1-1-1995

The Center for Human Modeling and Simulation

Norman I. Badler

University of Pennsylvania, badler@seas.upenn.edu

Dimitris Metaxas

University of Pennsylvania

Bonnie L. Webber

University of Pennsylvania, bonnie@inf.ed.ac.uk

Mark Steedman

University of Pennsylvania

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

Recommended Citation

Badler, N. I., Metaxas, D., Webber, B. L., & Steedman, M. (1995). The Center for Human Modeling and Simulation. Retrieved from <https://repository.upenn.edu/hms/72>

Copyright Massachusetts Institute of Technology. Reprinted from *Presence*, Volume 4, Issue 1, 1995, pages 81-96.

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/hms/72>
For more information, please contact repository@pobox.upenn.edu.

The Center for Human Modeling and Simulation

Abstract

The overall goals for the Center for Human Modeling and Simulation are the investigation of computer graphics modeling, animation, and rendering techniques. Major foci are in behavior-based animation of human movement, modeling through physics-based techniques, applications of control theory techniques to dynamic models, illumination models for image synthesis, and understanding the relationship between human movement, natural language, and communication.

Comments

Copyright Massachusetts Institute of Technology. Reprinted from *Presence*, Volume 4, Issue 1, 1995, pages 81-96.

The Center for Human Modeling and Simulation

Abstract

The overall goals of the Center for Human Modeling and Simulation are the investigation of computer graphics modeling, animation, and rendering techniques. Major foci are in behavior-based animation of human movement, modeling through physics-based techniques, applications of control theory techniques to dynamic models, illumination models for image synthesis, and understanding the relationship between human movement, natural language, and communication.

I Background

Our research on human body modeling and simulation has been done in the Computer Graphics Research Lab of the Computer and Information Science (CIS) Department at the University of Pennsylvania since approximately 1975. We have achieved international recognition for our research and specifically the *Jack*¹ software. Much of the Lab's work is reported in a recent book (Badler, Phillips, & Webber, 1993b). In January 1994, we became the *Center for Human Modeling and Simulation* (HMS), a formal University entity with a Director (Norman Badler) and an external Advisory Committee.

Government and industry support for our research efforts over the last 14 years has brought several millions of dollars of research funds into the University of Pennsylvania. Part of these funds came through the Army Research Office Center of Excellence in Artificial Intelligence at the University of Pennsylvania. More than half of the rest of the funding came from various Army, ARPA, NASA, Air Force, and Navy entities. The remainder of the funds came from industry grants,

through University support of graduate students and lab renovations, and from *Jack* license fees.

I.1 Staff

The Center has a full time Associate Director (Karen Carter), a Technology Transfer consultant (Dawn Becket), and three full time technical staff [John Granieri, Mike Hollick, and Welton (Tripp) Becket] to organize, execute, and oversee software development. Other computing systems support staff is available through the general research computing environment established within the CIS Department.

I.2 Educational Role

The Center provides a collegial and open atmosphere in which faculty, staff, and students cooperate and coordinate project work. Over 30 students are Ph.D. candidates in CIS: the Center educates a large fraction of the Ph.D.s in the CIS Department! Other students—at the Masters' level or from other University departments—are welcome and often active contributors. Undergraduate students also do independent study or Senior project work in the Center. All students taking the Computer Graphics courses learn to use *Jack* as the fundamental interaction and rendering system, and can utilize the film and video facilities of the Center.

I.3 Existing Facilities

The HMS Center occupies newly renovated, contiguous, and expanded facilities in the Moore Building.

Norman I. Badler, Dimitri Metaxas, Bonnie Webber, and Mark Steedman

Department of Computer & Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389

1. *Jack* is a registered trademark of The Trustees of The University of Pennsylvania.

Presence, Vol. 4, No. 1, Winter 1995, 81-96

© 1995 The Massachusetts Institute of Technology

Fifty years ago, this space housed the ENIAC computer! Currently the Center contains 24 Silicon Graphics workstations (Crimson VGXT, 4D-340VGX, Indigo Server, 2 Indigo Elan, Indigo Elan R4000, 3 Indigo XS24Z, 2 Indigo R4000 XS24Z, 2 4D70GT, 4D-50G, 5 4D20G, 5 4D30TG, 4D-70S) on Ethernet. Peripheral devices include a Tektronix color hardcopy unit, an HP color postscript printer, an Ascension Technologies Flock of [Big] Birds[™], an Apple LaserWriter, an Abekas A60 video disk recorder, an RGB Spectrum Videolink 1400AX, and video editors.

2 The HMS Center's Research Agenda

2.1 Foundations

Our research has built a solid underpinning of general techniques and theory:

- User interaction with 3D environments and especially articulated figures. The *Jack* user interface is one of its great attractions.
- Inverse kinematics running in real-time provides the underlying engine that permits manipulation of spatial constraints of a complex articulated figure, rather than requiring the user to manipulate individual joints directly.
- Reliance on open databases for human factors information. A user can, if desired, customize human capability, size, and characteristics to an application. The Army ANSUR data are the default human anthropometric database. The joint limit and strength data are mostly from NASA sources.
- Human behaviors are often goal-directed, and sets of parallel, interacting goals are the normal way to control movement.
- Balance maintenance triggers a number of natural human behaviors such as counterbalancing, stepping, turning, and gait modification.
- Much of the realism in human motion comes from strength and load distribution considerations. We have successfully demonstrated natural synthetic motions based on strength and load for lifting as well as locomotion tasks.
- Simulated sensors are needed to achieve realistic interaction with an environment. *Jack* includes self-collision avoidance as well as real-time environmental object avoidance.
- Human gesture and facial expressions are not arbitrary, but are shaped by the agent's communicative goals and intentions. We have linked an underlying planning and question-answering system with realistic human gestures and facial expressions as the queries and responses are uttered by a speech generation system. These mappings are based on cognitive science principles.
- Synthesizing known or novel motions in real-time is a prerequisite for embedding realistic human agents into Virtual Environments or Distributed Interactive Simulations. *Jack* is the first convincing demonstration of this capability in both the Army ISMS and the Navy TTES systems.
- Natural Language instructions have been largely ignored by the AI community as a research subject; we have changed that by considering the role, requirements, and impact of NL processing of instructions as animation specifications.

The *Jack* testbed provides a highly capable vehicle with which to carry out research in 3D human-computer interaction, human animation, and especially the link between language and action. An important goal of our work is to enable people to use Natural Language instructions to guide the behavior of semi-autonomous agents. These are agents with behaviors and skills of their own (from low-level behaviors such as walking, obstacle avoidance, and terrain navigation, to higher-level skills such as the ability to manipulate objects for particular ends) that a person could enlist in carrying out a task. The situation is analogous to that of the leader of a group and its team members: the leader may issue instructions and specify policies affecting how those instructions are to be carried out. The team members should carry out the orders in a situationally sensitive manner that is consistent with policy as well as their own low-level behavioral responses. The instructions themselves should reflect the leader's awareness of the team members' skills and behaviors. By linking Natural Lan-

guage to what is essentially real activity in a real world, we are able to experiment with theories concerning the situated interpretation of language in ways that theoreticians do not have access to.

2.2 Application Impact

We perform basic and applied research in human modeling and simulation. A number of particularly well-motivated application areas are

- human factors analysis,
- physical performance assessment,
- intelligent agents acting on instructions,
- maintenance/logistics technician,
- factory or assembly worker,
- dismounted, semi-autonomous soldier,
- virtual training simulator,
- medic or surgical assistant, and
- virtual human for surgical procedures.

There are technologically and economically justifiable uses for virtual human figures in these domains.

2.3 Past and Present

The first applications of *Jack* (and its predecessors) were in human factors analysis: fit, reach, accessibility, view, and comfort. These applications are evolving into systems for maintainability analysis such as DEPTH (USAF) and MAW (DMSO/University of Iowa). As *Jack*'s real-time behavior improved (and hardware became faster), real-time applications of the *Jack* human became viable: VE/VR avatars that could mimic an actual human subject's movements from simple sensor input (Army Research Lab and Sandia projects). Recently, precompiling complex behaviors such as walking, running, crawling, kneeling, and going prone has led to real-time DIS-controllable entities in the TTES (Navy) system.

2.4 The Near Future

Our research has been highlighted in three separate internal Department of Defense proposals (Army, Navy,

Air Force) to the Defense Modeling and Simulation Office (DMSO). Each proposal was concerned with different but compatible aspects of the creation of a real-time "dismounted soldier" for virtual, distributed, interactive environments. We will be turning the *Jack* software into the common human model for future interactive simulations, especially for training, logistics, and maintenance support activities.

Besides *Jack*'s immediate applicability to directing virtual agents in virtual environments, we can begin to create animated simulations from information-efficient task specifications. Our research will make significant contributions to theoretical and computational linguistics by studying

- How *instruments* are specified in Natural Language instructions and how agents recognize from these instructions and/or the world, ways in which instruments will support their intended activity;
- *Policy specification* and understanding, aimed at developing a framework for planning and sensing that enables policy to inform and/or constrain behavioral choice; and
- How instructions specify task-appropriate *sensing behavior*—what sensing modalities the agent should employ, what is being sensed for, and how often to perform sensing, since successful perception may require other physical activities to be performed as well.

On a recently funded HPCC project with the National Library of Medicine, we are working toward physiological models of human internal organs of the chest and abdomen by endowing deformable geometric models with biodynamic and functionally faithful properties. These models (designed by Dimitri Metaxas), in combination with an extended TraumAID system for penetrating trauma diagnosis and treatment planning (Clarke, Rymon, Webber, Hayward, Santora, Wagner, & Ruffin, 1993), will support surgical training and experimentation in virtual environments.

Finally, we are linking all of the above in an ARPA project (including Sandia National Labs, the Naval Postgraduate School, and the Medical College of Pennsylvania).

nia) called *MediSim*. This will be a system for training emergency medical technicians or medical corpsmen in initial casualty assessment, management, stabilization, and transport. It will utilize *Jack* as the simulated casualty—with appropriate injuries and physiological behavior responses (even including facial expressions of pain)—and as both a medical assistant and the trainee's virtual self or *avatar*. The intention is to use the trainee's instructions to trigger actions in the medical assistant; those actions will cause changes in the patient's state, and hopefully improve it. Not only will this system require and challenge *Jack*'s real-time performance capabilities, but it also should be executable in a semi-autonomous mode during DIS activities. The hope is that the semi-automated medical corpsman will respond to Global Positioning System and Personnel Status Monitor information on the simulated battlefield and thereby participate in the [virtual] reality of the battle by treating [virtual] casualties. The dual-use capacity to train civilian EMTs for accidents or natural disasters is also anticipated.

3 Research Directions of the Center

Current research in the Center can be viewed in several broad categories.

- **Simulated Humans in Three-Dimensional Virtual Environments:** The *Jack* software simulates the movements of one or more human figures: determining their mobility, seeing their viewable and reachable spaces, checking for collisions with other personnel or the environment, scaling the bodies to specific population statistics, modeling strength capabilities, and animating actual tasks to be carried out in that environment. A VR interface allows a virtual human. Program and socket interfaces permit other systems to drive *Jack* remotely.
- **Animation Control Techniques:** A variety of methods for specifying and controlling human figure animation are studied. Techniques include interactive systems, posture positioning and interpola-

tion, constraint systems, collision avoidance, strength-guided motion, and empirical motor-control models.

- **Locomotion and Navigation:** Locomotion and navigation through cluttered environments are crucial skills for an autonomous human agent. These capabilities are provided by kinematic and inverse dynamic locomotion models, forward dynamics and control, behavioral simulation, motion learning, and reactive planning.
- **Automatic Animation:** We study task description methodologies, task performance measurement, artificial intelligence models, and natural language directives. The language-related effort includes motion verb representations, semantics of verbs and adverbial modifiers, reactive planning, and formulation of executable animation primitives. Action simulation requires agent capability, resource models, tool use understanding, representation of function, knowledge-base partitioning, and object-specific reasoning. Gestural motions, facial expression, lip movement and coarticulation, head motion, eye movements, speech, and speech intonation are generated from a dialogue planner.
- **Physics-Based Modeling and Analysis:** This effort deals with the use of physics-based techniques for modeling, animation, and control of dynamic models. Techniques include modeling of complex dynamic objects through blending of parametric and spline or finite element models, systematic techniques for converting geometric degrees of freedom to physical degrees of freedom, dynamic animations through the Lagrange equations of motion, deformable model simulations, modeling of elastic or viscoelastic materials, finite element techniques, computationally efficient dynamic constraint algorithms, collision detection and impact force computation algorithms, fluid simulations, and control of dynamic models through the use of control theory algorithms.

- **Biomedical Modeling and Applications:** We use physics-based modeling for visualizations of human internal organs and their motions and deformations. A key notion in creating virtual surgical environments is the linkage of anatomical, physical, physiological, and functional representations. Simulated accident or battlefield casualties will provide virtual patients for medical technician training and evaluation.
- **Rendering Techniques for Complex Environments:** This research deals with the simulation of and user interaction with environments of geometric and illuminant complexity. We consider real-time rendering capabilities, novel natural light illumination models, and automatic and adaptive techniques for progressive refinement rendering.

The following sections describe individual research projects currently underway in the Center.

4 Simulated Humans in Three-Dimensional Virtual Environments

4.1 A Quick Introduction to *Jack*

Our *Jack* software supports the importation and graphic manipulation of externally-designed CAD models as a virtual world into which one or more anthropometrically valid human figures may be placed and animated. *Jack* supports the definition, positioning, animation, and human factors performance analysis of simulated human figures on Silicon Graphics workstations. Enhanced interactive control is provided by a set of natural behaviors such as multiple constraints, looking, reaching, balancing, collision avoiding, lifting, stepping, walking, grasping, and so on. Besides interactive specification through mouse, menu, and keyboard, *Jack* allows external control through direct sensing of the user's body position. *Jack* may also be controlled through command files, network sockets, or a simulation system that provides a convenient temporal and spatial parallel "programming language" for behaviors.

4.2 Integrating Virtual Humans in VR Applications: Michael Hollick and John Granieri

Michael Hollick and John Granieri have been working on several projects that involve the use of *Jack* as a part of large VR systems. Typically a nongraphic version of *Jack* is configured to communicate with the rest of the system via sockets or shared memory. Information about the human figures in the simulation is sent to the *Jack* process, which is responsible for realistic animation of the figures. Posture information is sent back to the rest of the system with each simulation frame.

One example of this use was demonstrated at INCOMMS '94 at Ft. Benning, Georgia. *Jack* was used as a part of the Army's Individual Soldier Mobility Simulator, along with software developed at the Naval Postgraduate School, University of Utah Center for Engineering Design, and Sarcos, Inc. In this system, a soldier was outfitted with a body suit to sense upper body joint angles, a head-mounted display (or projector screens), and a rifle equipped with a Polhemus tracker, and was seated on a pedal-based device that allowed him to "walk," feel the slope of the terrain, and turn in a convincing manner. *Jack* was used for three purposes:

- Upper body positioning: information from the sensor suit was processed and applied to the human figure.
- Posture changes: *Jack* generated animation sequences when the figure had to assume a different posture (standing, kneeling, prone, or dead).
- Locomotion: given the velocity and heading of the soldier, *Jack* generated realistic walking motions.

After the human figure was updated, each joint angle value was exported to the rendering systems.

A similar application is currently in progress with Sandia National Laboratories. One major difference is that upper body angles will be computed by *Jack*'s real-time inverse kinematics routines based on sensor data. Additionally, they are planning on experimenting with driving the locomotion from stationary walking. Semi-autonomous humans will also be included in the simula-

tion. These agents will react to the environment and the human participants, and will play several roles in different scenarios, such as a victim pinned under a tree, or a sleeping security guard.

Another application of *Jack* in VR is the TTES system being designed and built by the Navy (NAWCTSD). This system projects a soldier into a virtual environment, where he may engage hostile forces. The soldier stands in front of a large projection screen, which is his view into the environment. He has sensors on his head and gun. He locomotes through the environment by stepping on a resistive pad. The soldier may move his head, and the view frustum is updated accordingly. Essentially, both the hostiles and the soldier can move around the environment and engage each other. The hostiles are controlled via a DIS stream of commands coming from a Computer-Generated Forces (CGF) simulator written by the Institute for Simulation and Training at the University of Florida (Orlando). TTES filters and translates the DIS stream into a set of posture tokens that are passed to *Jack*. *Jack* then animates the human figures transitioning from one posture to another. *Jack* passes the joint angles back to TTES for animating an SGI Performer run-time articulated database of human geometry (created by Jonathan Crabtree). The connection is made through a double-buffered shared memory area. This allows typical updates of 6 human figures at 30 fps, with each human having 73 joints. The animation of the human is done via prerecording posture transitions at 30 fps, storing the joint angles, then playing back in real-time by selecting the appropriate frame from the recording, based on the real or wall-clock time of the simulation. TTES controls the global position of each human figure, using DIS dead-reckoning algorithms and information from the DIS stream. It also creates the necessary DIS Entity State Protocol Data Units (PDUs) to represent the real soldier, and sends them out to all other nodes on the network (i.e., to other TTES stations and the CGF system). It also performs the ballistics computation for firing the gun into the scene, and determining where the human figures get hit.

4.3 Controlling a Virtual Human Using Minimal Sensors: Michael Hollick and John Granieri

Jack's detailed human model can be driven by external sensors in real time (Badler, Hollick, & Granieri, 1993a). Using as few as 4 Ascension Technology Flock of Birds[®] 6-DOF sensors we are able to create a good approximation of a human operator's posture. By using both specific computational models of body sections (e.g., the spine) and a general inverse kinematics algorithm we can achieve a high degree of realism while minimally encumbering the operator. This allows more natural movement with fewer distractions and encumbrances than existing systems.

4.4 Jack LISP-API and PaT-Nets: Welton Becket

The *Jack* Lisp Application Programming Interface (Lisp-API) provides a consistent, well-documented, protected interface to *Jack* functionality for developers and users of *Jack* (Becket, 1994b). *Jack*-5.8 contains a version of the XLISP-STAT Common Lisp subset written by Luke Tierney of the University of Minnesota School of Statistics and David M. Betz. XLISP-STAT is time-sliced into the *Jack* process and allows general programming of *Jack* through approximately 150 access functions. Along with providing fast read/write access to *Jack* data such as figure positions and joint angles, the Lisp-API allows collecting data about joint torques, available strength, and other derived information in *Jack*. The Lisp-API is also available through Unix sockets, so a separate process can control *Jack* as well.

Parallel Transition Networks (PaT-Nets) define a language written on top of the Lisp-API allowing general sequencing of *Jack* commands with parallel, interacting automata. These nets run concurrently with *Jack* (in XLISP-STAT) and can schedule *Jack* events based on arbitrary conditions in *Jack*. PaT-Nets are currently being used as part of the gesture and facial animation (Casell, Pelachaud, Badler, Steedman, Achorn, Becket, Douville, Prevost, & Stone, 1994) project and the behavioral

walking and terrain reasoning projects (Ko, Reich, Becket, & Badler, 1994; Reich, Ko, Becket, & Badler, 1994).

4.5 Spreadsheet Anthropometric Scaling System (SASS) and X.SASS:

Francisco Azuola, Ann Song, and Susanna Wei

The Spreadsheet Anthropometric Scaling System, SASS, allows flexible interactive access to all anthropometric variables needed to define a computer-based human figure (NASA, 1978; Natick Labs, 1988; Grosso, Quach, & Badler, 1989; Grosso, Quach, Otani, Zhao, Wei, Ho, Lu, & Badler, 1987). SASS works as a relational spreadsheet supporting both body segment, joint, and attribute information. Segment dimensions and center of mass, joint degrees of freedom and movement limits, and joint strength data are supported. It connects related body dimensions and parameters, such as relevant segment lengths to stature, or segment mass to scale. Besides statistical data, SASS provides access to an anthropometric database containing data from "real-world" individuals.

SASS supplies the dimensions of each segment needed to define a virtual human figure, based upon population data supplied as input. The human model generated by SASS consists of 68 segments (body structures), of which 67 have a geometric representation. For each of these segments, there are three dimensions required, namely, length, width, and thickness. The articulation of the human model consists of 66 joints, (1 to 3 degrees of freedom per joint) for a total of 138 degrees of freedom. The program supplies the corresponding joint limits. Hands are included.

Francisco Azuola will be enhancing SASS to support the specification of body sizes via direct anthropometric measurements (sam's), as well as joint center to center segment x - y - z measurements (Azuola, Badler, Ho, Kakadiaris, Metaxas, & Ting, 1994). Supporting a set of direct body measurements will give SASS an intuitive and natural user interface.

Ann Song and Susanna Wei are building X.SASS:

SASS running under the X-windows programming environment (Ousterhout, 1993). Picture buttons and pull-down menus enhance the user interface. Multiple figures may also be simultaneously changed and rescaled.

4.6 Human Figure Modeling: Pei-Hwa Ho

Pei-Hwa Ho's work includes the building and modeling of computer graphics human figures that better reflect physiological attributes, based on anthropometric parameters. Specifically, his research looks at methodologies of measuring, scaling, and shaping existing human models in order to create models of different physiological and anthropometric attributes while maintaining visual realism. It is not too difficult to create an individual human model with an excellent physical appearance (they can be purchased from a number of commercial sources, such as Viewpoint DataLabs); the greater challenge lies in keeping them reasonable under anthropometric scaling such as accommodated in SASS.

4.7 Free Form Deformation: Bond-Jay Ting

Bond-Jay Ting's work involves modeling the human body by using free form deformations (Azuola et al., 1994). Free form deformations were originally used as a tool to model complex objects by deforming simpler primitives (Sederberg & Parry, 1986). They are also used to animate such objects. In our applications, the human torso is treated as one deformable segment. The deformation for the human spine is controlled by an attached free form deformation lattice. Changes to the lattice result in realistic deformations of the torso and internal organs.

5 Animation Control Techniques

5.1 Posture Interpolation: Rama Bindiganavale and Susanna Wei

Postures form a very important aspect of human figure simulation. Static postures such as sit, stand, su-

pine, or prone can be defined by the relative positioning of various parts of the body. Simulation of human motions requires suitable animation between these static postures. Unlike the standard techniques of key-frame animations, Rama Bindiganavale and Susanna Wei use a finite state machine to control the behaviorally reasonable transitions from any posture to a goal posture by finding the shortest path of the required predetermined motion sequences between the two (Badler, Bindiganavale, Granieri, Wei, & Zhao, 1994). If the motion sequence is not collision-free, then a collision avoidance strategy is invoked and the posture is changed to one that satisfies the required goal while respecting object and agent integrity. Such a posture transition network (formed from the Finite State Machine) has been successfully implemented for simulating the transitions between different key postures of a soldier such as Standing Stowed, Standing Fire, Kneeling Stowed, or Kneeling Fire. Bindiganavale and Wei are now working on generalizing this posture interpolation for any anthropometric figure. In conjunction with Xinmin Zhao, the posture transitions are being extended to consider strength, torque, and fatigue factors based on Philip Lee's original work (Lee, Wei, Zhao, & Badler, 1990).

5.2 Human Reach Trajectories: Hanns-Oskar Porr

Hanns-Oskar Porr's work includes a motion capture system that uses digitally sampled human reach trajectories. These data were collected by MOCO Inc. for an optimal reach space study for NASA. Four Ascension Technology Flock of Birds[®] sensors were attached to a human subject on the torso, upper arm, lower arm, and hand. The subject then reached for one of a large number of predetermined sites in his reachspace, and the movement was recorded. To re-create this motion using the sensor readings, the constraint system in *Jack* was used to force the computer-generated human figure into a posture that fits all the 6-DOF data readings.

5.3 Motion Planning: Xinmin Zhao

When a semi-autonomous agent is instructed to perform a task, e.g., lift a box and put it on a shelf, a sequence of motions has to be planned, either in advance or on-line. In general, this activity is called motion planning. To produce realistic motion sequences, a number of factors have to be considered.

1. Collision: motion planning has to ensure that the resulting motion sequence is free of collisions.
2. Strength: for a given load and the agent's strength, the motion sequence should be physically possible for the agent to perform.
3. Fatigue: the state of muscle fatigue will affect how the agent is performing a particular task.
4. Energy: we may wish to find the trajectory that minimizes the total energy expenditure of the agent.
5. Time: the amount of time allotted to the agent to perform the task will obviously affect the resulting motion sequence.

These factors are not independent. For example, the agent is able to exert maximum force/torque if its workspace is unrestricted, and it is able to exert less force/torque if its torso has to bend because of low clearance in the workspace. Xinmin Zhao's current work involves the study of the effects of these factors, especially the first two, on the agent's motion (Zhao & Badler, 1994).

5.4 Human Torque and Strength Analysis: Hyeongseok Ko

Hyeongseok Ko has developed a dynamic torque and strength analysis system. The Newton-Euler method is applied to a 97 degree-of-freedom human model to compute joint forces and torques in real-time. The torque is compared with recent NASA strength data (Pandya, Maida, Aldridge, Hasson, & Woolford, 1991) to check feasibility and motion comfort. Several visualization techniques are applied to validate and display the result.

Ko also developed an efficient real-time algorithm for animating human locomotion using inverse dynamics, balance, and comfort control. Balance is maintained by rotating or translating the pelvis and torso. Comfort control ensures that any excessive joint torque is distributed to other joints. The resulting motion is not only visually realistic but also dynamically sound. The effect of attached loads or external forces is simulated quite accurately.

6 Locomotion and Navigation

6.1 Human Locomotion: Hyeongseok Ko

Automatic generation of realistic walking animation, from the current stance to a goal, was developed by Hyeongseok Ko. The locomotion includes terrain navigation skills such as rhythmic curved path walking, non-rhythmic lateral or backward stepping, running, and transitions between walking and running for motion continuity (Ko & Badler, 1993a, 1993b). Running and the transitions to and from walking were developed with Byong Mok Oh. Locomotion attributes control the movement of the torso, pelvis, and swing leg, etc. They can be rhythmic to produce a visually realistic effect, or nonrhythmic to generate an intentional deviation from a normal gait, such as ducking under or stepping over obstacles. Development of other locomotion skills is in progress to navigate uneven surfaces, stairs, snow, ice, sand, mud, grass, water, swamp, etc. Virtual reality applications generate many demands for human locomotion. Our locomotion system has been used in developing an Individual Soldier Mobility Simulator—a large-scale virtual reality project for the U.S. Army to create a virtual combat environment to train U.S. infantry soldiers. It generated walking or running motion according to the locomotion speed and heading direction supplied from the virtual reality input device.

6.2 Animation and Control of Four Legged Animals: Evangelos Kokkevis

Evangelos Kokkevis' research currently focuses on

a new approach to generating natural looking motion of four legged animals. The technique used is based on the blending of established kinematic procedures with model-based dynamic control algorithms to achieve realistic animation. During the locomotion cycle, the body is supported by one or more legs that are pushing against the ground. Kinematics are used for the motion of the legs that are not in contact with the ground. The dynamic controller, given the required velocity and orientation of the animal, computes the forces and torques that should be applied by the rest of the legs in order to accomplish the motion. There are two available gait patterns, walking and trotting. A gait controller enforces the use of one of them depending on the current speed and is responsible for a smooth transition between the two when required. The advantages of his approach are real-time performance and automatic adaptability to environment conditions, such as locomotion in uneven terrain. The method has been tested with an anatomically accurate model of a dog based on data from Viewpoint DataLabs International.

6.3 Sensor-Based Navigation: Barry D. Reich

For us, sensor-based navigation means real-time terrain traversal by simulated human agents (Ko et al., 1994; Reich et al., 1994). Agents begin in arbitrary positions in an outdoor environment. During a simulation they walk to the goal, avoiding obstacles and each other, ducking under low branches, climbing over objects, avoiding difficult terrain where feasible, and squeezing through tight spaces where necessary. In addition, the agents attempt to evade detection by hostile agents by avoiding the sensory fields of any hostiles in the area.

An agent is not told how to reach its goal nor is off-line path-planning used. Instead, it is made aware of its environment through the use of a network of simulated sensors. These sensors acquire information on object size, type, and location, passageways, terrain type, exposure to hostile agents, and so on. Based on this information the path through the terrain is incrementally com-

puted, one step at a time. This allows the agent to react to moving obstacles, changing terrain, or unexpected events due to hostile agents or the effects of limited perception.

6.4 Behavioral Walking: Welton Becket

Behavioral walking uses behavioral or “reactive planning” techniques to navigate a simulated human or group of humans through a scene in order to get to a particular location without collisions (Becket & Badler, 1993). The behavioral walking relies on tight couplings of simulated sensors to simulated effectors rather than on static plans based on a static environment. As a result, the behavioral walking works in real time and works effectively with moving obstacles (such as other agents), unexpected obstacles, or unexpected opportunities.

The behavioral walking is controlled through a network of attraction and avoidance behaviors connected to simulated object proximity sensors or simulated sonar arrays. The attract and avoid behaviors collectively determine where each agent will step next—decisions about where to step next are made only at the beginning of each step.

Because the behavioral description is posed as a network of modules with floating point communication, low-level learning and numerical optimization techniques can be applied to aid the user in developing an adequate set of parameters for the behaviors. Becket (1994a) uses gradient search and proposes to use a genetic optimization routine for finding optimal weights given a user-defined fitness function. This use of optimization addresses one of the primary problems with designing behavioral or reactive systems—arbitration among competing behaviors to decide which gets control of the effectors rapidly becomes too complex for the agent designer. Also addressed in Becket (1994a) is the sequencing of behavioral weights given to optimize total reward using a variety of unsupervised reinforcement learning techniques such as Q-learning and classifier systems.

7 Automatic Animation

7.1 Gestures and Facial Animation: Justine Cassell, Catherine Pelachaud, Matthew Stone, Brett Achorn, Scott Prevost, and Brett Douville

This project is led by Justine Cassell, visiting us under the aegis of the National Science Foundation Visiting Professorship for Women Scholars. Realistic animation of human-like characters should be based on cognitive behavioral principles to maximize expressive realism. On the basis of the principles underlying movements performed during conversations between two people, this group has designed and implemented a system to *automatically* animate conversations between two human-like agents with appropriate and synchronized speech, intonation, facial expressions, and hand gestures (Cassell et al., 1994). The conversation is created by a dialogue planner that produces the text as well as the intonation of the utterances. The speaker/listener relationship, the text, and the intonation in turn drive facial expressions, lip motions, eye gaze, head motion, and arm gesture generators. The facial motion models are driven and synchronized by parallel transition networks. Coordinated arm, wrist, and hand motions are invoked to create semantically meaningful gestures.

7.2 Object Specific Reasoning: Libby Levison

Jack can animate a human figure through commands to reach for things, grasp them, pick them up, or walk along a path. Now suppose that one wanted to generate realistic simulations of the character performing different tasks. A task-level interface to this agent, which allows instructions to the agent, might have commands such as (**pickup jack glass**) or (**open betty door**). These high-level descriptions of action must somehow be mapped to the agent's control language: the command (**pickup jack glass**) must be converted to fully parameterized animation commands that specify details such as where on the glass to grasp, how high to lift the glass, and where to stand in relation to the glass.

Libby Levison is designing an intermediate planning

module, called the Object Specific Reasoner, to perform this conversion (Geib, Levison, & Moore, 1994; Levison & Badler, 1994). The mapping is done by considering action commands in terms of each agent's resources, the object attributes, the purpose governing the action, and the situation. Controlling animated figures requires specifying details of the desired motions that are rarely present in high-level commands. The Object Specific Reasoner will map a small set of high-level commands to the language of the animation system, providing the animation system with a task-level interface.

8 Physics-Based Modeling and Analysis

8.1 Active Part-Decomposition, Shape and Motion Estimation of Articulated Objects: A Physics-Based Approach: Ioannis Kakadiaris

Ioannis Kakadiaris has developed a novel, robust, integrated physics-based technique to reliably identify an articulated object's parts, joint locations, shape, and motion parameters (Kakadiaris, Metaxas, & Bajcsy, 1994; Azuola et al., 1994). To overcome the indeterminacy of part identification from a single image, a sequence of images of a moving articulated object is used instead. Initially, the program assumes that the object consists of a single part, and fits a deformable model to the given data using our physics-based framework (Metaxas & Terzopoulos, 1992, 1993). As the object attains new postures, decision criteria determine if and when to replace the initial model with two new models. These criteria are based on the model's state and the given data. The model fitting process uses a novel algorithm for assigning forces from the data to the two models, which allows partial overlap between them and determination of joint location. This algorithm is based on the theory of fuzzy clustering. The procedure is applied iteratively until all the object's moving parts are identified. Kalman filtering is employed to overcome noise and occlusion. We have successfully applied this technique to monocular image sequences of a moving robot arm, a human arm, and a human finger. Our algorithm has the following advantages, in that it

- couples the processes of segmentation, shape, and motion estimation,
- allows the reliable shape description of the real parts of an articulated object,
- estimates the true location of the joints between the parts,
- detects multiple joints, and
- obviates the need for markers or special equipment.

We are currently extending our approach by combining it with Jianmin Zhao's thesis work on articulated motion reconstruction (Zhao, 1993). The goal is to achieve complete three-dimensional shape and motion estimation of articulated figures based on multiple image sequences from a single camera.

8.2 Modeling Fluids: Nick Foster

Nick Foster is using direct numerical simulation of the Navier–Stokes equations to animate liquids with free surfaces. Since the model is based on a set of physically accurate equations, waves, splashing drops, and colliding surfaces behave realistically. Thus, Foster can render and animate a variety of fluid phenomena in an arbitrary 3D environment by simply defining a set of initial conditions and allowing the system to evolve. This is the start of a series of numerical physics-based models for fluids and possibly explosions.

8.3 Blended Shapes: Douglas DeCarlo

Doug DeCarlo has developed a new class of parameterized models based on the linear interpolation of two parameterized shapes using a blending function. This blending function describes how the shapes are combined. Using a small number of additional parameters, blending extends the coverage of shape primitives. For example, a bullet shape could be defined by a blend of a sphere and a cylinder. Blends between shapes such as a sphere and torus are also possible, even though they differ topologically. The result would be a shape that has a hole that can appear depending on the blending parameters.

These shapes are incorporated into a physics-based

Lagrangian dynamics framework that uses globally and locally deformable models (Metaxas & Terzopoulos, 1992). We can also perform shape morphing, by changing the parameters of an object over time. Since blended shapes have good shape coverage, it is possible to morph between a variety of interesting shapes, including shapes of differing genus. This morphing is tied in with the dynamics, so that forces can arise due to the morphing. These models can also be used in a physics-based shape estimation framework.

9 Biomedical Modeling and Applications

9.1 MediSim

We are just starting on MediSim: a system that extends virtual environments to represent simulated medical personnel interacting with simulated casualties. In conjunction with Sandia National Labs and the Naval Postgraduate School, this project will develop VE technology for planning, training, and evaluation of both medical corpsmen and civilian EMTs. Behaviors and behavioral control will be developed for the medical corpsmen that will enable their actions on the digital battlefield to conform to both military practice and medical protocols. Explicit representation of emergency care protocols will permit testing in either military or civilian conditions. From situationally appropriate injury models, a set of physical and behavioral manifestations in a simulated casualty will be determined and portrayed on a three-dimensional body. Through I-Port (individual virtual environment) interfaces and voice commands, these simulated casualties will be located, observed, assessed, triaged, stabilized, and either treated locally or evacuated for definitive management. Individual corpsmen may be trained and evaluated in realistic simulated battlefields or other emergency care environments populated by synthetic personnel—friendlies, hostiles, casualties—as well as other simulated entities, buildings, vehicles, and terrain. New medical and trauma care technologies (e.g., trauma pods) and doctrine (e.g., evacuation procedures) may be evaluated prior to deployment.

9.2 TrauMAP—Trauma Modeling of Anatomy and Physiology: Jonathan Kaye

VR systems require both realistic geometric and behavioral modeling. This means that objects in a virtual environment must look real, both in their visual presentation and their behavior. Systems that aim to model realistic human anatomy often focus on geometric aspects of the problem, such as accurately modeling shape, with lesser emphasis, if any, on the functional relationships among anatomical parts. Physiological simulations (in the form of differential equations), on the other hand, focus on functional relationships, with less consideration for the physical presence of objects and the physical space in which the processes take place.

TrauMAP models the immediate consequences of penetrating trauma (gunshot and stab wounds). The system ties together the spatial (geometric) properties of an anatomical object with the object's role as part of the appropriate physiological system. By associating the physical presence of a substance or object with its functional role, TrauMAP can infer simple effects of a structural change on physiology, and vice versa. This can be useful for detecting situations where independent physiological systems become dependent because they are physically adjacent to one another. Kaye represents anatomical parts accurately using a method that combines superquadric and finite element modeling (Metaxas & Terzopoulos, 1992). He models the significant physiological systems with qualitative models of behavior.

There are five main thrusts to this project:

1. Scaling deformable models of various major organs in the chest or abdominal cavity, where the scaling represents normal, diseased, or hypertrophied situations.
2. Allowing the scaled organs to move, interact and rearrange themselves into anatomically and physiologically realistic positions.
3. Mapping Visible Human datasets (or currently available imaging datasets) onto the deformation models to impart fine-grained structure and surgically relevant anatomy, such as major arteries, nerves, and support tissues.
4. Building a qualitative model of macroscopic organ

function correlated with spatial properties that represents, *inter alia*, notions of container, conduit, obstruction, leakage, distention, collapse, etc.

5. Using the qualitative functional model and the physical/spatial model to model candidate surgical procedures appropriate to training and experimentation in virtual simulated environments.

9.3 Organ Deformations: Jinah Park

Physics-based deformation models are applied to *in vivo* imaged organ data to reconstruct both shape and motion. Specific instances where such models can be applied are crash simulations, virtual surgery, blunt or penetrating trauma assessment, shape and motion recovery from CT, MRI, and/or SPAMM data, and recovery of human motion from video or range data. The SPAMM technique imprints a magnetic grid into tissue so that the motion of the grid intercepts can be detected and tracked even in the absence of observable anatomical features. Such information provides a means for assessing an athlete's fitness, for modeling and describing heart motion for disease diagnosis, and for estimating and simulating blood flow. Jinah Park has been able to model and quantify the shape and motion of the left ventricle using deformable models with a small number of "intuitive" parameters such as elongation and twisting. Based on the Lagrange equations of motion, these deformable meshes are fit to sequences of time-varying SPAMM data (Park, Metaxas, & Young, 1994).

10 Rendering Techniques for Complex Environments

10.1 Real-Time Rendering: Paul Diefenbach

Paul Diefenbach's area of research is in fast animation and rendering techniques. Past work includes using existing hardware pipelines to produce advanced rendering features such as refractions, reflections, shadows, and caustics. His current work involves automatic animation generation from storyboard sequences. This includes building an intelligent system for controlling character

animation as well as camera motion from a level higher than the typical keyframe level. These techniques could be used to greatly reduce traditional animation time from concept to product.

10.2 Photorealistic Animations: Jeff Nimeroff

Jeff Nimeroff's current work with Eero Simoncelli and Julie Dorsey centers around constructing photorealistic animations of naturally illuminated scenes (Nishita & Nakamae, 1986; CIE Technical Committee 4.2, 1973; IES Daylighting Committee, 1979) without the requirement of using a costly renderer on each frame (Nimeroff, Simoncelli, & Dorsey, 1994; Nimeroff, Simoncelli, Dorsey, & Badler, 1995).

More specifically Nimeroff is using steerable functions (Freeman & Adelson, 1991) to represent natural illumination for use in a "photorealistic animation via image interpolation" pipeline. The theory shows the conditions for generating photorealistic basis images and under what constraints these images can be interpolated to generate the scene under any possible natural lighting condition.

10.3 Radiosity: Min-Zhi Shao

Min-Zhi Shao's area of research is realistic image synthesis in computer graphics, particularly the radiosity method. Recently, he has developed a gathering and shooting overrelaxation radiosity method (Shao & Badler, 1993). In his algorithm, the conventional progressive refinement method is modified to make optimal use of the visibility information computed in each iteration: the light energy is first gathered from all visible patches and then immediately reshot to the environment together with the previously accumulated unshot light energy. Based on a concise record of the history of the unshot light energy distribution in the environment, a solid convergence speedup is achieved with little additional computation and memory usage. His future research is to develop and implement more efficient algorithms and techniques for rendering dynamic and complex environments.

Appendix A: Worldwide Active Jack Sites

- Anthropology Research Project, Inc., Yellow Springs, OH
- Battelle, Dayton, OH
- British Aerospace Dynamics, United Kingdom
- Caterpillar, Inc., Peoria, IL
- Children's Design Center, Paia, HI
- Computer Sciences Corp., San Diego, CA
- Deere & Co., Moline, IL
- FMC Corp., Santa Clara, CA
- GE Engines, Cincinnati, OH
- General Dynamics Land Systems, Sterling Heights, MI
- General Motors Corp., Warren, MI
- Goodwin Marcus Systems Ltd., United Kingdom
- Hughes Missile Systems Co., Tucson, AZ
- Hughes Training, Inc., MN
- Iowa State University, Ames, IA
- Israeli Air Force, Israel
- Israeli Ministry of Defense, Israel
- Japan Tech Services Corp., Tokyo, Japan
- Kimberly-Clark Corp., Neenah, WI
- MIT Media Lab, Cambridge, MA
- MOCO, Inc., Scituate, MA
- Martin Marietta, Moorestown, NJ
- McDonnell-Douglas Space Systems Co., Titusville, FL
- Micro Analysis and Design, Inc., Boulder, CO
- Ministry of Defense, United Kingdom
- NASA Ames Research Center, Moffett Field, CA
- NASA Marshall Space Flight Center, Huntsville, AL
- NASA Johnson Space Center, Houston, TX
- Naval Air Warfare Center, Training Systems Division, Orlando, FL
- NIOSH, Morgantown, WV
- NIOSH, Finland
- National Institute of Standards and Technology, Gaithersburg, MD
- Naval Surface Warfare Center, Bethesda, MD
- Naval Surface Warfare Center, Dahlgren, VA
- Ohio State University, Columbus, OH
- Sandia National Labs, Albuquerque, NM
- Synergy Integration Ltd., Israel
- Texas A&M University, College Station, TX
- Texas Woman's University, Houston, TX
- Tinker Air Force Base, Oklahoma City, OK
- Traces, Inc., Bala Cynwyd, PA
- U.S. Army Corps of Engineers, Vicksburg, MI
- U.S. Army Infantry School, Fort Benning, GA
- U.S. Army Natick RD & E Center, Natick, MA
- U.S. Army Night Vision & Electro Optics, Ft Belvoir, VA
- U.S. Army Research Institute, Alexandria, VA
- U.S. Army Research Laboratory, Aberdeen Proving Grounds, MD
- U.S. Army TACOM, Warren, MI
- U.S. Army Topographic Engineering Center, Alexandria, VA
- United Technologies Research Center, East Hartford, CT
- University of Central Florida, Orlando, FL
- University of Dayton Research Institute, Dayton, OH
- University of Delaware, Wilmington, DE
- University of Iowa, Iowa City, IA
- University of the Arts, Philadelphia, PA
- Vickers Defense Systems, United Kingdom
- Wright-Patterson Air Force Base, Dayton, OH

Acknowledgments

This research is or has been recently supported in part by ARO DAAL03-89-C-0031 including the ARO AI Center of Excellence at the University of Pennsylvania, Army Research Laboratory (Aberdeen), Army Natick Labs, NASA Ames Research Center, and Army Research Institute; Air Force DEPTH through Hughes Missile Systems F33615-91-C-0001; ARPA DAMD17-94-J-4486; Naval Air Warfare Center, Training Systems Division N61339-94-C-0005; Sandia Labs AG-6076; DMSO through the University of Iowa and DAAH04-94-G-0402; NASA KSC NAG10-0122; MOCO, Inc.; Robotics Research Harvesting, Inc.; National Library of Medicine N01LM-43551; NSF IRI91-17110, IRI93-09917, CISE CDA88-22719, VPW GER93-50179, graduate fellowships, and Instrumentation and Laboratory Improvement Program USE-9152503; National Defense Science and Engineering

Graduate Fellowship in Computer Science DAAL03-92-G-0342; Ascension Technologies, Inc.; Silicon Graphics, Inc.; and the Children's Hospital of Philadelphia.

References

- Azuola, F., Badler, N. I., Ho, P.-H., Kakadiaris, I., Metaxas, D., & Ting, B. (1994). Building anthropometry-based virtual human models. *Proc. IMAGE VII Conf.*, Tucson, AZ.
- Badler, N. I., Bindiganavale, R., Granieri, J. P., Wei, S., & Zhao, X. (1994). Posture interpolation with collision avoidance. *Proceedings of Computer Animation '94*, Geneva, Switzerland. New York: IEEE Computer Society Press.
- Badler, N. I., Hollick, M. J., & Granieri, J. (1993a). Real-time control of a virtual human using minimal sensors. *Presence* 2(1), 82-86.
- Badler, N. I., Phillips, C. B., & Webber, B. L. (1993b). *Simulating humans: Computer graphics, animation, and control*. New York: Oxford University Press.
- Becket, W. (1994a). *Discrete and continuous learning for behavioral control of simulated autonomous agents*. Ph.D. thesis, Dept. of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.
- Becket, W. (1994b). *The Jack Lisp API Version 1.1*. Technical Report MS-CIS-94-01, Dept. of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.
- Becket, W., & Badler, N. I. (1993). Integrated behavioral agent architecture. *The Third Conference on Computer Generated Forces and Behavior Representation*, Orlando, Florida, 57-68.
- Cassell, J., Pelachaud, C., Badler, N., Steedman, M., Achorn, B., Becket, W., Douville, B., Prevost, S., & Stone, M. (1994). Animated conversation: Rule-based generation of facial expression, gesture, and spoken intonation for multiple conversational agents. *Computer Graphics (Proc. SIGGRAPH)*, 413-420.
- CIE Technical Committee 4.2. (1973). Standardization of luminance distribution on clear skies. CIE Publication No. 22, Commission International de L'Eclairage, Paris, 1973.
- Clarke, J. R., Rymon, R., Webber, B., Hayward, C., Santora, T., Wagner, D., & Ruffin, A. (1993). The importance of planning in the provision of medical care. *Medical Decision Making*, 13(4), 383 (abstract).
- Freeman, W. T., & Adelson, E. H. (1991). The design and use of steerable filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(9), 891-906.
- Geib, C., Levison, L., & Moore, M. B. (1994). *Sodajack: An architecture for agents that search for and manipulate objects*. Technical Report MS-CIS-94-16, Dept. of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.
- Grosso, M., Quach, R., & Badler, N. I. (1989). Anthropometry for computer animated human figures. In N. Magnenat-Thalmann & D. Thalmann (Eds.), *State-of-the art in computer animation* (pp. 83-96). New York: Springer-Verlag.
- Grosso, M. R., Quach, R., Otani, E., Zhao, J., Wei, S., Ho, P.-H., Lu, J., & Badler, N. I. (1987). *Anthropometry for computer graphics human figures*. Technical Report MS-CIS-87-71, Dept. of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.
- IES Daylighting Committee. (1979). Recommended practice of daylighting. *Lighting Design and Application*, 9(2), 45-58.
- Kakadiaris, I. A., Metaxas, D., & Bajcsy, R. (1994). Active part-decomposition, shape and motion estimation of articulated objects: A physics-based approach. *IEEE Computer Vision and Pattern Recognition Conference*, Seattle, Washington.
- Ko, H., & Badler, N. I. (1993a). *Curved path human locomotion that handles anthropometrical variety*. Technical Report MS-CIS-93-13, Dept. of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.
- Ko, H., & Badler, N. I. (1993b). Intermittent non-rhythmic human stepping and locomotion. In *Proceedings of the First Pacific Conference on Computer Graphics and Applications, Pacific Graphics '93*, (pp. 283-300), Seoul. World Scientific.
- Ko, H., Reich, B. D., Becket, W., & Badler, N. I. (1994). Terrain navigation skills and reasoning. In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representations*, Orlando, Florida.
- Lee, P., Wei, S., Zhao, J., & Badler, N. I. (1990). Strength guided motion. *Computer Graphics*, 24(4), 253-262.
- Levison, L., & Badler, N. I. (1994). How animated agents perform tasks: Connecting planning and manipulation through object-specific reasoning. In *Working Notes from the Workshop on: Towards Physical Interaction and Manipulation*. AAAI Spring Symposium.
- Metaxas, D., & Terzopoulos, D. (1992). Dynamic deformation of solid primitives with constraints. *Computer Graphics*, Vol. 26(2), 309-312.
- Metaxas, D., & Terzopoulos, D. (1993). Shape and nonrigid motion estimation through physics-based synthesis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-15(6), 580-591.
- NASA. (1978). *The anthropometry source book*. Reference Publi-

- cation 1024, NASA Scientific and Technical Information Office.
- Natick Labs. (1988). *1988 Anthropometric survey of US Army personnel—Methods and summary statistics*. Natick Technical Report TR-89-044.
- Nimeroff, J. S., Simoncelli, E., & Dorsey, J. (1994). Fast re-rendering of naturally illuminated environments. *5th Annual Eurographics Symposium on Rendering*; Darmstadt Germany, June 1994, pp. 359–374.
- Nimeroff, J. S., Simoncelli, E., Dorsey, J., & Badler, N. I. (1995). Rendering spaces for architectural environments. *Presence*, in press, to appear in Presence 4.3.
- Nishita, T., & Nakamae, E. (1986). Continuous tone representation of three-dimensional objects illuminated by sky light. *Computer graphics (SIGGRAPH '86 Proceedings)*, 20(4), 125–132.
- Ousterhout, J. K. (1993). *TCL and the Tk toolkit*. New York: Addison Wesley.
- Pandya, A., Maida, J., Aldridge, A., Hasson, S., & Woolford, B. (1991). Development of an empirically based dynamic biomechanical strength model. *Space Operations Applications and Research Conference Proceedings*, 2, 438–444.
- Park, J., Metaxas, D., & Young, A. (1994). Deformable models with parameter functions: Application to heart-wall modeling. *Proceedings of the Computer Vision and Pattern Recognition Conference*, CVPR.
- Reich, B. D., Ko, H., Becket, W., & Badler, N. I. (1994). Terrain reasoning for human locomotion. *Proceedings of Computer Animation '94*, Geneva, Switzerland. IEEE Computer Society Press.
- Sederberg, T. W., & Parry, S. R. (1986). Free-form deformation of solid geometric models. *Computer Graphics*, 20(4), 151–160.
- Shao, M.-Z., & Badler, N. I. (1993). *A gathering and shooting progressive refinement radiosity method*. Technical Report MS-CIS-93-03, Dept. of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.
- Zhao, J. (1993). *Moving posture reconstruction from perspective projections of jointed figure motion*. Ph.D. thesis, Technical Report MS-CIS-93-74, Dept. of Computer and Information Science, University of Pennsylvania, Philadelphia, PA.
- Zhao, X., & Badler, N. I. (1993). Near real-time body awareness. First Bilkent Computer Graphics Conference, Ankara, Turkey. (*Computer-Aided Design*, in press.)