Artistically Based Computer Generation of Expressive Motion

Michael Neff* Eugene Fiume*

*Department of Computer Science
University of Toronto
neff|elf@dgp.toronto.edu

Abstract

Understanding how to create the right movement for a specific character is a significant open problem in computer animation. This same problem, however, has been at the heart of the work of performance artists for hundreds of years. In this work, we try to learn from the lessons contained in the performing arts literature so that we can apply them to creating rich, engaging, animated characters. Three classes of movement properties are identified: those relating to shape, those relating to transitions and those relating to timing. Computational models of some of these properties have been developed and are briefly detailed. A software framework is also presented that allows these properties to be applied independently to a character's movement. The resulting animation system allows the key aesthetic aspects of movement to be quickly varied and allows motion to be easily customized for a particular character.

1 Introduction

A talented actor, dancer or animator can take a few abstract notions of a character, and from those build a rich, consistent, concrete piece of movement that captures the essence of that character. This ability to take an abstract idea and give it concrete form is the heart of a performer's art. In the research field of computer animation, we have barely begun to understand how this mapping works: how an often vague description of a few character properties can be mapped to a complete, concrete style of movement. We can, however, gain insight into the process by studying the arts literature. This can help us to understand what aspects of movement have the most significant impact on the expressive content of that movement and what aspects are most useful to vary in order to develop various types of characters. In a sense, this is like building a vocabulary of fundamental movement parameters that can then be used at any time to define the movements of a particular character.

Automatically generated motion has generally lacked the expressiveness required to create engaging characters. For this reason, character systems often rely on an animator to provide either key poses or animation sequences as the basis for their animation engines. For example, in the AlphaWolf system of Tomlinson et al. (2002), an interactive wolf simulation demonstrated at SIGGRAPH 2001 by the Synthetic Characters Group from Media Lab, an animator crafted animation sequences at the extremes of a dominance scale and the system then interpolated between them. In approaches such as this, the manner in which a motion is modified remains a part of the animator's craft and is not under control of the character modeling system. A long term goal of our work is to provide the tools and vocabulary necessary for the character system to directly control how a motion is modified in order to capture the nuances of a specific character's movements.

Our work learns and applies lessons from the arts literature. Through studying the performance literature from theatre, animation, mime, theatre anthropology and movement theory, we have built a list of key movement properties that significantly affect the aesthetic aspects of a character's movement. These have been divided into three main categories: those that affect the character's shape or pose; those that affect transitions, namely how a character moves from pose to pose and other transient effects; and those that relate to timing. We have built computational models for many of these movement properties and developed a software framework that allows this extensible set of properties to be combined in order to generate animated sequences.

The focus of this work is on how to vary the manner in which a given movement is performed, rather than on what movement is performed. In other words, we are interested in defining the expressive range over which an action can be performed, and in finding generic ways to parameterize these ranges.

One of the major desired outcomes of this research has been the creation of better authoring tools for character animators. We hope this work can also play an important role in research on autonomous or semi-autonomous expressive characters. In these fields, it is important to be able to customize the movement style of various characters in order to match their mood and personalities. Our research not only elucidates the movement properties that are likely to be useful to vary, but also provides computational models for many of these properties, and a framework in which these properties can be combined. A long term goal of our research is to examine how these movement properties can be combined in various ways to create character sketches. Such sketches make use of the low-level movement properties to define a basic movement style for a character. Different requested actions will then be performed in a manner that is consistent with the character sketch.

2 Lessons from Computer Animation

Previous work in computer animation on expressive character motion can be divided into three categories: script based systems, model-based systems and data driven systems. These categories are distinguished by where the knowledge of what makes motion expressive lies. In script based systems, the user hand tunes expressive variations in motion using low-level constructs. The system does not encapsulate any knowledge of expressive movement. Model based methods try to capture some aspect of what makes motion expressive within a computational model, which can then be invoked by a user. The user is thus working at a higher or more abstract level. Data driven methods make use of motion capture data and the expressiveness is in the captured sequences. A user merely determines which data is used to create the final motion. Several model based approaches also use motion capture data, but the intent there is to build a model of a movement quality that will then be independent of the data.

A script based system was introduced by Perlin (1995) that allowed a user to interactively control an animated character. The actions of the character were specified ahead of time using script files that specified character poses and transition functions that could be used to move between these poses. The transition functions were either sinusoids or different frequency noise functions. Perlin and Goldberg (1996) extended this work to a general system for scripting interactive characters known as *Improv*. This system took a similar approach to generating animation, but also allowed behavioural scripts to be defined that determined what actions a character took.

Chi et al. (2000) use Laban's Effort-Shape movement analysis to define a fixed set of parameters that can be used to modify the style of a motion. This model-based approach is similar to our work in that their model is based on research in the arts literature, but we aim at a more open, extensible system and we target a larger set of movement properties.

Other model based approaches use captured data of human movement in order to develop a computational model. Work on expressive transforms (Unuma et al. (1995); Amaya et al. (1996)) has attempted to extract emotional content from a piece of captured motion. The extracted transforms can then be applied to other motions, in order to give them the same emotional 'feel'. Brand and Hertzmann (2000) provide for very high level editing of a motion's style. They learn a style, such as ballet or modern dance, from captured motion and can then apply this style to other movement sequences. Pullen and Bregler (2002) work at a similar high level, allowing an animator to specify key frames and then use a statistical model drawn from motion capture data to texture the key framed motion with the style of the captured data.

Rose et al. (1998) present a data driven system in which expressive motion is generated by interpolating within a large captured motion set. Actions, such as walking, are called verbs in their system. They capture the motion of each action of interest being performed in a large number of ways. These variations then define an "adverb" space. During playback, the user can specify an action and a location within the adverb space. The final motion is generated by interpolating between the captured motions. By modifying the location in adverb space, the user can change the style of the motion.

Bruderlin and Williams (1995) adjust captured motion by treating movement as a signal and adjusting the gain of various frequency bands of the signal. They suggest different bands capture different aesthetic qualities of the motion.

Our work is a hybrid between script based and model based systems. We build computational models of movement properties, but these are defined programmatically, potentially by the end user, and are highly configurable. Users can also build new properties which are composites of existing properties defined in the system. Unlike many previous approaches, the expressive properties represented in our system are all garnered from the arts literature and as such benefit from traditional practice. As well, a focus of our work is on how a varied and extensible set of movement properties can be combined. An additional advantage of our work is that we generally solve for expressive constraints at the same time as solving for hard constraints, such as having a character touch a certain location. Many of the above approaches operate by warping an existing piece of motion and can hence violate hard constraints in generating the final motion.

In previous work, we have presented a model for tension and relaxation in the human body and explored its expressive impact (Neff and Fiume (2002)). We have also presented a small set of simple yet powerful aesthetic edits that can be used to modify the feel of a motion sequence (Neff and Fiume (2003)).

3 Lessons from the Arts

One of the lessons taken from the arts literature is that performance movement is not the same as daily life movement. Eisenstein and Tretyakov (1996) argue that the purpose of stage movement is to infect the audience with emotion and according to Alberts (1997), it is generally when the attributes of an actor's movement are out of the ordinary that they will have the greatest significance for the audience.

Because of the need to communicate clearly with an audience, performance movement is based on two basic principles: *simplification* (Lawson (1957); Thomas and Johnston (1981); Lasseter (1987); Barba (1991b)) and *exaggeration* (Thomas and Johnston (1981); Lasseter (1987); Barba (1991a)). These two properties work together to clarify the meaning of a character's movement in the spectator's mind.

Simplification, also known as the "Virtue of Omission" (Barba (1991b)), works to bring focus to certain elements of a character's movement by eliminating extraneous movements. In mime, it is often necessary to slow down the pace of major actions to make them comprehensible (Lawson (1957)). In traditional animation, they preach the importance of having a character only do one thing at a time (Thomas and Johnston (1981); Lasseter (1987)). All this lends clarity to performance motion that is often lacking in the movements of daily life.

Once a movement has been simplified, it is exaggerated to ensure that its meaning is conveyed to the audience. Frank Thomas nicely summarizes the interplay of simplification and exaggeration as follows: "As artists, we need to find the essence of the emotion and the individual who is experiencing it. When these subtle differences have been found, we must emphasize them, build them up and at the same time, eliminate everything else that might appear contradictory or confusing." (Thomas, 1987, p.6). By so doing, an artist ensures that he/she is focusing the audience's attention on the key aspects of the movement and clear communication results.

We have organized our analysis of the arts literature into three main categories: shape, transition and timing. We will briefly examine a few of the most important properties from each category.

3.1 Shape

From the point of view of defining a character, one of the most crucial aspects of shape is stance. Stance is a combination of a character's posture and his/her balance point. Stance is one of the quickest indicators of both a character's overall personality and how he/she is feeling in a particular scene. One of the simplest views of posture, espoused by Alberts (1997), is as a measure of the overall level of tension in a character's body, or more specifically, to what degree the character expends energy to hold himself erect against the pull of gravity.

Barba (1991a) illustrate the importance of the curve adopted in the spine. The 'beauty line', seen perhaps most famously in the Venus de Milo and other ancient Greek statues, creates a large S-curve that snakes through a character's body and lends a very sensual appearance to the work. This provides a clear illustration of the role of spine curvature in the coronal plane. Spine curvature in the sagittal plane is most often related to a character's reaction to gravity, as suggested by Alberts. It can also be used to illustrate a character's reaction to an object by having a character either recoil away from an object or lean towards it.

Shawn (1963), in summarizing the work of the movement theorist, Delsarte, suggests that the part of the torso that a person habitually holds forward is a strong indicator of what kind of person they are. If they hold their chest high, this indicates self-respect and pride. If their abdomen is protruding, this indicates animality, sensuality and lack of bodily pride. A normal, balanced carriage will have the middle zone of the abdomen carried forward and the chest and abdomen withdrawn. This triad can be augmented by considering people who carry their head forward, normally indicating a mental or academic disposition.

Balance in stance refers to where a character's centre of mass is relative to his or her feet. Is the character's weight centred between the two feet or does the character have a lean? When we stand in daily life, we are never still, but rather are constantly making small adjustments, shifting our weight to the toes, heels, right side, left side, etc. Barba (1991b) argues that these movements should be modeled and amplified by the performer, underlying the expressive importance of balance adjustments. Indeed, Laban (1988) and Tarver and Bligh (1999) suggest it can be particularly powerful to situate a character near its balance limit. This makes the pose more precarious and acts to heighten the associated sense of tension for the audience.

Extent is another key shape property. Extent or extension refers to how far an action or gesture takes place from a character's body. It can be thought of as how much space a character is using while completing an action. If the arms are fully extended and straight out from the character's side, this would be maximal extension, whereas, if the hands were touching the torso and the elbows were held at the character's side, this would be minimal extension. Both Laban (1988) (also Tarver and Bligh (1999)) and Delsarte (Shawn (1963)) use the term *extension*. Alberts refers to this as *range*.

Laban refers to the area around a person's body as the kinesphere and defines three regions within it (Tarver and Bligh (1999)). The near region is anything within about ten inches of the character's body. This is the area for personal, intimate or perhaps nervous actions. The middle area is about two feet from the person's body and this is where daily activities take place, such as shaking hands. The area of far extent has the person extended to full reach. It is used for dramatic, extreme movements and in general is used more on stage than in daily life. Laban (1988) argues that people learn the space around their body and determine how it is most comfortable to perform an action.

The literature on body shape is rich and warrants considerably more investigation and, perhaps, critical scrutiny, than we are able to provide in this paper. Many writers deal with the meaning associated with different areas around the body. It is also common to associate meaning with different parts of the body. The sequencing of body shapes, or put differently, how shape varies over time, is also very important expressively.

3.2 Transitions

Transitions deal with how a character moves from shape to shape. They also deal with transient aspects of movement, such as the interplay of tension and relaxation.

Laban (1988) suggests that the flux of movement can flow continuously, be intermittently interrupted, yielding a trembling kind of movement, or stopped, yielding a pose. It is worth distinguishing between a motion that is paused (Laban (1988); Tarver and Bligh (1999)) or suspended (Alberts (1997)), and one that is stopped. According to Tarver and Bligh (1999), when a motion is paused, the actor maintains focus on the final destination and can continue the movement at any time. There is no perceptible loss of intensity nor break in intention. When a motion is stopped, the energy of that motion has been lost and the actor's focus is no longer on the completion of a motion. It cannot be seamlessly continued with the same intensity. This may seem to be an esoteric point, but it is quite easy to distinguish between a paused and stopped motion when observing a performer.

According to Laban, motion can be either complete or incomplete(Tarver and Bligh (1999)). Many actions in daily life are incomplete – they are stopped before they reach their natural conclusion. Being able to stop an action midway can have a powerful effect as it can be a strong indicator of a character's internal mental process. Consider a character that reaches out to comfort a former lover and then stops the motion part way through. This is a clear indication of the internal conflict the character is experiencing.

Lasseter (1987) and Thomas and Johnston (1981) describe how Disney animators found it effective to have the bulk of footage near extreme poses and less footage in between in order to emphasize these poses. They referred to this as slow in, slow out, indicating that most of the time was spent on the main poses. Interpolating splines are used to generate this result in computer animation (Lasseter (1987)), often going beyond simple ease-in, ease-out curves using tool such as tension, continuity and bias splines(Kochanek and Bartels (1984)).

Modifications to the transition envelope of a movement are very important expressively. Some movements, such as a tired character dropping his head, start slowly and accelerate as they proceed. Other movements, like a feinted lunge, start very quickly and slow near the end as the character stops the motion in a controlled way. Other motions, such as those common in tai chi, move at a consistent pace throughout. Further to this, some characters will have a light touch while others have a heavy touch. This is not necessarily a reflection of the physical mass of the character, but is rather controlled by how the character moves his mass. Very heavy characters may still move in a light way, and vice versa.

Arcs are another Disney animation principle (Thomas and Johnston (1981); Lasseter (1987)). The *arcs principle* claims that natural movements are normally not straight, but follow some form of arc.

The interplay of tension and relaxation is another widely cited movement property. See for example: Dorcy (1961); Laban (1988); Shawn (1963); Lawson (1957); Barba (1991a). There is a flow between tension and relaxation. There must first be relaxation in order for there to be tension and tension is followed again by relaxation. Shawn (1963) describes the inflow and outflow of energy that accompanies tension modulation. There may be different tension levels in different areas of a character's body. Furthermore, as Lawson (1957) describes, tension changes can take place through the entire body or a tiny part. They can occur suddenly or gradually and there can be spasmodic changes back and forth. Laban (1988) suggests a rise in tension can serve to accent a movement. Physical or emotional pain can be shown by spasmodic contractions of muscles, followed by relaxation as the pain eases, according to Shawn (1963).

3.3 Timing

The two main components of timing are tempo and rhythm (Lawson (1957); Laban (1988); Alberts (1997)). *Rhythm* refers to the overall beat structure or pattern of a set of movements. For instance, the pattern could be long, long, short, repeat. *Tempo* refers to the speed at which motions are completed. Tempo is independent of rhythm. For instance, a given rhythm could be performed with a fast or slow tempo.

Other terms used to define timing include duration, which is the length of time an action takes (Alberts (1997)) and speed, which is the rate at which we let movements follow one another (Laban (1988)).

The timing or speed at which an object moves reflects the size and weight of the object (Lasseter (1987); Shawn (1963)). Larger objects have more momentum and thus accelerate and move more slowly than small, light objects. In his "Law of Velocity", Delsarte extends this idea to include the relationship between the inner gravity of feeling or meaning and the pace and size of its expression (Taylor (1999); Shawn (1963)). Thus the deepest of feelings may be expressed with complete stillness (Taylor (1999)). "Emotions of profound and deeply serious import, then, require slow and large movement patterns; emotions that are petty, light, trivial, nervous etc. take on small and quick movement patterns" (Shawn, 1963, p.64). Extending the idea to character types suggests that "[m]ajesty and nobility of emotion and character can only be conveyed by broad, slow movements while the petty, small, nervous, fearful, irritated, shallow person and emotion is revealed by quick, small movements." (Shawn, 1963, p.65) This same phenomenon is often at play in film making when superheroes gifted with great speed are shown not moving more rapidly than those around them, but shown moving in slow motion to lend gravity to their deeds.

Lasseter (1987) suggests it is important to spend the correct amount of time on *anticipation* for an action, the action itself and then the reaction to the action. Timing can be used to indicate if a character is nervous, lethargic, excited or relaxed (Lasseter (1987)).

The use of *successions* is arguably the most important timing property. Successions deal with how a movement passes through the body. Rarely will every limb involved in a motion start and stop at the same time. As described by Shawn (1963), Delsarte defined two types of successions: true or normal successions and reverse succession. In a normal succession, a movement starts at the base of a character's torso and spreads out to the extremities. In a reverse succession, the movement starts at the extremities and moves in toward the centre of the character. Delsarte associated true successions with good, truth and beauty and reverse successions with evil, falsity and insincerity (Shawn (1963)). Shawn (1963), a pioneer in the medium, argues that the active use of successions was fundamental to the development of American modern dance.

The importance of successions was also known by traditional animators. Thomas and Johnson write: "Our most startling observation from films of people in motion was that almost all actions start with the hips; and, ordinarily, there is a drop – as if gravity were being used to get things going. From this move, there is usually a turn or tilt or a wind up, followed by a whiplash type of action as the rest of the body starts to follow through... . Any person starting to move from a still, standing position, whether to start walking or pick something up, always began the move with the hips."(Thomas and Johnston, 1981, p.72) This is consistent with a normal succession.

4 System Overview

Our animation system is focused on the creation of realistic, humanoid character motion for a standing figure. The actions we work with include gestures, posture changes and balance adjustments. Focussing on this limited set of motions allows us to explore the rich expressive range contained therein. The system can create motion either kinematically or using physics-based simulation for a 48 Degree of Freedom (DOF) skeleton. Dynamically simulated motion takes advantage of tension modelling and can produce higher quality results, but due to the computational cost of physical simulation, does not yet run as a real-time process. Kinematic generation is currently more appropriate for interactive applications.

To meet the goals of this work, it was necessary to develop a new software architecture. This architecture must achieve several ends. It must:

- allow motion input from multiple sources.
- allow different aspects of the motion to be varied separately.
- support the procedural definition of movement properties.
- provide a mechanism to smoothly combine all the movement properties that are active in the system.

The unifying concept underlying the architecture is the *movement property*, or *property* for short. Low-level properties vary a particular aspect of a movement, such as the amount of tension in a joint, the shape the torso takes during a reaching movement, or the transition function used to vary the envelope of a movement. Higher level properties may combine several low-level properties to vary several properties of a movement in a consistent way. Both high and low-level properties can be procedurally defined. Properties provide the handles by which an animator modifies a movement. They are used to both define and edit an action. Often a property will be bound to some underlying feature in the system that actually does the work of modifying the character's motion. For instance, the *Balance Shift* property allows an animator to adjust the desired *x-z* position of the character's centre of mass. This change is effected by the balance controller, which in conjunction with an IK algorithm, modifies the joint angles in the character's lower body to achieve the change. The property itself merely specifies the desired change.

The software architecture is shown in Figure 1. Animators, technical directors, or another computational process can interact with the system through three channels: the script, action descriptions, and direct animator edits. In the future, they will also be able to create character sketches that will automatically customize movement sequences for a given character.

The script contains a time-ordered list of actions. It defines what a character does. There are multiple tracks in the script allowing actions that affect different parts of the body to be superposed. The script can accept predefined actions or reactive controllers. Predefined actions are used for planned movements, such as reaching for an object. Reactive controllers actively change the character's behaviour in response to sensor input. Balance is the most significant reactive controller in the system. It actively modifies the angles in the lower body in order to achieve a desired balance point. This desired balance point can be moved over time to affect balance changes.

An action description is a small script file that defines a specific movement, such as a posture adjustment, a gesture or a reaching motion. The representation used for action descriptions is discussed in detail below. The basic philosophy of action descriptions is to attempt to find a minimal definition of a movement that then facilitates a wide range of expressive editing and customization.

Animator edits allow an animator or character controller to directly modify either an individual action or a series of actions that make up a movement sequence.

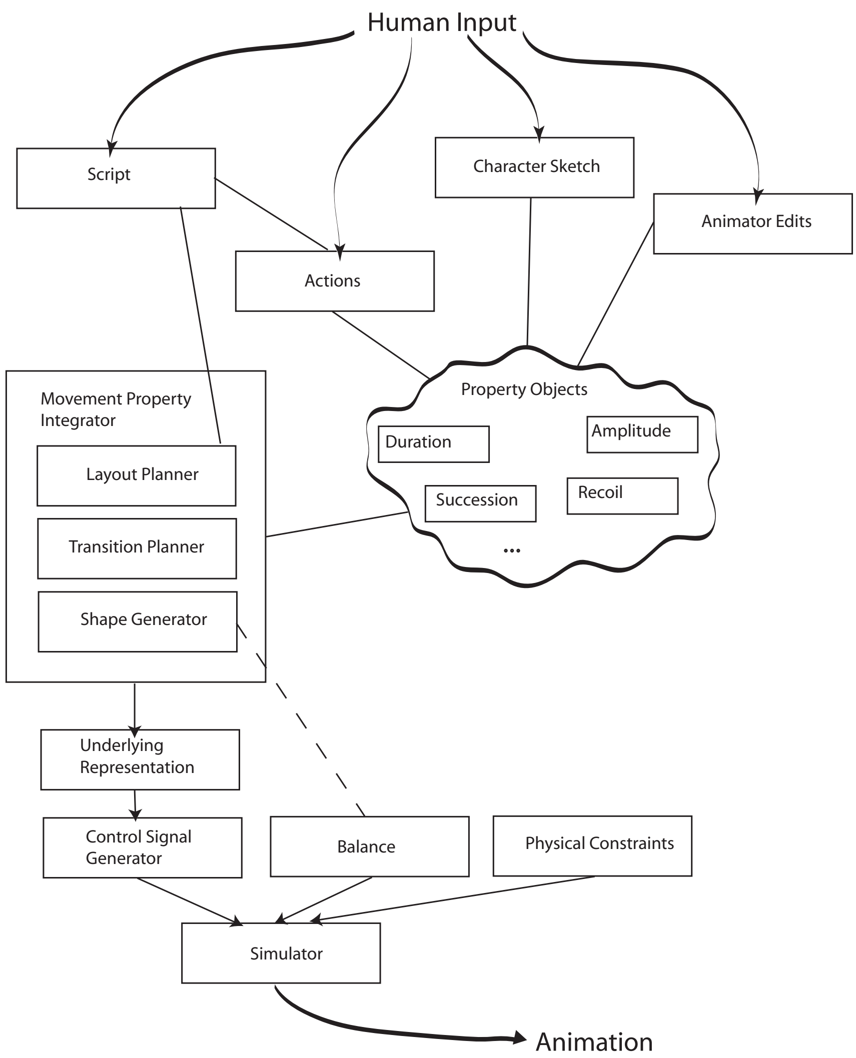

Figure 1: System architecture.

The edits work by either modifying the active movement properties or invoking additional movement properties.

The *character sketch* is used to specify a default set of tendencies for a given character. This functionality is currently being researched and will be added to future version of the system.

As much as possible, properties are defined to be orthogonal. For instance, modifying the shape of a character's torso is independent of the amount of tension in the torso and the timing used to move to this new shape. This orthogonality limits conflicts when various aspects of a character's movement are varied independently. Nonetheless, conflicts may occur. This can happen if two input sources, such as an animator edit and the character sketch, modify the same underlying property, or if composite properties modify a low-level property that is either directly modified or modified by a different composite property. It is the job of the *Movement Property Integrator* (MPI) to arbitrate between these conflicts in a consistent and predictable manner. Conflicts are currently dealt with by ordering the application of properties based on their source and time of application, but we are actively investigating more general conflict resolution schemes. The MPI is also responsible for mapping the movement properties to the low level system functions that are used to achieve them. Three planners are used to meet these goals, one for shape, one for timing and one for transitions. The MPI combines the movement properties in order to generate the *Underlying Representation* that will be used to drive the animation.

The Underlying Representation (URep) is used to generate the final either kinematic or dynamic animation. Animators do not need to be aware of this representation to create effective motion. All the movement properties active in the system are automatically transformed to the URep. The URep consists of a set of time-indexed tracks. There are two types of track. First, there is a track for each degree of freedom in the character's body which specifies the desired value for this DOF and other DOF related data. The second set of tracks contains control parameters that are not related directly to an individual DOF, such as the desired offset to the COM and the desired gaze direction. This data is used by processes that modify the character's movement as it is generated in order to meet these constraints.

Tracks are populated either with *Transition Elements* (TElements) or control parameters that are good for a single simulation time step. The attributes specified by TElements include a desired value of the DOF, the length of time it should take to reach that value, the transition function used to shape the envelope of the movement and desired tension changes. Planned actions are mapped down to a series of TElements by the Movement Property Integrator. Control params are generated at the beginning of each time step by reactive controllers, such as the balance controller. They are used to modify the character's movement in response to changes in the system state.

The *Control Signal Generator* (CSG) drives off the Underlying Representation to calculate whatever data is needed by the active simulator in order to generate the animation. With a kinematic simulator, the CSG simply determines the value of each DOF. With a dynamic simulator, the CSG must determine the two antagonistic joint gains and damping value required to achieve the desired DOF values. The details of the low-level muscle model are discussed in the Transition Properties section below.

The simulator is responsible for generating the final animation using all of the control data provided by the CSG. The Kinematic Simulator merely needs to update the transformation hierarchy based on the values of each DOF. The Dynamic Simulator must apply all forces acting on the character and then integrate the physical equations of motion in order to determine the new position of the character. A commercial package, SD/Fast (Hollars et al. (1994)), is used to generate the code that implements the equations of motion for the specific character.

4.1 Action Representation

The key role of the action representation is to facilitate editing of the motion. Action definitions attempt to capture the essence of a motion while specifying a minimal number of details. For instance, a reaching motion might be defined by a world space target the wrist must touch. A wave might require the arm to be in a particular area and the elbow angle to oscillate. Such minimal descriptions can be flushed out by other properties provided by the character sketch or animator edits. This allows a generic action to be customized for a particular character and situation.

Actions are essentially pose based, normally specifying particular poses at particular instants in time, although edits such as succession can break up this structure. Action descriptions are defined using both high and low-level properties. They include data on timing, some joint angles, external constraints such a world space point to touch and often composite movement properties with their parameters.

4.2 Timing Properties

Timing properties are applied directly to the time properties of the transition elements. Generally, they work to either shift the position of a transition element on its track or to scale it.

Succession is an example of a procedurally defined timing property. It operates by directly shifting the location of the transition elements in the underlying representation. A succession edit takes two parameters: whether the succession is normal or reverse and how much of a time offset (t) to use between the joints involved in the motion. The edit determines all of the transition elements it is being applied to and shifts their starting time based on where they are in the character's joint hierarchy. For instance, a normal succession would not modify the first joint in the spine, it would offset the next joint by t, the following joint by $2t$, and so on. The succession traces down all branches in parallel, for instance, modifying the start time of both collar bones, then both shoulders and then both elbows, etc.

4.3 Transition Properties

Two of the key transition properties in the system are transition functions and tension changes. Transition functions are used to warp the envelope of a motion. They are modelled using a cubic Hermite embedded in time and in space. This allows good control over the shape of the transition, including anticipatory and overshoot effects. A transition curve is associated with each transition element in the underlying representation. Each DOF can have its own transition curve or transition curves can be applied to joints, groups of joints, or an entire movement. Common transitions include: ease-in, ease-out; linear; fast start, slow end; slow start, fast end; backtrack before going forward (anticipation); and exceed final point and then return to it (overshoot).

Tension changes are used when the motion is dynamically simulated. Varying tension has a number of effects on a character's motion. Perhaps most obvious, the amount of tension in a character's body will effect how the character reacts to external forces, such as being shoved. One of the most important effects of tension changes is to warp the envelope of a motion. For instance, if a joint's tension is increased during a motion, the attached limb will move more quickly at the beginning of a motion and more slowly at the end. Low tension effects include overshoot, where a character moves slightly beyond its desired target before moving back to it, and pendular motion, for instance, when a character brings his arm down to his side. Tension also effects how forces travel through the body and how precisely the transition function is tracked. In general, modifying tension provides an effective way to differentiate between a character that moves in a loose, relaxed style and a character that moves in a more tight and rigid style.

Human muscle is organized in agonist-antagonist pairs. This simply means that muscles act in opposite directions around a joint. The final position of the limb will be the equilibrium point of the forces generated by these muscles and any other forces, such as gravity, acting on the limb. One theory of movement, the *equilibrium point hypothesis* proposed by Feldman (1966), suggests that humans move by directly varying the equilibrium point of their muscles. We developed a simple muscle model that is based on the two ideas of equilibrium point control and antagonistic muscle pairs. It consists of two opposing angular springs and a damper for each DOF.

The control law for the antagonistic actuator is

$$\tau = k_L(\theta_L - \theta) + k_H(\theta_H - \theta) - k_d\dot{\theta}, \qquad (1)$$

where τ is the torque generated, θ is the current angle of the DOF and $\dot{\theta}$ is its current velocity. θ_L and θ_H are the spring set points which serve as endpoints for the motion range of a DOF, k_L and k_H are the spring gains, and k_d is the gain on the damping term. The stiffness or tension of the joint is taken as the sum of the two spring gains, $k_L + k_H$. The desired angle is not explicitly shown in the actuator, but is controlled by the value of the two gains. This can be seen as a reparameterization of Proportional-Derivative (PD) control that is more appropriate for directly modeling tension changes. The gains that will achieve any desired angle are simply a linear function of each other.

In previous work using PD controllers, the gains of a PD controller were normally either carefully hand tuned to achieve a given behaviour or they were kept high in order to minimize positional error. The second case generates stiff characters. The first case requires much painstaking hand-tuning. Neither case leads to characters that can easily modify their tension. We avoid this by using equilibrium point control that takes external forces acting on the limb, mainly gravity, into account and adjusts spring gains to compensate for them. This allows us to achieve any desired angle while also varying the tension. A more detailed account of our work on modeling tension is given in Neff and Fiume (2002).

4.4 Shape Properties

Shape properties such as posture changes act through a shape generation module called the *body shape solver*. This module solves for a pose that satisfies the various shape properties acting on the character. The shape generator accepts a combination of hard world-space constraints, such as touching a particular world space point or maintaining a foot position; soft expressive constraints, such as maintaining a certain curve in the character's spine or keeping the character's elbows out from his side; and balance adjustments.

The body shape solver is customized for the human skeleton. It contains a combination of analytic and optimization based Inverse Kinematics routines. Balance control is based on a feedback strategy that uses the error in the character's centre of mass position to adjust the angles of the dominant ankle.

Other edits, such as the *extent* and *amplitude* edits, work by directly modifying the desired joint angles stored in the TElements of the Underlying Representation. More detail on these can be found in Neff and Fiume (2003).

5 Results

All animations discussed here can be found online at http://www.dgp.toronto.edu/people/neff

The system has been used to generate a wide range of animations involving reaching motions, gestures, posture changes and balance adjustments. To illustrate how various movement properties are layered together, consider a simple example of a character reaching for an object, say an apple, on her right. We want the reaching motion to have a sensual, languid feel. The initial action description specifies the reach target and uses default values for timing and the transition function. The result is a rather plain piece of animation. To create the sensual look desired, a shape property is invoked that causes the body shape solver to seek a beauty-line posture and adjust the balance point. The character's gaze direction is also modified to look toward the apple. A succession edit is applied to increase the sense of flow associated with the motion and give it a languid feel. The duration of the motion is extended and the transition functions are adjusted to make the motion smoother.

Now consider a second character that has a nervous temperament and is afraid to touch the apple, but must do so anyway. Starting from the same basic animation, we apply a different set of properties. First, we invoke a shape property that has the character recoil away from the apple while reaching for it. This time, we use a slight reverse succession to heighten the sense of unease. Timing is shortened and the transition function is varied to have the character jerk away from the object. Starting with the same basic action, we have created two animations with completely different feel. The initial reach solution, a frame from the sensual reach and a frame from the recoiling reach are shown in Figure 2

An animation was created of a very loose, relaxed character giving the direction "go left". Two versions of the animation were generated: a dynamic version that makes

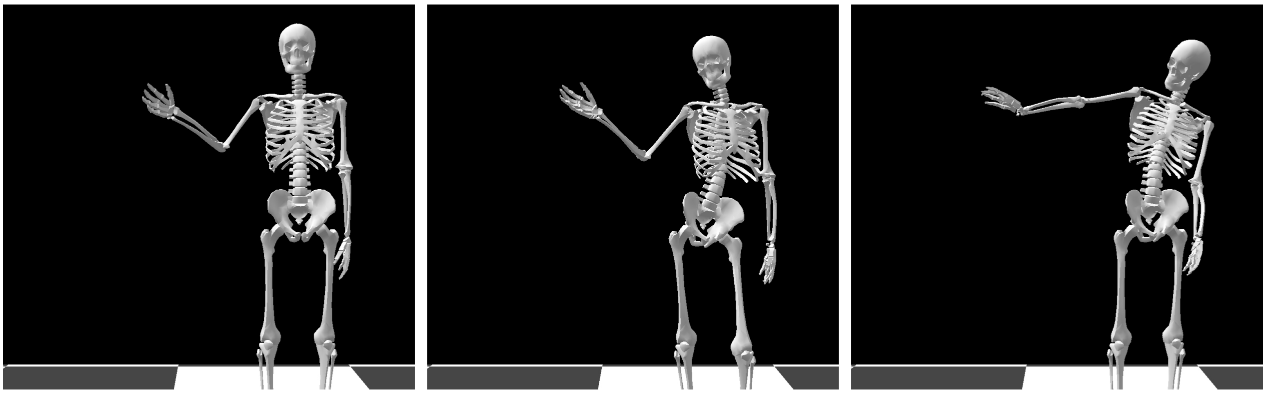

Figure 2: This figure shows a frame from each of three different reaching motions. The first is from the neutral, default motion. The second shows a more sensual character reaching for the object. The third shows a character who would rather avoid the object.

use of tension changes and a kinematic version for comparison purposes. Both animation sequences are based on the same set of actions and use the same poses. The dynamic animation with tension changes does a much better job of capturing the desired sense of looseness and flow. This is due to several effects resulting from tension modeling that only occur in the dynamic version: the loose wrists hang in an appropriate way and show the effect of momentum during movements, the arms have a slight bounce to them as they are brought to the character's side, and arm motions induce some secondary movement in the torso.

Figure 3 shows a 0.4 sec section from the animation. During this clip, the character has just brought his arms down to his sides such that his forearms are parallel to the ground. The top image shows five frames from the kinematic version of the animation. There is no movement during this period. The bottom image shows five frames from the dynamic version of the animation. The bounce in the spine, shoulders, elbows and wrists can be clearly seen. This additional motion greatly adds to both the sense of looseness capture by the animation and the realism of the clip.

The dynamic version of the animation is also smoother. This is because dynamic simulation acts much like a low-pass filter on the motion, leading to a smoother final output.

As a third example, we show several versions of an animation in which a character rises from a bow to make an emphatic two handed gesture. The intensity of the motion is increased by increasing the extent and adjusting the transition functions. A succession is again used to increase the sense of flow. Figure 4 shows three different versions of the motion with different extension.

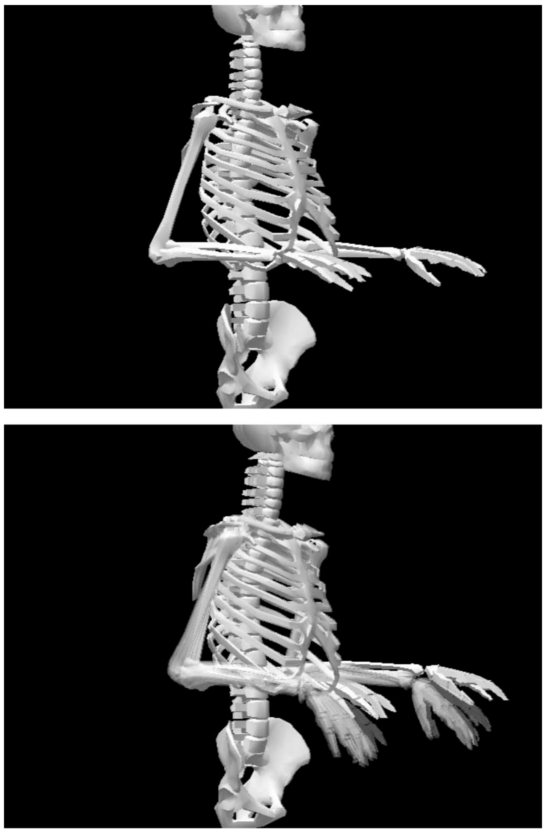

Figure 3: Comparison of kinematic and dynamic motion: The top image shows 0.4s from a kinematic version of the "go left" animation and the bottom image shows the same 0.4s from a dynamic version of the animation.

6 Conclusion and Discussion

Determining the correct movements for a specific character remains a challenging problem. The performing arts

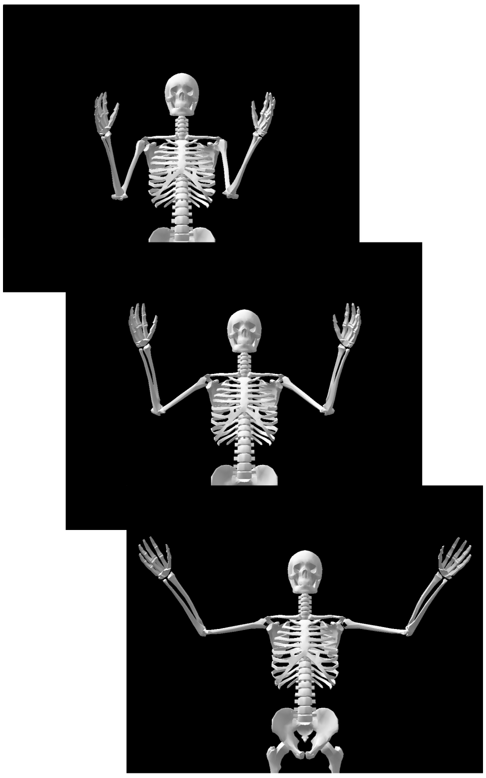

Figure 4: Three different extent versions of the same gesture.

literature, however, offers a rich source of information that can be applied to the task. We have tried to ground our computer animation tools in the lessons learned from the arts community. This helps to ensure that the problems we solve are the problems that need to be solved in order to effectively model expressive motion.

In this paper, we have summarized a number of key lessons from the arts community, providing a new organization of these ideas into the categories of *shape*, *transition* and *timing*. Computational models of a number of these properties were presented along with an architecture that allows various properties to be combined and new properties to be added to the system. The effectiveness of the approach was shown by demonstrating how motion could be quickly customized for a particular character.

Our work provides a vocabulary and set of computational models that can be used by high level character systems to modify the low-level properties of movement to match the personality and mood of a given character. Prescripted motion limits the actions a character can perform. Our hope is that the ideas contained here will be useful in developing character systems that can benefit from the flexibility of automatically generated motion while still offering expressive, engaging characters.

Given the difficulty of the problem, there is much more work to be done. Our research priorities are to develop the character sketch and provide a tighter set of rules for integrating the various movement properties. In addition, faster simulation techniques and more robust dynamic balance control would help make physics-based animation practical for real time applications.

Acknowledgements

Financial support for this work was provided by NSERC and the Department of Computer Science, University of Toronto. The software described here is built on top of DANCE (Dynamic Animation aNd Control Environment), developed by Victor Ng-Thow-Hing and Petros Faloutsos. DANCE is available at: http://www.cs.ucla.edu/magix/projects/dance/index.html Early on in this work, Prof. Stephen Johnson provided some very useful pointers into the arts literature.

References

David Alberts. *The Expressive Body: Physical Characterization for the Actor*. Heinemann, Portsmouth, N.H., 1997.

Kenji Amaya, Armin Bruderlin, and Tom Calvert. Emotion from motion. *Graphics Interface '96*, pages 222–229, May 1996. ISBN 0-9695338-5-3.

Eugenio Barba. Dilated body. In Eugenio Barba and Nicola Savarese, editors, *A Dictionary of Theatre Anthropology: The Secret Art of The Performer*. Routledge, London, 1991a.

Eugenio Barba. Theatre anthropolgoy. In Eugenio Barba and Nicola Savarese, editors, *A Dictionary of Theatre Anthropology: The Secret Art of The Performer*. Routledge, London, 1991b.

Matthew Brand and Aaron Hertzmann. Style machines. *Proceedings of SIGGRAPH 2000*, pages 183–192, July 2000. ISBN 1-58113-208-5.

Armin Bruderlin and Lance Williams. Motion signal processing. *Proceedings of SIGGRAPH 95*, pages 97–104, August 1995. ISBN 0-201-84776-0. Held in Los Angeles, California.

Diane M. Chi, Monica Costa, Liwei Zhao, and Norman I. Badler. The emote model for effort and shape. *Proceedings of SIGGRAPH 2000*, pages 173–182, July 2000. ISBN 1-58113-208-5.

Jean Dorcy. *The Mime*. Robert Speller and Sons, Publishers, Inc., 1961. Translated by Robert Speeler, Jr. and Pierre de Fontnouvelle.

Sergei Eisenstein and Sergei Tretyakov. Expressive movement. In Alma Law and Mel Gordon, editors, *Meyerhold, Eisenstein and Biomechanics: Actor Training in Revolutionary Russia*. McFarland and Company, Inc., Publishers, Jefferson, North Carolina, 1996.

A. G. Feldman. Functional tuning of the nervous system with control of movement or maintenance of a steady posture - ii. controllable parameters of the muscles. *Biophysics*, 11(3):565–578, 1966.

Michael G. Hollars, Dan E. Rosenthal, and Michael A. Sherman. *SD/FAST User's Manual*. Symbolic Dynamics Inc., 1994.

Doris H. U. Kochanek and Richard H. Bartels. Interpolating splines with local tension, continuity, and bias control. *Computer Graphics (Proceedings of SIGGRAPH 84)*, 18(3):33–41, July 1984. Held in Minneapolis, Minnesota.

Rudolf Laban. *The Mastery of Movement*. Northcote House, London, fourth edition, 1988. Revised by Lisa Ullman.

John Lasseter. Principles of traditional animation applied to 3d computer animation. *Computer Graphics (Proceedings of SIGGRAPH 87)*, 21(4):35–44, July 1987. Held in Anaheim, California.

Joan Lawson. *Mime: The Theory and Practice of Expressive Gesture With a Description of its Historical Development*. Sir Isaac Pitma and Sons Ltd., London, 1957. Drawings by Peter Revitt.

Michael Neff and Eugene Fiume. Modeling tension and relaxation for computer animation. In *ACM SIGGRAPH Symposium on Computer Animation*, pages 81–88, July 2002.

Michael Neff and Eugene Fiume. Aesthetic edits for character animation. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 239–244, July 2003.

Ken Perlin. Real time responsive animation with personality. *IEEE Transactions on Visualization and Computer Graphics*, 1(1):5–15, March 1995. ISSN 1077-2626.

Ken Perlin and Athomas Goldberg. Improv: A system for scripting interactive actors in virtual worlds. *Proceedings of SIGGRAPH 96*, pages 205–216, August 1996. ISBN 0-201-94800-1. Held in New Orleans, Louisiana.

Katherine Pullen and Christoph Bregler. Motion capture assisted animation: Texturing and synthesis. *ACM Transactions on Graphics*, 21(3):501–508, July 2002.

Charles Rose, Michael F. Cohen, and Bobby Bodenheimer. Verbs and adverbs: Multidimensional motion interpolation. *IEEE Computer Graphics and Applications*, 18(5):32–40, September - October 1998. ISSN 0272-1716.

Ted Shawn. *Every Little Movement: A Book about Francois Delsarte*. Dance Horizons, Inc., New York, second revised edition, 1963.

Jennifer Tarver and Kate Bligh. The physical art of the performer, November 1999. Workshop: A 30 hr. intensive introduction to Laban and Grotowski held at the Nightwood Theatre studio, Toronto.

George Taylor. Francois delsarte: A codification of nineteenth-century acting. *Theatre Research International*, 24(1):71–81, 1999.

Frank Thomas. The future of character animation by computer. In Bill Kroyer and Phillipe Bergeron, editors, *3D Character Animation By Computer*, volume 5 of *Siggraph 87 Course Notes*. ACM Siggraph, 1987.

Frank Thomas and Ollie Johnston. *The Illusion of Life: Disney Animation*. Abbeville Press, New York, 1981.

Bill Tomlinson, Marc Downie, Matt Berlin, Jesse Gray, Derek Lyons, Jennie Cochran, and Bruce Blumberg. Leashing the alphawolves: Mixing user direction with autonomous emotion in a pack of semi-autonomous virtual characters. In *ACM SIGGRAPH Symposium on Computer Animation*, pages 7–14, July 2002.

Munetoshi Unuma, Ken Anjyo, and Ryozo Takeuchi. Fourier principles for emotion-based human figure animation. *Proceedings of SIGGRAPH 95*, pages 91–96, August 1995. ISBN 0-201-84776-0. Held in Los Angeles, California.