d'interaction, on peut développer des méthodes de communication et d'utilisation des dispositifs pervasifs ou ubiquitaires plus conviviales. Ces méthodes sont analogues à la prise de parole tour à tour dans l'interaction et la communication humaine en groupe. La prise de parole tour à*

tour améliorer la capacité de l'utilisateur de conduire au premier plan le processus de la conversation. De même, les interfaces utilisateur attentives comblerent le fossé entre le premier plan et la périphérie de l'activité de l'utilisateur, en permettant à l'utilisateur de se déplacer lentement de l'un à l'autre. Les auteurs présentent un cadre permettant d'augmenter l'attention de l'utilisateur par des interfaces utilisateur attentives. Ils proposent cinq propriétés clés des systèmes attentifs: (1) sentir l'attention, (2) raisonner sur l'attention, (3) réguler les interactions, (4) communiquer l'attention et (5) augmenter l'attention. Les auteurs concluent par une discussion des problèmes posés par les interfaces utilisateur attentives quant au respect de la vie privée.

Mots-clés. Informatique de contexte – Informatique pervasive – Interfaces utilisateur attentives – Notification – Suivi du regard

It is our belief that the proliferation of ubiquitous digital devices necessitates a new way of thinking about human–computer interaction (HCI). Weiser (1991) said of ubiquitous computing: 'The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it.' Although we live in a ubiquitous computing age, with technologies that have interwoven themselves with our daily existence, the interface has far from disappeared. For many years our design and research efforts have focused on the development of computers as tools, extensions of analog devices such as paper, pencils and typewriters. While this view will continue to be extremely successful for years to come, we are now beginning to see limits to this approach. One of the main reasons for this is that, unlike traditional tools, computers are becoming increasingly *active* communicators. However, they are ill equipped to negotiate their communications with humans.



FIGURE 1

Email application with modal 'you have new mail' notification alert

Source: Qualcomm Corporation's Eudora, with permission.

Consider the example in Figure 1. An email tool brings up a modal dialog box to inform its user that a message has been received. Without any regard for the user's current activity, the dialog box pops up in the center of her screen. Only by clicking the OK button can the user continue her activities. This example points out a serious underlying flaw in user interfaces: the computer's lack of knowledge about the present activities of its user. Indeed the behavior of such devices could be described as socially inadequate.

HCI as multiparty dialog

As we evolve new relationships with the computing systems that surround us, there is a need to develop new design strategies. We have moved from many users sharing a single computer through a command line interface to a single person using one computer with a graphical user interface (GUI). Recently we have developed a multiparty relationship with our computers, one that causes existing channels of interaction to break down because:

- each user is surrounded by many active computing devices;
- these devices form part of a worldwide, connected network;
- users form part of a worldwide 'attention-seeking' community through these active devices.

Because of the ubiquity of active connected devices, users are now bombarded with interruptions from their Palm Pilots, BlackBerries, email programs, auction trackers, instant messaging tools and cell phones. Like the pop-up email example in Figure 1, the nature of interruptions is often acute, requiring immediate attention. As a consequence, user attention has become a limited resource, continually vied for by various devices, each claiming a high priority. We must design computers with channels to explicitly negotiate the volume and timing of communications with their user, pending the user's current needs. Our design strategy, which is to solve this problem by making interfaces more considerate (Gibbs, 2005) and less interruptive, rests upon the most striking parallel available: that of multiparty dialog in human group communication.

Taking turns for attention

We were motivated in part by work performed in the area of social psychology in view of understanding the regulation of human multiparty communication and attention. In human conversation, attention is inherently a limited

resource. Humans can only listen to and absorb the message of one person at a time (Cherry, 1953; Duncan, 1972). Thus we have developed two attentive mechanisms to focus on a single verbal message stream. (1) When there are many speakers, the ‘cocktail party phenomenon’ allows us to focus on the words of the one speaker we are interested in by attenuating speech from other individuals (Cherry, 1953). We can apply this approach to augment the attentive capacities of users. (2) However, a far more effective method for optimizing attention is to allow only one person to speak at a time, while the others remain silent. By using nonverbal cues to convey attention, humans achieve a remarkably efficient process of speaker exchange, or *turn-taking* (Duncan, 1972). Turn-taking provides a powerful metaphor for the regulation of communication with ubiquitous devices. So what information do humans use to determine when to speak or to yield the floor? According to Short and colleagues (1976), as many as eight cues may be used: completion of a grammatical clause; a socio-centric expression such as ‘you know’; a drawl on the final syllable; a shift in pitch at the end of the clause; a drop in loudness; termination of gestures; relaxation of body position and the resumption of eye contact with a listener. However, in group conversations, only one of these cues indicates to whom the speaker may be yielding the floor: eye contact (Vertegaal, 1999).

Eye contact points to targets of attention

Eye contact indicates with about 82% accuracy whether a person is being spoken to or listened to in four-person conversations (Vertegaal et al., 2001). When a speaker falls silent and looks at a listener, this is perceived as an invitation to take the floor. Vertegaal (1999) showed that, in triadic mediated conversations, the number of turns drops by 25% if eye contact is not conveyed. According to a recent study, 49% of the reason why someone speaks may be explained by the amount of eye contact with an interlocutor (Vertegaal & Ding, 2002). Humans use eye contact in the turn-taking process for four reasons:

1. Eye fixations indicate most reliably the target of a person’s attention, including their conversational attention (Argyle & Cook, 1976; Vertegaal et al., 2001).
2. The perception of eye contact increases arousal, which aids in proper allocation of brain resources and in regulating inter-personal relationships (Argyle & Cook, 1976).
3. Eye contact is a *nonverbal visual* signal, one that can be used to negotiate turns without interrupting the *verbal auditory* channel.

4. Eye contact allows them to observe the nonverbal responses, including the attentional focus, of others.

We have sought to implement similar characteristics in computing systems so as to allow them to communicate more sociably with their users. The gaze of the user, as an extra channel of input, seems an ideal candidate for ubiquitous devices to sense their users' attention. It may allow devices to determine whether a user is attending to them, or to another device or person. By tracking whether a user ignores or accepts requests for attention, interruptions can be made more subtly.

Designing windows and mice for the real world

Bellotti et al. (2002) posed five challenges for multiparty HCI. In this article, we hope to provide some suggestions for answering the first three:

1. How do I address one of many possible devices?
2. How do I know the system is ready and attending to my actions?
3. How do I specify a target for my actions?

How indeed do we move from GUI-style interactions where multiple entities are represented on a single computing device to interactions with many remote devices in the real world?

For one, it is important to note that many of the elements of GUIs were designed with attention in mind. According to Smith and colleagues (1982), windows provide a way to optimally allocate screen real estate to accommodate user task priorities. *Windows* represent foreground tasks at high resolution and occupy the bulk of display space in the center of vision. *Icons* represent peripheral tasks at low resolution in the periphery of the user's vision. *Pointers* allow users to communicate their focus of attention to graphic objects. By clicking icons to open windows, and by positioning, resizing and closing windows, users use their pointing device to manually manage their attention space. By control-clicking graphic objects, users indicate the target of menu commands. In clicking OK buttons, users acknowledge interruptions by alert boxes.

Figure 2 shows how we might extend the above GUI elements to interactions with ubiquitous remote devices, drawing parallels with the role of attention in human turn-taking. Windows and icons are supplanted by graceful increases and decreases of information resolution between devices in the foreground and background of user attention. Devices sense whether they are in the focus of user attention by observing presence and eye contact.

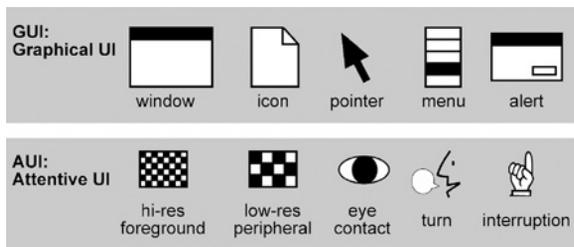


FIGURE 2
Equivalents of GUI elements in attentive UI

Menus and alerts are replaced by a negotiated turn-taking process between users and devices. Such characteristics and behaviors define an attentive user interface.

A framework for attentive user interfaces

Attentive user interfaces are interfaces that optimize their communication with users, such that information-processing resources of users and devices are dynamically allocated according to users' task priorities. This is achieved using measures and models of users' past, present and future attention for tasks. Five key properties of AUIs include:

1. *Sensing attention*: by tracking users' physical proximity, body orientation and eye fixations, interfaces can determine what device, person or task a user is most likely attending to.
2. *Reasoning about attention*: by statistically modeling the simple interactive behavior of users, interfaces can estimate user task prioritization.
3. *Communication of attention*: interfaces should make available information about the users' attention to other people and devices. Communication systems should convey to whom or what users are paying attention, and whether a user is available for communication.
4. *Gradual negotiation of turns*: like turn-taking, interfaces should determine the availability of the user for interruption by (a) checking the priority of their request; (b) progressively signaling this request via a peripheral channel; and (c) sensing user acknowledgment of the request before taking the foreground.
5. *Augmentation of focus*: the ultimate goal of all AUIs is to augment the attention of their users. Analogous to the cocktail party phenomenon,

AUIs may, for example, magnify information focused upon by the user and attenuate peripheral detail.
(Shell, Selker & Vertegaal, 2003)

Modern traffic-light systems provide an interesting parallel to an attentive user interface that augments the attention of all the users involved. They use presence sensors in the road surface to determine every vehicle's intent at the intersection, in effect *sensing the user's attention*. They are programmed with models that determine the priority of traffic on intersection roads with volume statistics, in effect allowing for *reasoning about attention*. Using peripheral displays, such as traffic lights, they *communicate the collective attention of drivers*. In so doing, they *negotiate turn-taking* on intersections to allow for smooth traffic flow.

Other related work on AUIs

Our work was also inspired by our interactions with a host of researchers, designers and media artists, as well as by the vision of ubiquitous computing (Weiser, 1991) and the seamless interaction created by considering foreground vs background in tangible user interfaces (Ishii & Ullmer, 1997). Since both of these paradigms are well known, we limit ourselves here to a discussion of existing attentive user interfaces, and discuss these examples as they relate to our framework.

Sensing attention: eye tracking as a tool in AUIs

Rick Bolt's Gaze-Orchestrated Dynamic Windows was one of the first true AUIs (Bolt, 1985). It simulated a composite of 40 simultaneously playing television episodes on one large display. All stereo soundtracks from the episodes were active, creating 'a kind of Cocktail Party Effect mélange of voices and sounds'. Via a pair of eye-tracking glasses, the system sensed the user's visual attention towards a particular image, turning off the soundtracks of all other episodes and zooming in to fill the screen with the image. Bolt's system demonstrated how a windowing system could be translated into a display with malleable resolution that exploits the dynamics of the user's visual attention. It shows the great potential of AUIs to augment attention by reducing information overload in congested audio-visual environments. Jacob (1995) and MAGIC (Zhai, Morimoto & Ihde, 1999) showed that eye tracking works best when it is applied to observe user attention rather than as a device for control. This is because, to the user, the

eyes are principally an input rather than an output organ. As a consequence, when the duration of an eye fixation on an on-screen object is used to issue commands, users may unintentionally trigger unwanted responses while looking (the ‘Midas Touch’ effect: Jacob, 1995). In a ‘non-command interface’ (Nielsen, 1993) version, instead of a user explicitly issuing commands, the computer observes user activity. The system then reasons about action, using a set of heuristics. In the classic game of Paddleball, the goal is to position a sliding paddle in the path of a moving ball using a joystick, which introduces an eye–hand coordination problem. In a non-command interface version of the game, the paddle location is given by the horizontal coordinate of a user’s on-screen gaze, communicating the visual attention of the user, and thus eliminating the game’s eye–hand coordination problem.

Reasoning about attention

Attentional interfaces (Horvitz, 1999; Horvitz, Jacobs & Hovel, 1999) are interfaces that use Bayesian reasoning to identify what channels to use and whether or not to notify a user. In the ‘Priorities’ system (Horvitz, Jacobs & Hovel, 1999), the delivery of email messages is prioritized using simple measures of user attention to a sender: the mean time and frequency with which the user responds to emails from that sender. Messages with a high-priority rating are forwarded to a user’s pager, while messages with low priority wait until the user checks them. Horvitz’s attentional interfaces are characterized by their ability to reason about user attention as a resource, rather than sense attention for a device.

Communication of attention

GAZE (Vertegaal, 1999) was one of the first attentive groupware systems. Using eye trackers, it communicates the visual attention of the participants during mediated group collaborations. GAZE treats awareness (Vertegaal et al., 1997) as an attentive phenomenon, and has been fundamental to a vision in which not just communication systems but all computing systems communicate attention. In GAZE-2 (Vertegaal et al., 2003), streaming media optimize bandwidth according to the user’s visual attention. Video images of users zoom on the basis of visual interest, and audio connections introduce an artificial cocktail party effect on the basis of the visual interest of the group.

Gradual negotiation of turns

Simple User Interest Tracker (*SUITOR*) (Maglio et al., 2000a) was one of the first attentive information systems. *SUITOR* provides a GUI architecture that tracks the attention of users through multiple channels, such as eye tracking, Web browsing and application use. It uses this to model the possible interest of the user, so as to present her with suggestions and Web links pertaining to her task. In order not to interfere with the user's foreground task, it displays all suggestions using a small ticker-tape display at the bottom of the screen. *SUITOR* shows the importance of modeling multiple channels of user behavior and demonstrates how to use a peripheral low-density display to avoid interrupting a user with information of which the relevance to her foreground task is not fully known.

Pong is a robot head that rotates to face users by tracking pupils with a camera located in its nose (Morimoto et al., 2000). *FRED* (Vertegaal et al., 2001) is an attentive embodied conversational system that uses multiple animated-head models to represent agents on a screen. Agents track eye contact with a user to determine when to take turns. Pong and *FRED* show how anthropomorphic cues from head and eye activity may be used to signal device attention to a user, and how speech engines can track eye contact to distinguish what entity a user is talking to. *FRED* shows how proximity cues may be used to move from foreground to peripheral display with malleable resolution. When the user stops talking and fixating at an agent, it looks away and shrinks to a corner of the screen. When users produce prolonged fixations at an agent and start talking, the agent makes eye contact and moves to the foreground of the display. Maglio et al. (2000b) and Oh et al. (2002) demonstrated that, when issuing spoken commands, users do indeed look at the individual devices that execute the associated tasks. This means eye-contact sensing can be used to open and close communication channels between users and remote devices, a principle known as 'look-to-talk'. EyeR (Selker, 2001) is a pair of tracking glasses designed for this purpose. By emitting and sensing infrared beams, these glasses detect when people orient their head towards another device or user with EyeR. EyeR does not sense eye position. It stimulated us to develop eye trackers suitable for look-to-talk: low-cost, calibration-free, long-range, wearable eye-contact sensors.

Augmentation of attention: less is more

Attentive focus through multi-resolution vision is a fundamental property of the human eye. The acuity of our retina is highest in a 2-degree region around the

visual axis, the fovea. Beyond 5 degrees, visual acuity drops into peripheral vision (Duchowski, 2003). Gaze-contingent displays update their images in between fixations to allow alignment of visual material with the position of the fovea, as reported by an eye tracker. Originally invented to study vision, reading and eye disease, they are now used to optimize graphics displays (McGonkie & Rayner, 1975; Duchowski, 2003). By matching the level-of-detail of 3D graphic card rendering with the resolution of the user's eye, virtual-reality display is improved (Murphy & Duchowski, 2001). This technology inspired our design for dynamic multi-resolution windows, discussed below.

With the move towards context-aware interfaces (Moran & Dourish, 2001), we are seeing increased use of attentive visualization in HCI. 'Focus + Context' (Baudisch, Good & Stewart, 2001) is a wall-sized low-resolution display with a high-resolution embedded display region. Users move graphic objects to the hi-res area for closer inspection without losing the context provided by peripheral vision. It is an elegant example of static multi-resolution windows. 'Popout Prism' (Suh et al., 2002) focuses user attention on search keywords found in a document by presenting keywords throughout a document in enlarged, colored boxes. Such attentive user interfaces are distinct from context-aware interfaces in that they focus on designing for attention.

Architects and designers, such as Mies van der Rohe (Carter & Mies van der Rohe, 1999), have long advocated focusing design resources in ways that provide synergies between manufacturing, human factors and aesthetic requirements. His adage 'Less is more' reflects the need to consider human attention in design. Many tools can be characterized as having been designed with attentive properties in mind. The thin blue lines that aid handwriting on paper are a good example. Since peripheral vision is least sensitive to blue detail, the lines are visible only when you need them (Duchowski, 2003). According to Goldhaber (1997), the Internet can be viewed as an economy of attention. Drawing analogies with human group communication, Goldhaber argues convincingly that buying and selling attention is its natural business model. Indeed advertising agencies sell page views, while the Google search engine ranks results by the number of outside links to a page. Our framework expands upon the basic principles outlined by these designers to create attention-aware systems that truly augment the user's attention and, in so doing, his or her intellect.

Creating effective attentive user interfaces

As a goal, attentive user interfaces emphasize the design of interactions such that they optimize the use of the user's attentive resources. At this point we

describe our efforts towards the development of a number of attentive user interface prototypes, along the categorization provided.

Sensing attention: the eye-contact sensor

With the design of eye-contact sensors, or ECS, we wanted to push attention sensing, in the form of eye tracking, beyond desktop use. Current desk-mounted eye trackers limit head motion of the user and do not track beyond a 60 cm distance (Duchowski, 2003), restricting the user to a localized area with little movement. In addition, head-mounted portable eye trackers are expensive, obtrusive and difficult to calibrate (Duchowski, 2003). To implement a system analogous to look-to-talk with ubiquitous computers, we needed a cheap ubiquitous input device that sensed eye contact only. The \$1000 eye-contact sensor consists of a camera that finds pupils within its field of view using computer vision (see Figure 3). A set of infrared LEDs is mounted around the camera lens. When flashed, these produce a bright pupil reflection (red-eye effect) in eyes within range. Another set of LEDs is mounted off-axis. Flashing these produces a similar image, with black pupils. By syncing the LEDs with the camera clock, a bright-and-dark pupil effect is produced in alternate fields of each video frame. A simple algorithm finds any eyes in front of a user by subtracting the even and odd fields of each frame (Morimoto et al., 2000). The LEDs also produce a reflection from the surface of the eyes. These appear near the center of the detected pupils when the onlooker is looking at the camera, allowing the detection of eye contact without any calibration. Eye-contact sensors stream information about the number and location of pupils, and whether these pupils are looking at the device over a wireless TCP/IP connection. When mounted on any ubiquitous device, the current prototype can sense eye contact with the device at up to 3–10 meters distance. By mounting multiple eye-contact sensors on a single ubiquitous device and by networking all eye-contact sensors in a room, eye fixations can be tracked with great accuracy throughout the user's environment.

Reasoning about turns: eyePliances and eyeReason

EyePliances are smart ubiquitous appliances with embedded attention sensors, designed to extend the existing concept of gradual turn-taking. Users interact with eyePliances through speech, keyboard, radio tags (Want et al., 1999) or manual interaction. Functionality in appliances is accessed through X10 home automation



FIGURE 3
Eye-contact sensor



FIGURE 4
AuraLamp eyePliance

software (X10, 2002) and wireless Internet connectivity. Figure 4 shows the simplest form of eyePliance, a light fixture appliance with an embedded eye-contact sensor (Mamuji et al., 2003). Using speech recognition, the light is turned on or off by saying 'On' or 'Off' while looking at its eye-contact sensor. Using eye-contact sensors as pointing devices for the real world eases problems of naming conventions for speech interaction and of juggling remote controls (Vertegaal et al., 2005). When users do use a remote or keyboard to control eyePliances, eye-contact sensing is used to determine the target of keyboard actions.

Figure 5 shows how eyePliances may function in a more complex, attention-sensitive environment that keeps track of the devices users are paying attention to, the preferred notification channels and prioritization of notifications. A personalized central server, called eyeReason, handles all remote interactions of a user with devices, including user notification by devices. Devices report to the server whether a user is working with them and what

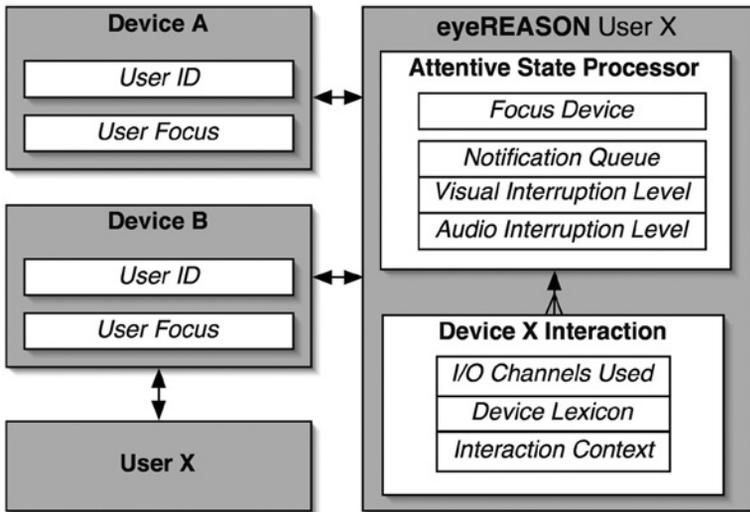


FIGURE 5
EyeReason architecture

that user's focus is by tracking manual interactions and eye contact with the device. Devices may use RFID tags (Want et al., 1999) to identify and detect users and objects. Any speech I/O (Input/Output) with a user is processed through a wireless headset by a speech recognition and production system on the server. As the user works with various devices, eyeReason switches its context to the lexicon of the focus device, sending commands to that device's I/O channels.

Communicating attention

As the GAZE systems showed (Vertegaal et al., 2003), AUIs can also communicate attention to others. AuraMirror is a media artwork, a video mirror that renders the virtual windows of attention through which we interact with other people. AuraMirror provides an ambient display which renders visualizations of virtual bubbles of attention, or 'attentive auras', that engulf groups of people during conversations and that distinguish subgroups in side conversations (Duncan, 1972). It unobtrusively communicates the negotiation of attention, and the effects of intrusion and interruption on this process. The mirror consists of a large plasma display mounted on a wall (see Figure 6).



FIGURE 6
AuraMirror showing merging of auras

This display reflects the world before it by displaying images from a video camera mounted on top of it. The images from the camera are also used to track the orientation of people standing in front of the mirror.

When two people standing in front of the mirror turn to look at each other, the virtual windows of attention between them are visualized through a merging of their auras. However, when people look at the mirror to see this, their auras break apart. It is now a permanent exhibit at the Ontario Science Center in Toronto, Canada.

Augmenting user attention

Attentive user interfaces can also enhance the user's cognitive processes, for example, by filtering out irrelevant information before it even reaches the brain. This allows users to focus their attentive resources differently, and potentially more effectively. The attentive headphones, for example, are a pair of noise-canceling headphones augmented with a microphone and an ECS (see Figure 7). Normally, noise-canceling headphones block out all noise around the user; however, this isolates the wearer from the attention of co-workers. The eye-contact sensors in the attentive headphones allow them to reason about when to turn off the noise cancellation, for instance when somebody makes eye contact with the wearer. The attentive headphones create a gradual, and consequently more natural, turn-taking effect in a social interaction than would otherwise be possible if auditory attention were blocked.



FIGURE 7
Attentive headphones

Attentive cubicles

The next step is to have all social interactions, including co-located ones, mediated by attention-aware systems. In office cubicle farms, where many users share the same workspace, problems of managing attention between co-workers are particularly acute. Our attentive cubicle system (Mamuji et al., 2004) addresses this problem by automatically mediating auditory and visual communications between co-workers on the basis of information about their social-geometrical relationships. Our prototype cubicle's walls were constructed using a special translucent material called Privacy Glass™ (see Figure 8). Privacy glass consists of a glass pane with an embedded layer of liquid crystals. When powered off, the crystals are aligned randomly, making the glass appear frosted and opaque. When a voltage is applied, the liquid crystals in the glass align, allowing light to go through the pane, thus rendering the glass transparent. When the privacy glass is opaque, cubicle workers cannot be seen by others, and are not distracted by visual stimuli from outside their cubicle. When the privacy glass is transparent, a cubicle worker can interact visually with workers on the other side of the cubicle wall. We augmented the privacy glass with a contact microphone to allow our system to detect the sounds of co-workers knocking on the pane. These knocks inform the system of a request for the attention of a person inside an opaque cubicle. To mediate auditory interactions, cubicle workers wear attentive headphones. Upon detection of a request for attention, these headphones automatically become transparent to sound from the outside world.

Users within our office environment are tracked by overhead cameras mounted in the ceiling. This allows the cubicle to detect co-location and co-

orientation of participants, as well as orientation towards objects of joint interest, such as whiteboards. For each tracked individual, the cubicle reports information about potential communication partners to that individual's eyeReason server. The eyeReason server controls the setting of the headset of the associated individual, as well as the transparency of the privacy walls of a cubicle entered by that individual.

Here is a scenario that illustrates how the system is used. Alex, the user, is busy finishing a report. Alex has a tight deadline, as he needs to have the report filed by the end of the day. While Alex is trying to focus on his writing, his colleague Jeff is discussing a design strategy with Laurie, a co-worker, in the next cubicle. All three individuals are wearing an attentive headset tracked by the system. The cubicle recognizes Laurie and Jeff are co-located and oriented towards each other, without any physical barriers between them. It reports each as a potential communication partner to the other person's eyeReason server. This causes their headphones to be set to transparent, allowing Jeff and Laurie to hear each other normally. At the same time, the cubicle detects that co-worker Alex is not co-located with anyone, and is oriented towards his computer. Alex's eyeReason server is notified that there are no apparent communication candidates, causing it to engage noise cancellation and render his cubicle's privacy glass opaque. When Jeff and Laurie require Alex's assistance, Jeff makes a request for Alex's attention by knocking on the cubicle's privacy glass. The request is forwarded to Alex's eyeReason server, which informs the cubicle to consider the wall between the two individuals removed. It also causes Alex's noise cancellation to be turned off temporarily, allowing him to hear the request. As Alex responds to the request, he orients himself to the source of the sound. The cubicle detects the co-orientation of Jeff and Alex. Alex's eyeReason server renders the privacy pane between Jeff and Alex translucent, allowing them to interact normally. After the conversation is completed, Jeff moves away from the cubicle wall, continuing his discussion with Laurie. Alex turns his attention back towards his computer system, causing the cubicle to conclude Alex and Jeff are no longer candidate members of the same social group. Alex's eyeReason server responds by turning on noise-cancellation in Alex's headset, and by rendering the privacy glass of his cubicle opaque again.

This scenario illustrates how entire rooms can be designed to balance social as well as privacy needs of co-workers in a dynamic fashion. The above scenario can also be applied to remote situations.

OverHear

An example of a remote surveillance interface that aims to augment the user's remote auditory attention is OverHear. It consists of an eye-tracking display



FIGURE 8

Attentive office cubicle prototype in opaque (upper) and transparent (lower) modes





FIGURE 9
OverHear remote surveillance interface with eye-tracking display



FIGURE 10
OverHear robotic shotgun microphone in close-up

showing a live audio and video feed obtained from a robotic directional microphone and Webcam at a remote public location (see Figure 9). When the user looks at a particular individual in the video stream, the directional microphone at the other end in the remote location (see Figure 10) will focus upon that person, allowing the user to hear that specific conversation. The OverHear interface simulates and enhances the natural cocktail party phenomenon by blocking out peripheral noise, creating a focus that augments the user's auditory attention in ways otherwise not possible. We anticipate – and are hopeful – that systems such as OverHear will eventually lead to new forms of hearing aids. Systems inspired by OverHear will allow the deaf to listen only to the person they look at, enhancing their hearing by means of an artificially managed cocktail party phenomenon.

Discussion

Throughout the process of designing attentive user interfaces, we came across many circumstances that helped us identify outstanding research questions. Among the concepts we explored, we found the metaphor of virtual windows of attention particularly inspiring. Whether in visual or auditory interactions with remote devices or people, users need to be supported by subtle cues that make up the virtual windows through which entities communicate with them. It is not sufficient to define such windows by the electronic channels through which interactions take place, because electronic channels do not delineate actual attention. By sensing user attention, devices may know when users are attending to them. By providing devices with a means of communicating their attention, users may know they are being attended to as well (Bellotti et al., 2002). This allows users and devices to establish the negotiation of joint interest that is characteristic of multiparty human turn-taking. We invite researchers and designers to further develop and improve upon the conceptual framework provided in this paper.

One of the technical problems we encountered is that of sensing attention for small or hidden devices. While physiological sensing technologies may address these issues, they are potentially invasive.

A second issue is the identification of users at a distance. While eye-contact sensors may one day be able to perform iris scanning, there are privacy implications that must be considered.

A third is that of prioritization of notifications. Can we trust automated services to rank and prioritize the information we receive?

We believe the most pressing issue relating to the sensing technologies presented in this paper is that of *privacy*. How do we safeguard the privacy of users when devices routinely sense, store and relay information about their identity, location, activities and communications with other people?

A Panopticon of attentive user interfaces?

One has only to examine the interface designs in movies such as *Minority report* to appreciate the dangers of incorporating sensing in ubiquitous computing interfaces. The paradox of attentive user interfaces is that sensing here serves to protect the user from unwanted intrusion on their privacy. Just as a secretary requires intimate knowledge of his or her manager in order to process requests in their place, surveillance of the user in AUIs serves to better filter information for the user. Like spam filters, they serve to filter information such that only solicited and pertinent information makes it through to our minds. However, few users of spam filters appear worried about the technological possibility of the developer of the filter monitoring their every email, for example, for marketing purposes (note that Gmail has already made the latter a reality). We conclude that every surveillance technology is a double-edged sword that is in no way different from any other technology that can be applied for good or evil.

Attentive user interfaces can either exacerbate the panoptic effect and widen the scope and net of surveillance, or give us an opportunity to take control of information privacy and manage our own data doubles. An attentive user interface is not only going to contain information as to what purchases we made and where we surf when a particular cookie is active. It will know when, how, where and what occupies our attention in both mediated and non-mediated worlds. It will know our buying preferences, our affiliations with organizations, and who our friends and foes are. It will know exactly what our priorities are, and how we like them sorted. AUIs will also have information about what is accessible to whom within our workplaces and households, including private files and secrets that we keep from each other within our family or work environments. Who we talk to and when we talk to them will be recorded to inform future prioritization of the caller. This could potentially be used to conduct surveillance on each other within the workplace or household. If we provide privacy controls on information within the context of use, restricting access to a resource can imply deceitfulness. In short, the surveillance required to facilitate such a system can produce accurate data doubles based on much more personal information than the current database and cookie-enabled systems, and could affect interpersonal dynamics greatly.

A simple and straightforward solution to this problem is to design AUIs to serve as locally encrypted surveillance systems that keep the information repository private. Using AUIs as personal and private 'diaries' of our activities allows users to benefit from improved communication among ubiquitous devices that respond to our attention without sacrificing informational

privacy. Like our secretaries, these systems can be our own personalized information facilitators that will accommodate our needs and preferences to improve our interaction with technology.

The sousveillance approach

Mann and colleagues (2003) discuss the use of 'sousveillance' as a strategy against surveillance. The idea here is that the watched can record the actions of the watcher, for example with the purpose of making this information public. As was demonstrated by the Abu Ghraib prison scandal, sousveillance technologies, such as offered by ubiquitous cameras, have the power to expose and disrupt illicit activities. If we are worried that the external communications provided by AUIs will be monitored by our service providers, our AUIs can implement automated sousveillance to thwart illicit use of the information. This can, for example, be implemented through automated querying of information that is contradictory to or outside of our real interests during offline time. Sousveillance offers a subtle opportunity for control of our external data double. When we are online, our personal AUI will make requests, actively engaging the useful resources, rather than passively being subjected to personalized direct-marketing information streams. If our AUIs are indeed stored and managed externally, the resources that pop up will certainly be relayed with more than our own interests in mind. AUIs and similar systems have great potential to gather both transactional and interactional information to construct externally managed personal profiles. Therefore, rather than being information facilitators that act solely on our behalf, they might be information gatekeepers with motives that differ from simply improving the quality of our lives with and among technology. This would produce a panoptic effect in which we would adjust our behavior on the assumption that we are always monitored by external interests.

We must be conscious of how access restrictions will affect our interpersonal relationships, by conducting experiments on prototypes, situated within the workplace or within a household. We should strive for ecological validity so we learn about how these systems will really be used not just how we think they will be used, and what the impact of their use might be. This very important research will require social scientists to conduct it properly. Designers must partner with social researchers to make use of the wealth of knowledge they possess, and social scientists must continue to take an active role in the development of new technology in order to influence and improve technologies that are changing the way we live.

Roel Vertegaal is one of Canada's foremost designers of highly futuristic interactive technologies. He is Associate Professor in Human-Computer Interaction at Queen's University in Kingston, Ontario, where he directs the Human Media Laboratory. He is also CEO of Xuuk, Inc., a startup that develops attention sensors for interactive, real-world viewing statistics gathering. *Author's address:* Roel Vertegaal, Director, Human Media Laboratory, Goodwin Hall 636, Queen's University, Kingston, ON, Canada K7L 3N6. [email: roel@cs.queensu.ca]

Jeffrey Shell is an alumnus of the Human Media Laboratory, and is currently a PhD student in Political Studies at John Hopkins University.

Note

1. This work was presented at the Cognitive Technologies seminar session 'Perceiving and being perceived in digital environments', held at the Fondation Maison des Sciences de l'Homme, Paris, 12 June 2006.

References

- Argyle, M. & Cook, M. (1976) *Gaze and mutual gaze*. London: Cambridge University Press.
- Baudisch, P., Good, N. & Stewart, P. (2001) 'Focus plus context screens: combining display technology with visualization techniques', in *Proceedings of UIST 2001*, pp. 31–40. Orlando, FL: UIST.
- Bellotti, V., Back, M., Edwards, K., Grinter, R., Henderson, A. & Lopes, C. (2002) 'Making sense of sensing systems: five questions for designers and researchers', in *Proceedings of CHI 2002*, pp. 415–22. Minneapolis, MN: ACM Press.
- Bolt, R. A. (1985) 'Conversing with computers', *Technology review* 88(2).
- Carter, P. & Mies van der Rohe, L. (1999) *Mies Van Der Rohe at work*. London: Phaidon Press.
- Cherry, C. (1953) 'Some experiments on the reception of speech with one and with two ears', *Journal of the acoustic society of America* 25: 975–9.
- Duchowski, A. (2003) *Eye tracking methodology: theory and practice*. Berlin: Springer-Verlag.
- Duncan, S. (1972) 'Some signals and rules for taking speaking turns in conversations', *Journal of personal and social psychology* 23: 283–92.
- Gibbs, W. (2005) 'Considerate computing', *Scientific American* 292(1): 54–61.
- Goldhaber, M. (1997) 'The attention economy and the net', *First Monday*, <http://www.firstmonday.dk>.
- Horvitz, E. (1999) 'Principles of mixed-initiative user interfaces', in *Proceedings of CHI'99*, pp. 159–66. Pittsburgh, PA: ACM Press.
- Horvitz, E., Jacobs, A. & Hovel, D. (1999) 'Attention-sensitive alerting', in *Proceedings of UAI'99*, pp. 305–13. Stockholm: Morgan Kaufmann.
- Ishii, H. & Ullmer, B. (1997) 'Tangible bits: towards seamless interfaces between people, bits, and atoms', in *Proceedings of CHI'97*, pp. 234–41. Atlanta, GA: ACM Press.

- Jacob, R. (1995) 'Eye tracking in advanced interface design', in W. Barfield & T. Furness (eds) *Virtual environments and advanced interface design*, pp. 258–88. New York: Oxford University Press.
- Maglio, P., Barrett, R., Campbell, C. & Selker, T. (2000a) 'SUITOR: an attentive information system', in *Proceedings of UIST 2000*, pp. 169–76. New Orleans.
- Maglio, P., Matlock, T., Campbell, C., Zhai, S. & Smith, B. (2000b) 'Gaze and speech in attentive user interfaces', in *Proceedings of the international conference on multimodal interfaces*. Berlin: Springer-Verlag.
- Mamuji, A., Vertegaal, R., Dickie, C., Sohn, C. & Danninger, M. (2004) 'Attentive office cubicles: mediating visual and auditory interactions between office co-workers', in *Video proceedings of Ubicomp 2004*.
- Mamuji, A., Vertegaal, R., Shell, J., Pham, T. & Sohn, C. (2003) 'AuraLamp: contextual speech recognition in an eye contact sensing light appliance', in *Adjunct Proceedings of Ubicomp 2003*.
- Mann, S., Nolan, J. & Wellman, B. (2003) 'Sousveillance: inventing and using wearable computing devices for data collection in surveillance environments', *Surveillance and society* 1(3): 331–55.
- McGonkie, G. W. & Rayner, K. (1975) 'The span of effective stimulus during a fixation in reading', *Perception and psychophysics* 17: 578–86.
- Moran, T. P. & Dourish, P., eds (2003) *Human–computer interaction* 16 (special issue on context-aware computing).
- Morimoto, C., et al. (2000) 'Pupil detection and tracking using multiple light sources', *Image and vision computing* 18.
- Murphy, H. & Duchowski, A. (2001) 'Gaze-contingent level of detail rendering', in *EuroGraphics 2001*.
- Nielsen, J. (1993) 'Noncommand user interfaces', *Communications of ACM* 36(4): 83–99.
- Oh, A., et al. (2002) 'Evaluating look-to-talk: a gaze-aware interface in a collaborative environment', in *Extended abstracts of CHI 2002*, pp. 650–1. Seattle, WA: ACM Press.
- Selker, T. (2001) 'Eye-R, a glasses-mounted eye motion detection interface', in *Extended abstracts of CHI 2001*. Seattle, WA: ACM Press.
- Shell, J. S., Selker, T. & Vertegaal, R. (2003) 'Interacting with groups of computers', *Communications of the ACM* 46(43): 40–6.
- Short, J., Williams, E. & Christie, B. (1976) *The social psychology of telecommunications*. London: John Wiley.
- Smith, D., Irby, C., et al. (1982) 'Designing the Star user interface', *BYTE* 7(4).
- Suh, B., Woodruff, A., Rosenholtz, R. & Glass, A. (2002) 'Popout Prism: adding perceptual principles to overview + detail document interfaces', in *Proceedings of CHI 2002*, pp. 251–8. Minneapolis, MN: ACM Press.
- Vertegaal, R. (1999) 'The GAZE groupware System: mediating joint attention in multiparty communication and collaboration', in *Proceedings of CHI '99*. Pittsburgh, PA: ACM Press.
- Vertegaal, R. & Ding, Y. (2002) 'Explaining effects of eye gaze on mediated group conversations: amount or synchronization?', in *Proceedings of CSCW 2002*. New Orleans, MI: ACM Press.
- Vertegaal, R., Cheng, D., Sohn, C. & Mamuji, A. (2005) 'Media eyePliances: using eye tracking for remote control focus selection of appliances', in *Extended abstracts of ACM CHI 2005*. Portland, OR: ACM Press.
- Vertegaal, R., Slagter, R., Van der Veer, G. & Nijholt, A. (2001) 'Eye gaze patterns in conversations: there is more to conversational agents than meets the eyes', in *Proceedings of CHI 2001*, pp. 301–8. Seattle, WA: ACM Press.

- Vertegaal, R., Velichkovsky, B. & Van der Veer, G. (1997) 'Catching the eye: management of joint attention in cooperative work', *SIGCHI Bulletin* 29(4).
- Vertegaal, R., Weevers, I., Sohn, C. & Cheung, C. (2003) 'GAZE-2: conveying eye contact in group videoconferencing using eye-controlled camera direction', in *Proceedings of CHI 2003*, Fort Lauderdale, TX: ACM Press
- Want, R., Fishkin, K., et al. (1999) 'Bridging physical and virtual worlds with electronic tags', in *Proceedings of CHI'99*, pp. 370–7. Pittsburgh, PA: ACM Press.
- Weiser, M. (1991) 'The computer for the 21st century', *Scientific American* 265(3): 94–104. X10 Home Automation, <http://www.x10.com> (2002).
- Zhai, S., Morimoto, C. & Ihde, S. (1999) 'Manual and Gaze Input Cascaded (MAGIC) pointing', in *Proceedings of CHI'99*, pp. 246–53. Pittsburgh, PA: ACM Press.