Newton's Nightmare: Reality Meets Faux Physics

Organizer
Dinesh K. Pai
University of British Columbia
Department of Computer Science
2366 Main Mall
Vancouver, British Columbia
V6T 1Z4 Canada
+1.604.822.8197
pai@cs.ubc.ca

Panelists
Dinesh K. Pai
University of British Columbia

Doug Roble
Digital Domain

Holly Rushmeier
IBM T. J. Watson Research Center

Richard Szeliski
Microsoft Research

Demetri Terzopoulos
New York University and University of Toronto

Reality – who needs it? Is computer graphics about building more and more accurate simulations of the real world, down to the last photon? Is computer graphics really hard physics dressed in Hollywood clothing? Or is reality, like, soooooo old fashioned? Is computer graphics now free from its shackles, free to create non-whatever realistic experiences, free to write its own laws, with no relation to reality? This panel sheds some light (real or imagined) on these complex questions.

Dinesh K. Pai

Computer graphics is indeed about reality, but reality as experienced by humans. We need models of reality but our needs are very different from, say, the needs of physics or engineering. I argue that:

1. New, creative applications in computer graphics need new types of models, but these still need to be rooted in reality, because human perceptual and cognitive systems evolved to cope with it. We need to model not only external physical systems, but also human systems that produce and consume the experience.

2. Traditional models of reality are based on the assumption that measuring the real world is a lot more expensive than simulating it. New and inexpensive sensors have changed the economics of measurement and hence of modeling, making radically different models possible.

3. All models of reality are wrong, but some are more wrong than others for a specific purpose.What matters is to clearly know the metric. Is interactive response more important for perception than accurate motion? Are we trying to convey the details of a real object on an e-commerce Web site, or are we trying to direct attention to the object's most important features?

Dinesh K. Pai, a professor of computer science at the University of British Columbia, received his PhD from Cornell University. His research interests span the areas of robotics, graphics, modeling, and simulation. His current interests are in interactive multimodal simulation of contact (including auditory and haptic displays) and acquiring multimodal models of everyday objects using automated measurement techniques.

Holly Rushmeier

Design applications where the image is not the end product but a means to predict what a physical design will look like require accurate simulation. Researchers have developed simulations of light transfer to compute the quantity of energy that would pass through each pixel. We can still develop better algorithms, but there are few major problems left in the simulation of light. There is no reason to consider quantum or relativistic effects. Our research challenges are now in psychophysics, understanding what features will have an impact on a human observer.

Many interesting effects have been developed by accident – by setting various parameters and seeing if the resulting image is pleasing. Our knowledge of how to simulate the physics of light gives us the ability to delibrately control rendering to create consistent, visually rich alternative environments. While talented artists have always done this with effort, we can now facilitate more experimentation. Rather than relying on accident, we should further exploit what we know about light and deliberately change the rules to provide powerful new tools.

Holly Rushmeier received her PhD degree in mechanical engineering from Cornell University and is now a research staff member at the IBM T.J. Watson Research Center. Her research interests include data visualization and realistic image synthesis.

Doug Roble

In the realm of visual effects, we are always trying to convince people to suspend their disbelief. Of course, there are the big effects (the asteroids and the spaceships and aliens), but audiences know that they aren't real, so digital artists can have fun and get away with non-real effects and graphics. When I reflect on the most gripping visual effects scenes, I think of small effects that fool me utterly: the cow getting hit by the car in "Oh Brother, Where Art Thou?" or Julia Roberts' car accident in "Erin Brockovich." These scenes are devastating in that computer graphics has been used to manipulate and simulate reality so well that you don't have to suspend disbelief. There was never a point where you disbelieved!

So, do we need to simulate reality? Of course! The more accurate the lighting, the fluid dynamics, the surface parameters, the modelling ... the more powerful an artist becomes. Just look at the trends in the effects industry: Years ago, particle systems were all the rage. Now every effects house is developing its own fluid dynamics package. For characters, an IK weighting system controlling a NURBS surface used to be good enough. Now we are all developing physically accurate bone/muscle/skin systems. Visual effects houses have all adopted computer-vision techniques to extract every last bit of information from the real world.

Can we mimic reality without accurately simulating it? Sure! That's what we've been doing for years, and you've seen the results on the movie screens. We need to continue to forge ahead with more detailed and accurate models so that the artists can produce the effects of the future.

Doug Roble is creative director of software at Digital Domain and Sketches and Applications Chair for SIGGRAPH 2002. He has been developing tools and doing research at Digital Domain since 1993. His computer vision system, "track," won a Technical

Achievement Award from the Academy of Motion Picture Arts and Sciences in 1998. He received his PhD in computer science from The Ohio State University in 1992.

Richard Szeliski
"Faux physics or no physics?" For many computer graphics applications, it is often sufficient to simply capture some real-world imagery, and then to manipulate it to get the desired effect. An early example of this was image morphing, where different video streams could be morphed or blended to get compelling transitions between different people or objects. More recently, image-based rendering has suggests that we can often approximate the 3D appearance of an object (and generate novel interactive views) by simply jumping (or interpolating) between different views. Current implementations of the "freeze frame" effect often do just that: jump between a densely spaced set of still images taken with cameras.

Of course, doing computer vision analysis (recovering the geometric side of the "physics") allows us to use fewer cameras or to get better interpolation results. Still images, however, are just a very narrow subset of what we want to synthesize in computer graphics. The temporal analog to image-based rendering is video-based rendering, where sample video clips can be manipulated to achieve novel synthetic video sequences. An early example of this was *video rewrite*, which manipulated (concatenated and blended) digitized lip motions to make a character say new speech. More recently, we have been working on *video textures*, which can synthesize realistic, novel, quasi-periodic motions (waterfalls, flames, swimming fish, talking heads) from sample video footage. Is this "data driven" or "machine learning" ("no-physics") approach the solution to everything? Obviously not.

For many (most?) applications, we will get more mileage by trying to understand (and then simulate) the actual physics (geometry, photometry, dynamics, behavior) of the phenomena we are modeling. For example, recovery of BRDF from multiple images is currently one of the hot areas in image-based modeling. It's just that a complete model is often very hard to achieve, both because of our limited understanding, and because the inverse estimation problems are often ill-posed. Judicious knowledge of when to "fake" aspects of the physics will always remain one of the hallmarks of successful application of computer graphics to complex phenomena.

Richard Szeliski is a senior researcher in the Vision-Based Modeling Group at Microsoft Research, where he is pursuing research in 3D computer vision, video scene analysis, and image-based rendering. His current focus is on constructing photorealistic 3D scene models from multiple images and video. He received a PhD in computer science from Carnegie Mellon University in 1988, and he has been at Microsoft Research since 1995.

Demetri Terzopoulos
My holy grail is a "reality emulator" as compelling as the one portrayed in "The Matrix." Although a multisensory computational simulation with such incredible fidelity (never mind all the exhilarating weirdness!) remains elusive, the trend in computer graphics is clear. With Moore's law on our side, researchers and practitioners alike are eagerly pursuing what might be characterized as the "Taylor series approximation to reality." I, for one, have found it intellectually stimulating to help establish some crucial, low-order terms of this approximation, which now epitomize the prominent physics-based and biology-based (artificial life) paradigms in CG modeling and animation. The endeavor of systematically augmenting the realism of CG models continues to excite me.

However, I also believe that we should explore alternatives to simulation. It behooves us to exploit the special computational structure of the brain, which after all is the client, ideally through a direct brain-machine interface, of our provisionally mythical reality emulator. The brain learns to perceive the raw reality of nature in certain ways and not in others. In this context, recent CG techniques such as the NeuroAnimator (SIGGRAPH 98) are provocative. They suggest that it should be possible to create a new breed of emulation algorithms that, through observation of reality by computational structures analogous to those found in the brain, can learn to mimic a wide variety of natural phenomena (physical dynamics in the case of the NeuroAnimator) with sufficient fidelity to render all residual errors imperceptible.

Demetri Terzopoulos holds the Lucy and Henry Moses Professorship in the Sciences at New York University and is professor of computer science and mathematics at NYU's Courant Institute. He is currently on leave from the University of Toronto, where he is professor of computer science and professor of electrical and computer engineering. He received his PhD degree from the Massachusetts Institute of Technology. He was elected a fellow of the IEEE, a fellow of the Canadian Institute for Advanced Research, a Steacie Fellow of the Natural Sciences and Engineering Research Council of Canada, and a Killam Fellow of the Canada Council for the Arts. Among his many awards are computer graphics honors from Ars Electronica, NICOGRAPH, and the International Digital Media Foundation.

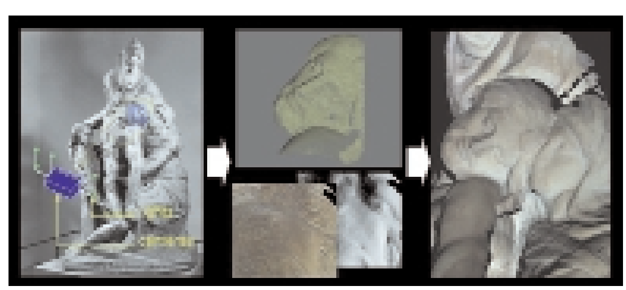

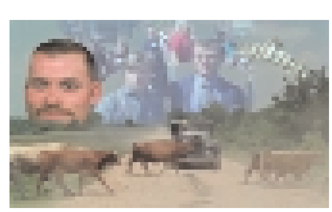

Real…(above)

…or not? (left)