



August 1999

Animation Control for Real-Time Virtual Humans

Norman I. Badler

University of Pennsylvania, badler@seas.upenn.edu

Martha Palmer

University of Pennsylvania

Ramamani Bindiganavale

University of Pennsylvania

Follow this and additional works at: <http://repository.upenn.edu/hms>

Recommended Citation

Badler, N. I., Palmer, M., & Bindiganavale, R. (1999). Animation Control for Real-Time Virtual Humans. Retrieved from <http://repository.upenn.edu/hms/111>

Postprint version. Copyright ACM, 1999. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *Communications of the ACM*, Volume 42, Issue 8, August 1999, pages 64-73.

Publisher URL: <http://doi.acm.org/10.1145/310930.310975>

This paper is posted at ScholarlyCommons. <http://repository.upenn.edu/hms/111>

For more information, please contact libraryrepository@pobox.upenn.edu.

Animation Control for Real-Time Virtual Humans

Abstract

The computation speed and control methods needed to portray 3D virtual humans suitable for interactive applications have improved dramatically in recent years. Real-time virtual humans show increasingly complex features along the dimensions of appearance, function, time, autonomy, and individuality. The virtual human architecture we've been developing at the University of Pennsylvania is representative of an emerging generation of such architectures and includes low-level motor skills, a mid-level parallel automata controller, and a high-level conceptual representation for driving virtual humans through complex tasks. The architecture—called Jack— provides a level of abstraction generic enough to encompass natural-language instruction representation as well as direct links from those instructions to animation control.

Comments

Postprint version. Copyright ACM, 1999. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *Communications of the ACM*, Volume 42, Issue 8, August 1999, pages 64-73.

Publisher URL: <http://doi.acm.org/10.1145/310930.310975>



☞ Marilyn on subway grate, from the film *Flashback*, 1990.

Image designed by Nadia Magnenat-Thalmann, University of Geneva, Switzerland, Daniel Thalmann, Swiss Federal Institute of Technology, Lausanne, Switzerland, and Benoit Lafleur.

[Norman I. Badler, Martha S. Palmer, and
Rama Bindiganavale]

WANT TO MAKE VIRTUAL HUMANS MORE HUMAN?
LET THEIR FLESH-AND-BLOOD COUNTERPARTS
ANIMATE THEIR ACTIONS AND INTENTIONS THROUGH
NATURAL-LANGUAGE INSTRUCTIONS.

Animation Control for REAL-TIME VIRTUAL HUMANS

The computation speed and control methods needed to portray 3D virtual humans suitable for interactive applications have improved dramatically in recent years. Real-time virtual humans show increasingly complex features along the dimensions of appearance, function, time, autonomy, and individuality. The virtual human architecture we've been developing at the University of Pennsylvania is representative of an emerging generation of such architectures and includes low-level motor skills, a mid-level parallel automata controller, and a high-level conceptual representation for driving virtual humans through

complex tasks. The architecture—called *Jack*—provides a level of abstraction generic enough to encompass natural-language instruction representation as well as direct links from those instructions to animation control.

Only 50 years ago, computers could barely compute useful mathematical functions. About 25 years ago, enthusiastic computer researchers were predicting that game-playing machines and

autonomous robots performing such surrogate functions as mining gold on asteroids were in our future. Today's truth lies somewhere in between. We have balanced our expectations of complete machine autonomy with a more rational view that machines should assist people in accomplishing meaningful, difficult, and often enormously complex tasks. When such tasks involve human interaction with the physical world, com-

Table 1. Requirements of representative virtual human applications.

Application	Appearance	Function	Time	Autonomy	Individuality
Cartoons	high	low	high	low	high
Games	high	low	low	medium	medium
Special Effects	high	low	high	low	medium
Medicine	high	high	medium	medium	medium
Ergonomics	medium	high	medium	medium	low
Education	medium	low	low	medium	medium
Tutoring	medium	low	medium	high	low
Military	medium	medium	low	medium	low

putational representations of the human body—virtual humans—can be used to escape the constraints of presence, safety, and even physicality.

Why are real-time virtual humans so difficult to construct? After all, anyone who can watch a movie can see marvelous synthetic animals, characters, and people. But they are typically created for a single scene or movie and are neither autonomous nor meant to engage in interactive communication with real people. What makes a virtual human human is not just a well-executed exterior design, but movements, reactions, self-motivated decision making, and interactions that appear “natural,” appropriate, and contextually sensitive. Virtual humans designed to be able to communicate with real people need uniquely human abilities to show us their actions, intentions, and feelings, building a bridge of empathy and understanding. Researchers in virtual human characters seek methods to create digital people that share our human time frame as they act, communicate, and serve our applications.

Still, many interactive and real-time applications already involve the portrayal of virtual humans, including:

- Engineering.** Analysis and simulation for virtual prototyping and simulation-based design.
- Virtual conferencing.** Teleconferencing, using virtual representations of participants to increase personal presence.
- Monitoring.** Acquiring, interpreting, and understanding shape and motion data related to human movement, performance, activities, and intent.
- Virtual environments.** Living and working in a virtual place for visualization, analysis, training, and even just the experience.
- Games.** Real-time characters with actions, alternatives, and personality for fun and profit.
- Training.** Skill development, team coordination, and decision making.
- Education.** Distance mentoring, interactive assistance, and personalized instruction.
- Military.** Simulated battlefield and peacekeeping

operations with individual participants. **Maintenance.** Designing for such human factors and ergonomics as ease of access, disassembly, repair, safety, tool clearance, and visibility.

Along with general industry-driven improvements in the underlying computer and graphical display technologies, virtual humans will enable quantum leaps in applications normally requiring personal and live human participation. The emerging MPEG-4 specification, for example, includes face- and body-animation parameters for real-time display synthesis.

Fidelity

Building models of virtual humans involves application-dependent notions of fidelity. For example, fidelity to human size, physical abilities, and joint and strength limits are essential to such applications as design evaluation. And in games, training, and military simulations, temporal fidelity in real-time behavior is even more important. Appreciating that different applications require different sorts of virtual fidelity prompts a number of questions as to what makes a virtual human “right”: What do you want to do with it? What do you want it to look like? What characteristics are important to the application’s success? and What type of interaction is most appropriate?

Different models of virtual-human development provide different gradations of fidelity; some are quite advanced in a particular narrow area but are more limited for other desirable features. In a general way, we can characterize the state of virtual-human modeling along at least five dimensions, each described in the following progressive order of feature refinement:

- Appearance.** 2D drawings, 3D wireframe, 3D polyhedra, curved surfaces, freeform deformations, accurate surfaces, muscles, fat, biomechanics, clothing, equipment, physiological effects, including perspiration, irritation, and injury.
- Function.** Cartoon, jointed skeleton, joint limits, strength limits, fatigue, hazards, injury, skills,

effects of loads and stressors, psychological models, cognitive models, roles, teaming.

Time. Off-line animation, interactive manipulation, real-time motion playback, parameterized motion synthesis, multiple agents, crowds, coordinated teams.

Autonomy. Drawing, scripting, interacting, reacting, making decisions, communicating, intending, taking initiative, leading.

Individuality. Generic character, hand-crafted character, cultural distinctions, personality, psychological-physiological profiles, gender and age, specific individual.

Different applications require human models that individually customize these dimensions (see Table 1). A model tuned for one application may be inad-

tem for creating, sizing, manipulating, and animating virtual humans. Our philosophy has yielded a particular virtual-human development model that pushes the five dimensions of virtual-human performance toward the more complex features. Here, we focus on the related architecture, which supports enhanced functions and autonomy, including control through textual—and eventually spoken—human natural-language instructions.

Other universities pursuing virtual human development include: the computer graphics laboratory at the Swiss Federal Institute of Technology in Lausanne, Georgia Institute of Technology, Massachusetts Institute of Technology Media Lab, New York University, the University of Geneva, the University of Southern California, and the University of Toronto. Companies include: ATR Japan, Credo, Engineering Animation, Extempo, Kinetix, Microsoft, Motion Factory, Phillips, Sony, and many others [3, 12].

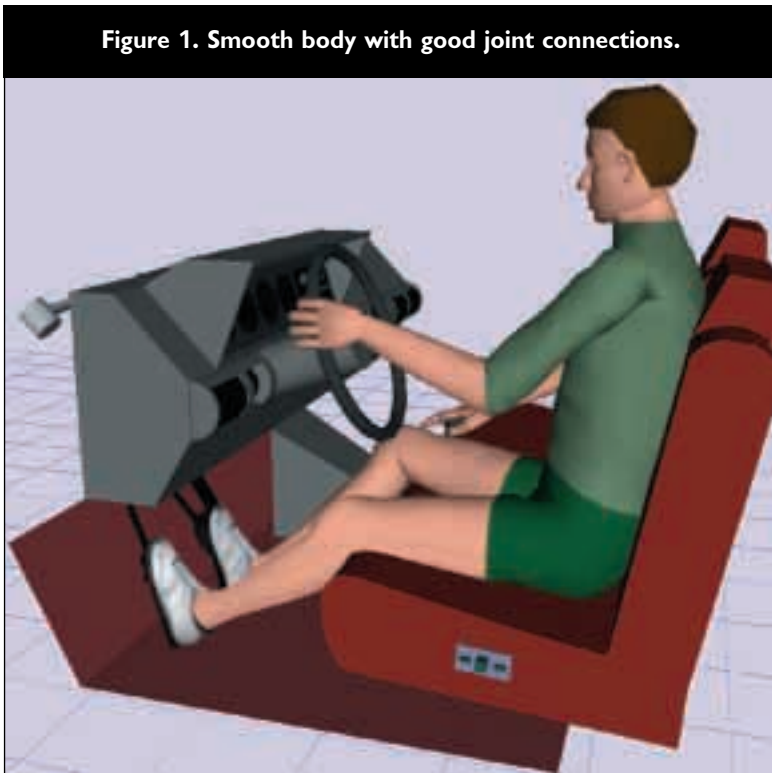
Levels of Architectural Control

Building a virtual human model that admits control from sources other than direct animator manipulations requires an architecture that supports higher-level expressions of movement. Although layered architectures for autonomous beings are not new, we have found that a particular set of architectural levels seems to provide efficient localization of control for both graphics and language requirements. A description of our multilevel architecture starts with typical graphics models and articulation structures, and includes various motor skills for endowing virtual humans with useful abilities. The higher architectural levels organize these skills with parallel automata, use a conceptual representation to describe the actions a virtual human can per-

form, and finally create links between natural language and action animation.

Graphical models. A typical virtual human model design consists of a geometric skin and an articulated skeleton. Usually modeled with polygons to optimize graphical display speed, a human body can be crafted manually or shaped more automatically from body segments digitized by laser scanners. The surface may be rigid or, more realistically, deformable during movement. Deformation demands additional modeling and computational loads. Clothes are desirable, though today, loose garments have to be animated

Figure 1. Smooth body with good joint connections.



equate for another. And many research and development efforts concentrate on refining one or more dimensions deeper into their special features. One challenge for commercial efforts is the construction of virtual human models with enough parameters to effectively support several application areas.

At the University of Pennsylvania, we have been researching and developing virtual human figures for more than 25 years [2]. Our framework is comprehensive and representative of a broad multiapplication approach to real-time virtual humans. The foundation for this research is Jack, our software sys-

offline due to computational complexity.

The skeletal structure is usually a hierarchy of joint rotation transformations. The body is moved by changing the joint angles and its global position and location. In sophisticated models, joint angle changes induce geometric modifications that keep joint surfaces smooth and mimic human musculature within a character's particular body segment (see Figure 1).

Real-time virtual humans controlled by real humans are called "avatars." Their joint angles and other location parameters are sensed by magnetic, optical, and video methods and converted to joint rotations and body pose. For movements not based on

Simultaneous execution also leads to the next level of our architecture's organization: parallel automata.

Parallel transition networks. Almost 20 years ago, we realized that human animation would require some model of parallel movement execution. But it wasn't until about 10 years ago that graphical workstations were finally powerful enough to support functional implementations of simulated parallelism. Our parallel programming model for virtual humans is called Parallel Transition Networks, or PaT-Nets. Other human animation systems, including Motion Factory's Motivate and New York University's Improv [9], have adopted similar paradigms with alternative syntactic structures. In general, net-

A VIRTUAL HUMAN SHOULD BE ABLE TO WALK, TALK, AND CHEW GUM AT THE SAME TIME.

live performance, computer programs have to generate the right sequences and combinations of parameters to create the desired movements' desired actions. Procedures for changing joint angles and body position are called motion generators, or motor skills.

Motor skills. Virtual human motor skills include:

- Playing a stored motion sequence that may have been synthesized by a procedure, captured from a live person, or scripted manually;
- Posture changes and balance adjustments;
- Reaching and other arm gestures;
- Grasping and other hand gestures;
- Locomoting, such as stepping, walking, running, and climbing;
- Looking and other eye and head gestures;
- Facial expressions, such as lip and eye movements;
- Physical force- and torque-induced movements, such as jumping, falling, and swinging; and
- Blending one movement into another, in sequence or in parallel.

Numerous methods help create each of these movements, but we want to allow several of them to be executed simultaneously. A virtual human should be able to walk, talk, and chew gum at the same time.

work nodes represent processes and arcs, which connect the nodes, and contain predicates, conditions, rules, and other functions that trigger transitions to other process nodes. Synchronization across processes or networks is made possible through message-passing or global variable blackboards to let one process know the state of another process.

The benefits of PaT-Nets derive not only from their parallel organization and execution of low-level motion generators, but from their conditional structure. Traditional animation tools use linear timelines on which actions are placed and ordered. A PaT-Net provides a nonlinear animation model, since movements can be triggered, modified, and stopped by transitions to other nodes. This type of nonlinear animation is a crucial step toward autonomous behavior, since conditional execution enables a virtual human's reactivity and decision making.

Providing a virtual human with humanlike reactions and decision-making skills is more complicated than just controlling its joint motions from captured or synthesized data. Simulated humanlike actions and decisions are how we convince the viewer of the character's skill and intelligence in negotiating its environment, interacting with its spatial situation, and engaging other agents. This level of perfor-

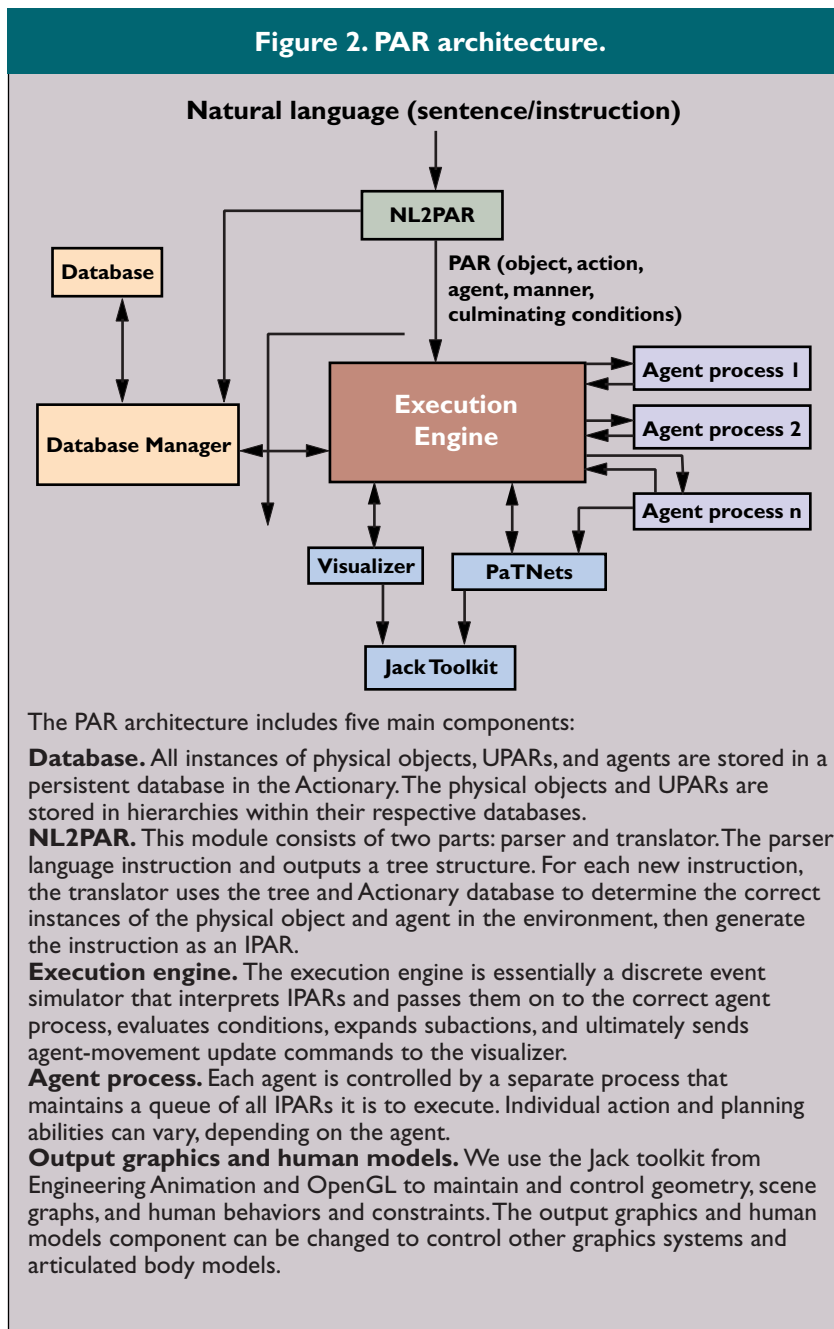
mance requires significant investment in action models that allow conditional execution. We have programmed a number of experimental systems to show how the PaT-Net architecture can be applied, including the game "Hide and Seek," two-person animated conversation [3], simulated emergency medical care [4], and the multiuser virtual world JackMOO [10].

PaT-Nets are effective but must be hand-coded in C++. No matter what artificial language we invent to describe human actions, it is not likely to represent exactly the way people conceptualize a particular situation. We therefore need a higher-level representation to capture additional information, parameters, and aspects of human action. We create such representations by incorporating natural-language semantics into our parameterized action representation.

Conceptual action representation. Even with a powerful set of motion generators and PaT-Nets to invoke them, we still have to provide effective and easily learned user interfaces to control, manipulate, and animate virtual humans. Interactive point-and-click tools (such as Maya from Alias | Wavefront, 3D StudioMax from Autodesk, and SoftImage from Avid), though usable and effective, require specialized training and animation skills and are fundamentally designed for off-line production. Such interfaces disconnect the human participant's instructions and actions from the avatar through a narrow communication channel of hand motions. A programming language or scripting interface, while powerful, is yet another off-line method requiring specialized programming expertise.

A relatively unexplored option is a natural-language-based interface, especially for expressing the intentions behind a character's motions. Perhaps not surprisingly, instructions for real people are given in natural language, augmented with graphical dia-

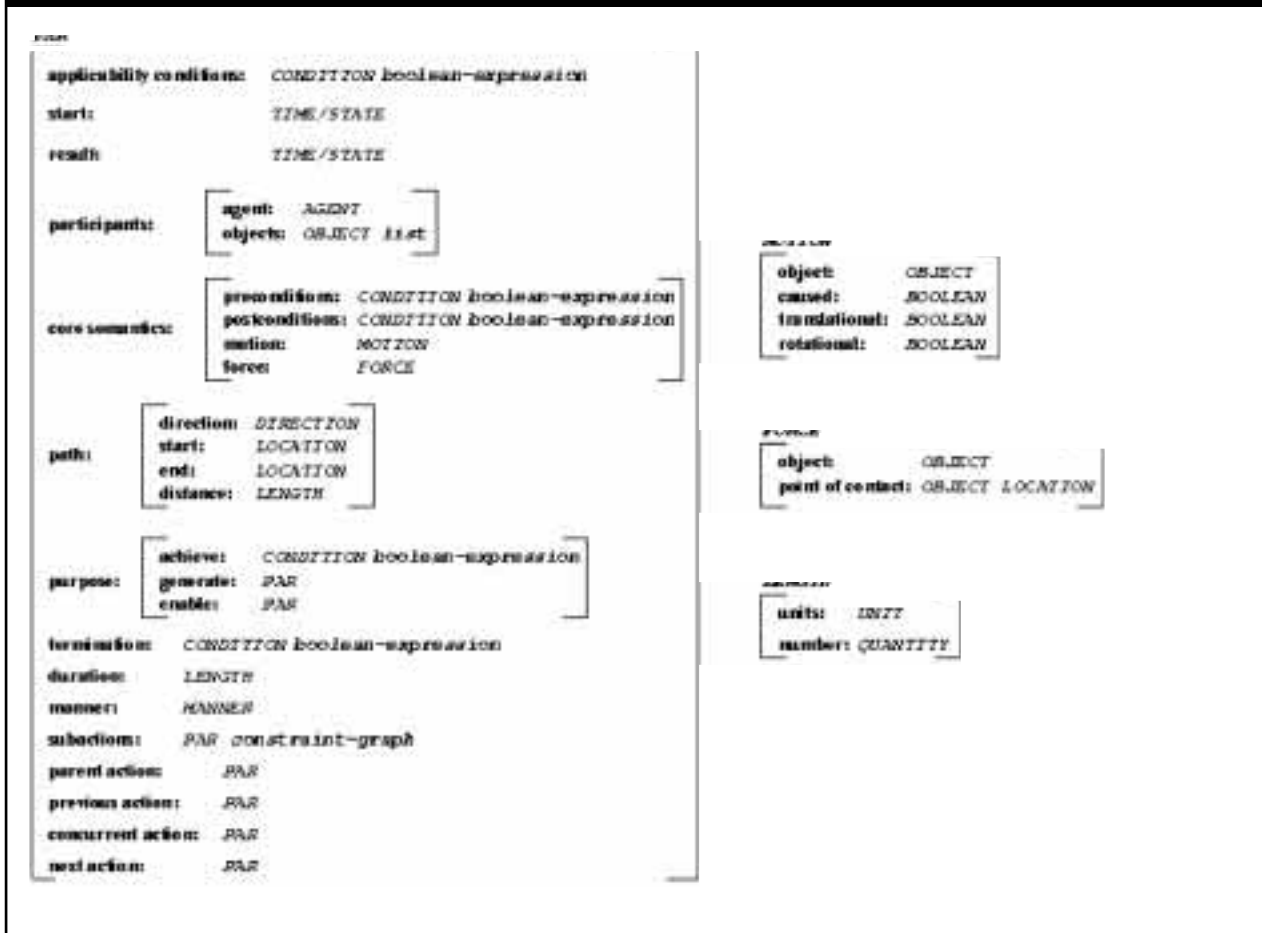
Figure 2. PAR architecture.



grams and, occasionally, animations. Recipes, instruction manuals, and interpersonal conversations can therefore use language as their medium for conveying process and action.

We are not advocating that animators throw away their tools, only that natural language offers a communication medium we all know and can use to formulate instructions for activating the behavior of virtual human characters. Some aspects of some actions are certainly difficult to express in natural language, but the availability of a language interpreter can bring the virtual human interface more in line with real interpersonal communication modes. Our

Figure 3. PAR template.



goal is to build smart avatars that understand what we tell them to do in the same way humans follow instructions. These smart avatars have to be able to process a natural-language instruction into a conceptual representation that can be used to control their actions. This representation is called a parameterized action representation, or PAR (see Figure 2).

The PAR has to specify the agent of the action, as well as any relevant objects and information about paths, locations, manners, and purposes for a particular action. There are linguistic constraints on how this information can be conveyed by the language; agents and objects tend to be verb arguments, paths are often prepositional phrases, and manners and purposes might be in additional clauses [8]. A parser maps the components of an instruction into the parameters or variables of the PAR, which is then linked directly to PaT-Nets executing the specified movement generators.

Natural language often describes actions at a high level, leaving out many of the details that have to be specified for animation, as discussed in a similar

approach in [7]. We use the example “Walk to the door and turn the handle slowly” to illustrate the function of the PAR. Whether or not the PAR system processes this instruction, there is nothing explicit in the linguistic representation about grasping the handle or which direction it will have to be turned, yet this information is necessary to the action’s actual visible performance. The PAR has to include information about applicability and preparatory and terminating conditions in order to fill in these gaps. It also has to be parameterized, because other details of the action depend on the PAR’s participants, including agents, objects, and other attributes.

The representation of the “handle” object lists the actions that object can perform and what state changes they cause. The number of steps it will take to get to the door depends on the agent’s size and starting location. Some of the parameters in a PAR template are shown in Figure 3 and are defined in the following ways:

Physical objects. These objects are referred to

Figure 4. Scene from Jack's MOOse Lodge.



within the PAR; each one has a graphical model and other properties. The walking action has an implicit floor as an object, while the turn action refers to the handle.

Agent. The agent executes the action. The user's avatar is the implied agent, and the walking and turning actions share the same agent. An agent has a specific personality and a set of actions it knows how to execute.

Start. This moment is the time or state in which the action begins.

Result. This is the state after the action is performed.

Applicability conditions. The conditions in this boolean expression must be true to perform the action. Conditions generally have to do with certain properties of the objects, the abilities of the agent, and other unchangeable or uncontrollable aspects of the environment. For "walk," one of the applicability conditions may be "Can the agent walk?" If conditions are not satisfied, the action cannot be executed.

Preparatory actions. These actions may have to be performed to enable the current action to proceed. In general, actions can involve the full power of motion planning to determine, perhaps, that a handle has to be grasped before it can be turned. The instructions are essentially goal requests, and the smart avatar must then figure out how (if possible) it can achieve them. We use hand-coded conditionals to test for likely (but generalized) situations and execute appropriate intermediate actions. Adding more general action planners is also possible, since the PAR represents goal states and supports a full graphical model of the current world state.

Subactions. Each action is organized into partially ordered or parallel substeps, called subactions. Actions described by PARs are ultimately executed as PaT-Nets.

Core semantics. These semantics represent an action's primary components of meaning and include preconditions, postconditions, motion, force, path, purpose, terminating conditions,

Figure 5. Virtual trainer for military checkpoints.



duration, and agent manner. For example, “walking” is a form of locomotion that results in a change of location. “Turning” requires a direction and an end point.

A PAR can appear as one of two different forms: uninstantiated PAR (UPAR) and instantiated PAR (IPAR): We store all instances of the UPAR, which contains default applicability conditions, preconditions, and execution steps, in a hierarchical database called the Actionary. Multiple entries are allowed, in the same way verbs have multiple contextual meanings. An IPAR is a UPAR instantiated with specific information on agent, physical object(s), manner, terminating conditions, and more. Any new information in an IPAR overrides the corresponding UPAR default. An IPAR can be created by the parser (one IPAR for each new instruction) or dynamically during execution, as in Figure 2.

A language interpreter promotes a language-centered view of action execution, augmented and elaborated by parameters modifying lower-level

motion synthesis. Although textual instructions can describe and trigger actions, details need not be communicated explicitly. The smart avatar PAR architecture interprets instruction semantics with motion generality and context sensitivity. In a prototype implementation of this architecture, called Jack’s MOOSE Lodge [10], four smart avatars are controlled by simple imperative instructions (see Figure 4). One agent, the waiter, is completely autonomous, serving drinks to the seated avatars when their glasses need filling. Another application runs a military checkpoint (see Figure 5).

Realistic Humanlike Movements

Given this architecture, do we see the emergence of realistic humanlike movements, actions, and decisions? Yes and no. We see complex activities and interactions. But we also know we’re not fooling anyone into thinking that these virtual humans are real. Some of this inability to mimic real human movements and interactions perfectly has to do with graphical appearance and motion details; real

humans readily identify synthetic movements. Motion captured from live performances is much more natural, but more difficult to alter and parameterize for reuse in other contexts.

One promising approach to natural movement is through a deeper look into physiological and cognitive models of behavior. For example, we have built a visual attention system for the virtual human that uses known perceptual and cognitive parameters to drive the movement of our characters' eyes (see Terzopoulos's "Artificial Life for Computer Graphics" in this issue). Visual attention is based on a queue of tasks and exogenous events that can occur arbitrarily [1]. Since attention is a resource, task performance degrades naturally as the environment becomes cluttered.

Another approach is to observe human movement and understand the qualitative parameters that shape performance. In the real world, the shaping of performance is a physical process; in our simulated worlds, assuming we choose the right controls, it may be modeled kinematically. That's why we implemented an interpretation of Laban's effort notation, which characterizes the qualitative rather than the quantitative aspects of movement, to create a parameterization of agent manner [1]. Effort elements are weight, space, time, and flow and can be combined and phrased to vary the performance of a given gesture.

Individualized Perceptions of Context

Within five years, virtual humans will have individual personalities, emotional states, and live conversations [11]. They will have roles, gender, culture, and situational awareness. They will have reactive, proactive, and decision-making behaviors for action execution [6]. But to do these things, they will need individualized perceptions of context. They will have to understand language so real humans can communicate with them as if they were real.

The future holds great promise for the virtual humans populating our virtual worlds. They will provide economic benefits by helping designers build more human-centered vehicles, equipment, assembly lines, manufacturing plants, and interactive systems. Virtual humans will enhance the presentation of information through training aids, virtual experiences, teaching, and mentoring. They will help save lives by providing surrogates for medical training, surgical planning, and remote telemedicine. They will be our avatars on the Internet, portraying ourselves to others—as we are, or perhaps as we wish to be. And they may help turn cyberspace into a real community. **G**

REFERENCES

1. Badler, N., Chi, D., and Chopra, S. Virtual human animation based on movement observation and cognitive behavior models. In *Proceedings of the Computer Animation Conference* (Geneva, Switzerland, May 8–10). IEEE Computer Society, Los Alamitos, Calif., 1999, pp. 128–137.
2. Badler, B., Phillips, C., and Webber, B. *Simulating Humans: Computer Graphics Animation and Control*. Oxford University Press, New York, 1993; see www.cis.upenn.edu/~badler/book/book.html.
3. Cassell, J., Pelachaud, C., Badler, N., Steedman, M., Achorn, B., Becket, W., Douville, B., Prevost, S., and Stone, M. Animated conversation: Rule-based generation of facial expression, gesture and spoken intonation for multiple conversational agents. In *Proceedings of Computer Graphics, Annual Conf. Series* (Orlando, Fla., July 24–29). ACM Press, New York, 1994, pp. 413–420.
4. Chi, D., Webber, B., Clarke, J., and Badler, N. Casualty modeling for real-time medical training. *Presence* 5, 4 (Fall 1995), 359–366.
5. Earnshaw, R., Magnenat-Thalmann, N., Terzopoulos, D., and Thalmann, D. Computer animation for virtual humans. *IEEE Comput. Graph. Appl.* 18, 5 (Sept.–Oct. 1998), 20–23.
6. Johnson, W., and Rickel, J. Steve: An animated pedagogical agent for procedural training in virtual environments. *SIGART Bulletin* 8, 1–4 (Fall 1997), 16–21.
7. Narayanan, S. Talking the talk is like walking the walk. In *Proceedings of the 19th Annual Conference of the Cognitive Science Society* (Palo Alto, Calif., Aug. 7–10 1997).
8. Palmer, M., Rosenzweig, J., and Schuler, W. Capturing motion verb generalizations with synchronous tag. In *Predicative Forms in NLP: Text, Speech, and Language Technology Series*, P. St. Dizier, Ed. Kluwer Press, Dordrecht, The Netherlands, 1998.
9. Perlin, K., and Goldberg, A. Improv: A system for scripting interactive actors in virtual worlds. In *Proceedings of ACM Computer Graphics, Annual Conference Series* (New Orleans, Aug. 4–9). ACM Press, New York, 1996, pp. 205–216.
10. Shi, J., Smith, T., Granieri, J., and Badler, B. Smart avatars in JackMOO. In *Proceedings of IEEE Virtual Reality'99 Conference* (Houston, Mar. 13–17). IEEE Computer Society Press, Los Alamitos, Calif., 1999, 156–163.
11. Thorisson, K. Real-time decision making in multimodel face-to-face communication. In *Proceedings of the 2nd International Conference on Autonomous Agents* (Minneapolis-St. Paul, May 10–13). ACM Press, New York, 1998, pp. 16–23.
12. Wilcox, S. *Web Developer.com Guide to 3D Avatars*. John Wiley & Sons, New York, 1998.

NORMAN I. BADLER (badler@central.cis.upenn.edu) is a professor of computer and information science in the Center for Human Modeling and Simulation in the Department of Computer and Information Science at the University of Pennsylvania, Philadelphia.

MARTHA S. PALMER (mpalmer@linc.cis.upenn.edu) is a visiting associate professor in the Center for Human Modeling and Simulation in the Department of Computer and Information Science at the University of Pennsylvania, Philadelphia.

RAMA BINDIGANAVALA (rama@graphics.cis.upenn.edu) is a Ph.D. student and systems programmer in the Center for Human Modeling and Simulation in the Department of Computer and Information Science at the University of Pennsylvania, Philadelphia.

This research is supported by the U.S. Air Force through Delivery Orders #8 and #17 on F41624-97-D-5002; Office of Naval Research (through the University of Houston) K-5-55043/3916-1552793, DURIP N0001497-1-0396, and AASERTs N00014-97-1-0603 and N0014-97-1-0605; Army Research Lab HRED DAAL01-97-M-0198; DARPA SBMDA-97-2951001; NSF IRI95-04372; NASA NRA NAG 5-3990; National Institute of Standards and Technology 60 NANB6D0149 and 60 NANB7D0058; Engineering Animation, Inc., SERI, Korea, and JustSystem, Inc., Japan.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.