

AUTODESK<sup>®</sup>  
MAYA<sup>®</sup> 8.5  
Now Available

[Find Out More >](#)

SUBSCRIBE > e-Newsletter

> Magazine

SEARCH

GO Advanced

[Home](#) | [Current Issue](#) | [Gallery](#) | [Portfolios](#) | [News](#) | [Contact Us](#) | [Job Board](#) | [CG Forum](#) | [Resources](#)

## Computer Graphics World - Issue Articles

v advertisements v

You are here: Computer Graphics World : CGW : Articles : Seeing with the Mind's Eye

 printer friendly format  email page

Issue Date: Volume: 24 Issue: 10 (October 2001)

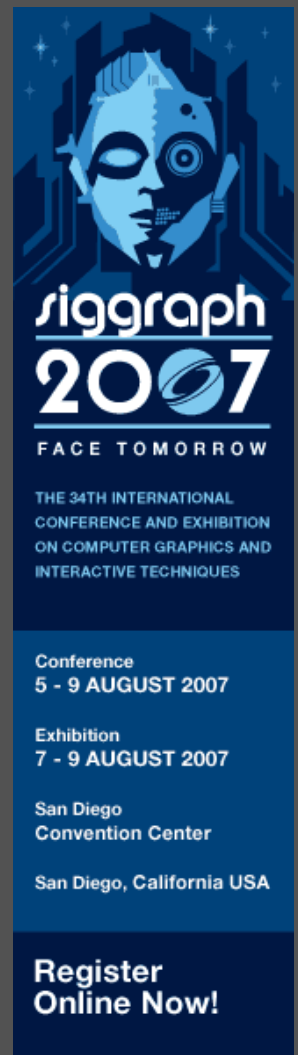

### Seeing with the Mind's Eye

By Diana Phillips Mahoney

The rapid-fire advances that have defined computer graphics over the past quarter century have filled our computer screens with more visual information than we know what to do with. Simulations that once required supercomputer power to calculate are now running on desktop workstations, and high-end animations and special effects for film, broadcast, and Internet applications are being churned out in minutes and hours rather than days, weeks, and months. But while Moore's law has delivered on its promise of ever-faster processors and greater storage and memory capacity, these advances have done little to help us grasp all aspects of the images on our computer screens. In fact, in some cases, the sheer volume of data that the new technologies allow us to generate actually undermines our ability to "see" important information.

In an effort not only to deal with the visual data overload, but also to benefit from it, a growing number of computer graphics researchers and practitioners are focusing their sights beyond the visual representations of data toward the study of how we perceive such information. The ultimate goal of these efforts is to generate images that complement human perceptual processes in order to communicate the visual data as effectively as possible.

Although perception issues have received a lot of attention in the visual computing community recently, they are not really new to computer graphics. In fact, says professor Jim Ferwerda, a researcher in the Computer Graphics Laboratory at Cornell University, "whether or not they realize it, animation and visualization designers use perceptually based stuff all of the time. The color spaces they use, the NTSC standard, JPEG, MPEG, display devices—all have roots in perception."



**siggraph**  
**2007**  
FACE TOMORROW

THE 34TH INTERNATIONAL  
CONFERENCE AND EXHIBITION  
ON COMPUTER GRAPHICS AND  
INTERACTIVE TECHNIQUES

Conference  
5 - 9 AUGUST 2007

Exhibition  
7 - 9 AUGUST 2007

San Diego  
Convention Center

San Diego, California USA

**Register  
Online Now!**

Lighter objects in a CG scene appear less glossy than darker ones (top) because of diffuse reflection. Using a vision model based on human perception, researchers at Cornell are able to generate images that appear similar in gloss despite their lightness



Come See  
For Yourself  
on  
November 14th  
[www.newtek.com/shows/selaunch](http://www.newtek.com/shows/selaunch)

What is new, however, is the appreciation of the benefits that can be had by incorporating tenets of perception theory into the computation and design of CG animation and visualization. And this appreciation comes from considering the problem not from a purely computer-science perspective, but from a multidisciplinary bridge that links the graphics community to physicists, vision researchers, and experimental psychologists.

Although the specific nature of the perception-based research and development efforts varies, a common thread is that most are seeking to incorporate a model of human perception into the visualization/animation design process in order to create "better" images. "Better" in this sense may mean images that are more physically realistic or data representations that promote an intuitive understanding of the numerical information being displayed. A secondary goal is to use the perception model to make processing more efficient. By understanding the limits of human vision and perception, designers can figure out what is necessary to show in order for the images to be effective, and they can also determine what computations can be avoided.




Salient structures in a volume-rendered abdominal image are highlighted using such illustration techniques as boundary and silhouette enhancement. This makes the visual information easier to understand, says Purdue researcher David Ebert.

Often, these two goals are synergistic. At Cornell, for example, Ferwerda and his colleagues use visual-perception information in their realistic image synthesis work. "To enhance the realism of images, we simulate the physics of light reflection. This is hugely computationally expensive, so we use what we know about perception to increase the efficiency of the algorithms-to avoid things that aren't going to be visible or important."

The first step in this process is to quantify human perception. "There's this huge body of literature extending over 150 years of experiments that tries to establish mathematical relationships between the physical properties of the world and the way we perceive those properties," says Ferwerda. "We take this information and use it to build computational models of vision and apply the models to various graphics problems."

*The human visual system operates differently in light and dark environments. To simulate this phenomenon, researchers at Cornell have built a vision model that maps the computed physical values of a scene to the display, accurately re-creating the variations that we perceive in the physical world. Here, the vision model was used to simulate the appearance of a street scene both under and outside of a bridge, and to depict the visible differences in colors and text viewed in daylight and moonlight.*


For example, Ferwerda and his colleagues have developed a vision model that maps the intensity values of a scene onto the display in different ways depending on whether the scene is supposed to show a day or a night view. "Vision operates differently in daytime and at night. To simulate the difference, we take the raw physical values that we've computed [for the scene] and put them through this vision model."



What's critical, says Ferwerda, is that the vision model does more than simulate the physics of light. "For a long time, people thought that if they simulated the physics of light reflection correctly or modeled the geometry correctly, they could produce images that were absolutely realistic. Well, it turns out that there's a lot more to realism than just reproducing light and geometry. Part of our mission is trying to understand what makes something realistic. What errors do we make that cause things to look fake?"

Because such perception simulation is so comprehensive, corners have to be cut in order to speed up the processing and visual representation of the information. An example of one such short cut is the implementation of perception-based camouflage techniques, such as masking. "In the psychology literature, there's information on such things as how one pattern in a scene can hide or obscure another. By modeling this phenomenon and putting it into the rendering algorithm, we can avoid computing more than we have to," says Ferwerda.

For example, the Cornell researchers have built a model that predicts how a texture map will hide some artifacts in an image, such as noise or banding. "This allows us to approximate certain things rather than compute them completely, because the system tells us the texture is going to hide the problem," says Ferwerda.



Traditional weather maps (top) are ill suited for the simultaneous display of multiple conditions, such as temperature, pressure, wind, and precipitation. To optimize the understanding of multidimensional relationships, researchers at North Carolina State

The perception issues being addressed in the graphics community are far from niche, application-dependent interests, says Ferwerda. "The perception research is a generally enabling endeavor, either in terms of allowing you to deal with bigger models or simulations, or increasing the performance of programs or making things more visually accurate. This cuts across all graphics domains."

Holly Rushmeier, a perception researcher at IBM's T.J. Watson Research Center, agrees. "In computer graphics and visualization, we are creating images for people to look at, so every application, every technique we develop, has perceptual issues to be resolved," she says.

Among the specific perception issues being considered by Rushmeier and her IBM colleagues are the use of human vision models to effectively map tone variations in an image, to determine the "allowable error" in global illumination calculations, and to create perceptually optimal geometric descriptions of objects.

Of these objectives, the tone-mapping problem is perhaps the most complex. "When we simulate an image of a real scene, the radiance we compute can have a dynamic range of 10,000:1. We will display the result on a device with a range of 100:1." As of yet, she says, there is no single solution for bridging this discrepancy. "It turns out you get different results if you use a method that tries to preserve the sense of brightness [being able to tell if a search light or dim bulb was used to light a scene, for example] versus one that considers visibility [the ability to discern the presence of an object only if you were able to do so if you were in the real environment]."

Perceptual models can help gauge the relative importance of these factors with respect to a given scene. A natural follow-on to this, says Rushmeier, is to use the perception information to determine the allowable error in an illumination simulation and thus prevent computing information that will have no value to the viewer. "In absolute terms, you may compute a pixel value with an estimated error of 30 percent, based on how it will be mapped to the display." Once the perceptible values have been mapped, she says, "you can stop your calculation." For the tone-mapping and error-bounding applications, the IBM researchers rely on vision models built by the image-processing community.

The third research area, the geometry problem, is an entirely different beast. The researchers are trying to quantify the relationship between physical attributes of an object, such as its weight, shape, or luminance, and its appearance. "We are doing experiments to understand which factors affect the way people judge object quality," says Rushmeier. With such information, a visualization can be designed to prioritize attributes that have a significant perceptual impact. This research is preliminary, however. "Many more experiments are needed before we can formulate explicit numerical models," she says.


Attention to perceptual influences is considered particularly important in scientific and information-visualization applications. "Because data comes in many flavors, such as numerical, categorical, field, sequence, image, graph, and text, and because there are many different dimensions, including color, shape, texture, and luminance, there are many possible ways of making pictures of data," says perceptual graphics researcher Bernice Rogowitz at IBM's T.J. Watson facility. "Understanding the mapping of data to visual dimensions can help people appreciate and understand the patterns and structures in the data." The notoriously difficult problem of creating the "right" mapping, she says, can be made easier by understanding how the human visual system processes these visual dimensions.

The data-mapping challenge is exacerbated by the increasing complexity of the datasets researchers are attempting to visualize. "A multidimensional dataset can contain many layers of information," says Christopher Healy, assistant professor and graphics researcher in the computer science department at North Carolina State University. Healy and his colleagues are developing perception-based methods to support rapid, accurate, effective visualization of these large, complex collections of data. "We use perception theory to help answer the question, 'how can we display some or all of this information simultaneously in a single display and at the same time ensure that viewers can explore, discover, analyze, and verify efficiently and effectively within their data?'"

Healy offers the example of weather maps to illustrate the challenge. "Most weather maps display a single weather condition, such as temperature via color, pressure via contour lines, winds via directed arrows, or precipitation via Doppler radar traces. Multiple conditions are displayed on multiple, separate maps." While this is fine for viewing conditions in isolation, he says, "it makes it difficult to identify trends or patterns that occur between conditions—for example, a cold low-pressure region with light rainfall and strong winds."

The North Carolina researchers are using an interdisciplinary approach to establish ways to show all of this information on a single map that allows the analysis of both individual conditions and the complex relationships

they may form with one another. "We begin by studying how the human visual system 'sees' the world around us. We then conduct controlled experiments to test our ability to distinguish visual features, such as color, texture, and motion, that we know are being detected in our visual system." The researchers test the experimental results by applying them to real-world visualization problems, including medical imaging, weather tracking, and e-commerce Web traces.



Illustrative enhancements (bottom) sharpen the features of an airflow visualization and make the separations and vortices more discernible. Researchers at Purdue analyze traditional art and design techniques and apply the perceptually successful methods t

The process of validating perceptual theory is a critical, but sometimes overlooked, step in developing practically useful models. This is because the theories were developed for simpler environments than those being depicted in most visualizations. "In the literature on human perception that we draw on, the research was conducted to study one dimension at a time [color, shape, or motion, for example]. We're taking those dimensions and putting them together to represent more complicated, multidimensional phenomena. That might not be entirely correct," says Ferwerda.

One way to validate the theoretical models is to work backward—that is, to determine what works and why in the physical world and try to re-create the success digitally. At Purdue University, computer science professor David Ebert and his colleagues have been analyzing perception-based technical and medical illustration techniques that have been used for hundreds of years and relating the techniques to the end results. "We then analyze the rendering process to determine how to create similar effects in the visualization pipeline," he says, with the ultimate goal of creating images that effectively convey information.

Ebert stresses that the goal of his group's research is not to create a photorealistic representation of a model, but rather "to enhance the appearance of an image to effectively convey information, whether the application involves flying through a storm in a motion-based ride or trying to find tumors in a medical image or flow features that make aircraft flight unstable."

While there are certainly compelling reasons to use perceptual knowledge as a way to filter huge amounts of data, there are risks inherent in doing so. "Whenever you remove information, you run the risk of taking away something that's important for viewers to see," says Healy.

There is also the risk of generating images with perceptible errors because the perception-based techniques and approximations might not be visually accurate, says Ebert. "These are the same problems that occur with visualization now. Transfer functions, segmentation, and rendering algorithms can all introduce error or create images that are not true representation of the data." Because of this, he says, "we need to make sure that we have error models and metrics and actually perform some user studies."

Also to be considered, says Rushmeier, is that the "correct" representation of data in many cases depends on what the user is trying to discover. There is no one-size-fits-all perceptual gauge. "In problems where there is

no natural 3D structure, such as business data, there are many ways to represent a dataset, because there are typically many variables and there can be an almost infinite number of possible combinations and methods for representing them."

Using an exhaustive visual search is a "terrible" way to try to gain understanding, says Rushmeier, so some type of filtering is needed to reduce the data, as long as the process is guided by domain knowledge. "The perceptual design should follow what you are looking for. If linear changes in pressure, for example, are important to you, a color map that represents pressure with a perceptually linear scale should be used." Achieving the most effective mapping from data to visual components, she says, "means using appropriate color scales, choosing appropriate geometric attributes, taking into account whether data is continuous or categorical, and so forth. We want to use what we know about perception to make sure our images help users see meaningful data structures and not filter them out."

In this vein, says Healy, researchers need to be careful how they build perception into their visual representations. "A poor choice can actually produce visualizations that actively mask information in the dataset," he says. "Our experiments are carefully designed to identify both the strengths and the limitations of the human visual system. Thus, when the visualizations are built, we can select combinations of visual features (colors, icons, intensity variations) that we know will complement one another and do a good job of showing viewers what's happening in their data." Careful feature combinations can be used to highlight important characteristics of the data as well as to display less important information in a way that is accessible, but does not mask the most salient properties.

In some cases, graphics researchers are forging new paths in perceptual science. "Perceptual models of color, shape, form, and motion can help inform our choice of parameters when representing data or graphical objects," says Rushmeier, as can literature on cognitive processes such as human memory and attention. But often there is no simple way of applying these results to the more complex CG environments being created, she says, "so we embark on new experiments to learn how people experience [complex visual information]." In this sense, she adds, we are providing new knowledge about human visual behavior.

While the need for perceptually based graphics tools seems obvious, and the demand for such capabilities is high, there are a number of obstacles to widespread application. And unlike many of today's pressing graphics problems, the chief challenge hampering perceptual graphics is not a computational one. Rather, says Rushmeier, "it is the complexity of human perception that we don't understand yet. There are models and results that can be used, but it isn't something someone outside of vision research can easily pick up and apply."

Healy agrees. "The use of perception per se is not technically challenging. Managing the data might require high-speed disk arrays and network connections, and certain techniques such as volume visualization may need dedicated hardware during the final rendering step, but deciding how to display the data is not computationally intensive." The main challenge, he says, is the lack of an easily accessible, unified body of information.

The first step in overcoming this obstacle is to foster collaboration among such groups as vision scientists, computer scientists, and artists, says Rushmeier. "Recently there have been workshops and conferences aimed at bridging this seemingly natural connection [the Siggraph 2001 Campfire on perception and graphics is one example], and a few schools have begun teaching perception as part of the graphics and visualization curriculum." Additionally, some research developments promise to break down some barriers.

For example, says Rogowiz, "one inhibitor to using perceptual models is that most of the perceptual studies are based on knowing the exact luminance and color values on a display screen. If it is not possible to calibrate the system-over the Internet, for instance-or if it is just too difficult, people won't do it, and may even eschew perceptual approaches for this reason." To address this, she says, "some of our perceptual work includes the development of color maps that provide faithful representations of data, and a new piece of work, called the Which Blair Project, that allows users to easily select a color map, even when their display systems are not calibrated."

Once perception-based techniques become more fully developed, they could easily be introduced into the rendering algorithms used in commercial animation and visualization packages, says Ebert. "In many ways, the commercial animation industry has been applying perception techniques for quite sometime. Animators know where they need detail in models and textures and where they don't based on the movement in a scene, just as on a movie set. If you carefully analyze a still from some animations, you may be surprised at how unrealistic some things really look."

On the visualization side, some of the perception work has already made its way into mainstream packages. The Pravda tool for selecting perceptually based color maps developed by IBM's T.J. Watson researchers is part of the open source Data Explorer visualization package, and the freely available global-illumination Radiance software has a tone-mapping capability.

These early efforts notwithstanding, what is needed to move perceptual considerations into mainstream graphics and visualization is access to the right tools and compelling applications. Most researchers agree that this day will come soon, as ongoing computational advances continue to give us more and more visual data to deal with. "We have greater access than ever to more complex, multidimensional visual information that users want to explore," says Ferwerda. "We have to develop tools to help them achieve this without overwhelming

their visual systems."

---

*Diana Phillips Mahoney is chief technology editor of Computer Graphics World.*

---

### Please post your comments

You must sign in to post

Username

Password

No Forum Poster account? [Sign up here.](#)

Latest News

[PNY Technologies To Debut DDR3 Memory Modules in Standard and Overclocked Versions](#)

6/1/2007

**Parsippany, N.J.** - Maker of OEM and aftermarket memory products, PNY Technologies Inc. will...

[PNY Technologies Introduces Two Graphics Cards Based on Nvidia GeForce 8 Series](#)

5/31/2007

**Parsippany, N.J.** - PNY Technologies Inc. has added two, new XLR8 graphics cards based on the Nvidia...

[On2 Technologies To Acquire Hantro Products, Maker of Embedded Video Technology](#)

5/31/2007

**Tarrytown, N.Y.** - On2 Technologies Inc., maker of video compression software and solutions, has entered...

more 

advertisement

v advertisement v



**NYU's Programs in Design,  
Digital Arts, and Film**  
"That's why I go to NYU."

[CLICK HERE FOR MORE](#)



[Home](#) | [About Us](#) | [Contact Us](#) | [Corporate Website](#) | [Privacy Policy](#) | [Site Map](#)

Copyright 2007: COP Communications, Inc., Glendale, CA; All Rights Reserved. | [Terms and Conditions](#) | [Webmaster](#)