

Nonphotorealistic Rendering



Adam Finkelstein and Lee Markosian
Princeton University

To state the obvious, the field of computer graphics is concerned with techniques for creating images via computer; and humans use images, fundamentally, to communicate visual information. It's perhaps slightly less obvious that photorealistic imagery might not be the best choice for communicating information about 3D scenes (real or virtual)—this decision depends on the intention behind the communication. For example, photographs are great for documenting a crime scene in which any small detail might turn out to be an important clue. Likewise, a real estate agent might show photographs of a house to potential buyers who want to see it in its current condition.

Let's change the premise slightly, however, to the case of an architect presenting a client with a preliminary design for a house not yet built. One study found that a photorealistic rendering of the house inhibits the dialog between architect and client.¹ The client might feel frustrated, perceiving (incorrectly) that the architect has unilaterally finalized the design without adequate input from the client. In contrast, an imprecise pencil sketch that omits many details suggests to the client that the design remains open to revision, facilitating the exchange of ideas (and keeping the client happy).

Other examples of the use of drawings instead of photographs abound. Owners' manuals typically employ schematic line drawings to illustrate device features, as these can more clearly delineate parts of interest (knobs, cables, spark plugs, and so on). For similar reasons, medical textbooks continue to use hand-drawn illustrations when the goal is to explain anatomical structures. (Photographs, on the other hand, are better for documenting particular symptoms—for example, a skin condition.) Whether explaining the structure of the solar system, an atom, or the inner workings of a volcano, hand-drawn illustrations can better communicate 3D structure, elide unimportant details, and emphasize important features. In the context of storytelling, writers often prefer hand-drawn images for their ability to convey an appropriate mood, or to depict imaginary worlds for which photographs aren't available. It's certainly true that simple, colorful drawings seem to hold a special appeal for children. For evidence, visit the children's section of your local library or bookstore, where hand-drawn images most likely outnumber photographic ones by a wide margin.

These arguments are well known within computer graphics. Indeed, they have fueled the emergence and rapid growth of the field of nonphotorealistic rendering (NPR), to which this special issue is dedicated. (For broader arguments in this vein or a survey of the field see the article by Durand² or either of two recent books on NPR.^{3,4}) One goal of this work is to develop algorithms for generating synthetic images that embody qualities of hand-drawn images: the selective emphasis of important features, suppression of unimportant details, and use of stylization and abstraction to suggest complex structures without resorting to literal representations. These qualities make NPR especially suited to applications where the purpose of the imagery is to explain, illustrate, or tell stories.

There is an altogether different reason why NPR might continue to expand its influence within 3D graphics and perhaps eventually grow to rival the importance of photorealistic rendering: realism is expensive. That is, we need a vast amount of detail to faithfully represent (and animate) a realistic natural scene. When a human designer creates that scene by hand, the process necessarily requires great time and effort. One strategy to avoid this inherent cost of realism is to capture details from the real world. Examples include image-based rendering, 3D scanning, and motion capture. But these strategies have limitations. For one, they require that the desired data be present in the real world (at a suitable scale). For many applications, including storytelling, this might not be the case. When scanning isn't an option, the alternative (currently) is to hire a team of trained experts and let them painstakingly model and animate the needed 3D content by hand. While this strategy works, it's only feasible for high-budget industries such as games and movies. For 3D graphics to reach new applications and attract new users, something must change. We believe NPR has the potential to solve the content creation problem because nonphotorealistic images can be far simpler to create by hand than photorealistic ones.

For that promise to be fulfilled, we must address many research challenges. This special issue presents five articles and a tutorial that focus on various aspects of NPR research. One article reviews algorithms for rendering silhouettes of 3D models; silhouettes are fundamental to most simple line-drawing styles. Another article pre-

sents a new fast algorithm for transferring texture from a source image to a destination image. While not strictly an NPR algorithm, this work enables style transfer, with applications in NPR. The remaining articles address various aspects of rendering 3D models with stylization. One presents a new technique for rendering volume data sets (for example, from medical scans) in the style of pen-and-ink illustrations. Another describes a new technique for rendering stylized highlights on 3D models when the goal is to produce an overall look resembling (or integrated with) traditional cel animation. The last article presents a new algorithm for rendering 3D models over multiple frames in the style of traditional stipple drawings. A key challenge is to generate stipples where needed, given the current viewing and illumination conditions, while maintaining temporal coherence. The tutorial presents a framework for understanding a variety of stroke-based rendering algorithms, which are central to NPR. This tutorial can be useful for anyone wishing to implement known NPR algorithms; it also identifies several open areas for research.

As this is a new field, many of the tough research challenges in NPR remain unsolved. Interactive techniques, which offer the opportunity for content creators to work more naturally, are just becoming possible due to recent advances in both algorithms and hardware. We have observed a burgeoning within the field and expect the expansion to continue. Several courses and many paper sessions have been offered at Siggraph, and a third biannual symposium devoted to this topic will be held next year: the International Symposium on Nonphotorealistic Animation and Rendering (NPAR 2004). Visit <http://www.npar.org> for more information. We hope this special issue offers inspiration for new researchers to contribute to the advance of the field. ■

Acknowledgments

We acknowledge the pioneering work of members of the NPR community who have established this as a field within computer graphics during the last decade. The articles selected for this special issue were chosen from among 19 submissions; we thank the authors of all of the submissions for sending us their latest research. Each article was reviewed by three or four experts whom we thank for their diligent readings and careful reviews. In addition, we appreciate the administrative support of Rebecca Davies at Princeton and Alkenia Winston at *IEEE Computer Graphics and Applications*. Finally, we thank Associate Editor-in-Chief Maureen Stone and the rest of the CG&A editorial board for the opportunity to organize this issue, and for their support and guidance.

References

1. J. Schumann et al., "Assessing the Effect of Nonphotorealistic Rendered Images in CAD," *Proc. ACM Computer Human Interaction (CHI 96)*, ACM Press, 1996, pp. 35-42.
2. F. Durand, "An Invitation to Discuss Computer Depiction," *Proc. Symp. Nonphotorealistic Animation and Rendering (NPAR 02)*, ACM Press, 2002, pp. 111-124.
3. B. Gooch and A. Gooch, *Non-Photorealistic Rendering*, A.K. Peters, 2001.
4. T. Strothotte and S. Schlechtweg, *Non-Photorealistic Computer Graphics: Modeling, Rendering, and Animation*, Morgan Kaufman, 2002.



Adam Finkelstein is an associate professor of computer science at Princeton University. His research interests include nonphotorealistic rendering, multiresolution representations, animation, and other applications of computer graphics in art. Finkelstein attended Swarthmore College as an undergraduate student, where he studied (occasionally) physics and computer science; he received a PhD in computer science from the University of Washington. Previously, he was a software engineer at TIBCO where he wrote software for people who trade stock.



Lee Markosian is joining the Department of Electrical Engineering and Computer Science at the University of Michigan as an assistant professor. His research interests include nonphotorealistic rendering and sketch-based modeling of 3D shapes. Markosian has a PhD in computer science from Brown University and worked as a postdoctoral research associate at Princeton University.

Readers may contact Adam Finkelstein at the Computer Science Dept., 35 Olden St., Princeton, NJ 08540; af@cs.Princeton.edu.

For further information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.