


Focusing on the Essential: Considering ATTENTION in Display Design

Attentive displays address the need for rendering power and computer display resolution.


The five examples presented here illustrate a common goal with very different approaches to achieving it.

Over the past years, industry has been addressing the resulting demand by offering displays of steadily increasing resolution, reaching resolutions of over nine million pixels (IBM T220 display). Although high resolution is desirable for a variety of applications, it results in an ongoing challenge for creators of rendering hardware, as the large number of pixels makes these displays especially hungry for computational resources.

Displaying computationally intense graphics, such as flight simulation or interactive scientific visualization, requires considerable rendering effort. It is important to note that when computing power is insufficient to support the task of the user, any benefits of large-screen technology to user productivity may be negated. In many cases, this issue can be addressed with parallel rendering hardware; a display system consisting of an array of projec-



tors, for example, is often driven by an array of PCs or a PC cluster. Parallel hardware, however, leads to substantially increased costs and space requirements. Furthermore, in the case of projector array-based displays, the increased need for rendering hardware is accompanied by the cost of the projector technology. So when people make decisions about display sizes and resolutions, not only do cost and space requirements for rendering





hardware and display play an important role—user productivity itself is at stake.

Several research projects have addressed the demand for rendering power and display resolution by taking the user's attentional focus into account. In this article, we will use the term “attentive displays” to refer to this class of techniques. Attentive displays address the demand for rendering power and display resolution, but their approach is different from the parallel display approach described here. Instead of requiring additional hardware, attentive displays make more out of available hardware by directing display and computation resources to where they count most. As we will illustrate, a single user can only be attending to a relatively small part of a display at a time (while multiple users can be attending to more than one location, presentations to groups are more likely to motivate more substantial computational resources). So, instead of rendering information at the same level of detail everywhere, these displays track the user's attention and render information in full detail only at the user's current focus of attention, while reducing information in peripheral areas. By shifting computational power

from peripheral regions to the region located in the user's focus of attention, attentive display systems can provide faster response times with higher subjective display quality than systems distributing their resources equally across the screen.

In order to achieve the focal effect, attentive displays need an add-on—a device informing the display about the user's current focus of attention. While other related approaches use models of attention based on properties of the displayed scene [5], most of the approaches we survey in this article use an eye tracker for this purpose. While eye tracking has long involved complex technology, recent technological progress in this area, as well as comparably moderate accuracy requirements, allow attentive displays to use relatively simple and inexpensive

Figure 1. Gaze-contingent display shows a scene from the movie *The Gladiator*. As the user focuses on the face of the shot's main character, all other display content is rendered at reduced resolution. This type of display can be used for gaze-contingent compression purposes or for the study of human visual perception—in this case the display is used to study glaucoma patients. (Original image © 2000 DreamWorks SKG and Universal Studios; gaze-contingent rendering and resolution map courtesy of Bill Geisler and Jeff Perry.)

By Patrick Baudisch, Doug DeCarlo, Andrew T. Duchowski, and Wilson S. Geisler

trackers. (For details, see Zhai's article in this section.) A general survey of eye tracking techniques can also be found elsewhere [4].

Attentive display prototypes have been applied to a variety of visually demanding tasks, spanning a wide range from driving simulators to advertisements and art [4]. Here, we review five different approaches to degrading the resolution of the screen in peripheral regions. The presented techniques cover a variety of different methods, encompassing customized display hardware and software and range from real-time animation to artistic applications. The first four of the presented displays aim to match the subjective quality of a non-degraded display. To prevent users from noticing the drop in peripheral resolution, that is, in order to design an *imperceptibly degraded* display, the size of the foveal regions of these displays is designed to at least match the extent of the user's *perceptual span* (see the sidebar, "The Perceptual Span"). In order to do so, these displays use a model of size and resolution of foveal and peripheral vision. The last of the five presented displays, however, does not try to obtain an imperceptibly degraded display. Instead, it quite noticeably removes image content to achieve a different effect; by presenting viewers with only the most important information, it aims at reducing the viewer's cognitive load.

Gaze-Contingent Displays

Our first example is a gaze-contingent display (GCD) [10]. GCDs degrade the resolution of peripheral image regions in order to reduce computational effort during image transmission, retrieval, or display. Figure 1 shows an example of a movie scene rendered using a GCD. As the user focuses on the face of the shot's main character, all other display content is rendered at reduced resolution, substantially reducing the rendering effort for this frame. As the movie plays, the high-resolution region moves with the user's focus of attention, so that the spot at the user's focus of attention is always rendered in high resolution. This effect is achieved by tracking the user's gaze with an eye tracker.

By compressing peripheral image information not resolvable by the user, GCDs help increase display speed. Applications include flight and driving simulators, virtual reality, infrared and indirect vision, remote piloting, robotics and automation, teleoperation, and telemedicine; image transmission and retrieval, and video conferencing [10]. In addition to these applications, GCDs have been invaluable for the purpose of studying perception, for example in order to obtain measurements of the human perceptual span such as those presented in the sidebar.

Designing an imperceptibly degraded GCD, that is, one indistinguishable from a full-resolution display, is desirable but difficult [10]. However, for cer-

The Perceptual Span

High-resolution color vision relies on cone photoreceptors located primarily in the fovea (the retina's small central region). Human ocular physiology places a limit on the range of the perceptual span: fine-grained visual acuity is limited by the dimension of the fovea to roughly 2° visual angle, or about 1%–3% of the field of view. A good rule of thumb to remember is the area of high resolution projected onto the fovea is about as large as one's thumbnail at arm's length. At a typical reading distance (30 cm), the foveal region is only about 2 cm in diameter—about as large as the red dot in this box.

Although the quality of the information extracted by the visual system declines with eccentricity, the useful field of view may extend considerably further, depending on the task. For tasks that require detailed discrimination, for example, when reading a newspaper, the useful field of view may cover only a fraction of the fovea. For readers of alphabetical orthographies, such

as English or French, for example, the span extends from 3–4 letters to the left of fixation to about 14–15 letter spaces to the right of fixation [9].

While reading tasks are comparably well understood, viewing tasks populate a much wider spectrum. Contrary to reading, there is no particular correct way to look at a picture. Context differences are generally at play and viewing behavior and eye movement patterns change as a function of the task. Viewing tasks range widely, from viewing art to performing visual search, such as in target detection, from driving to performing visual inspection, or from looking at advertisements to viewing a user interface. For tasks not requiring detailed discrimination, such as visual navigation or detection of large moving objects, the useful field of view may extend to a full 180°. For attentive display design, it is therefore important to consider the dominant viewing task for which the display will be used. **C**





Figure 2. Focus plus context screens complement a monitor-sized high-resolution area in the screen center with a large low-resolution space in the periphery. Courtesy of the Palo Alto Research Center.

tain tasks, such as visual search, the reduction in resolution may not necessarily interfere with user performance, even if the peripheral degradation of a GCD is quite noticeable. Given sufficient accuracy in tracking the user's eye movements and a sufficiently fast-moving foveal region, a GCD with two resolution regions can moderately degrade the peripheral region, while still producing search performance comparable to a full-resolution display—a foveal region of a 5° viewing angle was sufficient for this purpose [10].

Prior research of GCDs has mainly addressed lossy resolution compression of peripheral image regions. New research extends GCDs to support arbitrary resolution maps, which allow exploring two new aspects of GCDs [8]. First, this enhancement allows creating foveal regions of arbitrary shape and size with peripheral regions degraded by arbitrary means, for example, color or contrast reduction, not only resolution. The decoupling of resolution degradation from rendering allows the generation of high-quality images with minimal artifacts at real-time display frame rates [8]. Second, this GCD system allows the display of multiple foveal regions at the same time. Multiple foveal regions provide a suitable display strategy for future systems capable of predicting the user's next point of focus.

GCDs have been successfully deployed to save rendering effort [4]. However, while peripheral content is rendered in low resolution, the display hardware on which it is displayed is still the same resolution as any

other part of the screen surface. On large screens, where the greater part of the screen surface maps to the user's peripheral vision, this seems especially wasteful of the hardware. This aspect is addressed by focus plus context screens—an attempt to make better use of display hardware.

Focus Plus Context Screens

Focus plus context screens achieve a high-detail/low-detail effect by combining a wall-sized low-resolution display with an embedded high-resolution screen [1]. The installation shown in Figure 2 uses an LCD inset combined with projection for generating the low-resolution context. The shown version uses a fixed-position high-resolution focus screen; the iconic illustration at the bottom right shows where it is located. The inset shows the difference in resolutions between the focus and the context area. While the focus area offers enough resolution to allow users to see individual cars, the coarse pixels in the context area merely allow seeing larger objects, such as buildings.

In the example shown, the user is inspecting a specific neighborhood on a satellite image of San Francisco. If the user were using a regular-sized monitor showing the same level of detail as the shown setup, only the neighborhood of interest would be visible, without visual context. With residential areas looking very much alike, it would be difficult for the user to tell where the shown portion of the satellite image is located within the city, potentially disorienting the viewer. Adding the low-resolution context screen space brings the Bay bridge and the piers into view, providing additional landmarks that simplify orientation. When the user moves the mouse, the entire display content pans, which allows scrolling display content into the focus region in order to make it high resolution.

For tasks involving large maps or detailed chip designs, focus plus context screens were shown to allow users to work from 20%–35% faster than when using displays with the same number of pixels, but in homogeneous resolution or with multiple views. For an interactive driving simulation, users' error rates were only a third of those in a competing multiple-view setup [1].

In applications that continuously draw the user's attention to the focus area, as is the case in the driving simulation used in the experiment, focus plus context screens with a fixed-position focus succeed because the display's focus and context regions cover the user's foveal and peripheral vision the same way a corresponding high-resolution screen does. This makes this type of focus plus context screen, which can be built from comparably inexpensive off-the-shelf components, a cost-effective alternative to complex multiprojector high-resolution screens. By slaving the focus



Figure 3. Gaze-contingent spatial and temporal LOD modeling. As the viewer focuses outside the room at the left of the rendering (image at left, courtesy of David Luebke), scene objects located at the right side of the room are rendered using a lower level of spatial detail, indicated by larger triangles (overlaid). Collisions between L-shaped objects (image at right, courtesy of Carol O'Sullivan and John Dingliana) are calculated at a higher level of temporal detail if located within the user's current focus of attention.

display to the user's gaze, future versions may obtain high resolution wherever the user looks, thereby widening the applicability of focus plus context screens to applications where users continuously look around.

Real-Time 3D Graphics

Both GCDs and focus plus context screens degrade peripheral information by manipulating the image (that is, pixel) properties of the display. In computer graphics research, rendering speed is a primary concern. Interactive applications, such as virtual reality, demand high frame rates in order to satisfy real-time interaction and display. For complex scenes consisting of a large number of polygons, such as virtual terrain containing significant topological detail, or when using computationally expensive rendering techniques such as ray tracing or radiosity, achieving an acceptable combination of surface detail and frame rate requires a substantial hardware effort. Researchers are therefore exploring attentive UI techniques directing the bulk of system resources toward the scene components delivering the highest perceptual impact. One prominent example of an attentive 3D-rendering engine varies the Level of Detail (LOD) at which an object is drawn based on the user's gaze [6]. This way, unattended scene objects are modeled with fewer polygons, even when they are not distant in the scene. Gaze-contingent LOD reduction is similar to the GCDs in that both techniques reduce the complexity of the displayed image, however, unlike GCDs, graphical methods do so at the object geometry level, rather than at the image level.

Gaze-contingent LOD reduction was found to lead

to substantial performance improvements. In the example shown in Figure 3 (left), a reduction of the number of triangles by 70% still leads to an imperceptibly degraded display [6].

Gaze-contingent modeling has also been applied to real-time temporal resolution degradation [7]. The degradable collision handling mechanism shown in Figure 3 (right) evaluates object collisions inside the user's focus of attention with greater precision than collisions occurring in the user's periphery. The highlighted circle in the inset indicates the field of 4° visual angle inside which collisions are processed at greater precision. Saving processing time for collisions outside this area allows spending extra processing time on collisions in the user's focus of attention, which results in an overall improvement in the perception of the simulation.

Easily Perceived Displays

While the approaches described here *follow* the user's attention, attentive displays have also been used to *direct* the viewer's attention, for example, in the context of art. Artists have long been able to draw viewers' attention to specific artwork regions. Consider the painting shown in Figure 4. By controlling luminance, color contrasts, and depth cues, the painter is guiding the viewer's gaze toward the depictions of Christ and the kneeling woman.

The artist's success is evidenced in a recent large-scale eye-tracking study [11], which shows how only the main two figures in the image were fixated with the remainder of the image left largely unnoticed (see the inset). Neuroscientists such as Zeki [12] claim this lightens viewers perceptual burden, and enables them



Figure 4. Aggregated fixations from 131 subjects viewing Paolo Veronese's *Christ addressing a Kneeling Woman*. Subjects' gaze is drawn to the two main figures. (Original image © National Gallery, London, annotations © IBS, University of Derby, UK, courtesy of David Wooding.)

to look deeper into a piece of art, as the artist has left viewers with simpler visual inferences to make.

Work in the field of nonphotorealistic rendering [3] uses similar techniques to guide the viewer's attention and to allow computer generation of aesthetically pleasing images. Figure 5 shows an example. This system employs a perceptual model that works from gaze recordings from a single user (see the inset) to decide which parts of a photograph should be removed, as eye movement patterns are good indicators for what is important to the viewer [9]. Instead of blurring away detail where the user didn't look, the result is stylized using smooth black lines and colored regions. This produces a rendering that guides the viewers' attention to what the original user found important. This way, the incorporation of one viewer's gaze guides the attention of future viewers.

Conclusion

In this article, we presented five examples of attentive displays. All five techniques attempt to match the characteristics of computer displays to the characteristics of

human vision, namely its distinction between foveal vision and peripheral vision. They all try to make better use of limited rendering resources by tailoring display content to the affordances of human vision. The presented techniques differ, however, in which resource they try to preserve and in their adopted strategies for achieving this goal. Each of the discussed techniques falls onto a different point in the spectrum of attentive displays. GCDs, as well as the two presented 3D approaches, improve display frame rates and responsiveness given certain rendering hardware; focus plus context screens achieve better immersion and visual context with given display hardware; and nonphotorealistic rendering saves maybe the scarcest resource of all—the user's attention.

As rendering and display hardware continue to increase in power and decrease in cost, users will continue to see improved rendering quality on their computer screens. Desktop PCs will be able to display real-time graphics at a quality corresponding to that of today's movies—graphics that today require hours of offline rendering. As this happens, theater-quality graphics will have advanced another step, inching closer to the as-yet distant goal of *photorealism*. But in the future, we will also see more kinds of displays, in more places, and for more applications. The most effective use of these displays will seamlessly integrate



Figure 5. This gaze-based drawing was generated by transforming a photograph based on a user's fixations (inset). (From [3] © 2002 ACM, courtesy of Doug DeCarlo. Original photo courtesy of philip.greenspun.com.)

the requirements of the task and the needs of the user. Despite rapid technological progress, however, users will be limited by their current hardware configuration, no matter what state of advancement it is in. There will always be a desire to stay one step ahead of the current state of the art. The techniques described in this article will offer one possibility of doing so. These considerations suggest that attentive displays will be an enduring factor in the design of interactive computer systems. **C**

REFERENCES

1. Baudisch, P., Good, N., Bellotti, V., and Schraedley, P. Keeping things in context: A comparative evaluation of focus plus context screens, overviews, and zooming. In *Proceedings of CHI '02* (Minneapolis, MN, Apr. 2002). ACM, NY, 259–266.
2. Czerwinski, M., Tan, D.S., and Robertson, G.G. Women take a wider view. In *Proceedings of CHI '02* (Minneapolis, MN, Apr. 2002). ACM, NY, 195–202.
3. DeCarlo, D. and Santella, A. Stylization and abstraction of photographs. In *Proceedings of ACM SIGGRAPH 2002, ACM Transaction on Graphics* 21, 3 (2002). ACM, NY, 769–776.
4. Duchowski, A.T. *Eye Tracking Methodology: Theory & Practice*. Springer-Verlag, London, UK, 2003.
5. Horvitz, E. and Lengyel, J. Perception, attention, and resources: A decision-theoretic approach to graphics rendering. In *Proceedings of UAI '97* (San Francisco, CA, 1997). Morgan Kaufmann, San Francisco, CA, 238–249.
6. Luebke, D., Reddy, M., Cohen, J., Varshney, A., Watson, B., and Huebner, R. *Level of Detail for 3D Graphics*. Morgan-Kaufmann, San Francisco, CA, 2002.
7. O'Sullivan, C., Dingliana, J., and Howlett, S. Gaze-contingent algorithms for interactive graphics. *The Mind's Eyes: Cognitive and Applied Aspects of Eye Movement Research*. J. Hyönä, R. Radach, and H. Deubel, Eds. Elsevier Science, Oxford, UK, 2002.
8. Perry, J.S. and Geisler, W.S. Gaze-contingent real-time simulation of arbitrary visual fields. *Human Vision and Electronic Imaging* (San Jose, CA, 2002), SPIE.
9. Rayner, K. Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin* 124, 3 (1998), 372–422.
10. Reingold, E.M., Loschky, L.C., McConkie, G.W., and Stampe, D.M. Gaze-contingent multi-resolutional displays: An integrative review. *Human Factors* (2002). In press.
11. Wooding, D. Fixation maps: Quantifying eye-movement traces. In *Proceedings of ETRA '02* (New Orleans, LA, 2002). ACM, NY.
12. Zeki, S. *Inner Vision: An Exploration of Art and the Brain*. Oxford University Press, Oxford, UK, 1999.

PATRICK BAUDISCH (patrick.baudisch@acm.org) is a research scientist at the Large Display User Experience Group at Microsoft Research, Redmond, WA.

DOUG DECARLO (decarlo@cs.rutgers.edu) is an assistant professor of computer science with a joint appointment in cognitive science at Rutgers University, New Brunswick, NJ.

ANDREW DUCHOWSKI (andrewd@cs.clemson.edu) is an associate professor of computer science at Clemson University, Clemson, SC.

WILSON GEISLER (geisler@psy.utexas.edu) is a professor of psychology at the University of Texas at Austin, TX.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.