Autonomous Virtual Humans
ComS 572 Final Project

Tyler Streeter
12/11/03

Introduction
In a society focusing more on the autonomous control of technology, it is expected that virtual humans in 3D worlds have minds of their own. Surprisingly, almost every application of virtual humans uses pre-scripted motions or motion capturing, methods that are time-consuming and inflexible. My project (and future graduate research) concerns the creation of autonomous motor control systems for these virtual humans. I would like to find practical ways to enable them to perform high level tasks given appropriate commands. These self-controlled virtual people could be put into physically realistic virtual worlds and learn to balance themselves, walk, avoid obstacles, transport objects, and do higher level tasks by independently controlling their "virtual muscles." This would make the inclusion of 3D humans much more accessible to application programmers, and an enormous amount of time could be saved by not having to instruct each individual motion.

Among the influences that built my interest in this area are my previous class work, books, and articles. My class work in artificial intelligence, computer graphics, and virtual reality (described in my Previous Research Experience essay) has given me insight into what is possible today and what will be possible in the future. Staying current by reading nonfiction books and articles about science and technology has also inspired me greatly to want to make advances in the field.

The next section of this document outlines the current project implementation. Then I will explain why this type of research is important to society, what resources I will need to perform this research, references to similar work already done in this area, the results I achieved this semester, and future graduate research.

1. Current and Future Implementation
For this project I am using C++, OpenGL, and Open Dynamics Engine (for collision detection and simulated physics). I started by creating a 3D environment complete with simulated physics and incorporated a virtual person built from rigid bodies, joined together with hinges and ball-and-socket joints. (Note that the visual representation of the person can be incredibly realistic since it is separate from the simpler physics representation. At this time I am using a simple visual representation, though I plan to make the virtual humans look more realistic in future work.) After creating the body, I designed a neural network to control the body's motion. This network gets input values from every joint angle, and its outputs are sent as torque values to each joint in the body, effectively simulating the virtual person's muscles. To get the desired behavior without having to specify the neural network's weight values explicitly for each task, I am using a genetic algorithm to evolve the weights.

The fitness functions I have used so far were designed to achieve standing, jumping, and walking behaviors. The specific fitness functions are as follows. The fitness function for standing awards individuals for having a high average head height over the given time (10 seconds per person), and it penalizes individuals for having a large distance between their feet. (This penalty was implemented because the virtual humans would slide their feet apart to keep from falling over.) The jumping fitness function simply assigns scores based on the highest height of the head. Finally, the walking fitness function scores individuals based on how far forward the torso moves, the average head height, and the number of steps taken.

2. Why This is Important

Any kind of virtual environment that uses virtual humans would benefit from this kind of independent motor control. Military applications with simulated soldiers could be greatly enhanced if the soldiers could be given orders and then perform tasks independently. Psychologists wanting to study how humans learn from our surroundings could create simulations of controlled environments with virtual people. There has been much interest recently in creating software agents to assist people in all kinds of tasks, from simple personal assistants to company website tour guides, and each of these agents could be given 3D bodies with motor control if used in 3D environments. Virtual humans used in 3D video games would be natural beneficiaries. 3D meeting places used for business conferences, virtual chat rooms, or any other type of long-distance meeting would almost necessarily need this kind of control system if virtual humans were to be used.

My proposed method could even result in extremely complex and unexpected virtual human behavior with very little design work due to the nature of evolutionary algorithms. The beauty of such algorithms is that we don't need to know how a particular design works internally, just that it does what we want it to do [2]. It is a matter of controlled design vs. directed complexity, the latter giving us useful, robust solutions more efficiently when combined with natural selection.

3. Resources Needed

To do this research I will need software, computer hardware, and knowledge. I will use VRJuggler (the Virtual Reality Applications Center's open source VR software) to handle the interaction between software and VR hardware. I plan to extend a scene graph such as OpenGL Performer or OpenSG to help manage objects in the virtual worlds. I need to use a robust physics simulation package (Open Dynamics Engine seems to work well). To create realistic 3D models I will also need modeling software, such as 3D Studio Max or Maya. I will be able to do most of my work on a standard PC and monitor, but when I need to interact with the virtual humans in a more immersive environment, I may use a head-mounted display or six-sided cave display. I will need to keep learning more about artificial life, artificial intelligence, and computer graphics. I have learned the basics of these areas and will take relevant classes to expand that knowledge.

4. Similar Work

Some foundational work in this field was done by Sims [1] who evolved both creatures' control systems and their morphologies within a physically realistic 3D world. My work, in contrast, will focus on a static, complex morphology (namely, the human body) and will emphasize the development of increasingly complex control systems. Some research done by Klein [3] resulted in a 3D simulation/artificial life package called BREVE. This software is a general purpose tool to simulate artificial life in a 3D physical environment. Rather than using neural networks, however, it uses a simple sine function to control joint motions. More recently, Reil and Husbands [4] have been able to get a virtual human to walk using the evolved neural network method. A project called VLNET (Virtual Life Network) is a software system that allows users to collaborate in virtual worlds [5]. This system makes use of artificially intelligent characters (with predefined motion models) that can be assigned tasks to assist users.

5. Current Results

So far I have experimented with standing, jumping, and walking behaviors. The most visually appealing behavior is jumping, though it is the simplest to evolve. Usually after 50 generations (with population size at least 32), good jumping behavior emerges. The population's average jumping fitness soon stops increasing once the people have learned to maximize their leg power (Figure 1).

Figure 1: Jumping

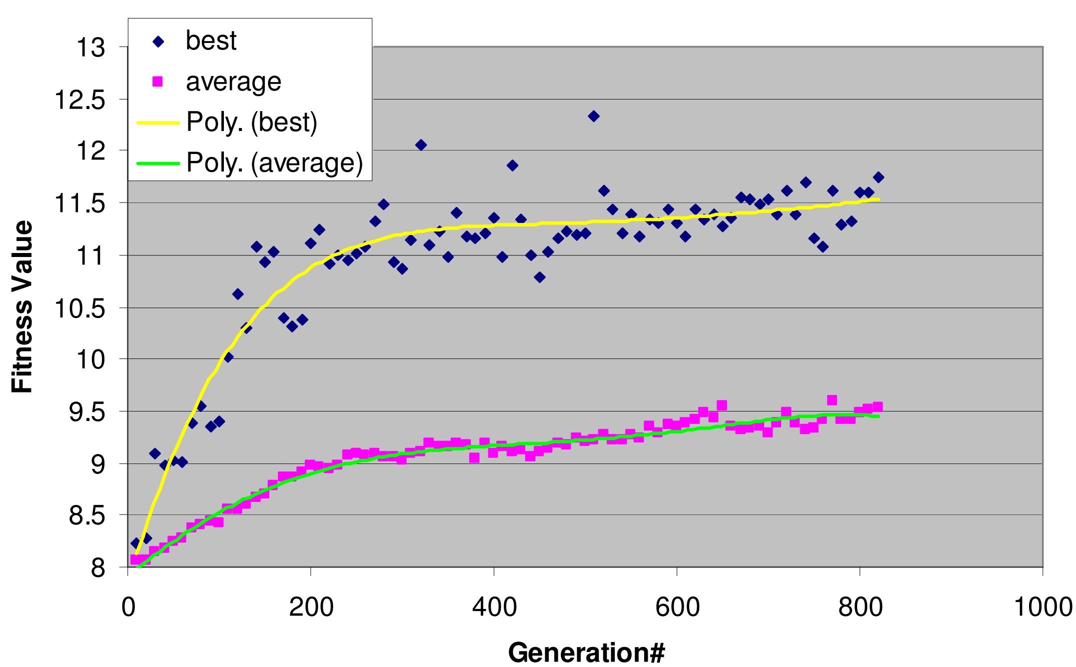

Standing behavior takes more time to elicit, but usually appears after several thousand generations. It usually appears quite suddenly, as evidenced by Figure 2. Many times the virtual humans waste a lot of energy holding themselves in convoluted upright positions while trying to stand. A future goal may be to evolve people that perform certain behaviors while minimizing the amount of energy used.

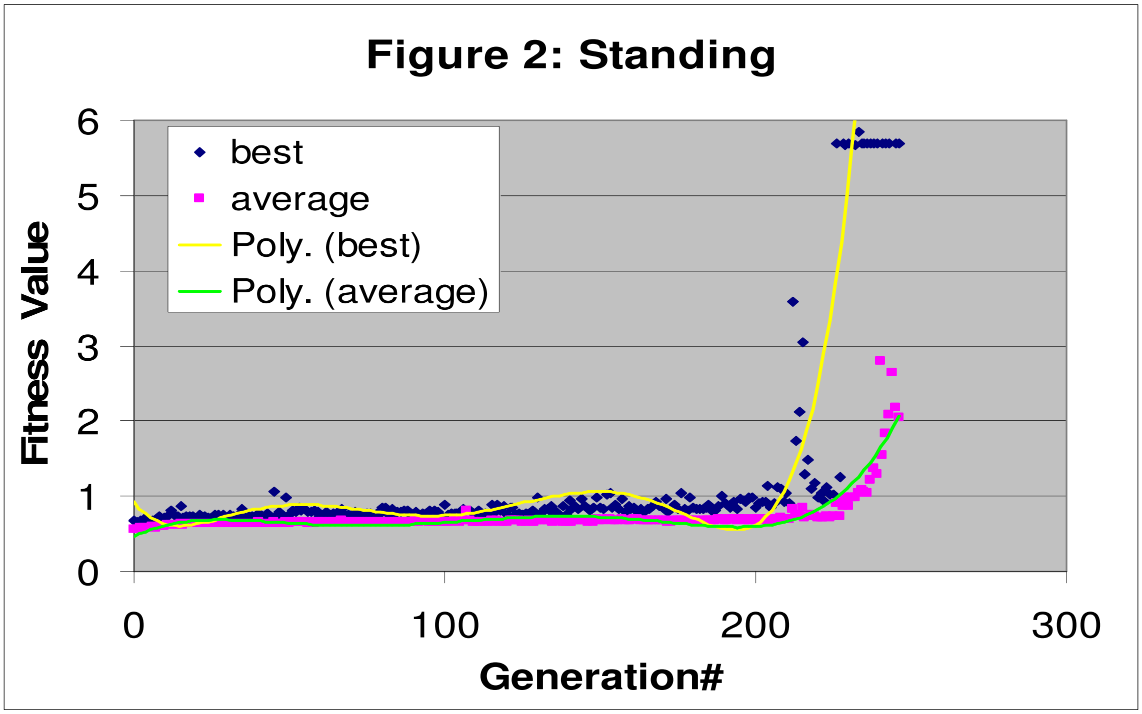

Walking behavior seemed to be the most difficult to evolve, at least with the current fitness function. Figure 3 shows a fairly good learning rate, though the behavior that emerges is still far from human-like walking. The people appear to do a series of hops rather than taking individual steps. Maybe awarding individuals that take slower steps could alleviate this problem.

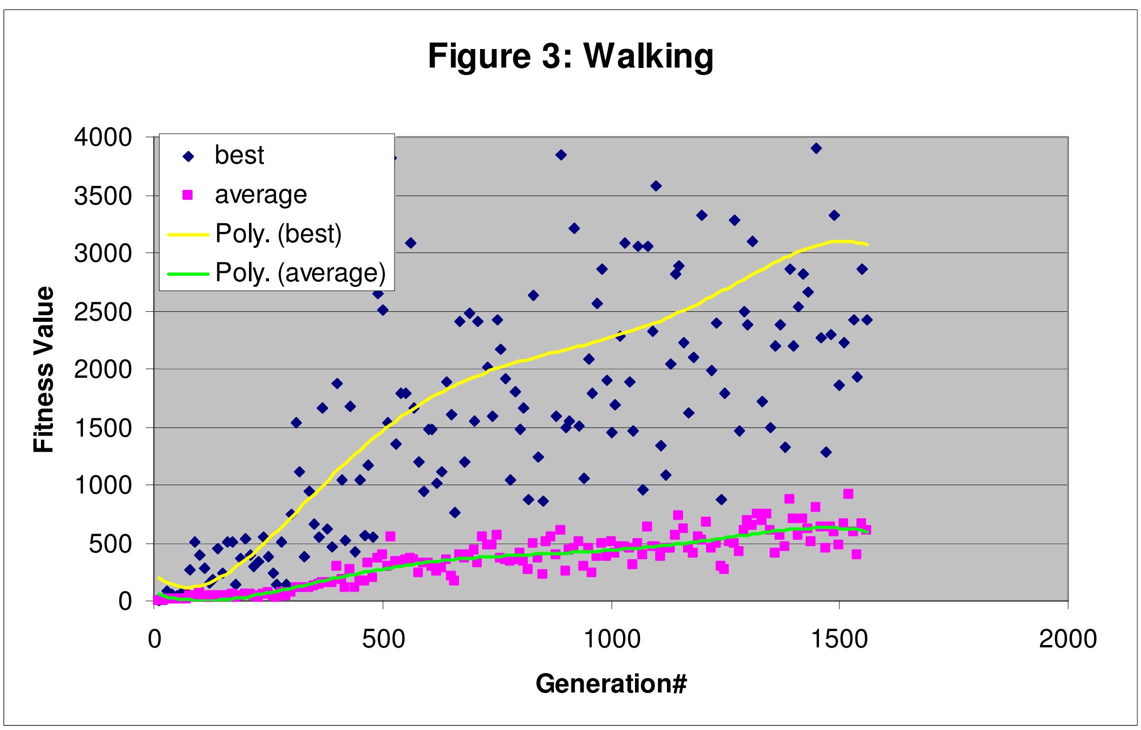

6. Future Graduate Research

Looking ahead, the questions I am asking are: How complex can the behaviors become? Will they be limited by computing power, or is there a fundamental problem with the evolved neural net approach? Are the desired behaviors even possible within the given search space of neural network weights? If so, what are the best ways to find the optimal set of weights? Would a recurrent neural net or simulated annealing method be more effective than the genetic algorithm? Once basic walking skills are evolved, could I use path-finding algorithms to help the virtual people travel through complicated environments?

To become better aware of their surroundings, virtual people will need good sensory inputs. I have found that thinking about how our bodies work helps me understand what senses virtual people will need. One without a sense of touch or sight, for instance, will have a difficult time learning to scale a mountain or open a door. To sense touch, the neural network could receive input signals when the body collides with external objects. Visual input could be implemented in various ways. One way might be simply to calculate the exact distance of other objects to the person's head (ray casting) to create a 3D model of the environment, although this seems more like sonar than sight. Another way might be to create an image of what the person can see and then use a computer vision algorithm to categorize objects and guess their distances. The same research could be used to simulate creatures besides humans. Additionally, these techniques could also be used to evolve physical robot control systems that could be implanted into real robots, as mentioned by Sims [1].

I plan to devise increasingly complex fitness functions to elicit more complex behaviors such as: staying balanced when pushed, walking across uneven terrain, carrying heavy/bulky objects, jumping over obstacles, climbing steep cliffs, and operating virtual machinery, all using simulated physics to avoid pre-scripted animations. More complex tasks will also require increasingly complicated virtual humans. For example, physically, as the tasks become more intricate, they will require jointed fingers to perform certain manipulations. To provide sufficient control of these more complex bodies, the neural networks used to manage their motions will need to become more complex. Perhaps even a hierarchy of neural networks could be used to allow a single virtual human to perform tasks with multiple steps.

Please see my project web site for visual demonstrations and other information:
http://www.vrac.iastate.edu/~streeter/avh.html

References

1. K. Sims. Evolving Virtual Creatures. ACM Computer Graphics (SIGGRAPH '94). pp. 15-22, 1994.

2. Bentley, P. Digital Biology. Simon & Schuster, 2002.

3. J. Klein. BREVE: a 3D Environment for the Simulation of Decentralized Systems and Artificial Life. Artificial Life VIII: The 8th International Conference on the Simulation and Synthesis of Living Systems. MIT Press, 2002.

4. T. Reil, Husbands. Evolution of Central Pattern Generators for Bipedal Walking in a Real-Time Physics Environment. *IEEE Transactions in Evolutionary Computation.* pp. 159-168, April 2002.

5. D. Thalmann, C. Babski, T. Capin, N. Magnenat Thalmann, I. S. Pandzic. Sharing VLNET worlds on the WEB. *Computer Networks & ISDN Systems.* Vol. 29, No. 14, pp. 1601-1610, 1997.