Autonomous Virtual Humans and Lower Animals: From Biomechanics to Intelligence

Demetri Terzopoulos
University of California, Los Angeles, USA

ABSTRACT

The confluence of virtual reality and artificial life, an emerging discipline that spans the computational and biological sciences, has yielded synthetic worlds inhabited by realistic artificial flora and fauna. Artificial animals are complex synthetic organisms that have functional, biomechanical bodies, perceptual sensors, and brains with locomotion, perception, behavior, learning, and cognition centers. Virtual humans and lower animals are of interest in computer graphics because they are self-animating graphical characters poised to dramatically advance the interactive game and motion picture industries even more so than have physics-based simulation technologies. More broadly, these biomimetic autonomous agents in their realistic virtual worlds also foster deeper computationally-oriented insights into natural living systems. Furthermore, they engender interesting applications in computer vision, sensor networks, archaeology, and other domains.

1. ARTIFICIAL LIFE MODELING

Human simulation has become an increasingly important topic for the interactive game and motion picture industries. We have taken an artificial life approach [1] to modeling humans that combines aspects of the fields of biomechanics, perception, ethology, machine learning, cognitive science, and sociology. This first part of the paper describes our comprehensive model of autonomous pedestrians (§ 1.1), outlines its application to computer vision and sensor network research (§ 1.2), and reviews our earlier work on the artificial life modeling of lower animals (§ 1.3).

1.1 Autonomous Pedestrians

We have demonstrated a reconstructed model of the original Pennsylvania (train) Station in New York City, richly populated with autonomous virtual pedestrians [2] (Fig. 1). In a departure from the literature on so-called "crowd simulation", we have developed a comprehensive human model capable of synthesizing a broad variety of activities in the large-scale indoor urban environment. Our virtual pedestrians are autonomous agents with functional bodies and brains. They perceive the virtual environment around them, analyze environmental situations, have natural reactive behaviors, and proactively plan their activities. We augment the environment with hierarchical data structures that efficiently support the perceptual queries influencing the behavioral responses of the autonomous pedestrians and also sustain their ability to plan their actions over local and global spatiotemporal scales.

Our artificial life simulation of pedestrians (Fig. 2) integrates motor, perceptual, behavioral, and, importantly, cognitive components, each of which we will review in the subsequent sections. Featuring innovations in these components, as well as in their combination, our model yields results of unprecedented fidelity and complexity for fully autonomous multi-human simulation in virtual public spaces.

1.1.1 Appearance and Locomotion

As an implementation of the appearance and motor levels of the character, we employ a human animation 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. Emulating the natural appearance and movement of human beings is a highly challenging problem and, not surprisingly, DI-Guy suffers from several limitations, mostly in its 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 conceal the details of the underlying kinematic control layer from our higher-level behavior modules, enabling the latter to be developed more or less independently.

1.1.2 Perception

In a highly dynamic virtual world, an autonomous intelligent character must have a keenly perceptive regard for its surroundings in order to interact with it effectively. The hierarchical world model is used extensively by pedestrians to perceive their environment, providing not only the raw sensed data (such as those obtained from perception maps), but also higher-level interpretations of perceived situations (such as those obtained from specialized objects) that are at least as important and useful to a pedestrian.

1.1.3 Behavior

The purpose of realistic behavioral modeling is to link perception to appropriate actions. We adopt a bottom-up strategy, 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

Cite as: Autonomous Humans and Lower Animals: From Biomechanics to Intelligence, Demetri Terzopoulos, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Arcade | Main Waiting Room | Concourses and Platforms

Figure 1: A virtual train station populated by self-animating virtual humans.

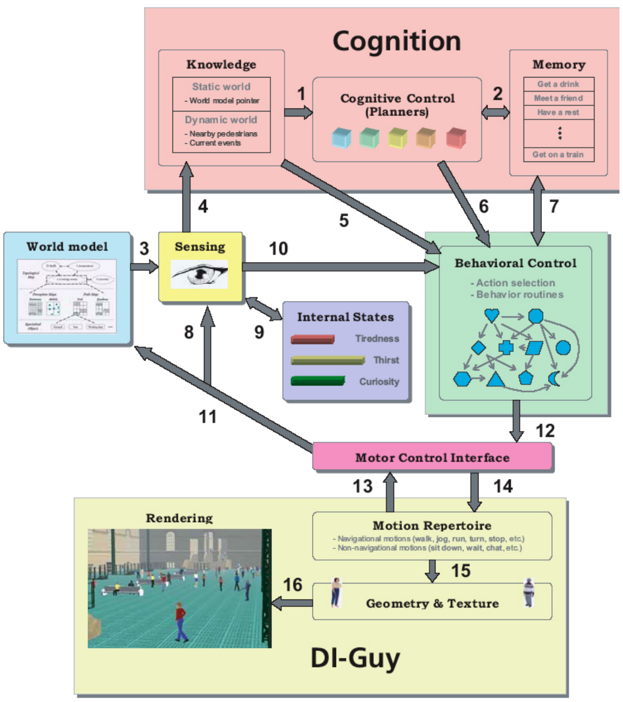

Figure 2: Autonomous pedestrian simulation model.

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 interesting, 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]. This probabilistic framework addresses complex social interactions between autonomous pedestrians in the presence of uncertainty.

1.1.4 Cognition

At the highest level of autonomous control, cognitive modeling [4] yields a deliberative, proactive autonomous human agent capable of applying knowledge to conceive and execute intermediate and long-term plans. A stack memory model enables a pedestrian to remember, update, and forget chains of goals. The stack couples the deliberative intelligence with the reactive behaviors, enabling a pedestrian to achieve its long-term goals. For example, a commuter can enter the station, with the 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 street 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.

1.2 A Reality Emulator

A *Reality Emulator*—i.e., 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—could be used in revolutionary ways with profound impact across multiple scientific disciplines [5]. Offering unprecedented predictive power, it would enable experimentation with complex and/or dangerous scenarios of significant social consequence that cannot be attempted in real life safely, repeatedly, or even at all. For example, the reality emulator could be used to investigate how a pathogen can be transmitted among pedestrians moving in a public space. Or it could simulate the reactions of crowds in crisis situations, such as terrorist attacks, in order to assess the likely consequences and plan appropriate emergency response missions. A realistic simulation of a fire in a virtual building occupied by virtual humans could help assess the visibility and accessibility of exits, as well as plan rescue and evacuation procedures as a function of the location and type of fire, smoke dispersal, ventilation, and so on. The fidelity of such simulations ultimately depends on a detailed modeling of the abilities and limitations of human bodies and brains under real-world conditions.

Our reality emulator can serve as an experimental, albeit reasonably realistic simulation environment for exploring important scientific, engineering, or social "what-if" questions. In particular, it has supported our research in computer vision and sensor networks [6].

Deploying a visual sensor network in the real world is a major undertaking whose cost can easily be prohibitive for most researchers interested in designing and experimenting with large-scale multi-camera surveillance systems. Moreover, privacy laws generally restrict the monitoring of people in public spaces for experimental purposes. A provocative alternative is to deploy networks of passive and active virtual surveillance cameras in our virtual train station, which is populated by autonomous pedestrians. The virtual cameras generate synthetic video feeds that model those generated by real surveillance cameras typically monitoring richly populated public spaces. Our simulation approach offers wonderful rapid prototyping opportunities with significantly greater flexibility during the design and evaluation cycle, thus expediting the engineering process. It has rapidly enabled us to develop novel control strategies for smart camera systems capable of carrying out persistent visual surveillance tasks automatically or with minimal human intervention [6]. The future of such advanced simulation-based approaches appears to be promising in other domains.

1.3 Autonomous Lower Animals

As a precursor to autonomous pedestrians, we developed several lower animals, among them artificial fishes [7]. Each fish is an autonomous agent with a deformable body actuated by internal muscles. The body includes eyes and a brain with motor, perception, behavior, and learning centers (Fig. 3). Through controlled muscle actions, artificial fishes swim through simulated water in accordance with hydrodynamic principles. Their articulate fins enable them

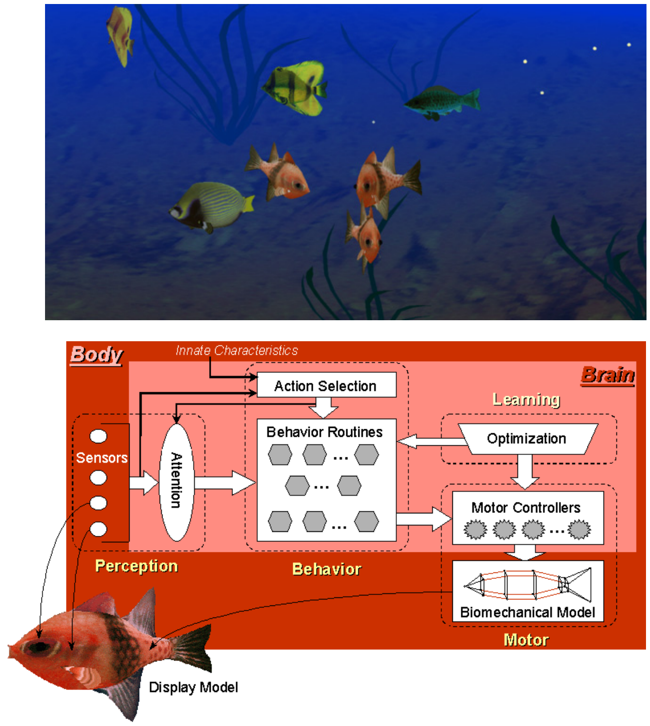

Figure 3: Top: Artificial fishes in their physics-based virtual world; the 3 reddish fish are engaged in mating behavior, the greenish fish is a predator hunting prey, the remaining 3 fishes are feeding on plankton (white dots). Bottom: Functional diagram of the artificial fish model illustrating the body and brain submodels.

to locomote and maneuver underwater. In accordance with their perceptual awareness of their virtual world, their ethological model arbitrates a repertoire of piscine behaviors, including collision avoidance, foraging, preying, schooling, and mating. Though rudimentary compared to those in real animals, the brains of artificial fishes are nonetheless able to learn how to swim and carry out perceptually guided motor tasks.

2. BIOMECHANICAL HUMAN SIMULATION AND CONTROL

In the remainder of the paper, we review our work on aspects of biomechanical human simulation and control.

2.1 Dynamic Human Characters

An ambitious goal in the area of physics-based computer animation is the creation of virtual agents that autonomously synthesize realistic human motions and possess a broad repertoire of lifelike motor skills. To this end, the control of dynamic, anthropomorphic figures subject to gravity and contact forces remains a difficult open problem. We have developed a *virtual stuntman*, a dynamic graphical character that possesses a nontrivial repertoire of lifelike motor skills [8]. The repertoire includes basic actions such as balance, protective stepping when balance is disturbed, protective arm reactions when falling, multiple ways of rising upright after a fall (Fig. 4), and several more vigorously dynamic motor skills. Our virtual stuntman is the prod-

Figure 4: A supine dynamic "virtual stuntman" rolls over and rises to an erect position, balancing in gravity.

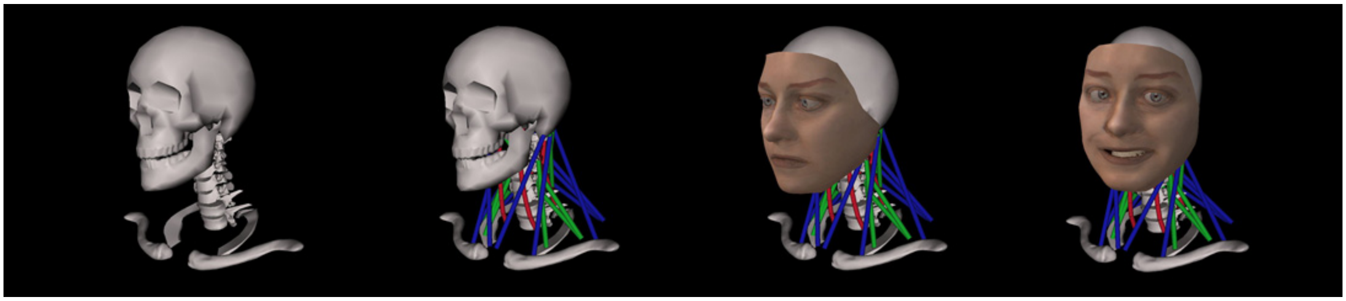

Figure 5: Biomechanical neck-head-face model with neuromuscular controller.

uct of a framework for integrating motor controllers, which among other ingredients includes an explicit model of pre-conditions; i.e., those regions of a dynamic figure's state space within which a given motor controller is applicable and expected to work properly. We have demonstrated controller composition with pre-conditions determined not only manually, but also automatically based on support vector machine (SVM) learning theory.

2.2 Simulating 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 [9] (Fig. 5). 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. 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 (the biomechanical face model described in [10]).[1]

[1]Extending this work, we have recently developed a comprehensive biomechanical model of the human upper body, confronting the combined challenge of modeling and controlling more or less all of the relevant articular bones (68) and muscles (814), as well as a finite element model that simulates the physics-based deformations of the soft tissues, including muscle bulging.

3. ACKNOWLEDGEMENTS

I thank my current and former PhD students whose work is reviewed in this paper and whose names are cited in the following list of references.

4. REFERENCES

[1] Terzopoulos, D., Artificial life for computer graphics. *Communications of the ACM*, **42**(8): 32–42, 1999.

[2] Shao, W., and Terzopoulos, D., Autonomous pedestrians. *Graphical Models*, **69**(5–6):246–274, 2007.

[3] Yu, Q., and Terzopoulos, D., 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, 2007.

[4] Funge, J., Tu, X., and Terzopoulos, D., Cognitive modeling: Knowledge, reasoning and planning for intelligent characters. In: Proc. ACM SIGGRAPH 99 Conf. *Computer Graphics* Proceedings, Annual Conference Series, Los Angeles, CA, 29–38, 1999.

[5] Terzopoulos, D., Perceptive agents and systems in virtual reality. In *Proc. 10th ACM Symposium on Virtual Reality Software and Technology (VRST 2003)*, Osaka, Japan, 1–3, 2003.

[6] Qureshi, F.Z. and Terzopoulos, D., Surveillance camera scheduling: A virtual vision approach. *ACM Multimedia Systems*, **12**:269–283, 2006. Special Issue on "Multimedia Surveillance Systems".

[7] Terzopoulos, D., Tu, X., and Grzeszczuk, R., Artificial fishes: Autonomous locomotion, perception, behavior, and learning in a simulated physical world. Artificial Life, **1**(4):327–351, 1994.

[8] Faloutsos, P., van de Panne, M., and Terzopoulos, D., Composable controllers for physics-based character animation. In: Proc. ACM SIGGRAPH 01 Conf. *Computer Graphics* Proceedings, Annual Conference Series, Los Angeles, CA, 151–160, 2001.

[9] Lee, S.-H., and Terzopoulos, D. Heads up! Biomechanical modeling and neuromuscular control of the neck. *ACM Transactions on Graphics*, **25**(3): 1188–1198, 2006. *Proc. ACM SIGGRAPH 06 Conf.,* Boston, MA.

[10] Lee, Y., Terzopoulos, D., and Waters, K., Realistic modeling for facial animation. In: Proc. ACM SIGGRAPH 95 Conf. *Computer Graphics* Proceedings, Annual Conference Series, Los Angeles, CA, 55–62, 1995.