



11-7-2001

Pedestrians: Creating Agent Behaviors through Statistical Analysis of Observation Data

Koji Ashida
University of Pennsylvania

Seung-Joo Lee
University of Pennsylvania

Jan Allbeck
University of Pennsylvania, allbeck@seas.upenn.edu

Harold Sun
University of Pennsylvania

Norman I. Badler
University of Pennsylvania, badler@seas.upenn.edu

See next page for additional authors

Copyright 2001 IEEE. Reprinted from *Computer Animation 2001 : Proceedings of the 14th Conference on Computer Animation*, pages 84-92.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

This paper is posted at Scholarly Commons. http://repository.upenn.edu/cis_papers/214
For more information, please contact repository@pobox.upenn.edu.

Pedestrians: Creating Agent Behaviors through Statistical Analysis of Observation Data

Abstract

Creating a complex virtual environment with human inhabitants that behave as we would expect real humans to behave is a difficult and time consuming task. Time must be spent to construct the environment, to create human figures, to create animations for the agents' actions, and to create controls for the agents' behaviors, such as scripts, plans, and decision-makers. Often work done for one virtual environment must be completely replicated for another. The creation of robust, procedural actions that can be ported from one simulation to another would ease the creation of new virtual environments. As walking is useful in many different virtual environments, the creation of natural looking walking is important. In this paper we present a system for producing more natural looking walking by incorporating actions for the upper body. We aim to provide a tool that authors of virtual environments can use to add realism to their characters without effort.

Comments

Copyright 2001 IEEE. Reprinted from *Computer Animation 2001 : Proceedings of the 14th Conference on Computer Animation*, pages 84-92.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Author(s)

Koji Ashida, Seung-Joo Lee, Jan Allbeck, Harold Sun, Norman I. Badler, and Dimitris Metaxas

Pedestrians: Creating Agent Behaviors through Statistical Analysis of Observation Data

Koji Ashida Seung-Joo Lee Jan M. Allbeck Harold Sun Norman I. Badler
Dimitris Metaxas
Center for Human Modeling & Simulation
University of Pennsylvania
ashida@graphics.cis.upenn.edu

Abstract

Creating a complex virtual environment with human inhabitants that behave as we would expect real humans to behave is a difficult and time consuming task. Time must be spent to construct the environment, to create human figures, to create animations for the agents' actions, and to create controls for the agents' behaviors, such as scripts, plans, and decision-makers. Often work done for one virtual environment must be completely replicated for another. The creation of robust, procedural actions that can be ported from one simulation to another would ease the creation of new virtual environments. As walking is useful in many different virtual environments, the creation of natural looking walking is important. In this paper we present a system for producing more natural looking walking by incorporating actions for the upper body. We aim to provide a tool that authors of virtual environments can use to add realism to their characters without effort.

1 Introduction

Creating a complex virtual environment with human inhabitants that behave as we would expect real humans to behave is a difficult and timing consuming task. Time must be spent to construct the environment, to create human figures, to create animations for the agents' actions, and to create controls for the agents' behaviors, such as scripts, plans, and decision-makers. Often work done for one virtual environment must be completely replicated for another. The creation of robust, procedural actions that can be ported from one simulation to another would ease the creation of new virtual environments. As walking is useful in many different virtual environments, the creation of natural looking walking is important. In this paper, we are not concentrating on creating state of the art leg motion for walking. Our

contribution to locomotion research is the addition of more natural motion of the upper body during walking.

When people walk, their upper bodies rarely follow the classic mirroring the legs pattern found in most animations of walking. Rather, the arms are nearly always engaged, either due to holding a package, gesturing while participating in conversation, or just touching the body. Also, the torso and head should move periodically as the figure glances at the ground and at attention-demanding people, objects, and motions in the environment. Our hypothesis is that the addition of these "subconscious actions" will add a sense of life to otherwise repetitive walking.

We define a subconscious action to be any action that a person is not entirely conscious he or she is performing (though they could become conscious of it). We contrast subconscious actions with conscious or planned actions, which are performed as a step toward a goal. Unlike planned actions, our subconscious actions have no inherent semantic value and do not progress the agent toward a goal. In other words, adding our subconscious actions to a simulation does not alter the semantic content of the simulation like an action, such as, clapping might.

We have, however, come to realize that if our subconscious actions are always performed in the same manner, they may detract from the consistency of individual characters [1]. For example, a very tired or very sad character would be unlikely to perform actions in the same manner as a very energetic character would. In order to preserve the consistency of the characters, we have linked high level agent model parameters, such as tiredness and sadness, with the low level animation parameters such as position and velocity.

Similarly, it seems that some character traits influence the probability of a character performing some subconscious actions. For example, a fearful character is less likely to perform behaviors such as scratching [9]. Through observation, we have created a statistical distribution for our



Figure 1. Pedestrians

subconscious actions. This distribution represents the frequency of occurrences of subconscious actions for an average person. We can then alter the distribution in accordance with agent model parameters to customize the distribution to the characters.

2 Related Work

Using motion capture may be the most effective way to adding natural looking walking to virtual environments, but it is not flexible enough to reuse in different environments. Many different research groups have explored the generation of locomotion. Most methods for generating locomotion can be classified into kinematics [4, 5], dynamics, which are slower and more complex [6, 17, 12, 15], motion editing [11, 16, 21], and combinations of these methods [14, 10, 2, 24]. These existing walking systems have different features including, flexibility (procedural models and figure size independence), parameterization (stylistic walking), curved path locomotion, and the ability to handle uneven terrain. However, all of these systems have concentrated on the movement of the lower body, and nearly ignored the upper body. As a result, the leg movement may appear natural, but the character still appears lifeless.

Similar to the work of Rose, Bodenheimer, and Cohen [23], we use joint angle space to interpolate between animations in order to express more variations in an action. These interpolations are based on higher level parameters, such as emotion. Our work, however, also interpolates the frequency with which our subconscious actions are performed. Also, the weight of their interpolants, such as the amount of limping is designed by animators. The weight of our interpolants is based on statistical observation data and agent model parameters.

Perlin and Goldberg’s Improv system combined scripted actions with stochastically generated motions to achieve controllability and believability at the same time[20]. Although our aim is somewhat similar to theirs, there are a couple of crucial differences. First, their method for avoiding the mechanical appearance of character movements is adding coherent noise to limb motions. We choose instead to provide low level parameters that control aspects of the motion, such as velocity, range of motion, and orientation. The values of these low level parameters are then set in accordance with character properties such as emotion to produce natural-looking motion that reflects the emotional state of the character. Second, the Improv system is centered around a scripting language. This language provides a mechanism for authors to create decision rules containing probabilistic choices between actions. Thus, the stochastic nature of the character’s behavior must be scripted by a simulation’s authors. While our system also allows authors to provide probabilities for actions, our system works without this sort of author intervention. In our system, subconscious actions are chosen stochastically based on psychological and physiological aspects of the agent. Therefore, our system does not require authors to be familiar with human behavior patterns in order to create realistic animations.

The EMOTE system of Chi et al [7] provides artists a method of incorporating movement manner specified in higher level abstract concepts into any pre-existing motions. We also provide a way for artists to control animations with higher level parameters. However, our goal is to off load as much of the work as possible from animators. Towards this end, we built a system to create those higher level parameters according to a distribution so that a series of human agents can be created automatically. From another perspective, our work is different from that of Chi et al in that they modified motion trajectories, whereas we add extra motions to animation sequences.

In [18], Musse and Thalmann have reported their method of crowd simulation with various levels of autonomy. They incorporated rule-based behaviors with programmed and user controlled agents in the same scenario. This added controllability gives the perception of intelligence where necessary. Our focus is more on naturalness. To this end, the quality of an individual’s walking and its variation among the members of a group is more important than being able to give specific commands to a single agent or group of agents. Our controllers provide a method to change the qualities of individuals. For example, through graphical user interfaces, we can control how “depressed” a character is. Musse and Thalmann also can control high level parameters such as emotion, which in turn effect lower level parameters such as walking style. In their system other actions, such as applauding, are either scripted or activated through

action points that are manually placed in the environment. Our system uses stochastic processes which are influenced by emotional state to evoke actions in a less structured, pre-defined manner. Because our actions are stochastically driven, our agents are better able to express emotion. Both the frequency of the actions and the motion quality of the actions are effected by emotional state.

Funge et al [8] used cognitive modeling to create intelligent characters. While we also try to make our system smart enough to create an illusion of life without much effort on the animator side, the focus of our work is very different from theirs. Funge et al concentrated on creating consistent, believable high level behaviors, instead of natural-looking human actions.

We hypothesize that the combination of our system with those cited above would create even better simulations. In particular, the addition of our subconscious actions and motion variations would create a more natural crowd simulations. In our work we have not examined high level behaviors of crowds as others have [22, 19]. We have not focused on flocking behaviors, navigation, or sociological aspects of crowds, but the stochastic addition of subconscious actions to crowd simulations would add needed variation, interest, and realism.

3 Stochastic Behavior Modeling

3.1 Subconscious Actions

As stated in the introduction, our aim is to increase realism in simulations by adding stochastically modeled subconscious actions. These are not actions that require high level planning or progress a story-line. These are actions that we, as humans, perform on a daily basis without necessarily even being aware of. They are similar to idling behaviors found in some games, but our actions are based on a distribution and linked to character traits, whereas in most games idling behaviors are based only on the amount of time the character has been idle.

Subconscious actions are not, at their very root, stochastically driven. They likely have a definable cause and if one were omniscient, one could point to the source of each action. However, from the various levels of explanation for a subconscious action, including the neuron level, the nerve level, the organ level (i.e. skin itch), and the cognitive level (i.e. emotions), it is clear that modeling the underlying causes of subconscious actions is a formidable task. Rather than model each and all of these levels, we instead model the generation of subconscious actions as a stochastic process whose distribution can be modified based on a number of factors, including dry skin and emotions.



Figure 2. A frame from the video used to obtain our stochastic data.

3.2 Statistical Analysis of Observation Data

In order to gather data with which to form our distribution of subconscious actions, we video-taped people as they walked along a section of sidewalk on the University of Pennsylvania campus. Each of the 33 subjects was taped for 35 seconds¹. We then studied the video tape and determined the type of actions we wished to model in our system. In this, the first version of our system, we chose; head touches, body touches, head scratches, body scratches, yawning, looking around, looking down, and gesturing. As we are not attempting to model communication in our system and have no basis for linking speech acts with gesturing, we simply created two different gesture actions for added variability in the simulation.

Next we counted the occurrences of these actions for each subject and averaged them over the 33 subjects, creating a representation of the actions for an average person². We then created a distribution of these actions over time based on an average action duration of 3 seconds. At every frame of the virtual environment simulation, this distribution is passed to the Agent Model to be customized to the character.

3.3 Agent Model

Each character in the virtual environment has a parameterized Agent Model. Including an Agent Model in the system adds to the consistency of the behaviors. Not everyone performs the same subconscious actions or performs them

¹We taped each subject for 35 seconds, because it was the longest time during which the subjects were not occluded by any trees or buildings.

²We are not claiming that our sampling was done in a principled way in order capture data for an average person. It is clearly biased toward a university community. Different distributions could be constructed for different environments.



Figure 3. Sad Walking

at the same rates. In fact even the same person exhibits different actions, such as nervous mannerisms, with different frequencies in different situations.

We can not determine cognitive or behavioral influences, such as emotions, boredom level, or tiredness from a 35 second video clip, so the ways in which the Agent Model parameters alter the subconscious action distribution are empirically derived.

In our system the parameters of the Agent Model, including emotions, physiological factors, and role, can be set by a user before the simulation starts, set randomly, or set through graphical user interfaces during the simulation. Some of the Agent Model parameters change as the simulation progresses. For example, the longer the agent stays awake, the more tired it will become.

The parameters of the Agent Model affect the actions of the agents in two ways. First, they effect the frequency with which it performs the subconscious action. For example, boredom increases the frequency of restless behaviors such as scratching and touching. Similarly, fear increases the frequency of looking around as though the agent were visually scanning the environment for danger. Role also effects the frequency of many of the subconscious actions, including the gestures. Gestures are much more likely when the character is in a speaking role. At every frame in the simulation, the Agent Model updates the distribution of subconscious actions to reflect its current state.

Second, the parameters of the agent model effect the style or manner in which the character's actions are performed. The Agent Model can set parameters of our basic walking (walking with out subconscious actions) to produce

walking that reflects the cognitive state of the agent. For example, we can produce sad walking(See Figure 3) where the agent's body has a more closed shape and the speed of the walking is slower. Furthermore, the Agent Model parameters effect the style of the subconscious actions. One example is handedness. The handedness of the agent effects which hand and arm is used for an action. The predominant hand is preferred, but if that hand is already being utilized, then the other hand will be used (See Section 4.1 for more details).

4 Simulation System

The task of animating such subconscious actions required the design of a new animation architecture. To understand why, consider how online animations are usually generated. Typically, animation data, in the form of joint angle sets for the whole body, are created using an off-line procedure and associated with some action label, such as "Open door." When this action is required, the animation data set is played back on a figure.

This approach commandeers the entire body for the duration of the animation. However, if multiple actions are to take place, then animation data sets cannot take over the whole body, as this prevents other actions from being performed simultaneously. Body segments and joints which are not necessary for the action should not be used and should be left for use by other actions.

Furthermore, subconscious action can occur at any time. Scratching your nose might need to occur while driving, walking, or lying down. This implies that a more flexible notion of action is needed for animation. Animating based on joint angle data sets will be unable to handle many situations in which the subconscious actions will be called.

Finally, subconscious actions are often satisfiable using any available resource. For example, if you are writing using your right hand and need to scratch your nose, you might use your left hand, or the back of the pen in your right hand. This implies that a resource manager for the body will be necessary to determine which joint/segment resources can be used to satisfy an action request. The fact that the resources for performing a subconscious action are known only at run-time (since subconscious actions are generated stochastically), reinforces the previous observation that actions must be very flexible in order to work successfully with different resources, as well as in different situations.

Figure 4 shows the architecture of our system. After the distribution of subconscious actions has been altered by the Agent Model, the Stochastic Behavior unit chooses an action according to the distribution. It then instantiates a PAR (Parameterized Action Representation) [3] for that action. PARs contain high level information such as the agents and object participating in the action, as well as, low level in-




Figure 4. System Architecture

formation such as which end effector(s) is preferred for the action and how strongly it is preferred. The Stochastic Behavior unit uses the parameters of the Agent Model and objects of the environment to set parameters of the actions. For example, if scratching the head is chosen as an action, the Stochastic Behavior unit will consult the Agent Model to determine the handedness of the agent. It then stores this information along with the degree of handedness (in this case low because scratching can be performed with either hand) in the PAR, and places the PAR on the subconscious action queue. Similarly, if looking around is chosen as an action, the Stochastic Behavior unit will provide the action with objects in the environment to look at.

There is also a queue of Planned Actions. Actions on this queue are also represented as PARs, and can be scripted ahead of time, come from a planner or cognitive unit, or placed on the queue dynamically during the simulation through a graphical user interface.

4.1 Action Execution

The Action Execution unit is responsible for managing the body resources (i.e. arms, legs, torso, and head) of the agent and translating the high level parameters of the Agent Model into low level parameters for the Animation system. It examines both action queues and determines which ac-

tions from each queue can be executed given the available resources. Preference is given to actions on the Planned Actions queue. The Action Execution unit then provides additional information such as goal positions and calls the Animation System with all of the necessary information.

The Action Execution unit notifies the Agent Model when an action is being executed. This allows the Agent Model to reset some of its properties when necessary. For example, itch grows over time, increasing its urgency. When the itch is finally scratched, the parameter is reset to its default value. Finally, the Animation System then uses inverse kinematics to animate the character. The Animations System notifies the Action Execution unit when an action has been completed, so that the Action Execution unit can use the body resources that are now free.

We observed that people often initiate subconscious actions concurrently with planned and walking actions or even other subconscious actions. Those actions are not always synchronized to one another, and they do not always end in the same order as they started. When a human starts an action, he or she unconsciously selects parts of the body which are not engaged in any other actions at the time. If necessary he or she would wait until some of those parts are freed before launching the action, or he or she might give up and forget about the action he or she was about to start.

The Action Execution unit examines the queues of actions to find actions with higher priorities that can be executed by the agent given the available body resources.

We enumerated body parts as resources ($\mathcal{R} = \{r_i | 0 \leq i < N\}$ where N is the number of such resources). Currently, we have *Head*, *Left Arm*, *Right Arm*, and *Upper Torso* as the resources. The legs are reserved for walking. Each action (both subconscious and planned) is defined with the set of resources it requires. In other words, for action A_i , there is $R_i \subseteq \mathcal{R}$ such that:

$$R_i = \{r_j | \forall j A_i \text{ uses } r_j\}$$

At any moment in time, an agent has a set of free resources. We let $F \subseteq \mathcal{R}$ denote this set. Action A_i can be launched when $R_i \subseteq F$. Many actions are defined to use any one of the set of resources. For example, the “Scratch Head” action is able to use either $\{LeftArm Head\}$ or $\{RightArm Head\}$.

The queued actions are also associated with priorities and actions with higher priority have precedence for execution. As soon as an action has been chosen to run, the free resource set F is updated, and this action is pulled out from the action queue. It should be noted that in contrast to the work of Perlin and Goldberg[20], we choose between the timing of actions, not the actions themselves. Planned actions have higher priority, but the urgency of subconscious action increase if they are not performed (like an itch that is not scratched).

4.2 Animation System

This module considers the parameters in the PARs, such as objects and handedness. These parameters either decide the trajectory of an action or the manner in which the resources are chosen.

In our system, an *Action* is defined to be a sequence of goal positions of one or more end effectors along with their intervals. Using inverse kinematics, the goal positions are linearly interpolated in joint angle space to achieve continuous motions.

In order to make those actions work in situations along with other actions concurrently, we define the goal positions to be relative to the coordinate frame of the relevant body segment. For example, some goal positions of the “Scratch Head” motion are defined relative to the coordinate frame of the head segment. This simple scheme works for our purposes, allowing the head to move, while it is being scratched.

Another issue that must be addressed in our system is the blending and stitching together of disparate actions. This becomes much more difficult in light of the unknown state of the body, due to the stochastic nature of the subconscious actions, and in light of our desire not to return to a “body neutral” state after each action. Returning to a known state not only does not work when actions can overlap, but is undesirable because of the unnatural look it gives.

We blend motions in a fashion similar to that of Perlin and Goldberg[20]. In order to transition from one action to another, each required joint gradually changes the percentage of the actions it is performing. As example, take transitioning from arm swing to scratching head. At the beginning of the transition the arm joints are performing 100 percent arm swing and no scratching. They gradually start lowering the amount of arm swing and increasing the amount of scratching until it is eventually fifty percent of each and finally 100 percent scratching.

Subconscious arm motions are generated using inverse kinematics (IK). Goal positions for the IK are generated inside of our Action Execution module. Scratching the tip of a nose is an example of a subconscious arm motion. In order to animate this action, the Action Execution module finds two relative positions on the nose surface from a pre-defined site on the character’s face. The IK then moves the specified end effector to the goal positions. Other goal positions are given from higher levels of this system because they can be different depending on the current environment or time. For example, an object could be given as a goal position for an attention action.

One of the key issues in animating simultaneous actions is how to generate the trajectories for the end effectors, so that the actions are satisfied no matter what other actions are occurring. For example, the trajectory for scratching the

face is very different when the figure is upright and walking than when it is bent over a table writing. Because our system checks goal position at every cycle of the animation, that goal position does not need to be static. Going back to our previous example of scratching the nose, at every cycle our system finds the two goal positions from the pre-defined site, which may have changed its global position since the last frame, if for example the agent is moving its head. The IK then uses these newly found goal positions to move the end effectors. Thus, it can handle dynamic trajectories.

The Animation System uses IKAN[25] inverse kinematics and Transom Jack[13] to animate the actions and move the human figure.

4.2.1 Parameterization of Animation

We have applied our ideas about stochastic modeling of subconscious actions to creating more natural looking walking. Therefore, we have divided walking into upper and lower body movements. Most research on locomotion concentrates mostly on the lower body movement and simply adds synchronized arm swing to the upper body. We believe that movement of upper body adds significant reality and naturalness to walking animation. Unlike legs during walking, the upper body in general does not create a closed chain of joints. This gives us greater variability, which makes adding affect to the walking simpler. However, the movement of upper body does need to be closely synchronized with the lower body motions, and any changes happening to the lower body motions need to be reflected to the upper body. We describe our system of controlling upper body in section 4.2.2. This upper body controller is responsible for spine, arm, and neck motions. In addition we need to control the nuances of locomotion. We need a walking animation which allows us to control its parameters in real time. Our solution to this problem is described in section 4.2.3. This leg controller is responsible for movement of all the joints from the pelvis to the feet. The Agent Model and Action Execution modules can set parameters of the actions for both upper and lower body movements.

4.2.2 Upper Body Movements

Unlike the leg motion which plays the main role in the functionality of forwarded movement, upper body movement has more room for affect to show through. We focused on the movement of upper body which can convey this secondary information. In order to create an upper body movement which is controllable enough to generate different styles of walking, we employed a simple kinematic, procedural method.

In figure 5, we show the structure of our human figure. We control 11 degrees of freedom (DOFs) in this structure to create the upper body movement for walking. The DOFs




Figure 5. Joint Structure of the Upper Body

we use are associated with elbow joint (1 DOF), shoulder joint (2 DOFs), clavicle joint (2 DOFs), upper $i + 1$ spine joints (3 DOFs), and lower $N - i$ spine joints (3 DOFs) where N is the total number of spine joints of the figure (including the waist) and i is the index of specially designated spine joint.

In order to cancel out the rotational movement of the pelvis that is created during walking, the lower $N - i$ spine joints move counter to the pelvis rotation. The necessary rotation of 3 DOF is evenly distributed in those $N - 1$ joints so that the upper body does not rotate in any direction above the i th joint because of the pelvis movement. Selection of i effects the naturalness of the upper body movement. From empirical study, we found $i = 13$ to be a good setting for our figure with $N = 17$ spine joints.

Upper body and leg movements are synchronized using a common time parameter t . All the rest of DOFs are controlled by sine curves with this t as a parameter.

$$R_j = A_j \sin(a_j t + b_j)^{p_j} + B_j$$

where $0 \leq j < 11$ and R_j is the rotation angle of the joint for time instance t . Controls of symmetric joints are achieved using the same sine curve with phase shift of π .

$$R_j^l = A_j \sin(a_j t + b_j)^{p_j} + B_j$$

$$R_j^r = A_j \sin(a_j(t + \pi) + b_j)^{p_j} + B_j$$

Those 5×11 parameters (A_j , a_j , b_j , p_j , and B_j with $0 \leq j < 11$) define a style of upper body movement. Changing these parameters, changes the walking style of a figure. Each simulation begins with default (average) values for these parameters. In order to create different styles of walking for different agents, we have created a graphical tool to interactively adjust these parameters. Instead of creating this parameter set for each agent, we have constructed




Figure 6. Trajectory of Toe During Walking.




Figure 7. Distance Checking

the parameter set associated with emotions of agents. For example, we have created a parameter set for “Angry” walking. We used weighted barycenter of those parameters with agent’s emotion parameters as the weights. This created a walking style which reflects the agent’s emotional state.

4.2.3 Lower Body Movements

Like the arm movement, walking in our system is also animated by using inverse kinematics. The walking has pre-stored key coordinates for the goals used by the leg inverse kinematic system. The root foot is non-swinging foot that stays in contact with the ground. The goal position of the swinging leg is calculated based on the distance from the root foot to the center of the pelvis (See Figure 7).

An intermediate goal position and orientation of the swinging leg is provided as a parameter of the walking. The goal position is first calculated relative to the pelvis. It is then transformed into the global coordinates necessary for the inverse kinematics. The style of walking can be varied simply by changing one key coordinate for the intermediate goal. This coordinate is used to set the maximum height of the swing leg and the position of this maximum height along the leg’s stride path (See Figure 6).

During a walking cycle, a time t , where $-1 \leq t \leq 1$, is maintained. This time value represents the position of the legs in the walking cycle. For example, $t = 0$, represents the

left foot in full contact with the floor, and $t = -1$ or 1 represents the right foot in full contact with the floor. Note that because this is a repeating cycle, -1 is equivalent to 1 . This time t is then used to coordinate the upper body movements with the leg movements, such as the arm swing with the leg strides. Our upper body movements and subconscious actions can be used with any walking system that provides this timing information, because as we have mentioned in the previous section, our animation system uses the lower spine joints to compensate for the pelvis rotations.

5 Evaluations

In order to evaluate our system, we performed an informal study in which we showed various animation clips to 6 people. The majority of the participants (83%) found the walking with subconscious actions to be more natural than walking without subconscious actions. The participant that found plain walking more natural indicated that she was distracted by movement in the hips. All of the participants correctly identified our sad walking clip as sad. Some (66%), however, also described our worried walking as sad. Most of the participants were unable to identify angry walking. We believe adjustments in the parameters and an angry facial expression is needed to more clearly depict anger. Additionally, most (66%) of participants correctly identified happiness.

6 Conclusions and Future Work

In this paper, we have presented a system for producing more natural looking walking by incorporating actions for the upper body. Our aim is to provide a tool that authors of virtual environments can use to add realism to their characters without effort. Our system models subconscious actions with a statistical distribution and alters this distribution to customize the actions to traits of the characters in the virtual environment. Additionally, the actions themselves, including the walking, are parameterized allowing their performance to be consistent with the cognitive, social, physiological, and emotional state of the agent.

We found that our work is particularly well suited for the simulation of pedestrians (See Figure 1). With our system, generating a lot of agents with individualized walking actions is quite simple. All that is necessary is the setting of the Agent Model parameters. If desired the parameters of the agent could even be set stochastically.

There are, however, many areas where our work could be expanded. First, we are currently using a simple method for generating walking. We plan to add additional parameters to the walking.

Secondly, we have observed that there is phrasing in subconscious actions. For example, a person may scratch mul-

iple locations on their head in the same phrase of motion. Currently, this would only occur in our system if the various scratches had similar priorities. When one subconscious action is marked for animation, we could check the urgency of others that could be phrased with it. For example, if scratching the nose is about to execute and the value of itch on the chin is above a certain threshold, we could phrase them together.

Additionally, we would like to add more subconscious actions to our system. We have noticed that available objects seem to promote some subconscious behaviors. For example, a watch or a pair of glasses are often adjusted while walking. More emotive actions, including a facial animation system, and a component, such as EMOTE [7], for expressive gestures would also contribute realism to our system.

Lastly, we plan to more fully integrate an agent model with our system. The cognition of our current agent model is limited. One area of further exploration is the linkage of plans and subconscious actions. An agent's plans effect among other things its emotions and physiological state. These type of states are reflected in the frequency and style of subconscious actions, but the link between the plans and these states is currently not in place.

References

- [1] J. M. Allbeck and N. I. Badler. Consistent communication with control. In *Workshop on Representing, Annotating, and Evaluating Non-Verbal and Verbal Communicative Acts to Achieve Contextual Embodied Agents at Autonomous Agents*, 2001. To appear.
- [2] F. Anderson and M. Pandy. A dynamic optimization solution for one complete cycle of human gait. In *International Society of Biomechanics XVII Congress*, page 381, 1999.
- [3] R. Bindiganavale, W. Schuler, J. Allbeck, N. Badler, A. Joshi, and M. Palmer. Dynamically altering agent behaviors using natural language instructions. In *Autonomous Agents 2000*, 2000.
- [4] R. Boulic, N. Magnenat-Thalmann, and D. Thalmann. A global human walking model with real-time kinematic personification. *The Visual Computer*, 6:344–358, 1990.
- [5] A. Bruderlin and T. Calvert. Knowledge-driven, interactive animation of human running. In W. A. Davis and R. Bartels, editors, *Graphics Interface '96*, pages 213–221. Canadian Information Processing Society, Canadian Human-Computer Communications Society, May 1996.
- [6] A. Bruderlin and T. W. Calvert. Goal-directed, dynamic animation of human walking. In J. Lane, editor, *SIGGRAPH '89*, volume 23, pages 233–242, July 1989.
- [7] D. Chi, M. Costa, L. Zhao, and N. Badler. The EMOTE Model for Effort and Shape. In *SIGGRAPH '00*, pages 173–182, 2000.
- [8] J. Funge, X. Tu, and D. Terzopoulos. Cognitive modeling: Knowledge, reasoning, and planning for intelligent characters. In *SIGGRAPH '99*, pages 29–38, 1999.

- [9] P. E. Gallher. Individual differences in nonverbal behavior: Dimensions of style. *Journal of Personality and Social Psychology*, 63(1):133–145, 1992.
- [10] M. Girard and A. Maciejewski. Computational modeling for the computer animation of legged figures. *ACM Computer Graphics*, 19(3):263–270, 1985.
- [11] M. Gleicher. Motion editing with space-time constraints. In *SIGGRAPH '97*, pages 139–148, 1997.
- [12] J. Hodgins, W. Wooten, D. Brogan, and J. O'Brien. Animating human athletics. In *SIGGRAPH '95*, pages 71–78, 1995.
- [13] T. Jack. <http://www.transom.com/>.
- [14] H. Ko and N. Badler. Animating human locomotion in real-time using inverse dynamics, balance and comfort control. *IEEE Computer Graphics and Applications*, 16(2):50–59, March 1996.
- [15] J. F. Laszlo, M. van de Panne, and E. Fiume. Limit cycle control and its application to the animation of balancing and walking. In H. Rushmeier, editor, *SIGGRAPH '96*, pages 155–162. Addison Wesley, Aug. 1996.
- [16] J. Lee and S. Shin. A hierarchical approach to interactive motion editing for human-like figures. In *SIGGRAPH '99*, pages 39–47, 1999.
- [17] M. McKenna and D. Zelter. Dynamic simulation of autonomous legged locomotion. In *SIGGRAPH '90*, pages 29–38, 1990.
- [18] S. Musse and D. Thalmann. Hierarchical model for real-time simulation of virtual human crowds. *IEEE Trans. on Visualization and Computer Graphics*, 7(2):152–164, 2001.
- [19] S. R. Musse and D. Thalmann. A model of human crowd behavior: Group inter-relationship and collision detection analysis. In *Proc. Workshop of Computer Animation and Simulation of Eurographics '97*, Budapest, Hungary, 1997.
- [20] K. Perlin and A. Goldberg. Improv: A system for scripting interactive actors in virtual worlds. In *SIGGRAPH '96*, pages 205–216. Addison Wesley, 1996.
- [21] Z. Popovic and A. Witkin. Physically based motion transformation. In *SIGGRAPH '99*, pages 11–20, 1999.
- [22] C. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proc. SIGGRAPH '87*, volume 21, July 1987.
- [23] C. Rose, M. F. Cohen, and B. Bodenheimer. Verbs and adverbs: Multidimensional motion interpolation using radial basis functions. *IEEE Computer Graphics and Applications*, 18(5):32–40, 1998.
- [24] H. Sun, A. Goswami, D. Metaxas, and J. Bruckner. Cyclogram planarity is preserved in upward slope walking. In *International Society of Biomechanics XVII Congress*, page 514, 1999.
- [25] D. Tolani, A. Goswami, and N. Badler. Real-time inverse kinematics techniques for anthropomorphic limbs. *Graphical Models*, 62:353–388, 2000. <http://hms.upenn.edu/software/ik/>.