

Methods for Exploring Expressive Stance

Michael Neff*, Eugene Fiume

Department of Computer Science, University of Toronto

Abstract

The postures a character adopts over time are a key expressive aspect of its movement. While IK tools help a character achieve positioning constraints, there are few tools that help an animator with the expressive aspects of a character's poses. Three aspects are required in good pose design: achieving a set of world space constraints, finding a body shape that reflects the character's inner state and personality, and making adjustments to balance that act to strengthen the pose and also maintain realism. This is routinely done in the performing arts, but is uncommon in computer graphics. Our system combines all three components within a single *body shape solver*. The system combines feedback based balance control with a hybrid IK system that utilizes optimization based and analytic IK components. The IK system has been carefully designed to allow direct control over various aesthetically important aspects of body shape, such as the type of curve in the spine and the relationship between the collar bones. The system allows for both low-level control and for higher level *shape sets* to be defined and used. Shape sets allow an animator to use a single scalar to vary a character's pose within a specified shape class, providing an intuitive parameterization of a posture. Changing a shape set allows an animator to quickly experiment with different posture options for a movement sequence, supporting rapid exploration of the aesthetic space.

Key words: computer animation, character animation, human modeling, inverse kinematics, stance, expressive control, style, animation tools

1 Introduction

When we first meet someone, one of the most significant indicators of their character is how they stand. A person's posture or stance provides a strong

* Corresponding author.

Email addresses: `neff@dgpc.toronto.edu` (Michael Neff),
`elf@dgpc.toronto.edu` (Eugene Fiume).




Fig. 1. Four different reach postures satisfying the same constraint. The last three use shape sets to modify the expressiveness of the stance.

guide to whether the person is energetic or lethargic, shy or arrogant, nervous or relaxed. It gives us a powerful first impression that is often stressed in drama, dance and literature.

While posture is often seen as a reflection of a character’s inner state, it is also an indicator of how a character feels about its environment. We often show our disdain for or interest in an object by how we adjust our stance relative to that object. Put a large plate of brussel sprouts in front of an average eight year old and watch him recoil from the plate. This is an unambiguous expression of his feelings toward brussel sprouts. Similarly, when someone takes interest in an object, he will often orient himself toward it and lean closer.

Stance adjustments also play a key role when we reach out to touch objects. Again, the posture adjustment accompanying the reaching motion can indicate how the character feels about the object she is reaching for. If a character is given something she really does not want, she will often pull back from it with her body, even while reaching out to take it with her hand. A sultry, sensual person will move her body in a completely different manner while reaching for an object than a tired, nervous, old man. We are in fact quite attuned to the motions performed by one character that mock or parody the stereotypical motions of another.

Such posture adjustments are not generally the main actions that define a movement sequence. They are, however, often the movements that convey the most information and clarify the intentions of a character. They tell us how the character feels about himself in relation to other objects and people around him, and indicate his thoughts.

This work takes an integrated approach to pose modeling, combining balance adjustments, world space reach constraints and aesthetic body shape constraints in a single *body shape solver*. We demonstrate the approach by focusing on the areas introduced above: posture changes of a free standing character, posture changes relative to world space objects, and posture changes accompanying reaching motions.

A significant focus of this work is to allow animators to easily explore an aesthetically meaningful posture space. To this end, we introduce both a low-level and a high-level interface. Animators can interact with our system by directly modifying the low level parameters of the shape solver. These parameters have been designed based on lessons from the arts literature to relate to key expressive aspects of posture. Alternately, they can work more efficiently at a high level of abstraction by selecting a *shape set*. Shape sets are procedurally defined and encapsulate a particular aesthetically meaningful range of poses. Example shape sets will be shown for recoil motions and a posture range. Shape sets make use of the parameters in the body shape solver and provide animators with a simple interface consisting of a scalar intensity value and possibly one or more world space constraints. Switching shape sets allows an animator to efficiently explore the range of expressive postures, and the animator can always switch to the low-level interface to fine tune a given posture. Shape sets are fully extensible and new shape sets can be developed as needed during a production.

This paper proceeds by first reviewing the relevant background from both the performance and the computer science literature and then presents our algorithm for solving body shape. This algorithm features feedback based balance adjustment and a pose solver that uses both analytic and optimization based IK routines. A set of shape classes is presented to illustrate how higher level interaction is achieved, followed by a review of our results and conclusion.

2 Background

2.1 Expressive Value of Stance

The controls built into our body shape solver are guided by which aspects of movement are expressively salient. We take our cue from the theatre and dance literature, which as art forms carry a deep cultural understanding of the relationship between body posture and meaning. We will review some of the most salient postural devices here.

Stance can be divided into two separate but coupled aspects: adjustments made to body shape and adjustments made to balance.

Balance is central to expressive movement. Barba suggests that “[t]he characteristic most common to actors and dancers from different cultures and times is the abandonment of daily balance in favour of a ‘precarious’ or extra-daily ‘balance’” [1, p.34]. These balance adjustments do not normally convey a particular meaning, but serve to intensify motions and emphasize a performer’s presence [1]. While computer animation algorithms often try to ensure a char-



Fig. 2. Venus de Milo

acter is well balanced, being near the edge of balance or slightly off balance actually greatly heightens the excitement of a performer's movements [2].

Body shape refers to both the overall contour of the body and more specific features, such as how a character holds her shoulders. Alberts suggests that posture is a combination of two components, the level of tension displayed in the body and the overall body position [3]. Body position includes standing, leaning, kneeling, sitting and lying down. He proposes the following posture scale: hunched, stooped, slumped, drooped, slouched, sagging, tired, relaxed, straight, upright, uptight, erect and over-erect (at attention).

If a person leans their torso away from an object, this indicates repulsion. Leaning it toward an object indicates attraction. There is an exception to this in a certain form of sensual admiration that has a character lean away from the object of interest and cock his head slightly back towards it. [4].

The part of the torso that a person habitually holds forward is a clear 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 [4]. This triad can be augmented by considering people who carry their head forward, normally indicating a mental or academic disposition.

In general, Laban suggests there are three principal components of trunk movement: rotational movement about the length of the spine; "pincer-like" curling from one or both ends of the trunk and "bulge-like" shifting of the central area of the trunk out of its regular position [2].

The *S curve* or *Beauty Line* involves the legs, torso and neck in making a large S curve with the entire body, generally in the coronal plane. It is a key pose in Indian dance and was also prominent in ancient Greek sculpture, seen perhaps most famously in the Venus de Milo (Figure 2) [5].

Delsarte suggests that raised shoulders act to strengthen an action[4], while “[t]he elbow approaches the body by reason of humility, and moves outward, away from the body, to express pride, arrogance, assertion of the will” [4]. Moving the shoulders forward or backward also plays an important role in closing or opening the chest.

Similar findings are reported in the experimental psychology literature which studies the relationship between body movements and emotions. For example, leaning towards a subject is viewed as positive and leaning back has negative connotations [6]. Contracted poses featuring hunched shoulders, a forward leaning trunk and sunken chest have a sad or negative connotation [6,7], whereas expansive motions are found with angry or happy emotions [7]. An erect posture seldom accompanies shame, sadness or boredom, which tend to feature more collapsed postures[8]. Lifting the shoulders is typical of joy and hot anger, and moving the shoulders forward is more frequent with disgust, despair and fear[8]. Body movement and postural activity is an important indicator of the quantity of emotion, as well as quality[8].

While it is reassuring that the psychology literature also recognizes the emotional impact of these movements, it is important to note that our system does not enforce particular bindings for particular emotions. We feel this choice is best left to the animator as deviations from the norm are often what make a character interesting. This freedom allows an animator to create characters with behavioural idiosyncrasies or characters that “go against type”.

2.2 Computational Models

Research on inverse kinematics (IK) has been a significant focus in character animation for at least twenty years. Space limitations preclude a representative selection from the vast literature on this topic, so we offer a brief summary instead. The *IK problem* is to solve for a set of joint angles that will allow a kinematic chain to reach a given point in world space. Three main approaches have emerged to solve the IK problem. The first approach relates an instantaneous change in joint angles to an instantaneous change in the position of the end effector: $\dot{x} = J(\theta)\dot{\theta}$ using the Jacobian J . This relationship can be inverted by calculating the pseudo-inverse of J . By locally linearizing this equation and taking a series of small steps, the joint angles can be solved for. An example of this approach is [9]. The second approach to IK uses optimization to solve for the joint angles in order to satisfy the active constraints in the system (e.g., [10]). The third approach uses a dense sampling of motions and interpolation to achieve desired constraints (e.g., [11] and with reaching motions [12]). A variation on this approach uses learning techniques to generate poses close to those in a sample set[13]. We use a hybrid analytic and optimization approach. Lee and Shin [14] also use a hybrid IK solver, but their solver is not designed

to meet the same aesthetic constraints.

While inverse kinematics tools are useful for positioning a character to achieve a world space constraint, they provide only a limited ability to explore the widely varied set of expressively interesting postures available while meeting that constraint. Many IK systems do support secondary objectives, but these are often simple objectives, aimed at ensuring that the motion appears ‘natural’. One common secondary objective is to maintain joint angles near a desired value, often their centre [9,15]. Interpolation methods will generate the same posture as is in the sample data and it is difficult to vary substantially from this. Our work encourages an animator to actively explore the expressive range of posture. By selecting different *shape sets*, an animator can quickly experiment with various postures that each convey a different meaning, rather than being confined to a single “natural” pose.

A number of systems have explored balance control. Phillips and Badler use an optimization based IK system to allow animators to interactively control the balance point of a character[16]. Wooten uses a feedback method to control dynamic balance which is inspirational in our work[17].

Boulic et al. present an impressive system that combines inverse kinetics for balance control and inverse kinematics for reach constraints[18]. Their solutions take the form of the pseudo-inverse strategy for inverse kinematics as discussed above. They cascade the solutions, using the IK solution as a constraint in the null space of the inverse kinetics solution. A similar approach is taken by Aydin and Nakjima [19], who extend the work with a force distribution scheme to deal with different centres of support. Baerlocher and Boulic present an IK system that supports constraints at an arbitrary number of priority levels[20]. Position and orientation constraints are supported for arbitrary end effectors, including the centre of mass. A look-at orientation constraint is also supported.

Recent balance work has used the *Zero Moment Point* (ZMP), a measure of dynamic balance, to correct balance errors in motion sequences, ensuring that balance is always maintained [21,22].

Becheiraz and Thalmann present a system in which autonomous characters interact by changing their posture in response to other character’s posture changes.[23]

Other work has focused on aesthetic aspects of motion. Most closely related to our problem, the EMOTE system presented by Chi et al. includes a shape component which varies torso shape by interpolating between key poses in the horizontal, vertical and sagittal planes[24]. Other systems also make expressive adjustments to character pose, although this may not be their main focus. For example, the *Verb and Adverb* system by Rose et al. samples motions

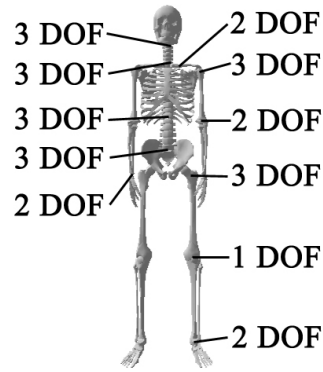


Fig. 3. The degrees of freedom for the full skeleton.

done in a number of expressively different ways and then interpolates between samples to generate a continuous range of expressive variation[25]. Brand and Hertzmann employ machine learning to acquire a movement style and apply it to other movement samples[26]. Pullen and Bregler use the high frequency components of a sample motion to “texture” another piece of motion[27].

3 System Overview

Animators can use the shape algorithms described here to solve for character poses that are used to define an animation sequence. An animator can interact with the system by either using an interactive interface that exposes all the low level parameters available in the system, or by invoking a shape set, which encapsulates the low level parameters. In the latter case, the animator simply selects a shape set of interest and provides an intensity value and possibly world space object locations. The low level interface provides maximal control and is useful if the animator wants to fine tune a motion or design a new shape set. Shape sets, on the other hand, are very fast to use and allow an animator to quickly explore a broad range of stance possibilities. Both shape sets and the low level parameters provided in the system are described in detail below.

Our skeleton contains 48 degrees of freedom: six DOFs for the position and orientation of the root and 42 rotational DOFs for the joints, as illustrated in Figure 3. We use a right hand coordinate system in which Y is up, Z is forward and X increases to the left.

4 Balance Adjustment

Balance control is based on a feedback strategy and works in conjunction with the lower body IK algorithm described below. The support polygon is the convex hull of the monitor points attached to the character’s feet. The

balance algorithm attempts to keep the projection of the character's centre of mass (COM) onto the ground plane at a particular location within the support polygon. Provided the character's momentum is low, she will remain balanced as long as the COM projects into the support polygon. The desired location of the projection can be changed either by a shape set or directly by the animator, allowing for balance adjustments.

One foot is set to be the dominant foot and balance control is actuated through this foot. The dominant foot can be changed arbitrarily during a simulation, but should be the foot that remains on the ground if the character lifts a foot. The current X and Z error in the location of the COM projection is fed back into the Z and X axes of the ankle in order to minimize the error. The other joint angles in the lower body are then calculated by the IK algorithm in order to keep the other foot rooted.

A similar feedback approach has previously been used in control of dynamic characters by Wooten [17], but their work fed the error into the desired value of both ankles and both hips. If this approach is used to control the actual angles directly instead of the desired values, it could lead to errors in the lower body joint angles, depending on the width of the stance and location of the balance point. In the dynamic case, ground reaction forces will prevent floor penetration and maintain foot contact. When these values are used directly in kinematic control, however, errors in joint angles can cause foot sliding and floor penetration. These errors are avoided in our work by controlling only one ankle and using our IK system to solve for the remaining angles in the lower body to maintain the position of the other foot.

Feedback based balance control is used when generating an animation. As discussed below, calculating a pose makes use of an optimization routine that must update the character's balance. For this routine, a simple search is performed to directly find the correct X and Z ankle offsets rather than using the feedback routine which adjusts these values over multiple time steps.

5 Body Shape Solver

Our shape algorithm will solve for a character's pose whether or not reaching constraints have been specified. Foot placements are always maintained and world space targets can be specified for either or both wrists. If no wrist targets are specified, the system returns the pose that satisfies the character's shape constraints.

Our algorithm and implementation is specific to the human body. This allows us to construct the controls that we find most desirable for making expressive adjustments to human pose and also allows us to take advantage of our

knowledge of the human form. An expressive animation system for general creatures is a related but separate goal.

Our algorithm breaks the body into four parts which are dealt with separately: the lower body (legs and pelvis), torso, arms and head. This breakdown is the same as is used when analyzing movement in mime [28], and is based on the different aesthetic contributions made by each of the four zones. Analytic methods are used for the lower body, arms and head. An optimization method is used for the torso and to help meet reach constraints. Each of the four sections is described in detail below.

5.1 *Lower Body*

The input to the lower body algorithm is the location and orientation of each foot, desired knee angles, balance adjustments to the root ankle from the balance routine, and desired pelvis X , Y and Z orientations. Default values are used for any parameter that is not specified. The pelvic Y orientation is relative to the line that passes through the centre of the two ankles. Note that while the knee angles are part of the skeleton state, the rest of the values are constraints the algorithm will try to achieve. Balance point, stance width, pelvic orientation and knee bends are the key aesthetic parameters for the lower body and they are all directly controllable in this model.

In brief, the algorithm starts with the position and orientation of each foot and then calculates the centre of each of the two hip joints. Starting at the root ankle the root hip position is calculated based on the ankle adjustments from the balance routine and the requested knee angle. It then solves for the free hip position in order to keep the free ankle in position and achieve the other pelvic and knee constraints. The skeleton is then subsequently fit to these points, maintaining the foot orientations, which requires calculating the orientation of both legs in order to set the hip angles. A detailed explanation of the lower body IK algorithm is given in the appendix.

5.2 *Torso*

The torso is comprised of two components: the spine, including the neck, and the collar bones. The spine is parameterized with three values and the collar bones with two. We analyze spine shape in terms of deformations in the coronal plane, deformations in the sagittal plane and the amount of twist or transverse motion. These planes are shown in Figure 4. Sagittal and coronal configurations belong to one of five classes depending on the shape they generate: large S, large C, small S, small C and straight. The straight class indicates no deformation in that plane. Large S and C make the shape of the




Fig. 4. The coronal, sagittal and transverse planes are displayed as they relate to the human skeleton.




Fig. 5. Emphasized for illustrative purposes, the five coronal shape classes for the spine are shown from left to right: Large C, Small C, Large S, Small S, Straight. Additional weight shifts have not been added.

letter including the lower body. Small S and C make the shape totally within the spine. These are illustrated in Figure 5 for the coronal plane.

The configuration of the spine is determined by a coronal shape class and amplitude, a sagittal shape class and amplitude, and a transverse amplitude. Each shape class defines a relative weight for each DOF in the spine that is part of the affected plane. The final DOF value is calculated as the product of the weight and the amplitude. For example, the DOFs that control movement in the sagittal plane are calculated as follows if the LargeS shape class is active:

$$AbX = sagittalAmplitude * LargeSWeight_{AbX} \quad (1)$$

$$ChestX = sagittalAmplitude * LargeSWeight_{ChestX} \quad (2)$$

$$NeckX = sagittalAmplitude * LargeSWeight_{NeckX} \quad (3)$$

$$HeadX = sagittalAmplitude * LargeSWeight_{HeadX} \quad (4)$$

The weight values are summarized in the appendix. For the transverse DOFs, the amplitude is evenly distributed amongst the spinal joints.

These shape classes provide sufficient control to achieve the spinal properties outlined in the background section, namely curling, shifting of the central torso, Albert's posture range, the beauty line and control over what part of the body is leading a pose. For instance, the large C class naturally lends to hunched behaviour, whereas the S classes can be used to thrust the chest out or generate the beauty line.

The collar bones consist of Y and Z rotations. In each of the Y and Z dimensions, the two collars are set to either align with or oppose each other. If they are aligned that means that both shoulders move in the same direction; opposed indicates that the shoulders move in opposite directions. A shrug in which both shoulders are raised is an example of an aligned motion. Collar movement is characterized by a Y and a Z amplitude and Y and Z states indicating whether the motion is aligned or opposed.

The collar movement works in conjunction with the spinal movement to generate the desired torso pose. Aligned backward movement of the collar bones opens up the torso, while moving the collars forward closes the chest. Aligned movement upwards raises the shoulders, indicating tension or apprehension. Conversely, downward movement is more relaxed or dejected. When the collars are opposed, they act as single rigid unit. This is particularly useful to emphasize a transverse twist in the spine by rotating in the same direction or to oppose the twist by acting in the opposite direction.

We also found it occasionally useful to allow an offset to be calculated for the X angle of the abdomen. This occurs when a particular amplitude gives the desired spinal shape, but due to the location of a world space object, you want to rotate the entire torso closer or further from the object.

An animator, either directly through an interface or by invoking a shape set, specifies which shape class to use for each dimension of the collarbones and the sagittal and coronal components of the spine. An optimization routine is then used to solve for the coronal, sagittal and transverse amplitude; the X and Z amplitude of the collar bones; and if it is included, the abdominal offset. The shape classes and the amplitudes operate together to determine the value of each DOF in the torso.

The optimization routine is necessary to resolve two types of conflicts: conflicts between desired internal body values and world space reach constraints; and conflicts over the desired X abdomen value, which is controlled by both the sagittal class and the abdominal offset. The latter of these is more significant and motivates the use of optimization.

A penalty function is defined for each amplitude using four cubic Hermite curves as shown in Figure 6. A desired value is specified for each parameter and two curves on either side of this value increase the penalty as the parameter




Fig. 6. The penalty function consists of desired point (shown by the filled circle that lies on the parameter value axis) and four cubic Hermite curve segments that increase the penalty as the parameter moves away from the desired point.

varies from this set point. The penalty functions have three desirable features: they are smooth; they can have flat basins when any value within a range is acceptable or steep sides when a certain value is desired; and they can easily be made asymmetrical which allows solutions on one side of the desired point to be favoured. An additional penalty function is included for each wrist constraint that is active. These penalty functions specify a desired distance from the shoulder of the reaching arm to the world space constraint. The objective function to be minimized is defined as the weighted sum of penalty functions f :

$$obj = \sum_i w_i f_i . \quad (5)$$

Additional penalty functions can be added to enforce other desired world space constraints. The penalty functions indicate how easily the different aspects of torso motion can be adjusted in order to achieve desired reach constraints, as well as specifying the importance of the reach constraints themselves and their desired accuracy.

We use a simple direction set optimization procedure that performs well for our application and does not require the computation of derivatives. A *line* minimization minimizes a one dimensional function. *Direction set methods* operate in n -dimensional space. They find a minimum by performing a sequence of line minimizations on n -dimensional direction vectors [29]. Different strategies can be used to determine the best set of directions. The dimensions of our space correspond to the amplitude terms discussed above. We use the simplest set of search directions: the unit vectors corresponding to each amplitude.

Our algorithm uses Brent's Method to perform line minimization. This uses inverse parabolic interpolation to speed the search for the minimum. Details and code are available in [29]. The main loop of the optimization procedure performs a line minimization for each amplitude term in sequence. This is done iteratively until the objective is within tolerance of zero or the objective stops improving, indicating a local minimum.

In our experience, the optimization procedure generally converges in a small number of iterations, rarely taking more than five. This indicates that the shape of the function being optimized is reasonably well aligned with the search directions. If this were not the case, an alternate set of directions for line search would perform better, but since our iteration count is low, we have not investigated more complicated search strategies.

5.3 *Arms*

If there is a wrist constraint on the arm, we solve for the elbow angle given the location of the constraint and the location of the shoulder. A rotation around the vector from the shoulder to the world space constraint is then applied to the arm triangle. This controls how far the elbow is from the chest, which as discussed is expressively important. The amount of rotation is parameterized and is either specified by the animator or the active shape set.

If there is no wrist constraint, we simply determine the elbow angle and rotate the arm out from the body based on a parameter provided by the animator or shape set.

5.4 *Head*

The key contribution of the head is to indicate gaze direction. When desired, we solve for the head orientation in order to have the character look at a specified world space point.

Gaze involves a combination of head movement and eye movement. It can be unnatural for a character to turn his head to completely align with a gaze direction when in real life the gaze would be achieved by a combination of head and eye movement. This is particularly true if the character is indifferent, casual or frightened. The system includes a weight that allows head position to be varied from the rest position to being fully aligned with the gaze direction.

6 **Shape Sets**

The controls in the body shape solver were carefully designed to allow animators to directly alter the most aesthetically important aspects of a character's pose. In practice, the interactive control palette allows an animator to arrive quite quickly at a desired pose. Nonetheless, there are still many significant parameters to specify and it is often beneficial to allow an animator to more

quickly explore the shape space. *Shape sets* are procedurally defined and encapsulate a particular class of body shapes. They take a small set of parameters, generally the name of the desired shape, an intensity value, and if required, one or more world space points that are needed to solve for the shape.

In our system, a particular pose is defined by a set of states and a set of numeric values. States include the items discussed above, such as the coronal class for the spine. Examples of numeric values include the coronal amplitude of the spine, the X dimension balance offset and the orientation of the elbow relative to the spine. The state set defines the interpolation space in which an artist is working. Different state sets lend themselves to different aesthetic meanings, providing a particular tunnel through pose space. Working within a shape set allows an animator to more quickly explore a particular pose range and switching shape sets allows an artist to more efficiently explore the total range of pose space.

The numeric values can be viewed as defining a shape vector. If all the numeric values are zero, for any state set, the character will stand in a neutral pose with his arms at his side. By interpolating from the origin to a given shape vector and then extrapolating beyond it, one obtains a continuous range of “intensities” of a pose. This is the basic idea behind the shape sets and the intensity parameter. Shape sets are based on one or more shape vectors and state sets and the intensity parameter interpolates between these. The interpolation can be between the origin and a given vector or it can be between two or more vectors.

Simple interpolation between vectors is not always sufficient, and for this reason shape sets are defined procedurally. Often certain aspects of shape will relate to the location of the world space object the character is reacting to. If a person is recoiling from a large spider, the direction of the recoil will depend on the location of the spider. By procedurally defining shape sets, the directional aspects of a pose can be varied based upon the world space constraints specified by an animator. In other circumstances, some aspects of shape will change at a different pace than others, making regular interpolation inappropriate. Consider a posture range from an over-erect army officer standing at attention, to someone badly hunched over with their knees slightly bent. There may be a continuous interpolation of the spine and collarbones between the two poses, but the legs will be straight over most of this range and the knee bend should only be introduced near the end of the interpolation range. This is straightforward to do with a procedural representation. Furthermore, shape sets facilitate the use of different interpolation functions for different parameters, allowing different parts of the body to change at different rates over an interpolation range. The interpolation functions can also be biased to emphasize a certain area within interpolation space.

Shape vectors have other uses as well. By making small perturbations to the state vector, the pose of the character can be varied while still keeping the character in the same general area of “shape space”. This makes it easy to have a character maintain a consistent set of poses without having the character always hold the *same* pose. Negating and reducing the amplitude of a shape vector allows a character to move backwards to wind up for a motion before moving forwards. This is an important anticipatory effect.

6.1 Building Shape Sets

To create a shape set, a user will normally begin by experimenting with the low-level interface. Through this, one or more appropriate sets of states and corresponding shape vectors will be determined. These are tied together procedurally in a shape set by specifying what interpolation functions will be used to move between them as the user varies the input parameters: an intensity scalar and possibly a world space location. The intensity value, denoted i , is normally defined over $[0, 1]$ or $[-1, 1]$.

As a rule of thumb, the shape vector components are typically divided into two categories: those parameters that have a world space dependency and those parameters that are primarily used for modifying the character’s posture. The latter are varied based on the intensity value and the former are varied based on the world space location, and possibly the intensity value as well. Parameters with a world space dependency are often primarily used for orienting the character. The pelvic Y and transverse twists, weight shifts and knee bends can all be used for orientation. Arm length, collar twists and in some cases, spinal deformation, can have a world space dependency as well.

Shape sets are designed to operate over a particular working range. For example, reach sets will generally operate well for reach targets that are a comfortable distance in front of or beside the character, but may not gracefully handle outlying target locations.

This discussion can be made more concrete by considering the example of the *keen reach* shape set. This shape set uses a short arm constraint in order to pull the character close to a given world space target and adjusts the character’s body to show intense interest. An example generated with this shape set is shown as the fourth posture in Figure 1. The generated poses involve a forward lean, so the working range of the shape set includes targets in a reasonable range in front of the character. An example set of poses generated with the shape set is shown in Figure 7, illustrating the shape sets coverage.

The parameters of this shape set are summarized in Table 1. Notice that the world space location is used only to determine the angle θ in the horizontal




Fig. 7. Examples poses for various reach locations generated with the *keen reach* shape set. The intensity for each pose is 0.7. The targets are shown as white boxes. The vertical spacing between targets in each row is 20 cm. The horizontal target spacing is -50 cm, -30 cm, 0 cm, 30 cm and 50 cm.

plane that the target makes with the forward vector. This in turn is used to modify the pelvic and spinal twists to orient the character towards the desired location. The intensity value will determine the desired shape vector. The final values of the shape vector are generated by the optimization process as it works to satisfy the reach constraint on the character's right arm. The reach constraint plays a significant role in deforming the torso. Since the gaze factor is 1, the character will look directly at the target. The impact of changing the intensity value is shown in Figure 8.

If a target is outside of the shape set's working range, errors can occur. Figure 9 shows two types of error. The system does not deal with collision detection and hence the character can interpenetrate itself for certain configurations. This is shown on the left for a reach target that is too close to the character. In the top row of Figure 7, the reach targets for the middle three examples were moved out to avoid this target. For the middle example, the target was moved 10 cm and it was moved 5cm for the examples to the right and the

Table 1

State sets and desired shape vector values for the *keen reach* shape set. The desired value for any parameter not shown is zero. $i \in [0, 1]$ is the intensity value. If the world space location is projected onto the horizontal plane with its origin at the centre of the character, θ is the angle that point is away from the axis coming directly forward from the character. All angles are in radians.

<i>Shape Set Parameters for Keen Reach</i>			
CorClass	LargeS	CorAmp	0
SagClass	LargeC	SagAmp	$3.7i$
TranAmp	$(5.7\theta + 0.15)i$	PelvisY	$(0.7\theta + 0.05)i$
RArmLength	$1.0 - (.7i)$	LArmLength	$1.0 - (.57i)$
GazeFactor	1.0	RArmAngle	$0.5i$
Collar Z	oppose	Collar Y	oppose





Fig. 8. The *keen reach* shape set is used with intensity values 0.1, 0.4 and 0.7 respectively. The reach target is shown as a white square.

left. The other targets in the figure are at the same depth. A second type of error occurs when the reach constraint is too difficult to satisfy and the optimization procedure forces the character into an unnatural pose in order to achieve it. This is shown on the right of Figure 9. This example is for the target in the lower left of Figure 7. The error was corrected in that pose by reducing the arm length constraint by ten percent. That was the only pose in the figure that required such an adjustment.

The optimization process will not automatically change the lower body posture. This was a design decision because it was felt that lower body posture changes are part of a higher decision process. Figure 10 shows the *keen reach* shape set being applied to a low target. By default, the character will remain standing and will bend over to reach the target. The shape set can be augmented to have the character automatically crouch for targets below a certain height. Both examples are shown in the figure. Note that neither option is necessarily right nor wrong, but the choice to crouch is likely related to the personality and health of the character.



(a) Interpenetration Error (b) Unnatural Pose Error

Fig. 9. This figure shows two potential types of error. Because there is no collision detection, interpenetration can occur. If the requested posture and reach constraint are not reasonable, unnatural poses may result.

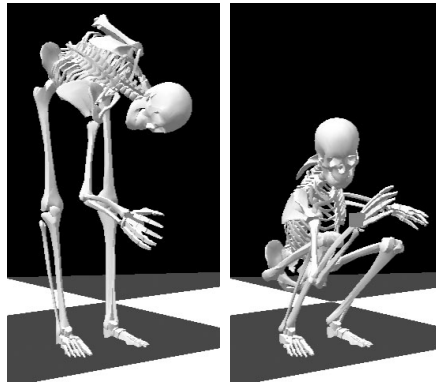


Fig. 10. The *keen reach* shape set is applied to a low target. By default, the character will remain standing. An augmented version of the shape set has the character automatically crouch for low targets.

Finally, it should be noted that at any point, the animator can flip back to the low-level interface to fine-tune the solution determined from the shape set. This allows shape sets to be used for rapid search, followed by a refinement phase to achieve the needed pose.

7 Results

All animations discussed in this paper are available online[30]. To emphasize the role played by shape, we have not changed the timing, succession, transition envelope or other expressively important aspects of the motion. A simple




Fig. 11. A continuous range of postures.

ease-in, ease-out curve is used for all transitions. A system that makes use of our shape solver and also deals with other expressive aspects of motion is described in [31].

To demonstrate the correctness of the balance algorithm and lower body IK routine, we have the character complete a complicated set of motions including a crouch, pelvic Y twist and balance adjustments. This sequence is performed with both a narrow and a wide stance.

The posture range proposed by Alberts is shown in Figure 11. The parameters that define the two extreme poses in the sequence are described in Table 2. The interpolation scheme will move from one parameter set to the other, passing through the neutral zero pose. State changes are made at the neutral pose. The spine's sagittal shape class changes from a SmallS, used for the over-erect posture, to a LargeC used for the hunch as the sagittal amplitude passes through zero. It is also here that the old man's bent knees and arms begin to be included in the interpolation. Similarly, for the beauty line posture displayed in the second frame of Figure 1, the knee that is bent changes as the amplitude passes from positive to negative. These changes are easy to effect with a procedural implementation of the interpolation scheme. Alberts' posture scale and the beauty line posture represent two posture shape sets. Other potential postures include a slovenly posture in which the pelvis is thrust forward and straight spine postures where the abdomen angle is varied, but the spine is held straight.

Four poses generated by different version of the recoil shape set are shown in Figure 12. These show different reactions to an object, represented by the red dot. The top two poses show different forms of recoil, indicating apprehension towards the object. The bottom two poses show interest in the object.

Table 2

The two extreme postures in Albert's postures range, denoted here as "soldier" and "old man". The desired value for any parameter not shown is zero.

<i>Extreme Posture Vectors</i>		
	Soldier	Old Man
SagClass	SmallS	LargeC
SagAmp	2.2	2.6
Collar Y	align	align
Collar Y Amp	-1.7	1.1
Collar Z	align	align
Collar Z Amp	0.9	-1.2
Ab X	-0.7	0
Arm Length	0 (straight)	0.08
Arm Angle	4 degrees	0
CentreOfMassZ	0.3	-0.2
Knee Bend	0	0.9




Fig. 12. Reactions to an object, indicated by the dot.




Fig. 13. Four different intensities of the strong recoil shape set.

It should be noted that the four poses are not simply linear interpolations within a pose space, but each is structurally different. The upper left recoil uses Large C shape classes for the coronal and sagittal planes, opposes the collar movement in both dimensions and adds a weight shift and pelvic Y rotation. The recoil in the upper right uses a Large S coronal class and a Small C sagittal class. The collars are opposed in the Y dimension, aligned in the Z dimension and the weight shift and pelvic Y rotation are increased. The lower left interest shape set uses Large S coronal and sagittal classes with aligned collars in each dimension. Finally, the lower left interest class uses a Large C for the coronal and sagittal shape classes and the collars are aligned in each dimension.

Four different intensities of the strong recoil (upper right in Figure 12) are shown in Figure 13. Table 3 summarizes the parameters that define the shape set. By selecting a shape set and providing a scale parameter, an animator can quickly experiment with both the type of reaction and the vehemence of that reaction.

Figure 1 shows four different reaching poses for the same target and initial character position. The first pose is the system's default solution. While it is reasonably "natural", and not surprising for an IK system, it communicates no information about how the character is feeling and hence does nothing to help the animator tell his story. The three subsequent poses are generated based on shape sets. They all feature different relationships with the object and convey more information about the character's state. This makes them

Table 3

State sets and desired shape vector values for the *strong recoil* shape set. The meaning of θ is the same as above. $f(\theta)$ is an interpolating function that biases the values away from the origin (currently $\pm\sqrt{fabs()}$).

<i>Shape Set Parameters for Strong Recoil</i>			
CorClass	LargeS	CorAmp	$3.1if(\theta)$
SagClass	SmallC	SagAmp	$2.7i$
TranAmp	$-3.9if(\theta)$	PelvisY	$0.7if(\theta)$
Left Knee	$0.55 - .05f(\theta)$	Right Knee	$0.55 + .05f(\theta)$
RArmLength	$1.0 - i(0.54 - .18f(\theta))$	CentreOfMassX	$-0.86if(\theta)$
LArmLength	$1.0 - i(0.54 - .18f(\theta))$	CentreOfMassZ	$-0.43i$
Collar Z	align	Collar Y	oppose
Collar Z amp	$2.4i$	Collar Y amp	$-1.7i$
Ab X	$-2.4i$		

far more useful in an animation sequence. It also illustrates how an animator can quickly explore pose space by switching shape sets.

Consider the second, “beauty-line” pose in Figure 1. It is unlikely that such a pose would be found with a traditional IK or optimization algorithm. The joint angles are intentionally and continuously varied from their central value. Furthermore, in a single kinematic chain, some limbs are rotated toward a goal while others are rotated away from it. While not “optimal”, such movements have a clear expressive impact.

8 Conclusion

This paper presents an integrated model for body shape, balance and reach constraints. The controls have been carefully chosen to align with the aesthetically salient aspects of pose. We have much to learn from the performing arts on how to communicate emotion and intent through variation in posture. The concept of shape sets provides an intuitive, high level interface for animators that is informed by the performing arts. One of the main hopes for this work is that it will make it easier for animators to actively explore the different postures, and different meanings, available when solving for a pose.

8.1 *Evaluation*

Experience with the system has revealed a number of strengths and weaknesses. The low-level controls were carefully chosen based on the arts literature to provide an aesthetically meaningful parameterization. In the practice of an experienced user, they appear to meet this task and allow desired poses to be rapidly created. The parameterization performs a degree of freedom reduction on pose space, simplifying the user's task. Shape sets offer a means for very rapid exploration and experimentation. Since shape sets are defined using the low-level parameters in the system, a pose arrived at through a shape set can be further refined by simply propagating its parameters into the low-level interface. The two level interface thus supports both exploration and refinement. The shape representation is extensible and customizable, allowing the tool to grow over time to better meet an individual animator's needs.

The disadvantages include the limited working range of shape sets and that the system does not prevent self-penetration. Some labour is required to create new shape sets. More complicated shape sets could be built to cover a very wide range of movements, but at the expense of additional labour. Occasionally, the weights in the optimization system need to be adjusted to achieve a desired pose. This is quite rare and arises when an arm constraint strongly opposes the body shape constraints.

User tests of the system are worthwhile and await completion of a full interface and access to a suitable group of test subjects. There are several potential user communities with whom evaluation would be beneficial. Experienced animators have a productive workflow and may resist tools that espouse a different, possibly disruptive, workflow. People who are new to animation but have a strong background in movement and the arts are a natural target for a system such as this. They are likely to understand the basis of the model and do not have the bias of an established workflow. For people who are both new to animation and new to understanding movement, the system may play a pedagogical role, helping them to learn which pose adjustments are meaningful. An ideal testing arrangement would occur over an extended production as users should both be involved in designing new shape sets and using existing shape sets and the low-level interface.

8.2 *Future Work*

We have only begun to explore the expressive uses of balance. Greater use could be made of changes in stance width and the greater range of balance adjustments used in dance would be interesting to explore. It would be worthwhile to extend our lower body IK routine to allow characters to go up on




Fig. 14. i) The location of the free foot and root hip are known and the location of the free hip must be solved for. ii) The initial estimate of the hip vector intersects the plane at the point marked with an X. This point is moved to the closest point on the circle of valid hip locations.

their toes. Adding in hand and other world space constraints would also be worthwhile.

The large number of matrix operations used in the lower body IK routine suggests it might be amenable to implementation in graphics hardware.

Acknowledgements

Financial support for this work was provided by Natural Sciences and Engineering Research Council of Canada. The software discussed in this paper was built on top of the DANCE framework, developed by Victor Ng-Thow-Hing and Petros Faloutsos. We would like to thank the anonymous reviewers for their comments; in particular, one reviewer provided useful pointers into the psychology literature.

Appendix: Lower Body IK Algorithm

The input to the lower body algorithm is the location of each foot, desired knee angles, balance adjustments to the root ankle from the balance routine, and desired pelvis X , Y and Z orientations. The goal is to calculate the left and right hip locations and all the angles in the lower body.

The desired angle of the root knee is taken as the actual value for this joint. The length of the vector from the root ankle to the root hip is solved for using the knee angle and thigh and shin lengths. The X and Z offsets provided by the balance algorithm are applied to this vector to determine the root hip location relative to the location of the root foot.

Consider the triangle shown in Figure 14 made up of the hip length (side B), the line from the free foot to the root hip (side A) and the line from the free foot to the free hip (side C). The location of the root hip and the free foot are known and we want to determine the location of the free hip. In general, it is not possible to satisfy the pelvic constraints and the knee constraint simultaneously. The animator can specify how these constraints are to be traded off. We will consider the case here where the algorithm tries to exactly achieve the knee angle and minimize the error in the requested pelvic values, which is the compromise we normally use in practice. Other compromises involve straightforward variations of this solution.

The free knee angle determines the distance between the free foot and free hip, or the length of side C. There is a plane that is perpendicular to the line A and passes through the free hip location. All potential values for the free hip lie on a circle in this plane. The circle is centred where line C penetrates the plane and its radius is the height of the triangle, taking C as the base.

To determine the free hip location, the desired pelvis Y and Z rotations are first applied to the hip vector. The vector is then (extended and) intersected with the plane. In extreme poses, it is possible for the hip vector to lie in or close to the plane. Intersection is not well behaved at this point, so we instead project the end point of the vector onto the plane and use this as the intersection point. The intersection point is moved to the closest point on the circle of possible values, as shown in Figure 14, and this is used as the location of the free hip. Once the free hip location is known, the angles of the free ankle can be calculated.

The final step is to determine the angles of the two 3-DOF hips. Given the frame of the root thigh, a rotation can be determined that will align the hip with the calculated location for the hip vector. This defines two degrees of freedom. The third degree of freedom is a rotation around the hip vector that controls the pelvic X rotation. We compute a view matrix that looks along the hip vector and has the pelvis's Y axis as its up vector. The pelvic X rotation is used to determine a desired up vector. This is projected onto a plane perpendicular to the hip vector and the projected location is used to determine the necessary rotation that must be applied around the hip vector. Quaternions have proved a useful representation for doing these calculations and the final result can be converted to whatever representation is necessary for the animation system.

The process for determining the free hip degrees of freedom is similar. The frame of the pelvis and the location of free thigh are now known so a rotation that will align the thigh with its desired location can be calculated. Again, this uses two degrees of freedom, the third being the amount of rotation around the thigh. This rotation is calculated in order to ensure that the shin and foot

Table 4
Sagittal Weights

<i>Sagittal Weights</i>				
<i>Class</i>	<i>Degree of Freedom</i>			
	AbX	ChestX	NeckX	HeadX
Large S	.15	-.1	-.12	-.1
Large C	.15	.15	.15	.1
Small S	.15	-.2	.06	-.1
Small C	-.15	.3	.15	.1
Straight	0	0	0	0

Table 5
Coronal Weights

<i>Coronal Weights</i>				
<i>Class</i>	<i>Degree of Freedom</i>			
	AbZ	ChestZ	NeckZ	HeadZ
Large S	.08	-.1	-.12	-.06
Large C	.15	.15	.12	.06
Small S	.15	-.2	.12	.06
Small C	-.15	.3	.12	.06
Straight	0	0	0	0

are correctly located. This is done using a variation of the projection technique above.

Appendix: Shape Class Weights

The coronal and sagittal shape classes specify a weight for each DOF of the spine in their respective planes. Final DOF values are obtained by multiplying the coronal and sagittal amplitudes by the weight factors associated with each DOF. The weight factors are summarized in Tables 4 and 5.

References

- [1] E. Barba, Theatre anthropology: First hypothesis, in: E. Barba, N. Savarese (Eds.), A Dictionary of Theatre Anthropology: The Secret Art of The

- Performer, Routledge, London, 1991.
- [2] R. Laban, *The Mastery of Movement*, 4th Edition, Northcote House, London, 1988, revised by Lisa Ullman.
 - [3] D. Alberts, *The Expressive Body: Physical Characterization for the Actor*, Heinemann, Portsmouth, N.H., 1997.
 - [4] T. Shawn, *Every Little Movement: A Book about Francois Delsarte*, second revised Edition, Dance Horizons, Inc., New York, 1963.
 - [5] E. Barba, N. Savarese, *A Dictionary of Theatre Anthropology: The Secret Art of The Performer*, Routledge, London, 1991.
 - [6] A. Mehrabian, Significance of posture and position in the communication of attitude and status relationships, *Psychological Bulletin* 71 (5).
 - [7] J. Montepare, E. Koff, D. Zeitchik, M. Albert, The use of body movements and gestures as cues to emotions in younger and older adults, *Journal of Nonverbal Behavior* 23 (2) (1999) 133–152.
 - [8] H. G. Wallbott, Bodily expression of emotion, *European Journal of Social Psychology* 28 (1998) 879–896.
 - [9] M. Girard, A. A. Maciejewski, Computational modeling for the computer animation of legged figures, *Computer Graphics (Proceedings of SIGGRAPH 85)* 19 (3) (July 1985) 263–270.
 - [10] J. Zhao, N. I. Badler, Inverse kinematics positioning using nonlinear programming for highly articulated figures, *ACM Transactions on Graphics* 13 (4) (1994) 313–336.
 - [11] C. F. Rose III, P.-P. J. Sloan, M. F. Cohen, Artist-directed inverse-kinematics using radial basis function interpolation, *Computer Graphics Forum* 20 (3) (2001) 239–250.
 - [12] D. Wiley, J. Hahn, Interpolation synthesis of articulated figure motion, *IEEE Computer Graphic and Applications* 17 (6) (1997) 39–45.
 - [13] K. Grochow, S. L. Martin, A. Hertzmann, Z. Popović, Style-based inverse kinematics, *ACM Transactions on Graphics* 23 (3) (2004) 522–531.
 - [14] J. Lee, S. Y. Shin, A hierarchical approach to interactive motion editing for human-like figures, in: *Proceedings of SIGGRAPH 99*, 1999, pp. 39–48.
 - [15] K. Yamane, Y. Nakamura, Natural motion animation through constraining and deconstraining at will, *IEEE Transactions on Visualization and Computer Graphics* 9 (3) (2003) 352–360.
 - [16] C. B. Phillips, N. I. Badler, Interactive behaviors for bipedal articulated figures, *Computer Graphics (Proceedings of SIGGRAPH 91)* 25 (4) (July 1991) 359–362.

- [17] W. L. Wooten, Simulation of leaping, tumbling, landing, and balancing humans, Ph.D. dissertation., Georgia Institute of Technology (1998).
- [18] R. Boulic, R. Mas-Sanso, D. Thalmann, Complex character positioning based on a compatible flow model of multiple supports, *IEEE Transactions on Visualization and Computer Graphics* 3 (3) (1997) 245–261.
- [19] Y. Aydin, M. Nakajima, Balance control and mass centre adjustment of articulated figures in interactive environments, *The Visual Computer* 15 (3) (1999) 113–123.
- [20] P. Baerlocher, R. Boulic, An inverse kinematics architecture enforcing an arbitrary number of strict priority levels, *The Visual Computer* 20 (6) (2004) 402–417.
- [21] S. Tak, O. young Song, H.-S. Ko, Motion balance filtering, *Computer Graphics Forum* 19 (3) (2000) 437–446.
- [22] M. Oshita, A. Makinouchi, A dynamic motion control technique for human-like articulated figures, *Computer Graphics Forum* 20 (3) (2001) 192–202.
- [23] P. Becheiraz, D. Thalmann, A model of nonverbal communication and interpersonal relationship between virtual actors, in: *Computer Animation 96*, IEEE Computer Society Press, 1996, pp. 58–67.
- [24] D. M. Chi, M. Costa, L. Zhao, N. I. Badler, The emote model for effort and shape, *Proceedings of SIGGRAPH 2000 (July 2000)* 173–182.
- [25] C. Rose, M. F. Cohen, B. Bodenheimer, Verbs and adverbs: Multidimensional motion interpolation, *IEEE Computer Graphics and Applications* 18 (5) (September - October 1998) 32–40.
- [26] M. Brand, A. Hertzmann, Style machines, *Proceedings of SIGGRAPH 2000 (July 2000)* 183–192.
- [27] K. Pullen, C. Bregler, Motion capture assisted animation: Texturing and synthesis, *ACM Transactions on Graphics* 21 (3) (2002) 501–508.
- [28] J. 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.
- [29] W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd Edition, Cambridge University Press, 1992.
- [30] M. Neff, <http://www.dgp.toronto.edu/people/neff>.
- [31] M. Neff, E. Fiume, Artistically based computer generation of expressive motion, in: *Proceedings of the AISB 2004 Symposium on Language, Speech and Gesture for Expressive Characters*, 2004, pp. 29–39.