# Real Time Virtual Humans

*Norman I. Badler, Rama Bindiganavale, Juliet Bourne, Jan Allbeck,*
*Jianping Shi, and Martha Palmer*
Center for Human Modeling and Simulation
Department of Computer and Information Science
University of Pennsylvania Philadelphia, PA 19104-6389

## Abstract

The last few years have seen great maturation in the computation speed and control methods needed to portray 3D virtual humans suitable for real interactive applications. Various dimensions of real-time virtual humans are considered, such as appearance and movement, autonomous action, and skills such as gesture, attention, and locomotion. A virtual human architecture includes low level motor skills, mid-level PaT-Net parallel finite-state machine controller, and a high level conceptual action representation that can be used to drive virtual humans through complex tasks. This structure offers a deep connection between natural language instructions and animation control.

## 1 Virtual Humans

Only fifty years ago, computers were barely able to compute useful mathematical functions. Twenty-five years ago, enthusiastic computer researchers were predicting that all sorts of human tasks from game-playing to automatic robots that travel and communicate with us would be 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 to accomplish meaningful, difficult, and often enormously complex tasks. When those tasks involve human interaction with the physical world, computational representations of the human body can be used to escape the constraints of presence, safety, and even physicality.

*Virtual humans* are computer models of people that can be used

- as substitutes for ``the real thing'' in *ergonomic* evaluations of computer-based designs for vehicles, work areas, machine tools, assembly lines, etc., *prior to the actual construction of those spaces*;
- for *embedding real-time representations of ourselves or other live participants* into virtual environments.

Recent improvements in computation speed and control methods have allowed the portrayal of 3D humans suitable for interactive and real-time applications. These include:

- **Engineering**: Analysis and simulation for virtual prototyping and simulation-based design.
- **Virtual-Conferencing**: Efficient tele-conferencing using virtual representations of participants to reduce transmission bandwidth requirements.
- **Interaction**: Real-time graphical bodies inhabiting virtual worlds.
- **Monitoring**: Acquiring, interpreting, and understanding shape and motion data on human movement, performance, activities, or intent.

- **Virtual Environments**: Living and working in a virtual place for visualization, analysis, training, or just the experience.
- **Games**: Real-time characters with actions and personality for fun and profit.
- **Training**: Skill development, team coordination, and decision-making.
- **Education**: Distance mentoring, interactive assistance, and personalized instruction.
- **Military**: Battlefield simulation with individual participants, team training, and peace-keeping operations.
- **Design/Maintenance**: Design for access, ease of repair, safety, tool clearance, visibility, and hazard avoidance.

Besides general industry-driven improvements in the underlying computer and graphical display technologies, virtual humans will enable quantum leaps in applications requiring personal and live participation.

In building models of virtual humans, there are varying notions of *virtual fidelity*. Understandably, these are application dependent. For example, fidelity to human size, capabilities, and joint and strength limits are essential to some applications such as design evaluation; whereas in games, training, and military simulations, temporal fidelity (real-time behavior) is essential. Understanding that different applications require different sorts of virtual fidelity leads to the question *of 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 success of the application?

There are gradations of fidelity in the models: some models are very advanced in a narrow area but lack other desirable features.

In a very general way, we can characterize the state of virtual human modeling along at least five dimensions, each with a wide range of realizations. Some significant datapoints along each one are listed below:

1. Appearance: 2D drawings > 3D wireframe > 3D polyhedra > curved surfaces > freeform deformations > accurate surfaces > muscles, fat > biomechanics > clothing, equipment > physiological effects (perspiration, irritation, injury)
2. Function: cartoon > jointed skeleton > joint limits > strength limits > fatigue > hazards > injury > skills > effects of loads and stressors > psychological models > cognitive models > roles > teaming
3. Time: off-line animation > interactive manipulation > real-time motion playback > parameterized motion synthesis > multiple agents > crowds > coordinated teams
4. Autonomy: drawing > scripting > interacting > reacting > making decisions > communicating > intending > taking initiative > leading
5. Individuality: generic character > hand-crafted character > cultural distinctions > personality > psychological-physiological profiles > gender and age > specific individual

Different applications require specialized human models that individually optimize character, performance, intelligence, and so on. Many research and development efforts concentrate on pushing the envelope of one or more dimensions toward the right.

If the need demands it, the *appearance* of increasingly accurate physiologically- and biomechanically-grounded human models may be obtained. We can create virtual humans with *functional* limitations that go beyond cartoons into instantiations of known human factors data. Animated virtual humans can be created in human *time* scales through motion capture or computer synthesis. Virtual humans are also beginning to exhibit *autonomy* and intelligence as they react and make decisions in novel, changing environments rather than being forced into fixed movements. Finally, rather several efforts are underway to create characters with *individuality* and *personality* who react to and interact with other real or virtual people[1] [2] [3] [4] [5] [6].

Across various applications, different capabilities are required as shown in Table 1. A model that is tuned for one application may not be adequate for another. An interesting challenge is be build virtual human models with enough parameters to provide effective support cross several application areas.

| 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** |
| Medical | **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** |

**Table 1: Comparing Applications for Virtual Humans**

We have been very actively engaged in research and development of virtual human figures for over 25 years[7]. Our interest in human simulation is not unique, and others have well-established efforts that complement our own, for example[8] [9] [10] [11] [12], The framework for our research is a system called *Jack*®[*]. Our philosophy has led to a particular realization of a virtual human model that pushes the above five dimensions toward the more complex features. In particular, here we will look at various aspects of each of the dimensions above, primarily working toward enhanced function and autonomy.

Why are real time virtual humans difficult to construct? After all, anyone who goes to the movies can see marvelous synthetic characters but they have been created typically for one scene or one movie and are not meant to be re-used (except possibly by the animator -- and certainly not by the viewer). The difference lies in the *interactivity* and *autonomy* of virtual humans. What makes a virtual human *human* is not just a well-executed exterior design but movements, reactions, and decision-making which appear natural, appropriate, and context-sensitive. Communication by and with

---

[*] Jack® is now the basis of a commercial software product distributed by Engineering Animation, Inc.

virtual humans gives them a uniquely human capability: they can let us know their intentions, goals, and feelings thus building a bridge of empathy and understanding. Ultimately we should be able to communicate with virtual humans through all our natural human modalities just as if they, too, were real.

## 2 Levels of Control

Animating virtual humans may be accomplished through a variety of means. To build a model that admits control from other than direct animator manipulations, however, requires an architecture to support higher-level expressions of movement. While layered architectures for autonomous beings are not new[13], we have found that a particular set[14] of levels seems to provide an efficient localization of control with sympathies to both graphics and language requirements. We examine this multi-level architecture, starting with a brief description of typical graphics models and articulation structure. We then examine various motor skills that empower virtual humans with useful capabilities. We organize these skills with parallel automata at the next level. The highest level uses a conceptual representation to describe actions and allows linkage between natural languages and action animation.

### *2.1 Graphical Models*

A typical virtual human model consists of a geometric skin and an articulated skeleton. Usually modeled with polygons to optimize graphical display speed, a human body may be manually crafted or more automatically shaped from body segments digitized by laser scanners. The surface may be rigid or, more realistically, deformable during movement. The latter accrues additional modeling and computational loads. Animated clothes are a desirable addition, but presently must be done offline[15] [16].

The skeletal structure is usually a hierarchy of joint rotation transformations. The body is moved by changing the joint angles and the global position and location of the body. In sophisticated models (Figure 1), joint angle changes induce geometric modifications that keep joint surfaces smooth and mimic human musculature within the body segment[17] [18]

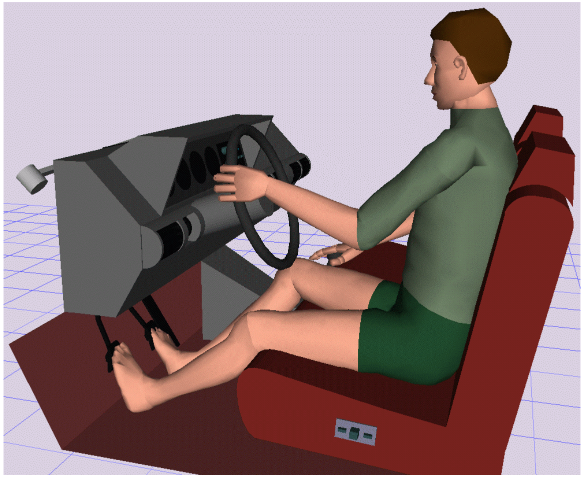

**Figure 1. Smooth Body (by Bond-Jay Ting).**

Animated virtual humans may be controlled by real people, in which case they are called *avatars*. The joint angles and other location parameters are sensed by magnetic, optical, or video methods, and converted to rotations for the virtual body. For a purely synthetic figure, computer programs must generate the right sequences and combinations of parameters to create the desired movements. Procedures to change joint angles and body position are called *motion generators* or *motor skills*.

## 2.2 Motor Skills

Typical virtual human motor skills include:

- Playing a stored motion sequence; this may have been synthesized by a procedure, captured from a live person, or manually scripted.
- Posture changes and balance adjustments.
- Reaching (and other arm gestures).
- Grasping (and other hand gestures).
- Locomoting (stepping, walking, running, climbing).
- Looking (and other head gestures).
- Facial expressions.
- Physical force- or torque-induced movements (jumping, falling, swinging).
- Blending (coarticulating) one movement into the next one.

Numerous methods exist for each of these; a comprehensive survey is beyond our scope. What is important here is that several of these activities may be executed simultaneously: a virtual human should be able to *walk, talk, and chew gum*. This leads to the next level of architectural organization: *Parallel Transition Networks*.

## 2.3 Parallel Transition Networks

Two decades ago we realized that human animation would require some model of parallel movement execution. About a decade ago[19] graphical workstations became fast enough to support feasible implementations of simulated parallelism. Our model for a parallel virtual machine that animates graphical models are called Parallel Transition Networks or PaT-Nets. Other human animation systems have adopted similar paradigms. In general, network nodes represent processes and arcs contain predicates, conditions, rules, or other functions that cause transitions to other process nodes. Synchronization across processes or networks is effected through message-passing or global variable blackboards.

The benefits of PaT-Nets accrue not only from their parallel organization and execution of low level motor skills, but also from their conditional structure. Traditional animation tools use linear time-lines on which actions are placed and ordered. A PaT-Net provides a *non-linear* animation model, since movements can be triggered, modified, or stopped by transition to other nodes. This is the first crucial step toward autonomous behavior since conditional execution enables reactivity and decision-making capabilities.

Providing a virtual human with human-like reactions and decision-making is more complicated than just controlling its joint motions from captured or synthesized data. Here is where we need to 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 performance requires significant investment in non-linear action models. Through numerous experimental systems we have shown how the PaT-Net architecture can be applied: games such as *Hide and Seek*[20], two person animated *conversation (Gesture Jack)*[3], simulated emergency medical care (*MediSim*)[21], a real-time animated *Jack Presenter*[22 23], and multi-user *JackMOO*[24] virtual worlds.

PaT-Nets are effective but must be hand coded in Lisp or C++. No matter what artificial language we invent to describe human actions, it is not likely to be just the way people conceptualize the situation[*]. We therefore need a higher level, conceptual representation to capture additional information, parameters, and aspects of human action. We do this by drawing on natural language semantic concepts.

## 2.4 Conceptual Action Representation

Even with a powerful set of motion generators and PaT-Nets to invoke them, a challenge remains to provide effective and easily learned user interfaces to control, manipulate and animate virtual humans. Interactive point and click systems (such as *Jack* and numerous other animation production toolsets) work now, but with a cost in user learning and menu traversal. Such interfaces decouple the human participant's instructions and actions from the avatar through a narrow and *ad hoc* communication channel of hand motions. A direct programming interface, while powerful, is still an off-line method that moreover requires specialized computer programming understanding and expertise. The option that remains is a natural language-based interface.

Perhaps not surprisingly, instructions for people are given in natural language augmented with graphical diagrams and occasionally, animations. Recipes, instruction manuals, and interpersonal conversations use language as the medium for conveying process and action[7 25 26]. The key to linking language and animation lies in constructing *Smart Avatars* that *understand what we tell them to do*. This requires a conceptual representation of actions, objects, and agents which is simultaneously suitable for execution (simulation) as well as natural language expression. We call this architectural level the *Parameterized Action Representation* or PAR. It must drive a simulation (in a context of a given set of objects and agents), and yet support the enormous range of expression, nuance, and manner offered by language[27]. The PAR gives a high level description of an action that is also directly linked to PaT-Nets which execute movements. A PAR is *parameterized* because an action depends on its participants (agents, objects, and other attributes) for the details of how it is accomplished. A PAR includes *applicability* and *preparatory* conditions that have to be satisfied before the action is actually executed. The action is finished when the *terminating* conditions are satisfied. Some of the PAR slots are described below:

- **Physical Objects**: the list of objects referred to within the PAR. Each physical object has a graphical model and other properties.

---

[*] Discussions with Bonnie Webber led to this observation.

- **Agent**: the agent who will be executing the action. Here, the user's *avatar* is the implied agent. An agent is a special type of object and has additional capabilities such as a set of actions it knows how to execute.
- **Start**: the time or state in which the action begins.
- **Result**: the time or state after the action is performed.
- **Applicability Conditions**: a boolean expression of conditions (conditions conjoined with logical **and**s and **or**s) which must hold (be true) in order for the action to be appropriate to perform. These 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. Unlike the preconditions (see below), it would be impossible or impractical to try to satisfy the applicability conditions as sub-goals before performing the action. For *walk* one of the applicability conditions may be: *Can the agent walk?* If not, conditions are not satisfied and the action is aborted. Going across the street requires that the agent be mobile and self-propelled in some fashion. Applicability conditions may also replace an action with a more specific one: opening the door might be specialized to a sliding action if that is what this particular door calls for.
- **Subactions**: the breakdown of the action into partially-ordered or parallel sub-steps. It is a collection of actions connected in a graph structure which indicates the temporal relationships (if any) between the actions (e.g. whether two actions are to be done sequentially, in parallel, *etc.*). Actions ground out as PaT-Nets. Thus a PAR can either describe a complex action or a primitive action. A complex action can list a number of sub-actions that may need to be executed in sequence, parallel, or a combination of both. A primitive action is a PaT-Net. Parameters pass from PAR to PaT-Net to motion process.

  In general, preparatory actions or applicability conditions may involve the full power of motion planning. The commands, after all, are essentially goal requests[28] and the smart avatar must then figure out how (if at all) it can achieve them. Presently we use PaT-Nets with hand coded conditionals to test for likely (but generalized) situations and execute appropriate intermediate actions. Adding more general actions planners is possible since the PAR represents goal states and supports a full graphical model of the current world state[20].
- **Core Semantics**: the primary components of meaning of the action and includes Preconditions, Postconditions, Motion, Force, Path, Purpose, Terminating Conditions, Duration, and Agent Manner.

A PAR appears in two different forms:

- **UPAR(Uninstantiated PAR)**: We store all instances of the uninitialized PAR in a database (called the *Actionary*) in a hierarchical tree. A UPAR contains default applicability conditions, preconditions, and execution steps. This is the heart of the Actionary. Multiple entries are allowed: just as verbs have multiple contextual meanings. *Go to bed* means much more than *go to the door* because it entails preparatory (and possibly) optional actions such as *undressing* and *lying down when at the bed*.
- **IPAR (Instantiated PAR)**: An IPAR is a UPAR instantiated with specific information on agent, physical object(s), manner, terminating

conditions, etc. 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 can be created dynamically during execution.

## 2.5 Architecture

Figure 2 shows the architecture of the PAR system.

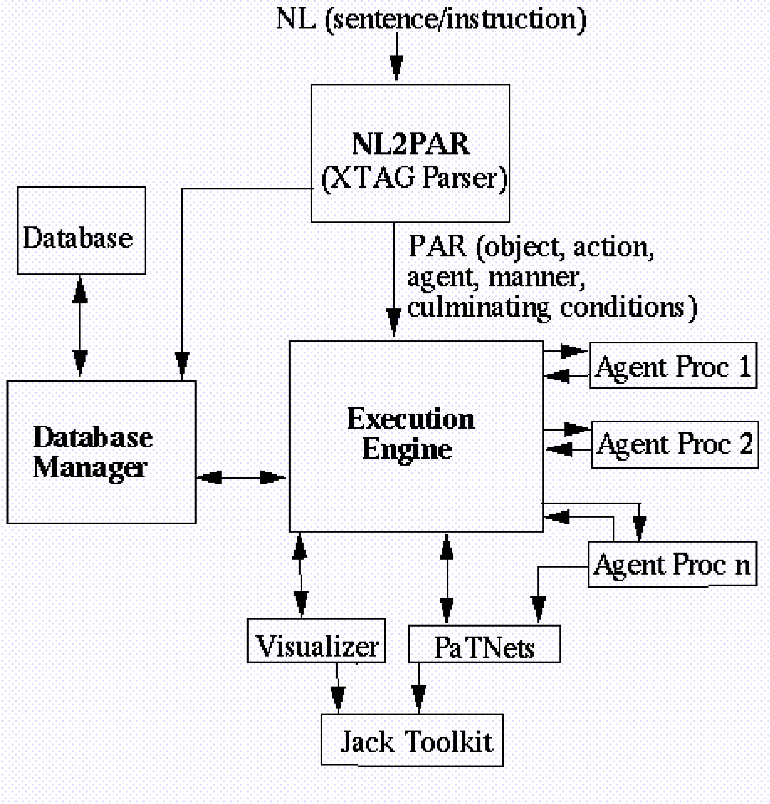

**Figure 2. PAR Architecture**

- **NL2PAR**: This module consists of two parts: parser and translator. The parser takes a natural language instruction and outputs a tree structure. For each new instruction, the translator uses the tree and Actionary database to first determine the correct instances of the physical object and agent in the environment. It then generates the instruction as an IPAR.
- **Database**: All instances of physical objects, UPARs, and agents are stored in a persistent database contained in the Actionary. The physical objects and UPARs are stored in hierarchies within their respective databases.
- **Execution Engine**: The execution engine is the main controller for the agent actions. It accepts a PAR from the NL2PAR module, passes it on to the correct agent process, evaluates conditions, expands PARs if necessary, and ultimately sends agent movement update commands to the visualizer.

- **Agent Process:** Each agent is controlled by a separate process, which maintains a queue of all IPARs it is to execute. Individual action capabilities and planning abilities may vary across agents.
- **Output Graphics and Human Models:** We use the EAI/Transom Jack toolkit and OpenGL to maintain and control the actual geometry, scene graphs, and human behaviors and constraints. This component may be easily changed to control other articulated body models.

A language interpreter promotes a *language-centered* view of action execution, but augmented and elaborated by parameters modifying lower-level motion synthesis. Although textual instructions can describe and trigger actions, details need not be explicitly communicated. The smart avatar PAR architecture interprets the semantics of instructions for both motion generality and environmental context-sensitivity. In a prototype implementation of this architecture, called Jack's MOOse Lodge[24], four smart avatars are controlled by simple imperative instructions (Figure 3). One agent, the waiter, is completely autonomous and serves drinks to seated avatars when their glasses need filling.

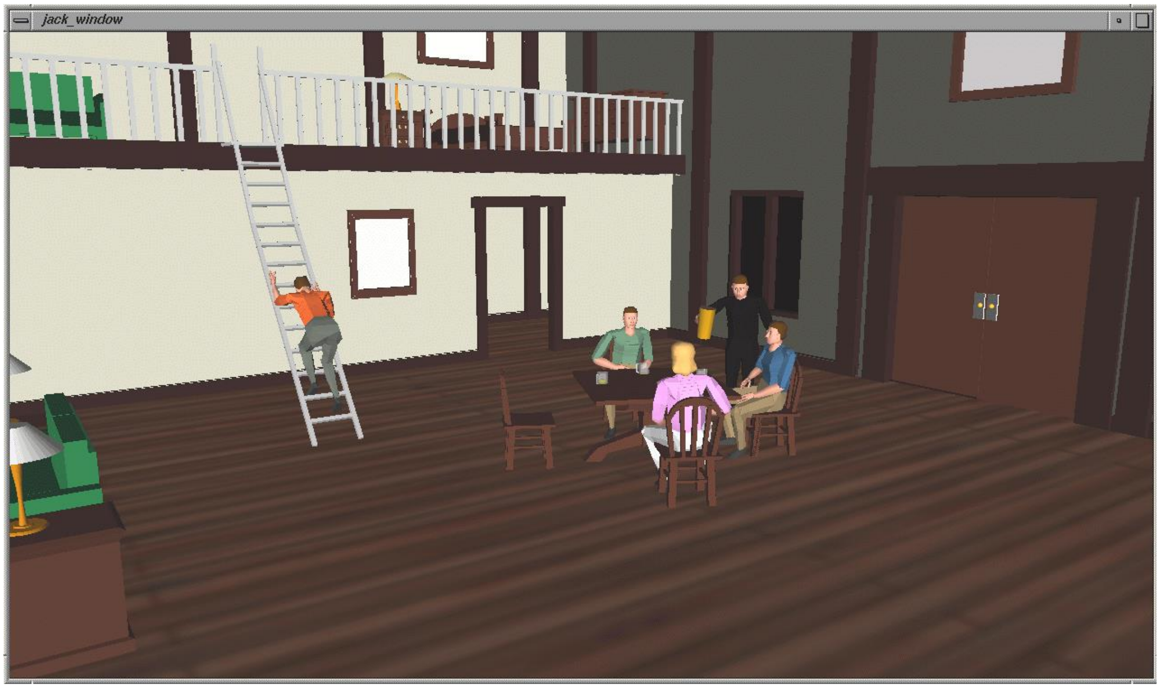

**Figure 3. Jack's MOOse Lodge.**

## 3 Discussion

This exposition has described virtual human modeling and control, with an emphasis on real-time motion and language-based interfaces. In particular, we discussed such issues as appearance and motion, autonomous action, and motor skills. A PaT-Net parallel finite-state machine controller can be used to drive virtual humans through complex tasks.

We next described a first version of a Parameterized Action Representation. The PAR is meant to be the intermediate structure between natural language instructions with complex semantics and task execution by a virtual human agent. An algorithm for interpreting PARs within an object-oriented system has been implemented.

We have established a role for language in action modeling. Linguistic classifications have helped us by identifying typical properties and modifiers of animate agents, such as the dimensions along which agent behavior can vary. In addition, linguistic analysis can help identify typical actions of animate agents and typical modifiers for their actions. Basing an agent and action ontology on linguistic evidence and movement models ensures extensibility. However, the development of the virtual human model from the bottom-up assures that a rich set of necessary capabilities are present.

Given this architecture, do we see the emergence of realistic human-like movements, actions, and decisions? Yes and no. On the positive side, we see complex activities and interactions. On the negative side, we're not fooling anyone into thinking that the virtual humans are real. While some of this has to do with graphical appearance, synthetic movements are still easy to pick out. Motion captured from live performances is much more natural, but harder to alter and parameterize for re-use in other contexts.

One approach to natural movement that offers some promise is to look deeper into physiological and cognitive models of behavior. For example, we have built an attention system for the virtual human that uses known perceptual and cognitive parameters to drive the movement of the eyes. Attention is based on a queue of tasks and exogenous events that may impinge arbitrarily. Since attention is a resource, as the environment becomes cluttered, task performance naturally degrades[29]. Attention can also predict re-appearance of temporarily occluded objects.

Another approach is to observe human movement and understanding the parameters that shape performance. In the real world this is a physical process; in our simulated world it may be modeled kinematically if we choose the right controls. We have implemented[30] an interpretation of Laban's Effort notation to have a parameterization of agent manner. The Effort elements are Weight, Space, Time, and Flow; they may be combined and phrased to effect the performance of a given set of key poses for a character's arms, hands, and body.

Soon virtual humans will have individual personalities, emotional states, and live conversations[31]. They will have roles, gender, culture, and situation awareness[32]. They will have reactive, proactive, and decision-making behaviors for action execution[33]. They will need to have individualized perceptions of context. They must understand language so that we may communicate with them as if they were real.

The future holds great promise for the virtual humans who will populate our virtual worlds. They will provide economic benefits by helping designers early in the product design phases to produce 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. And Virtual humans will help save lives by providing surrogates for medical training, surgical planning, and remote telemedicine. They will be our avatars on the Internet and will portray ourselves to others, perhaps as we are or perhaps as we wish to be. They may help turn cyberspace into a real, or rather virtual, community.

## 4 Acknowledgments

Many students, staff, and colleagues in the Center for Human Modeling and Simulation make this effort possible. Additional information and contributors may be found through **http://hms.upenn.edu/**.

This research is partially supported by U.S.~Air Force through Delivery Orders #8 and #17 on F41624-97-D-5002; Office of Naval Research (through Univ. 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 SB-MDA-97-2951001; NSF IRI95-04372; NASA NRA NAG 5-3990; National Institute of Standards and Technology 60 NANB6D0149 and 60 NANB7D0058; SERI Korea, and JustSystem Japan.

## 5 References

[1] J. Bates. The role of emotion in believable agents. *Comm. of the ACM*, 37(7), pages 122—125, 1994.

[2] J. Bates, A. Loyall, and W. Reilly. Integrating reactivity, goals, and emotion in a broad agent. *In Proc. of the 14th Annual Conf. of the Cognitive Science Society*, pages 696—701, Hillsdale, NJ, 1992. Lawrence Erlbaum.

[3] J. Cassell, C. Pelachaud, N. Badler, M. Steedman, B. Achorn, W. Becket, B. Douville, S. Prevost, and M. Stone. Animated conversation: Rule-based generation of facial expression, gesture and spoken intonation for multiple conversational agents. In *Computer Graphics*, Annual Conf. Series, pages 413—420. ACM, 1994.

[4] P. Maes, T. Darrell, B. Blumberg, and A. Pentland. The ALIVE system: Full-body interaction with autonomous agents. In (N. Magnenat-Thalmann and D. Thalmann, editors) Computer Animation, pages 11—18. IEEE Computer Society Press, Los Alamitos, CA, 1995.

[5] K. Perlin and A. Goldberg. Improv: A system for scripting interactive actors in virtual worlds. In *ACM Computer Graphics*, Annual Conf. Series, pages 205—216, 1996.

[6] D. Rousseau and B. Hayes-Roth. Personality in synthetic agents. Technical Report KSL-96-21, Stanford Knowledge Systems Laboratory, 1996.

[7] N. Badler, C. Phillips, and B. Webber. Simulating Humans: Computer Graphics Animation and Control. Oxford University Press, New York, NY, 1993.

[8] R. Earnshaw, N. Magnenat-Thalmann, D. Terzopoulos, and D. Thalmann. Computer animation for virtual humans. *IEEE Computer Graphics and Applications*, 18(5), pages 20—23, Sept.-Oct. 1998.

[9] S. K. Wilcox. Web Developer's Guide to 3D Avatars. Wiley, New York, 1998.

[10] J. Hodgins, W. Wooten, D. Brogan, and J. O'Brien. Animating human athletics. In *ACM Computer Graphics*, Annual Conf. Series, pages 71—78, 1995.

[11] M. Cavazza, R. Earnshaw, N. Magnenat-Thalmann, and D. Thalmann. Motion control of virtual humans. *IEEE Computer Graphics and Applications*, 18(5), pages 24—31, Sept.-Oct. 1998.

[12] K. Perlin. Real time responsive animation with personality. IEEE Trans. on Visualization and Computer Graphics, 1(1), pages 5—15, 1995.

[13] R. Brooks. A robot that walks: Emergent behaviors from a carefully evolved network. Neural Computation, 1(2), 1989.

[14] D. Zeltzer. Task-level graphical simulation: Abstraction, representation, and control. In N. Badler, B. Barsky, and D. Zeltzer, editors, Making Them Move: Mechanics, Control, and Animation of Articulated Figures, pages 3—33, Morgan-Kaufmann, San Francisco, 1990.

[15] M. Carignan, Y. Yang, N. Magnenat-Thalmann, and D. Thalmann. Dressing animated synthetic actors with complex deformable clothes. ACM Computer Graphics, Proc. SIGGRAPH '92, pages 99—104, July 1992.

[16] D. Baraff and A. Witkin. Large steps in cloth simulation. ACM Computer Graphics, Proc. SIGGRAPH '98, pages 43—54, July 1998.

[17] J. Wilhelms and A. van Gelder. Anatomically-based modeling. ACM Computer Graphics, Proc. SIGGRAPH '97, pages 173—180, July 1997.

[18] B.-J. Ting. Real time human model design. PhD thesis, CIS, University of Pennsylvania, 1998.

[19] N. Badler and S. Smoliar. Digital representations of human movement. *ACM Computing Surveys*, 11(1), pages 19—38, 1979.

[20] T. Trias, S. Chopra, B. Reich, M. Moore, N. Badler, B. Webber, and C. Geib. Decision networks for integrating the behaviors of virtual agents and avatars. In *Proceedings of Virtual Reality International Symposium*, 1996.

[21] D. Chi, B. Webber, J. Clarke, and N. Badler. Casualty modeling for real-time medical training. *Presence*, 5(4), pages 359—366, 1995.

[22] T. Noma and N. Badler. A virtual human presenter. In IJCAI '97 Workshop on Animated Interface Agents, Nagoya, Japan, 1997.

[23] L. Zhao and N. Badler. Gesticulation behaviors for virtual humans. Proc. Pacific Graphics, pages 161-168, 1998.

[24] J. Shi, T. J. Smith, J. Granieri, and N. Badler. Smart avatars in JackMOO. In *IEEE Virtual Reality Conf.*, 1999.

[25] N. Badler, B. Webber, J. Kalita, and J. Esakov. Animation from instructions. In N. Badler, B. Barsky, and D. Zeltzer, editors, Making Them Move: Mechanics, Control, and Animation of Articulated Figures, pages 51—93. Morgan-Kaufmann, San Francisco, 1990.

[26] B. Webber, N. Badler, B. Di Eugenio, C. Geib, L. Levison, and M. Moore. Instructions, intentions and expectations. *Artificial Intelligence J.*, 73, pages 253—269, 1995.

[27] N. Badler, B. Webber, M. Palmer, T. Noma, M. Stone, J. Rosenzweig, S. Chopra, K. Stanley, J. Bourne, and B. Di Eugenio. Final report to Air Force HRGA regarding feasibility of natural language text generation from task networks for use in automatic generation of Technical Orders from DEPTH simulations. Technical report, CIS, University of Pennsylvania, 1997.

[28] N. Badler, B. Webber, W. Becket, C. Geib, M. Moore, C. Pelachaud, B. Reich, and M. Stone. Planning for animation. In N. Magnenat-Thalmann and D. Thalmann, editors, Computer Animation. Prentice-Hall, 1996.

[29] S. Chopra. Where to look? Automating some visual attending behaviors of virtual human characters. Ph.D. Dissertation, CIS, University of Pennsylvania, 1999.

[30] D. Chi. Animating expressivity through Effort elements. Ph.D. Dissertation, CIS, University of Pennsylvania, 1999.

[31] K. Thorisson. Real-time decision making in multimodal face-to-face communication. *Proc. Second Annual Conf. on Autonomous Agents*, ACM, 1998.

[32] J. Allbeck and N. Badler. Avatars a lá snow crash. In *Proc. Computer Animation*. IEEE Press, 1998.

[33] W. L. Johnson and J. Rickel. Steve: An animated pedagogical agent for procedural training in virtual environments. *SIGART Bulletin*, 8(1-4), pages 16—21, 1997.