



Artificial Life and Biomechanical Simulation of Humans

Demetri Terzopoulos
University of California, Los Angeles
www.cs.ucla.edu/~dt

Abstract

I review our recent work on human simulation, specifically: (1) an artificial life framework that addresses the problem of emulating the rich complexity of human appearance and activity in urban environments, resulting in autonomous pedestrian models, and (2) a comprehensive biomechanical model of the human body that confronts the combined challenge of modeling and controlling more or less all of the relevant articular bones and muscles, as well as simulating the physics-based deformations of the soft tissues.

1. Introduction

The simulation of humans is a “grand challenge” that confronts us with many difficult problems spanning a breadth of topics from biomechanics to artificial intelligence. My research group has engaged in this broad scientific endeavor over the past decade. This paper presents a two-part review of our work on human simulation.

First, I will describe an artificial life framework [1] that addresses the problem of emulating the rich complexity of pedestrians in urban environments. Our approach integrates motor, perceptual, behavioral, and cognitive components within a comprehensive model of pedestrians as highly capable autonomous agents that can perform a variety of natural activities. Featuring innovations in each of these components, as well as in their combination, our agent model yields results of unprecedented fidelity and complexity for fully autonomous multi-human simulation in an extensive urban environment taking the form of a well-populated virtual train station.

Second, because of the complexity of the human body, its detailed biomechanical modeling has not received adequate attention. I will describe a comprehensive biomechanical model that confronts the combined challenge of modeling and controlling more or less all of the relevant articular bones and muscles, as well as simulating the physics-based deformations of the soft tissues. As a precursor to our comprehensive biomechanical model, I will also review our prior effort in simulating the highly important head-neck-face complex. This model includes a biologically inspired, neuromuscular controller. Using machine learning techniques, the neural networks comprising the controller are trained offline to efficiently generate the online control signals necessary to synthesize autonomous movements for the behavioral animation of the human head and face.

2. Autonomous Pedestrians

We have tackled the challenging problem of emulating the rich complexity of pedestrians in urban environments [2]. In a departure from the substantial literature on so-called “crowd simulation”, we have developed a decentralized, comprehensive model of pedestrians as autonomous *individuals* capable of a broad variety of activities in large-scale synthetic urban spaces. Analogous to real humans, our autonomous pedestrians perceive the virtual environment around them, analyze environmental situations, exhibit natural reactive behaviors, and proactively plan their activities. We represent the environment using hierarchical data structures, which efficiently support the perceptual queries that influence the behavioral responses of the autonomous pedestrians and sustain their ability to plan their actions over local and global spatiotemporal scales.

Our artificial life approach integrates motor, perceptual, behavioral, and, importantly, cognitive components, each of which I will review in the subsequent sections.

2.1. Appearance and Locomotion

As an implementation of the appearance and motor levels of the character, we employ a human animation software package called *DI-Guy*, which is commercially available from Boston Dynamics Inc. *DI-Guy* provides a variety of textured geometric human models together with a set of basic motor skills, such as strolling, walking, jogging, sitting, etc. *DI-Guy* is intended for manually scripted animation. Emulating the natural appearance and movement of humans is a highly challenging problem and, not surprisingly, *DI-Guy* suffers from several limitations, mostly in the kinematic control of human motions. To ameliorate the visual defects, we have customized the motions of *DI-Guy* characters and have implemented a motor control interface to hide the details of the underlying *DI-Guy* kinematic layer from our higher-level behavior modules, enabling the latter to be developed more or less independently.

2.2. Perception

In a highly dynamic virtual world, an autonomous intelligent character must have a keenly perceptive regard for the external world in order to interact with it effectively. The hierarchical world model is used extensively by each autonomous pedestrian to perceive its environment, providing not only the raw sensed data, but also



Fig. 1 A large-scale virtual train station populated by self-animating virtual humans.

higher-level interpretations of perceived situations that are important to a pedestrian. The details of the pedestrian perception model are given in [2].

2.3. Behavior

The purpose of realistic behavioral modeling is to link perception to appropriate actions. We adopt a bottom-up strategy [1], which uses primitive reactive behaviors as building blocks that in turn support more complex motivational behaviors, all controlled by an action selection mechanism.

At the lowest level, we developed six key reactive behavior routines that cover almost all of the obstacle avoidance situations that a pedestrian can encounter. The first two are for static obstacle avoidance, the next three are for avoiding mobile objects (mostly other pedestrians), and the last one is for avoiding both. Given that a pedestrian possesses a set of motor skills, such as standing still, moving forward, turning in several directions, speeding up and slowing down, etc., these routines are responsible for initiating, terminating, and sequencing the motor skills on a short time scale guided by sensory stimuli and internal percepts. The routines are activated in an optimized sequential manner, giving each the opportunity to alter the currently active motor control command (speed, turning angle, etc.).

While the reactive behaviors enable pedestrians to move around freely, almost always avoiding collisions, navigational behaviors enable them to go where they desire. We developed several such routines—*passageway selection*, *passageway navigation*, *perception guided navigation*, *arrival-at-a-target navigation*, etc.—to address issues, such as the speed and scale of online path planning, the realism of actual paths taken, and pedestrian flow control through and around bottlenecks. Furthermore, to make our pedestrians more capable, we have augmented their behavior repertoires with a set of non-navigational, motivational routines, such as *select an unoccupied seat and sit*, *approach a performance and watch*, *queue at ticketing areas and purchase a ticket*, etc.

An action selection mechanism triggers appropriate behaviors in response to perceived combinations of external situations and internal affective needs represented by the mental state. For example, in a pedestrian whose thirst exceeds a predetermined threshold, behaviors will be triggered, usually through online planning, to locate a vending machine, approach it, queue if necessary, and finally purchase a drink. In case more than one need awaits fulfillment, the most important need ranked by the action selection mechanism receives the highest priority. Once a need is fulfilled, the value of the associated mental state variable decreases asymptotically back to its nominal value. We instantiate different classes of pedestrians suitable for a train station environment, each class having a specialized action selection mechanism, including commuters, tourists, law enforcement officers, and buskers.

We have recently developed a decision network framework for behavioral human animation [3], leading to advanced strategies for action selection. Combining probability, decision, and graph theories, our probabilistic framework addresses complex social interactions between autonomous pedestrians in the presence of uncertainty. It enables behavior modeling and intelligent action selection subject to a multitude of internal and external factors in the presence of uncertain knowledge, yielding autonomous characters that can make nontrivial interpretations of social situations and arrive at rational decisions dependent upon multiple considerations. We have demonstrated our decision network framework in behavioral animation scenarios involving interacting autonomous pedestrians, including an elaborate emergency response simulation.

2.4. Cognition

At the highest level of autonomous control, a cognitive model yields a deliberative autonomous human agent capable of applying knowledge to conceive and execute intermediate and long-term plans. A stack memory model enables a pedestrian to “memorize”, “update”, and “forget” chains of goals. The stack couples the deliberative




Fig. 2 Musculoskeletal neck-head-face model with 72 Hill-type muscle actuators and a neuromuscular controller.




Fig. 3 Head-Eye gaze behavior (left); the model gazing at a visual target in different directions. Autonomous behavior-based interaction between three face-head-neck systems (right).

intelligence with the reactive behaviors, enabling a pedestrian to achieve its goals. For example, a commuter can enter the station, with the long-term goal of catching a particular train at a specific time. The cognitive model divides this complex goal into simpler intermediate goals, which may involve navigating to the ticket purchase areas to buy a ticket (which may involve waiting in line), navigating to the concourse area, possibly purchasing a drink if thirsty, sitting down to take a rest if tired, watching a performance if interested, meeting a friend, and/or navigating to the correct stairs and descending to the proper train platform when the time comes to board a train.

3. Biomechanically Simulating and Controlling the Neck-Head-Face Complex

The neck has a complex anatomical structure and it plays an important role in supporting the head atop the cervical spine, while generating the controlled head movements that are essential to so many aspects of human behavior. We have developed a biomechanical model of the human head-neck system that emulates the relevant anatomy [4] (Fig. 2). Characterized by appropriate kinematic redundancy (7 cervical vertebrae coupled by 3-DOF joints) and muscle actuator redundancy (72 neck muscles arranged in 3 muscle layers), our model presents a challenging motor control problem, even for the relatively simple task of balancing the mass of the head atop the cervical column in gravity.

We have developed a neuromuscular control model for the neck that emulates the relevant biological motor control mechanisms. Incorporating low-level reflex and high-level voluntary sub-controllers, our hierarchical

controller provides input motor signals to the numerous muscle actuators. In addition to head pose and movement, it controls the coactivation of mutually opposed neck muscles to regulate the stiffness of the head-neck multibody system. Taking a machine learning approach, the neural networks within our neuromuscular controller are trained offline to efficiently generate the online pose and tone control signals necessary to synthesize a variety of autonomous movements for the behavioral animation of the human head and face (Fig. 3).

Biomechanically simulating the human face presents its own deep challenges. We have been working on this problem for approximately 20 years, and employ an improved version of the biomechanical face model described in [5]. Conceptually, the face model decomposes hierarchically into several levels of abstraction related to the behavioral control of facial expression, the anatomy of facial muscle structures, the histology and biomechanics of facial tissues, as well as facial geometry and appearance. Like our biomechanical model of the neck, the face model is muscle-driven. Its 44 facial muscles are arranged in an anatomically consistent manner within the bottom layer of a synthetic facial soft tissue. The tissue is modeled as a lattice of uniaxial viscoelastic units assembled into multilayered prismatic elements with epidermal, dermal, sub-cutaneous fatty tissue, fascia, and muscle layers. The elements enforce volume preservation constraints and model contact response against the skull substrate. Expressive facial tissue deformations result from numerically simulating the physical response of the element assembly to the stresses induced by appropriately coordinated facial muscle contractions. The face simulation runs at real-time, interactive rates on a PC.




Fig. 4 Our comprehensive human body simulation is characterized by the biomechanical modeling of the relevant musculoskeletal tissues. The skeleton with 75 bones is actuated by 846 muscles (left). The motion of the skeleton and the activation level of each muscle deforms the inner soft tissue (center) and, hence, the outer skin (right).




Fig. 5 An inverse dynamics motor controller drives the musculoskeletal system to track a sequence of target poses.

4. Comprehensive Biomechanical Simulation of the Human Body

Extending the above work, we have recently been developing a comprehensive biomechanical model of the human body [6], confronting the combined challenge of modeling and controlling more or less all of the relevant articular bones and muscles, as well as simulating the physics-based deformations of the soft tissues, including muscle bulging (Fig. 4).

In particular, we have created a physics-based skeletal model that consists of *75 bones* and *165 DOFs* (degrees of freedom), with each vertebral bone and most ribs having independent DOFs. To be properly actuated and controlled, our detailed bone model requires a comparable level of detail with respect to muscle modeling. We incorporate a staggering *846 muscles*, which are modeled as piecewise line segment Hill-type force actuators. We have also developed an associated physics-based animation controller that computes accelerations to drive the musculoskeletal system toward a sequence of preset target key poses, and then computes the required activation signal for each muscle through inverse dynamics.

Our volumetric human body model incorporates de-

tailed skin geometry, as well as the active muscle tissues, passive soft tissues, and skeletal substructure. Driven by the muscle activation inputs and resulting skeletal motion, a companion simulation of a volumetric, finite element model of the soft tissue introduces the visual richness of more detailed 3D models of the musculature. Specifically, we achieve robust and efficient simulation of soft tissue deformation within the finite element framework by decoupling the visualization geometry from the simulation geometry. A total of *354,000 Body-Centered-Cubic (BCC) tetrahedra* are simulated to create detailed deformation of the embedded high-resolution surfaces of the skin and each of the muscles.

Figure 5 demonstrates biomechanically simulated arm flexing motions with dumbbell loads.

5. Applications and Future Work

The work reviewed in this paper was initially motivated by applications in computer graphics for the motion picture and interactive game industries. However, we have also pursued other promising applications of our autonomous pedestrian simulation, one in archaeology, the

other in computer vision.

First, we have applied our pedestrian simulator in the context of virtual archaeology for the purposes of visualizing social life in reconstructed archaeological sites. Specifically, we employed a reconstruction of the Great Temple archaeological site of ancient Petra, with which we were able to visualize and test hypotheses about human activity in the Great Temple theater [7].

Second, our prototype simulator has enabled our *Virtual Vision* paradigm in computer vision and sensor networks [8]. In particular, it has facilitated our ability to design large-scale visual sensor networks in virtual reality, develop multi-camera coordination and scheduling algorithms to control them, and experiment with these systems on commodity personal computers. Our approach offers wonderful rapid prototyping opportunities with significantly greater flexibility during the design and evaluation cycle, thus expediting the engineering process.

In view of the complexity of human simulation, our work is inevitably incomplete. For example, our biomimetic modeling approach heightens the need to model skeletal joints more accurately. To this end, since the elementary joints conventionally used in physics simulation cannot produce the complex movement patterns of biological joints, we have introduced a new joint model, called “spline joints”, that can better emulate biological joints [9]. Beyond such low-level technical issues, among the biggest high-level challenges for us going forward is to tackle the neuromuscular control problem for the comprehensive biomechanical body model and, subsequently, to integrate the model within our overarching artificial life framework, leading to fully biomechanically-simulated autonomous pedestrians.

My long-term objective is a computer simulated world that approaches the complexity and realism of the real world, inhabited by virtual humans that look, move, and behave like real humans. Such a *Reality Emulator* could be used in revolutionary ways with profound impact across multiple scientific disciplines.

Acknowledgements

The work reviewed herein was made possible through research grants from DARPA and the NSF, and from NSERC Canada. I proudly acknowledge the outstanding contributions that my current postdoctoral fellows and former doctoral students have made to this research effort. Wei Shao developed the autonomous pedestrian simulator and also applied it to virtual archaeology. Using the simulator, Qinxin Yu developed the decision network framework for human behavioral animation. Faisal Qureshi also employed the simulator to develop our *Virtual Vision* paradigm for computer vision research. Sung-Hee Lee developed the neck-head-face simulator, as well as the incomparable biomechanical simulation of the entire human body in collaboration with Eftychios Sifakis who developed and implemented the associated finite-element soft-tissue model. Yuencheng Lee developed the biomechanical model of the human face.

References

- [1] TERZOPOULOS, D., 1999. Artificial life for computer graphics. *Communications of the ACM*, **42**(8):32–42.
- [2] SHAO, W., AND TERZOPOULOS, D., 2007. Autonomous pedestrians. *Graphical Models*, **69**(5–6):246–274.
- [3] YU, Q., AND TERZOPOULOS, D., 2007. A decision network framework for the behavioral animation of virtual humans. *Proc. ACM SIGGRAPH/EG Symposium on Computer Animation (SCA 2007)*, San Diego, CA, 119–128.
- [4] LEE, S.-H., AND TERZOPOULOS, D. 2006. Heads up! Biomechanical modeling and neuromuscular control of the neck. *ACM Transactions on Graphics*, **25**(3), 1188–1198. *Proc. ACM SIGGRAPH 06 Conference*, Boston, MA, August.
- [5] LEE, Y., TERZOPOULOS, D., AND WATERS, K. 1995. Realistic modeling for facial animation. In *Proceedings of ACM SIGGRAPH 95*, Computer Graphics Proceedings, Annual Conference Series, 55–62.
- [6] LEE, S.-H., SIFAKIS, E., AND TERZOPOULOS, D. 2009. Comprehensive Biomechanical Modeling and Simulation of the Upper Body. *ACM Transactions on Graphics*, to appear.
- [7] SHAO, W., AND TERZOPOULOS, D., 2006. Populating reconstructed archaeological sites with autonomous virtual humans. *Lecture Notes in Artificial Intelligence, Vol. 4133*, J. Gratch et al. (eds.), Springer-Verlag, Berlin, 2006, 420–433.
- [8] QURESHI, F.Z. AND TERZOPOULOS, D., 2008. Smart camera networks in virtual reality. *Proceedings of the IEEE*, **96**(10):1640–1656.
- [9] LEE, S.-H., AND TERZOPOULOS, D. 2008. Spline joints for multibody dynamics. *ACM Transactions on Graphics*, **27**(3), 22:1–8. *Proc. ACM SIGGRAPH 08 Conference*, Los Angeles, CA, August.



Demetri Terzopoulos

Demetri Terzopoulos (PhD '84 MIT) is the Chancellor's Professor of Computer Science at UCLA. He is a Fellow of the ACM, a Fellow of the IEEE, a Fellow of the Royal Society of Canada, and a member of the European Academy of Sciences. Professor Terzopoulos is one of the most highly cited authors in engineering and computer science, with more than 300 scholarly publications, including several volumes, spanning computer graphics, computer vision, medical imaging, computer-aided design, and artificial intelligence/life. His numerous awards include an Academy Award for Technical Achievement from the Academy of Motion Picture Arts and Sciences for his pioneering research on physics-based computer animation, and the inaugural Computer Vision Significant Researcher Award from the IEEE for his pioneering and sustained research on deformable models and their applications.