Demetri Terzopoulos

BIOLOGICAL AND EVOLUTIONARY MODELS GIVE LIFE TO SELF-ANIMATING GRAPHICAL CHARACTERS WITH BODIES, BRAINS, BEHAVIOR, PERCEPTION, LEARNING, AND COGNITION.

ARTIFICIAL LIFE
for Computer Graphics

Computer graphics modeling for image synthesis, animation, and virtual reality has advanced dramatically over the past decade, revolutionizing the motion picture, interactive game, and multimedia industries. The field has advanced from first-generation, purely geometric models to more elaborate physics-based models. We can now simulate and animate a variety of real-world, physical objects with stunning realism. What's next?

Graphics researchers are exploring a new frontier—a world of objects of much greater complexity than those typically accessible through geometric and physical modeling alone—objects that are *alive*. The modeling and simulation of living systems for computer graphics resonates with an emerging field of scientific inquiry called artificial life, or ALife, a discipline that transcends the traditional boundaries of computer science and biological science.[1] The synergy between computer graphics and artificial life now defines the leading edge of advanced graphics modeling.

Artificial life concepts play a central role in the construction of advanced graphics models for image synthesis, animation, multimedia, and vir-

[1]For a readable introduction and historical perspective, see S. Levy's *Artificial Life*, Vintage Books, New York, 1992.

ALife

Cognitive

Evolutionary

Behavioral

Physical

Kinematic

Geometric

An early milestone in computer graphics modeling for animation was combining geometric models and (direct and inverse) kinematic models to simplify keyframing. In the mid-1980s, researchers began incorporating physical principles to develop physical models for animating particles, rigid bodies, deformable solids, fluids, and gases, helping generate copious quantities of realistic motion through dynamic simulation. In the context of physics-based modeling, biomechanical modeling employs simulated physics to automate the lifelike animation of animals with biological tissue and internal muscle actuators. More recently, research in behavioral modeling has progressed toward self-animating characters that react appropriately to perceived environmental stimuli. At the apex of the modeling pyramid, cognitive modeling has emerged as the use of artificial intelligence techniques, including knowledge representation, reasoning, and planning, to produce graphical characters with some level of deliberative intelligence and free will. ALife modeling spans the biomechanical-to-cognitive layers of the pyramid and investigates the possibility of applying evolutionary models to evolve graphical characters or aspects of their bodies and brains.

tual reality. New graphics models have taken bold steps toward the realistic emulation of a variety of living things—including plants and animals—from lower organisms all the way up the evolutionary ladder to humans. Typically, these models take complex forms and inhabit virtual worlds in which they are subject to the laws of physics. Consequently, they often incorporate state-of-the-art geometric and physics-based modeling techniques. But more significantly, these models must also simulate many of the natural processes that uniquely characterize living systems—including birth and death, growth and development, natural selection, evolution, perception, locomotion, manipulation, adaptive behavior, learning, and intelligence.

Here, I explore an exciting and highly interdisciplinary area of computer graphics that offers a wealth of provocative research problems and great commercial potential. The challenge is to develop and deploy sophisticated graphics models that are self-creating, self-evolving, self-controlling, and self-animating, by simulating the natural mechanisms fundamental to life.

Signs of ALife

Artificial life for computer graphics has spawned several principal avenues of research and development. The artificial

Table 1. Examples of behavioral animation in the movies.

Year	Title	Director	Producer
1987	*Stanley and Stella in: Breaking the Ice* (short)	L. Malone	Symbolics
1988	*Behave* (short)	R. Allen	R. Allen
1989	*The Little Death* (short)	M. Elson	Symbolics
1992	*Batman Returns*	T. Burton	Warner Brothers
1993	*Cliffhanger*	R. Harlin	Carolco Pictures
1994	*The Lion King*	R. Allers, R. Minkoff	Walt Disney Productions
1996	*From Dusk Till Dawn*	R. Rodriguez	Miramax Films
1996	*The Hunchback of Notre Dame*	G. Trousdale, K. Wise	Walt Disney Productions
1997	*Hercules*	R. Clements, J. Musker	Walt Disney Productions
1997	*Spawn*	M. Dippé	New Line Cinema
1997	*Starship Troopers*	P. Verhoeven	Tristar Pictures
1998	*Mulan*	T. Bancroft, B. Cook	Walt Disney Productions
1998	*Antz*	E. Darnell, L. Guterman, T. Johnson	Dreamworks SKG/PDI
1998	*A Bug's Life*	J. Lasseter, A. Stanton	Walt Disney Productions/Pixar
1998	*The Prince of Egypt*	B. Chapman, S. Hickner, S. Wells	Dreamworks SKG
1999	*Star Wars: Episode 1—The Phantom Menace*	G. Lucas	Lucasfilm

Figure 2. Artificial plants.
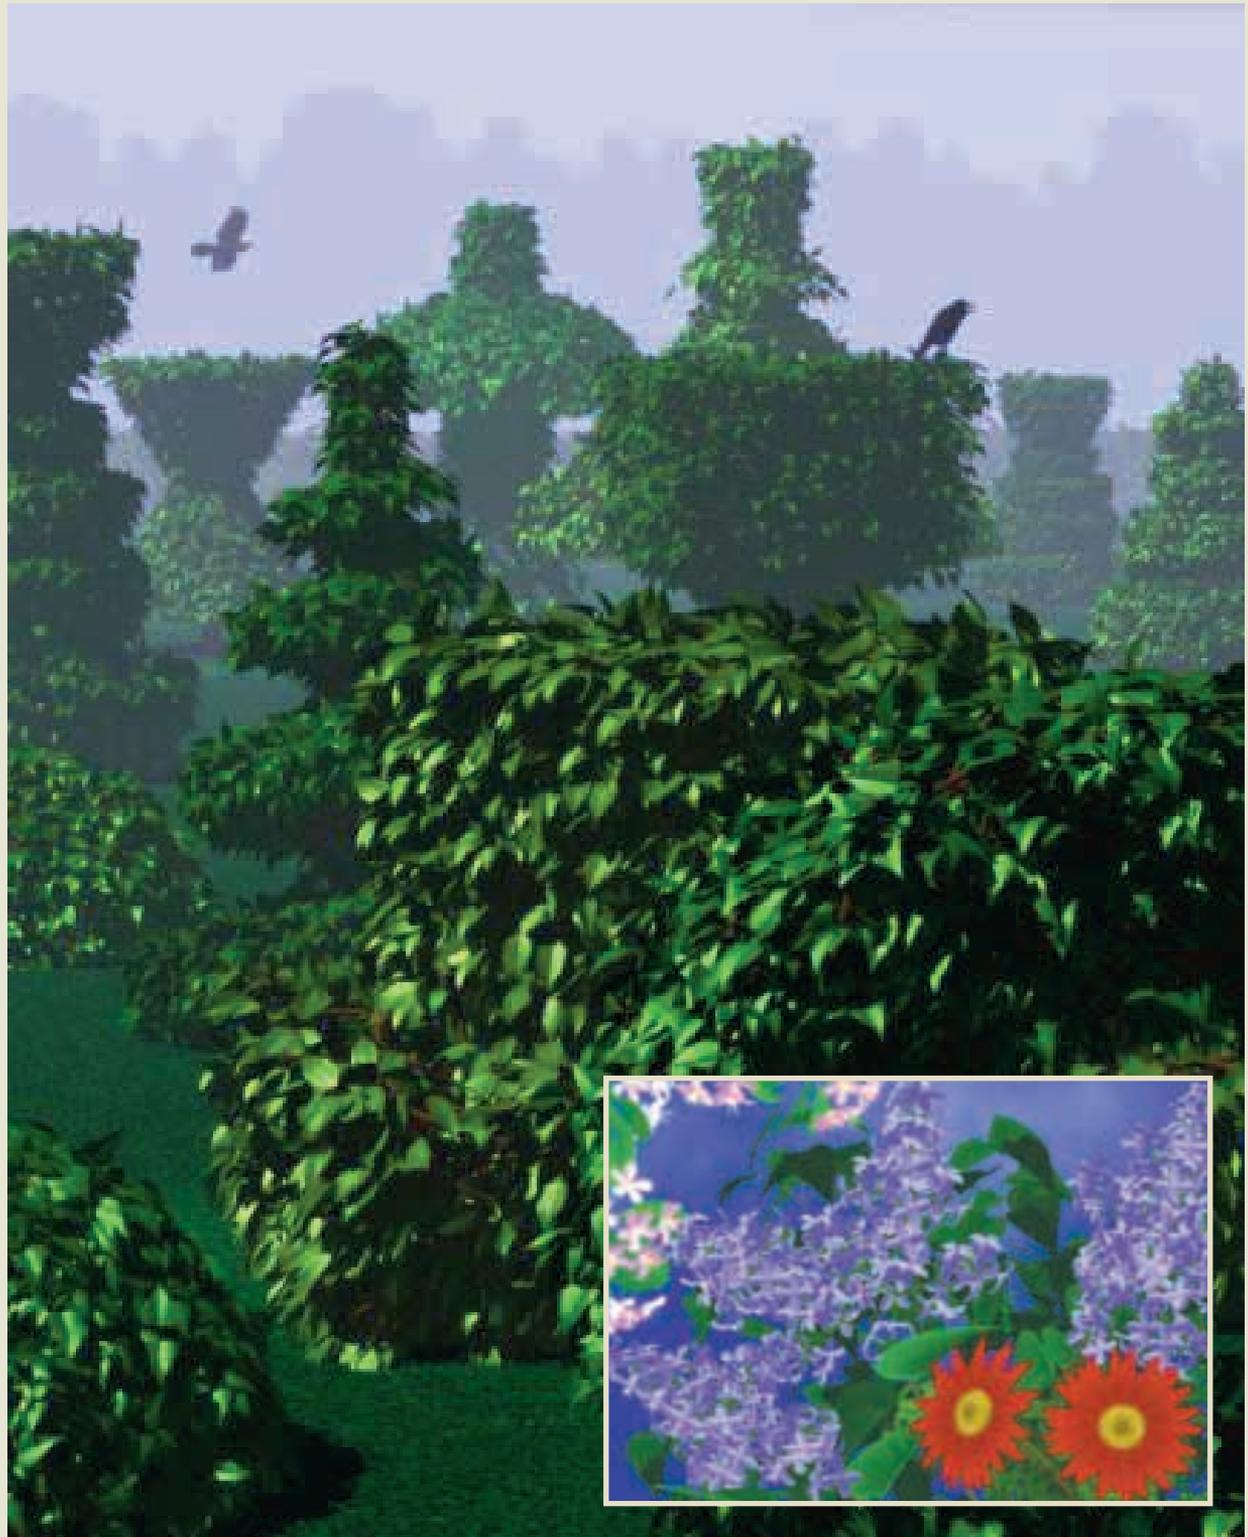
A synthetic model of the topiary garden at Levens Hall, England, by
R. Mëch, P. Prusinkiewicz, and M. James. "Garden of L" (inset) by P. Prusinkiewicz,
F. Fracchia, J. Hanan, and D. Fowler; see www.cpsc.ucalgary.ca/~pwp

Figure 3. Evolutionary art.
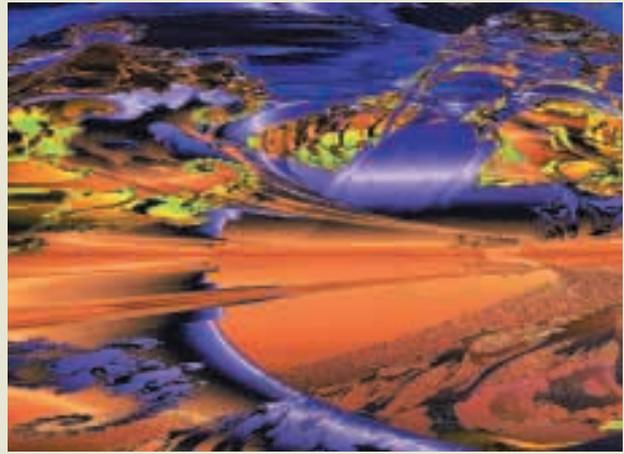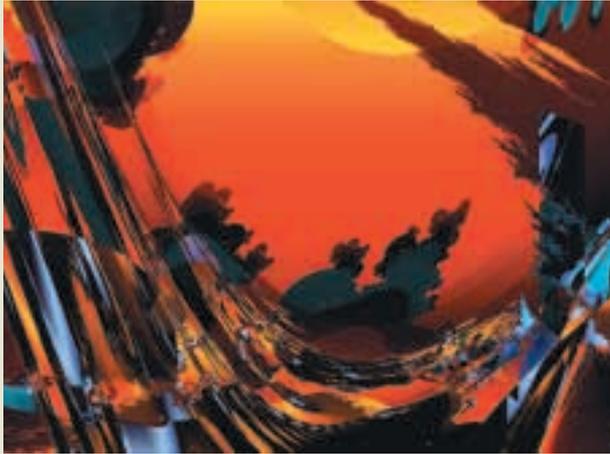
"In the Beginning" (left); "The High Plains of Kilimanjaro" (right).

life approach has proved especially effective for advanced animation (see Figure 1). Techniques are now available for realistically modeling and animating plants, animals, and humans. Behavioral modeling is a major trend in the motion picture special effects industry. The relentless increase in computational power is drawing the attention of researchers and practitioners to synthetic characters for interactive games. Moreover, artificial evolution is establishing itself as a powerful technique for image synthesis and potentially for model synthesis.

Artificial plants. Formalisms inspired by biological development processes have been used to grow highly complex and realistic graphics models of plants. Przemyslaw Prusinkiewicz, of the University of Calgary, is famous for his work in this area [5]. Lindenmayer systems (L-systems), a formal language framework introduced in the 1960s by biologist Aristid Lindenmayer as a theoretical framework for studying the development of simple multicellular organisms, has more recently been applied by Prusinkiewicz and others to the study of morphogenesis in higher plants with impressive results (see Figure 2). L-systems can realistically model branching and flowering patterns, as well as the propagation of the internal hormonal signals controlling plant growth and development.

Geometric and stochastic plant models expressed using L-systems were also recently proposed by Prusinkiewicz and his students for simulating the interaction between a developing plant and its environment, including light, nutrients, and mechanical obstacles. These environmentally sensitive L-systems appropriately regard the plant as a living organism in continual interaction with its environment. Prusinkiewicz has extended them to include a model

of the response of plants to pruning, which yields realistic synthetic images of sculptured plants found in topiary gardens, such as the Levens Hall garden in England, which inspired the synthetic image shown in Figure 2.

Artificial evolution. At Thinking Machines in 1991, Karl Sims pioneered a fascinating computer graphics approach inspired by theories of natural evolutionary processes [7]. Artificial evolution, a form of digital Darwinism, allows complex virtual entities to be created without need for detailed design and assembly. For example, the digital images in Figure 3 by "evolutionary artist" Steven Rooke are the result of artificial evolution. Artificial evolution evolves complex genetic codes (genotypes) that specify the computational procedures for automatically growing entities (phenotypes) useful in graphics and animation. Fortunately, graphics practitioners do not have to understand these codes; they can simply specify the subjective desirability of phenotypes as the entities evolve. The computer does most of the work, applying the principle of "survival of the fittest." The software instantiates populations of individual phenotypes from a variety of genotypes, sexually reproduces new individuals through the combination of genotypes (subject to occasional mutations), and terminates undesirable individuals unfit for survival.

Sims has shown that artificial evolution holds great promise for evolving many graphical entities, in addition to synthetic image art, including 3D sculptures, virtual plants, and virtual creatures.

Behavioral modeling and animation. In 1987, Craig Reynolds, then at Symbolics, Inc., published his landmark "boids" experiment, bridging the gap between artificial life and computer animation [6].

Figure 4. Artificial humans.

A face model, "cloned" from a real person, synthesizes realistic expressions. The subject is scanned by a Cyberware 3D color digitizer; algorithms convert the data into an anatomically correct synthetic face, including a deformable (finite element) facial tissue model with two dozen embedded contractile actuators modeling the muscles of facial expression, supported by a skull substructure with articulated jaw [4]. The virtual human models in the playground scene incorporate dynamically controlled mechanical models of articulated skeletons that can jog, use a swing, and more.
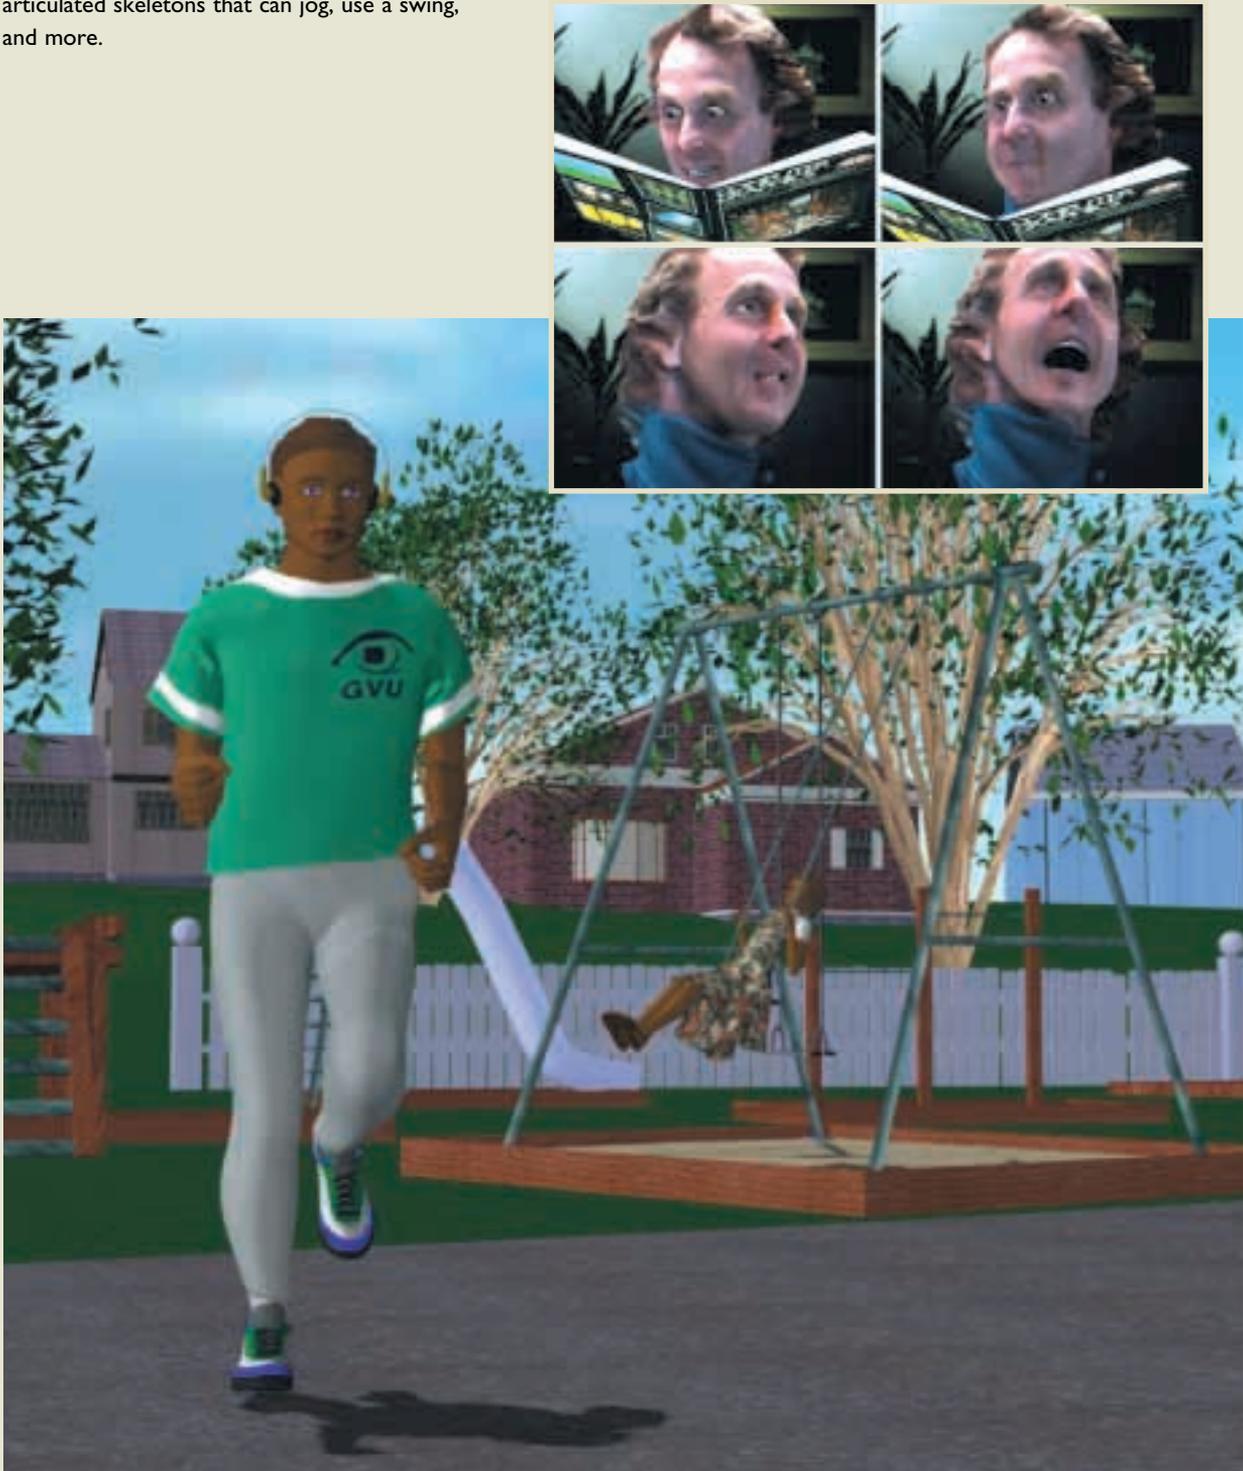
Figure 5. Interactive synthetic characters.

Interactive characters in the virtual environment *Swamped* (left). A "norn" speaks out in the interactive game *Creatures 2* (middle). *Demosaurus Rex* (right), an experimental interactive game engine with "cognitively enabled" prehistoric characters, developed at Intel and Angel Studios [2].

IMAGE COURTESY B. BLUMBERG, MIT MEDIA LAB; SEE WWW.MEDIA.MIT.EDU/PEOPLE/~BRUCE

IMAGE COURTESY T. SIMPSON, CYBERLIFE, CAMBRIDGE, ENGLAND; SEE WWW.CYBERLIFE.TECHNOLOGY.UK

His work on flocking showed how complex animations can emerge with minimal effort on the part of the animator through models of how characters should behave. These models take the form of behavioral rules governing the interaction of multiple autonomous agents capable of locomotion and perception within a virtual world. Behavioral modeling and animation now represents a major trend in computer graphics, and the technique has been applied extensively in the motion picture industry (see Table 1). Behavioral animation techniques, first demonstrated by Reynolds in the groundbreaking 1987 animated short film *Stanley and Stella*, have been used to create such feature film special effects as animated flocks of bats in *Batman Returns,* herds of wildebeests in *The Lion King,* and crowd scenes of epic proportions in *Mulan.* Behavioral techniques have also been popular for controlling multiple animated characters in interactive games and multimedia applications.

Artificial animals. A comprehensive artificial life approach to the realistic modeling of animals for animation and virtual reality is emerging, along with convincing results [9]. Its key ingredients are functional models of animal bodies and brains. Functional body modeling involves simulating the physics of the animal in its world, the use of biomechanics for locomo-

tion, and the operation of active sensory organs, such as eyes. Functional brain modeling involves emulating the information processing in biological brain centers responsible for motor control, perception, behavior, learning, and, in higher animals, cognition. I revisit the topic of artificial animals in greater depth later in the article, referring to my research group's work on artificial marine animals.

Artificial humans. For the typical computer graphics animator, the most important animal species is *Homo sapiens,* and significant effort has been invested in modeling and animating the most highly advanced living system known. In particular, the human face has attracted much attention, from Frederic Parke's pioneering work at the University of Utah in the 1970s to recent work on the biomechanical and anatomical modeling of faces [4] (see Figure 4). Daniel Thalmann at the Swiss Federal Institute of Technology in Lausanne and Nadia Magnenat-Thalmann at the University of Geneva in Switzerland have for years championed the grand challenge of modeling and animating virtual humans (see [8], ligwww.epfl.ch, and www.miralab.unige.ch). The Thalmanns have recently investigated the increasingly important role of sensory perception in human modeling, equipping virtual humans with visual, tactile, and auditory sensors to

make them aware of their environments. Sensory awareness supports such human behavior as visually directed locomotion, object manipulation, and response to sounds and utterances.

Their demonstrations include sensor-based navigation, walking on irregular terrain, grasping, game playing, behavior of virtual human crowds, and more. They have also explored communication among virtual humans and among real and virtual humans, including the compositing of interactive virtual humans into real scenes. Other researchers in the virtual human area are Norman Badler and Dimitri Metaxas at the University of Pennsylvania (see Badler et al.'s "Animation Control for Real-Time Virtual Humans" in this issue) and Jessica Hodgins at the Georgia Institute of Technology, whose impressive research involves biomechanical modeling and motor control of human locomotion for animation.

Interactive synthetic characters. Forthcoming computer-based interactive entertainment will captivate users with lifelike graphical characters in rich environments, an application for which artificial life techniques are well suited. Key concepts from artificial life are finding application in home entertainment, such as in the computer game *Creatures* (see Figure 5). *Creatures* allows users to interact with cute

autonomous agents called "norns," whose behavior is controlled by genetically specified neural networks and biochemistry. The commercial success of this title, with more than a million copies sold worldwide, reflects the relationships many users are willing to form with believable artificial life characters (see www.creatures.co.uk).

Increasingly potent artificial life modeling techniques for graphical characters are being explored in university and corporate research labs. Notably, Bruce Blumberg of the MIT Media Lab has developed prototype systems, such as the Artificial Life Interactive Virtual Environment (ALIVE) system, which enables full-body interaction between human participants and graphical worlds inhabited by engaging artificial life forms [1]. These characters have their own motivations and can sense and interpret the actions of other characters, as well as the human participants, responding to them in real time. His most recent project, which involves similar goals, is *Swamped* (see Figure 5).

The apex of the computer graphics modeling pyramid in Figure 1 addresses human cognitive functionality, an area that has challenged researchers in artificial intelligence since the field's earliest days. John Funge is pioneering the use in computer games and

Imagine a virtual marine world inhabited by a variety of realistic fishes. In the presence of underwater currents, the fishes employ their muscles and fins to swim gracefully around immobile obstacles and among moving aquatic plants and other fishes. They autonomously explore their dynamic world in search of food (see the figure). Large, hungry predators stalk smaller prey in this deceptively peaceful habitat. The sight of predators compels prey fishes to take evasive action. When a dangerous predator appears in the distance, similar species of prey form schools to improve their chances of survival. As the predator nears a school, the fishes scatter in terror. A chase ensues in which the predator selects and consumes victims until satiated. Some species seem untroubled by predators, finding comfortable niches and feeding on floating plankton when they get hungry. Driven by healthy libidos, they perform enticing courtship dances to attract mates.

From the perspective of computer animation, artificial fishes are not just highly realistic graphical puppets like the dinosaurs in the feature film *Jurassic Park* whose moves were painstakingly plotted by highly skilled human animators. Rather, these self-animating artificial animals have "eyes" that see their virtual world and "brains" that autonomously govern their actions. They swim, forage, eat, and mate entirely on their own. Given such artificial life models, computer animators can begin to play a role less like that of graphical model puppeteers and more like that of (National Geographic Society) nature cinematographers. **C**

Artificial fishes in their physics-based virtual world as it appears to an underwater observer. The three reddish fish are engaged in mating behavior, the greenish fish is a predator, and the remaining fishes are feeding on plankton (white dots) among seaweed growing from the ocean bed and swaying in the underwater currents. A predator shark stalks a school of prey fish. Functional diagram (right) of the artificial fish, showing body, brain, and display submodels [9].

Display Model
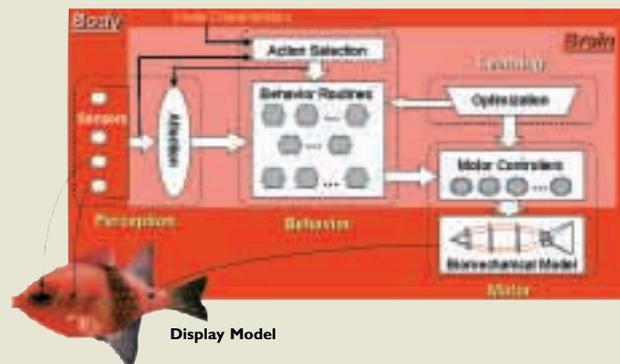
animation of hardcore artificial intelligence techniques, such as logic-based knowledge representation, reasoning, and planning [2]. This form of "animal logic" is a feature of *Demosaurus Rex,* an experimental interactive game environment, illustrated in Figure 5. Systematic cognitive modeling of this sort will ultimately lead to self-animating virtual humans and other graphical characters intelligent enough to be directed like human actors.

Lifelike Autonomous Agents

A particularly sophisticated artificial life model for computer animation was developed at the University of Toronto by Xiaoyuan Tu [9] (see the sidebar "Artificial Fishes"). The artificial fish is an autonomous agent with a deformable body actuated by internal muscles. The body also includes eyes and a brain with motor, perception, behavior, and learning centers, as shown in the sidebar figure. Through controlled muscle actions, artificial fishes swim through simulated water in accordance with hydrodynamic principles. Their articulate fins enable them to locomote, maintain balance, and maneuver in the water. Thus, the *functional* artificial fish model captures not just 3D shape and appearance in the form of a conventional computer graphics display model. More significantly, it also captures the basic physics of the animal (biomechanics) and its environment (hydrodynamics), as well as the function of the animal's brain. In accordance with its perceptual awareness of the virtual world, the brain of an artificial fish arbitrates a repertoire of piscine behaviors, including collision avoidance, foraging, preying, schooling, and mating. Though these artificial brains are rudimentary compared to the biological brains in real animals, they can also learn basic motor functions and carry out perceptually guided motor tasks, as shown by Radek Grzeszczuk while at the University of Toronto [3].

Artificial fish display models, such as the one in the sidebar, should capture the form and appearance of real fish with reasonable visual fidelity. To this end, we convert photographs of real fish into 3D spline surface body models using an interactive image-based modeling approach. The shapes and textures of the fish bodies are extracted from digitized photographs through computer vision techniques.

The motor system is the fish's dynamic model, including its muscle actuators and motor controllers. The biomechanical body model produces realistic piscine locomotion using only 23 lumped masses and 91 viscoelastic elements interconnected to maintain structural integrity under muscle-actuated deformation. Elements running longitudinally along the body function as actively contractile muscles. Artificial

fishes locomote like natural fishes—by autonomously contracting their muscles in a coordinated manner. As the body flexes, it displaces virtual fluid, producing thrust-inducing reaction forces that propel the fish forward. The mechanics are governed by systems of Lagrangian equations of motion (69 equations per fish) driven by hydrodynamic forces. The coupled, second-order ordinary differential equations are continually integrated through time by a numerical simulator (employing a stable, implicit Euler time-integration method). The model achieves a good compromise between realism and computational efficiency, while permitting the design of motor controllers using data gleaned from the literature on fish biomechanics.

A set of motor controllers in the motor center of the artificial fish's brain coordinates muscle actions to carry out specific motor functions, such as swimming forward, turning left, and turning right. Additional motor controllers coordinate the actions of the pectoral fins, enabling the neutrally buoyant artificial fish to pitch, roll, and yaw its body in order to navigate freely in its 3D world.

Artificial fishes are aware of their world through sensory perception, relying on a set of onboard virtual sensors to gather information about the dynamic environment. To achieve natural sensorimotor behaviors, it is necessary to model not only the abilities but also the limitations of animal perception systems. The perception center of the brain includes an attention mechanism that allows the artificial fish to sense the world in a task-specific way. For example, the artificial fish attends to sensory information about nearby food sources when foraging. (A biomimetic approach to perception based on computational vision was developed by Tamer Rabie at the University of Toronto; see www.cs.toronto.edu/~dt/animat-vision.)

The behavior center of an artificial fish's brain mediates between its perception system and its motor system. A set of innate characteristics determines a fish's (static) genetic legacy, such as whether it is male or female, predator or prey, and more. A (kinetic) mental state includes variables representing hunger, fear, and libido, whose values depend on sensory inputs. The behavioral repertoire of an artificial fish includes such primitive, reflexive behavior routines as obstacle avoidance, as well as more sophisticated motivational behavior routines, such as schooling and mating, whose activation depends on the fish's kinetic mental state. An appropriately curtailed (piscene) cognitive capacity stems from the fish's action-selection mechanism.

At each simulation time step, action selection entails combining the innate characteristics, the

mental state, and the incoming stream of the fish's sensory information to generate sensible, survival-dependent goals, such as avoiding obstacles and predators, hunting and feeding on prey, and courting potential mates. A limited, single-item behavior memory reduces dithering, giving goals persistence to improve the robustness of the more prolonged foraging, schooling, and mating behaviors.

Though rudimentary compared to those in real animals, the brains of artificial fishes are also able to learn. For example, an artificial fish can learn to loco-mote through practice and sensory reinforcement. Coordinated muscle contractions that produce forward movements are remembered. These partial successes then form the basis for the fish's subsequent improvement in its swimming technique. More formally, the learning center of an artificial fish's brain comprises a set of optimization-based motor learning algorithms that can discover and perfect muscle controllers capable of producing efficient locomotion. These artificial animals can also learn to accomplish higher-level sensorimotor tasks, including the delightful sort of leaping stunts trained dolphins perform at aquatic theme parks [3].

A large-scale virtual seaquarium implemented by Qinxin Yu at the University of Toronto was installed at SMART Toronto (www.sto.org), February 1998. This real-time version of the artificial fishes virtual reality simulation enables users to pilot a virtual submarine through the synthetic undersea world. Participants are fascinated by the lifelike marine animals they observe through the virtual submarine's huge, panoramic "window." The simulation runs in a Trimension Reality Theater (see www.trimension-inc.com), which combines an SGI eight-processor (R10000) Onyx2 system and multichannel PRO-DAS projection technology from U.K.-based SEOS Displays. The system features three SGI InfiniteReality graphics pipelines, each feeding video to an overhead-mounted projector. The system simulates, animates, and renders the 3D virtual marine world at a sustainable rate greater than 30 frames per second. The three projectors produce a seamless display of approximately 4000×1000-pixel resolution across an 18-foot by 8-foot curved screen.

Conclusion

Artificial life is playing an increasingly important role across computer graphics, including image synthesis, modeling, animation, interactive games, multimedia, and virtual reality. A fast-growing technological corpus is available for realistically modeling and animating objects that are alive. The techniques emulate phenomena fundamental to bio-logical organisms, such as birth, growth, and death, biomechanics and motor control, awareness, behavior, intelligence, and even evolution.

The exciting challenge to graphics researchers and practitioners is to develop and creatively deploy ever more sophisticated graphics models—that are self-creating, self-controlling, self-animating, and self-evolving—by simulating the mechanisms of life. That an area of such interdisciplinary reach can emerge within computer graphics is a testament to the richness and scope of the field. **C**

REFERENCES

1. Blumberg, B., and Galyean, T. Multi-level direction of autonomous creatures for real-time virtual environments. In *Proceedings of SIGGRAPH'95* (Los Angeles, Aug. 6–11). ACM Press, New York, 1995, pp. 47–54.
2. Funge, J., Tu, X., and Terzopoulos, D. Cognitive modeling: Knowledge, reasoning, and planning for intelligent characters. In *Proceedings of SIGGRAPH'99* (Los Angeles, Aug. 8–13). ACM Press, New York, 1999; see also Funge, J. *AI for Games and Animation: A Cognitive Modeling Approach*, A.K. Peters, Wellesley, Mass., 1999.
3. Grzeszczuk, R., and Terzopoulos, D. Automated learning of muscle-actuated locomotion through control abstraction. In *Proceedings of SIGGRAPH'95* (Los Angeles, Aug. 6–11). ACM Press, New York, 1995, pp. 63–70.
4. Lee, Y., Terzopoulos, D., and Waters, K. Realistic modeling for facial animation. In *Proceedings of SIGGRAPH'95* (Los Angeles, Aug. 6–11). ACM Press, New York, 1995, pp. 55–62.
5. Mëch, R., and Prusinkiewicz, P. Visual models of plants interacting with their environment. In *Proceedings of SIGGRAPH'96* (New Orleans, Aug. 4–9). ACM Press, New York, pp 397–410; see also Prusinkiewicz, P., and Lindenmayer, A. *The Algorithmic Beauty of Plants*, Springer-Verlag, Berlin, 1990.
6. Reynolds, C. Flocks, herds, and schools: A distributed behavioral model. In *Proceedings of SIGGRAPH'87* (Anaheim, Calif., July 27–31). ACM Press, New York, pp. 25–34.
7. Sims, K. Artificial evolution for computer graphics. In *Proceedings of SIGGRAPH'91* (Las Vegas, July 28–Aug. 2). ACM Press, New York, 1991, pp. 319–328; see also Evolving virtual creatures. In *Proceedings of SIGGRAPH'94* (Orlando, Fla., July 24–29). ACM Press, New York, 1994, pp. 15–22.
8. Thalmann D. The artificial life of virtual humans. In *Artificial Life for Graphics, Animation, Multimedia, and Virtual Reality (SIGGRAPH'98 Course Notes)*, D. Terzopoulos, Ed. ACM Press, New York, 1998; see also H. Noser et al. Playing games through the virtual life network. In *Artificial Life V*, C. Langton and K. Shimohara, Eds., MIT Press, Cambridge, Mass., 1996, pp. 135–142.
9. Tu, X., and Terzopoulos, D. Artificial fishes: Physics, locomotion, perception, behavior. In *Proceedings of SIGGRAPH'94* (Orlando, Fla., July 24–29). ACM Press, New York, 1994, pp. 43–50; see also Tu, X. *Artificial Animals for Computer Animation*. ACM Distinguished Dissertation Series, Springer-Verlag, Berlin, 1999.

DEMETRI TERZOPOULOS (dt@cs.toronto.edu) is a professor of computer science and a professor of electrical and computer engineering at the University of Toronto, where he directs the Visual Modeling Group (see www.cs.toronto.edu/~dt).