Computer Animation for Virtual Humans

Rae Earnshaw
University of Bradford

Nadia Magnenat-Thalmann
University of Geneva

Demetri Terzopoulos
University of Toronto and Intel

Daniel Thalmann
Swiss Federal Institute of Technology

Advances in computer animation techniques have spurred increasing levels of realism and movement in virtual characters that closely mimic physical reality. Increases in computational power and control methods enable the creation of 3D virtual humans for real-time interactive applications.¹ Artificial intelligence techniques and autonomous agents give computer-generated characters a life of their own and let them interact with other characters in virtual worlds. Developments and advances in networking and virtual reality (VR) let multiple participants share virtual worlds and interact with applications or each other.

Controlling actions and behavior

High-level control procedures make it possible to give behaviors to computer-generated characters that make them appear “intelligent”—that is, they interact with other characters with similar properties and respond to environmental situations in a meaningful and constructive way. Such scenarios have the potential of receiving script information as input and producing computer-generated sequences as output. Application areas include production animation and interactive computer games. In addition, researchers are currently investigating ways of having virtual humans perform complex tasks reliably.¹

Shared environments

Computer-supported collaborative work (CSCW) often involves interaction and discussion about computer-generated information such as models, simulations, annotations, and data accessed in shared virtual environments (VEs). Representations of users by computer-generated characters (avatars) facilitate communication and interaction. An interesting question arises as to what form such avatars should take to best promote life-like and interesting behaviors that mirror the owner, and invoke meaningful and creative responses from other avatars' owners in the virtual world. A shared experience in an artificial computer-generated world implies, in some sense, a belief that the world is real

(that is, the suspension of disbelief). It's clear from research to date that creating environments that look real and believable is easier than creating moving characters that look real. Increasing the characters' fidelity doesn't necessarily increase the feeling that their world is real. Engaging users in the tasks required appears to be the first step toward making the interface transparent and enhancing the relationship with other objects or users in the virtual world. Computer-generated games such as “Doom” and the SimNet tank interface² both get the user to concentrate on task performance at an early stage. Pausch et al.³ also reported similar results.

Distributed VEs

Avatars and agents have an interesting relationship. An agent personalizes information. The presence of avatars and agents in the same environment seems a fruitful area for further work. Current evidence suggests that avatars link the user to the virtual world very well initially. But from then on, less sophisticated representations suffice to convey information and facilitate communication—except in application domains where the framework is just as important as the action (for example, when playing tennis in a public forum). However, the tennis players themselves could operate on more basic physical models and representations, since they're concentrating on the task rather than the framework, or the event as a whole. This is probably one reason why computer games succeed.

A second aspect of the rapid rate of change is the increasing degree of real-time control passed on to the user or viewer by giving them access to new forms of interactive content. A third aspect is the increasing importance and prominence of the Internet and the facilitation of distributed VEs that Web technology provides. We're thus seeing convergence of content creation and technology delivery as well as a migration of infrastructure technologies down to the Internet.^{5,6} Both these trends increase the relevance and importance of tools and techniques for realistic modeling and movement of human-like characters to populate scenes or represent human users in geographically dispersed places.

Virtual Humans—Challenges and Opportunities in the US

Current work on virtual humans in the US includes the following projects.

For an overview of distributed VR, visit <http://ece.uwaterloo.ca/~broehl/distrib.html>. Also, Virtual Personalities in Pittsburgh (<http://web.vperson.com/>) and Haptik in Santa Cruz, California (<http://www.haptik.com/>) are two companies that develop humanoid characters with natural language response. In addition, Protozoa in San Francisco and Medialab Studio in Los Angeles have developed systems to produce 3D animated characters that provide real-time editing and control. They have been used to let 3D characters interact with, and react to, live performers in real time. For more information go to <http://www.protozoa.com> and <http://www.medialab.com>.

Facial modeling and animation

Scott King of Ohio State University gives an overview of facial animation with a bibliography of references at <http://www.cis.ohio-state.edu/~sking/FacialAnimation.html>. The Perceptual Science Laboratory at the University of California at Santa Cruz is another useful resource for facial animation (<http://mambo.ucsc.edu/psl/fan.html>). For examples of facial analysis, visit <http://mambo.ucsc.edu/psl/fanl.html>. Boston Dynamics (<http://www.bdi.com>) in Cambridge, Massachusetts developed the DI-Guy system for interactive humans such as dancers, pedestrians, athletes, and soldiers.

Jack and TJ Toolkit

Jack is a 3D interactive environment for controlling articulated figures developed by Norman Badler at the University of Pennsylvania and is available from Transom Technologies. It features a detailed human model and includes realistic behavioral controls and task animation. For more information visit <http://www.cis.upenn.edu/~hms/jack.html>.

Modeling and Simulation—Linking Entertainment and Defense

Michael Zyda of the Naval Postgraduate School in Monterey, California chaired a committee on modeling and simulation to investigate opportunities for collaboration between the defense and entertainment research communities. The US Department of Defense has funded hardware, networks, and simulation environments, and the entertainment industry has funded the development of games. Though these communities differ in their motivations, objectives, and cultures, they share a common interest in modeling and simulation. In entertainment, modeling and simulation technology is a key component of a \$30 billion annual market for video games, location-based entertainment, theme parks, and films. In defense, modeling and simulation provide a cost-effective means of conducting joint training; developing new doctrine, tactics, and operational plans; assessing battlefield conditions; and evaluating new and upgraded systems.

Visit <http://www.nap.edu/readingroom/books/modeling/> for a report of a workshop that discussed mutual areas of interest and potential for greater collaboration. During the workshop, participants identified one area of shortfall—a deficit of talented researchers with cross-disciplinary skills in areas such as modeling, simulation, VEs, electronic storytelling, and content production. The NPSNet Research Group has more information and a list of publications at <http://www-npsnet.cs.nps.navy.mil/npsnet/publications.html>.

The Improv Project

This project, at New York University's Media Research Lab, is building technologies to produce distributed responsive VEs in which human-directed avatars and computer-controlled agents interact with each other through a combination of procedural animation and behavioral scripting techniques.

Multimodal interaction paradigms being explored combine Improv with speech and gesture recognition in conjunction with various forms of presentation, including 2D and 3D display. Currently implemented as a set of Java classes, communicating with 2D and 3D graphical environments such as the Virtual Reality Modeling Language (VRML) 2.0, Improv supports a network distributed responsive world within standard Web browsers.

For more information visit <http://mrl.nyu.edu/> or see Perlin and Goldberg's paper "Improv: A System for Scripting Interactive Actors in Virtual Worlds" (*Proc. Siggraph 96*, ACM Press, New York, 1996, pp. 205-216).

Synthetic Characters Group

Bruce Blumberg of the MIT Media Laboratory is developing an ethologically inspired architecture for building autonomous animated creatures that live in virtual 3D environments. They can sense the world, the state of their internal goals, and motivations, and decide what to do next based on this information and the set of actions they can perform. For more information visit <http://bruce.www.media.mit.edu/people/bruce/>.

Natick Research Development and Engineering Center

Simulating and modeling strategic maneuvers involving infantry demand increased fine granularity. At the Natick Research Development and Engineering Center, Natick, Massachusetts, virtual humans are being used to supply the level of detail required. For more information visit <http://www.metavr.com/USArmySoldierSystemsCommand.htm>.

Virtual Humans Conference

This year's Virtual Humans Conference took place 16-17 June in Los Angeles. Key issues considered were international standards for virtual humans (VH3, based on the VRML 97 standard), authoring tools, dialog synthesis, autonomous humanoids, and intellectual property law as applied to digital actors. For more information visit <http://www.vrnews.com/eventsvh3main.html>.

Networked VE Projects in Europe

A number of pan-European projects funded by the European Commission under the Advanced Communications Technologies Services (ACTS) program (<http://www.at.infowin.org/ACTS/>) in the area of telepresence and shared VEs are currently aggregating their results into guidelines for usability, collaboration, teleoperation, augmented reality, and standards. These projects include Collaborative Integrated Communications for Construction (CICC), Coven, Distributed Video Production (DVP), Maestro, Midstep, Resolv, Tapestry, USInacts, and Visinet. Work from these projects emphasizes that presence is more than observing realistic looking models that move with the user, but involves subtle relationships among the user, the interface, the task, the user's involvement in the task, and more importantly, the emotional responses provided by human-like objects in the virtual world. This area needs more research.

At the same time and under the same program, a number of projects concerned with content generation and new forms of broadcasting (Interact, Mirage, Momusys, Emphasis, and Vista) are exploiting the convergence taking place between technology domains and content-generation domains. Virtual interactive studio and television applications (the Vista project) that use networked graphical supercomputers bring content creators, technologists, and broadcasters together to let content and interaction be generated by high-end computing systems. Virtual and real elements can be integrated in real time. Home users can interact via the telephone and Internet, giving a high-end and low-end spectrum of applications support. For more information visit <http://www.edc.nl/vista/>.

Another project, Vpark (<http://www1.iis.gr/vpark/>), is creating a virtual amusement park to integrate a number of novel applications based on distributed VEs. This will extend the VLNet⁴ shared VE systems to support heterogeneous network configurations, greater functionality for navigation, object manipulation, gesture activation, representation of emotion, and development of user interfaces for defining face and body behaviors.

The European Commission Framework 5 program is moving the center of gravity from the technology to the user of the technology, with increasing emphasis on content generation tools and facilities (<http://europa.eu.int/comm/dg12/press/1998/pr1302en.html>).

space of adverbs such as emotional characteristics or physical traits. Radial basis functions and linear regression are used to map a desired point in adverb space to the appropriate combination of basis motions. At run-time, the motion is controlled by a set of parameters called "adverbs" and through a graph of motions (such as walking or running) called "verbs." The graph defines the possible transitions between verbs and how they must be performed. Verbs, adverbs, and verb graphs are defined offline in an authoring system. User annotations place example basis motions along dimensions such as "happiness," or more generally at some point in the adverb space. During a transition between two graph nodes, only a simple blending is performed due to real-time constraints. The authoring system permits the definition of kinematic constraints, allowing, for example, a hand to hold on to a lever during a particular time period (via standard inverse kinematic techniques).

Moccozet et al. describe an innovative interactive animation system for building and simulating real-time virtual humans. The system emphasizes aspects of modeling and deformation that increase the realism of virtual humans' appearance. Two applications illustrate the system's usability and performance. The first, virtual tennis, allows two virtual humans to play a game of tennis judged by an autonomous virtual referee. In the second, CyberDance, a real choreographer is linked via sensors to a metallic robot. A further sequence links a real dancer to a virtual one.

The article by Brogan, Metoyer, and Hodgins describes two VEs showing novel uses of dynamically simulated characters. The first is a border collie environment, and the second, an Olympic bicycle race. Both examples use dynamically simulated, animated characters in networked VEs, and thus let the user interact intuitively with responsive characters. The article presents a real-time solution with 16 dynamically controlled characters. The system architecture for integrating various components to give the required real-time performance is also a significant contribution. Such an environment can test the hypothesis of whether the generation of complex and interesting behaviors in response to real-time user actions facilitates the user's involvement in the scenarios being simulated.

Eisert and Girod present a technique for analyzing video sequences of people's heads and faces. The rigid movement and deformation of the face are estimated from the sequence by combining optical flow techniques with a synthetic 3D model of the person. This leads to a robust and linear algorithm that estimates facial animation parameters with low computational complexity. A multiresolution framework overcomes the restriction of small object motion. A head model constrains the motion and deformation in the face to a set of facial animation parameters defined by the MPEG-4 video standard. This enables a description of both global and local 3D head motion as a function of the unknown facial parameters to be obtained.

The way forward

This issue presents significant and important developments of computer animation for virtual humans, par-

This issue

This special issue features five articles on computer animation for virtual humans. The first is a survey of virtual humans and the techniques that control the face and body. The article also covers higher level interfaces that allow direct speech input and an examination of issues associated with real-time control. This is particularly important in avatar rehearsal scenarios for animation production, where the director requires characters to interact in real time during the production. In cases where the director shouts "Stop" or "Move now," the real-time constraints are considerable. To provide instantaneous response requires a behavioral model for the characters more sophisticated than currently available.

The article by Rose, Bodenheimer, and Cohen presents a technique for interpolating between basis motions derived from annotated motion-capture data or traditional animation. The interpolation is defined over a

ticularly in the context of networked environments with distributed users. These developments have great potential as technologies converge and tools for content creation become increasingly synergetic with those for shared environments and interaction. Content scripts need high-level tools for translation into life-like and realistic behaviors of computer-generated characters capable of emotional responses (just as real actors do). In turn, this will engage users and achieve the same levels of satisfaction and enablement in shared applications as users currently do with entertainment applications. ■

Acknowledgments

We acknowledge input received from David Leever, chair of the European Commission special interest group on distributed environments (SID) Chain on Telepresence and Shared Virtual Environments. Work in progress and current documents may be found at http://www.infowin.org/acts/analysys/concertation/chains/si/home/ch_sid/. This Web site also contains a proposed Reference Model for Telepresence and Shared Virtual Environments.

A list of Virtual Human Web pointers may be found at <http://www.cis.upenn.edu/~badler/vhlist.html/>.

References

1. N.I. Badler, "Real-Time Virtual Humans," *Proc. Pacific Graphics 97*, IEEE CS Press, Los Alamitos, Calif., 1997, pp. 4-13, <http://www.cis.upenn.edu/~badler/paperlist.html>.
2. A. Pope, *The SimNet Network and Protocols*, BBN Report No. 7102, BBN Systems and Technologies, Cambridge, Mass., 1989.
3. R. Pausch et al., "Disney's Aladdin: First Steps Toward Storytelling in Virtual Reality," *Proc. Siggraph 96*, ACM Press, New York, 1996, pp. 193-203.
4. T.K. Capin et al., "Virtual Human Representation and Communication in VNet," *IEEE CG&A*, Vol. 17, No. 2, March/April 1997, pp. 42-53.
5. *The Internet in 3D: Information, Images and Interaction*, R.A. Earnshaw and J.A. Vince, eds., Academic Press, San Diego, 1997, p. 392.
6. *Virtual Worlds on the Internet*, J.A. Vince and R.A. Earnshaw, eds., IEEE CS Press, Los Alamitos, Calif., 1998, p. 360.

Rae Earnshaw is professor and head of electronic imaging and media communications at the University of Bradford, UK. His research interests include imaging, graphics, visualization, animation, multimedia, virtual reality, art, design, and the convergence of computing, telephony, media, and broadcasting. He obtained his PhD in computer science from the University of Leeds. He is a member of the editorial boards of *The Visual Computer*, *IEEE Computer Graphics and Applications*, and *The Journal of Visualization and Computer Animation*, and managing editor of *Virtual Reality*, vice-president of the *Computer Graphics Society*, chair of the *British Computer Society Computer*

Graphics and Displays Group, and a fellow of the *British Computer Society*. He is a member of *ACM*, *IEEE*, and *Eurographics*. Contact: University of Bradford, Bradford BD7 1DP, UK, e-mail r.a.earnshaw@bradford.ac.uk.

Nadia Magnenat-Thalmann has researched virtual humans for more than 20 years. She studied psychology, biology, and chemistry at the University of Geneva and obtained her PhD in computer science (cum laude) in 1977. In 1989 she founded *Miralab*, an interdisciplinary creative research laboratory at the University of Geneva. Some recent awards for her work include the 1992 Moebius Prize for the best multimedia system awarded by the European Community, "Best Paper" at the British Computer Graphics Society congress in 1993, to the Brussels Film Academy for her work in virtual worlds in 1993, and election to the Swiss Academy of Technical Sciences in 1997. She is president of the *Computer Graphics Society* and chair of the *IFIP Working Group 5.10* in computer graphics and virtual worlds.

Demetri Terzopoulos is a professor of computer science and electrical engineering at the University of Toronto and heads the computer animation research group at Intel. He received his PhD from MIT. He has held fellowships from the *Natural Sciences and Engineering Research Council of Canada* and the *Canadian Institute for Advanced Research*. In 1998 he was named a Killam Fellow of the Canada Council for the Arts. He has written extensively about computer vision and graphics, medical imaging, computer-aided design, artificial intelligence, and artificial life. He has received awards from the *IEEE*, *AAAI*, *Nicograph*, *Ars Electronica*, *International Digital Media Foundation*, *Canadian Image Processing and Pattern Recognition Society*, and the University of Toronto. He was program chair of the 1998 conference on *Computer Vision and Pattern Recognition (CVPR 98)*.

Daniel Thalmann researches real-time virtual humans in virtual reality, networked virtual environments, artificial life, and multimedia at the Swiss Federal Institute of Technology (Ecole Polytechnique Fédérale de Lausanne—EPFL). He received a diploma in nuclear physics in 1970, a certificate in statistics and computer science in 1972, and a PhD in computer science (cum laude) in 1977 from the University of Geneva. He is co-editor-in-chief of the *Journal of Visualization and Computer Animation*, member of the editorial board of *The Visual Computer*, *CADDM Journal* (China Engineering Society), and *Computer Graphics (Russia)*. He is co-chair of the *Eurographics Working Group on Computer Simulation and Animation* and a member of the executive board of the *Computer Graphics Society*.