[DRAFT] : Psychologically inspired anticipation and dynamic response for impacts to the head and upper body

Ronald Metoyer[†], Victor Zordan[‡], Benjamin Hermens[†], Chun-Chi Wu[‡], and Marc Soriano[‡]

Oregon State University[†]
University of California, Riverside[‡]

Abstract

We present a psychology-inspired approach for generating a character's anticipation of and response to an impending head or upper body impact. Protective anticipatory movement is built upon several actions that have been identified in the psychology literature as response mechanisms in monkeys and in humans. These actions are parameterized by a model of the approaching object (the threat) and are defined as procedural rules. We present a hybrid forward and inverse kinematic blending technique to guide the character to the pose that results from these rules while maintaining properties of a balanced posture as well as characteristics of the behavior just prior to the interaction. In our case, these characteristics are determined by a motion capture sequence. We combine our anticipation model with a physically-based dynamic response to produce animations where a character anticipates an impact before collision and reacts to the contact, physically, after the collision. We present a variety of examples including threats that vary in approach direction, size and speed.

Index Terms

motion capture, animation, procedural, inverse-kinematics, dynamics, physical simulation.

I. INTRODUCTION

Responsiveness of characters is an important problem in computer animation and has been receiving increasing interest based on recent publication trends. This problem is challenging because pre-recorded or pre-generated motions do not afford environment-specific reactions that will lead to believable interactions primarily due to the large space of possible reactions. Thus, to make characters respond to a particular situation, a motion synthesis system requires a dedicated method for generating or modifying character behavior 'on the fly' based on the current conditions. In recent years, researchers have proposed several means for combining motion capture animation with physical simulation to produce motion that is compelling and lifelike while also physically correct. However, in much of this work and in many interactive games, the character generally appears to be oblivious to the impending contact. In contrast, we present a solution inspired by psychological findings for generating a humanlike response

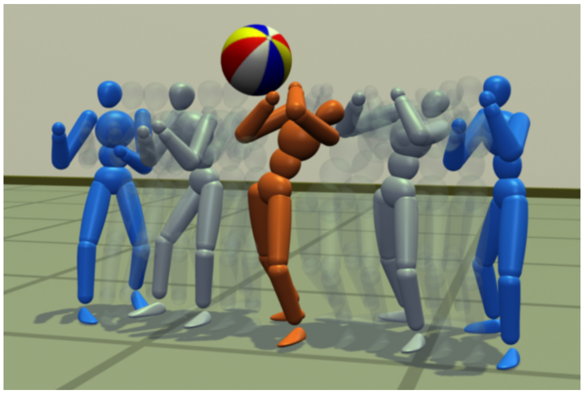

Fig. 1. An example anticipation pose. The blue poses represent motion capture frames, gray poses represent interpolation frames, and the orange pose represents the character's pose at impact.

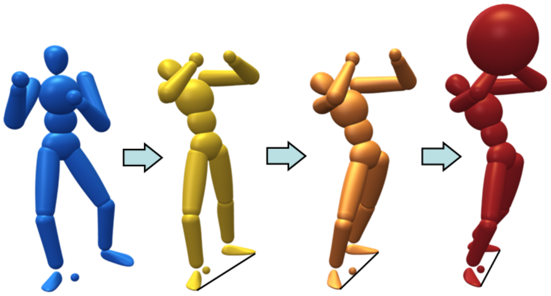

Fig. 2. The steps of our technique illustrated by snapshots taken at the instant of impact. From left to right: in blue, the pose from the original motion at the time of impact (note that the left foot is not planted); in yellow, the result of executing the rules; in orange, the balanced final anticipation pose; and in red, the collision and dynamic simulation. The front edge of the support polygon and the ground plane projection of the COM are shown for balance comparison.

in anticipation of an approaching interaction - that is, the response that occurs *before* contact takes place - with the ultimate goal of creating a character that can anticipate an interaction as well as react to it physically. We focus on impacts to the head and upper body.

Our basic approach generates anticipatory response by interpolating from a motion capture sequence to a procedurally generated ‘anticipation pose’ and then subsequently back to motion capture animation (see Figure 1). The simplifying assumption that makes this reasonable is that a single, sustained pose encapsulates the desired anticipatory response and, as we will see, this assumption is supported in part by several psychological studies on the subject. The first step in the process is to compute this pose that anticipates the impending impact of the object we call *the threat*. We model the threat with a small set of controllable parameters that guide the characteristics of the anticipation pose. We derive these characteristics from behaviors described in the psychology literature and encode them as rules that change the character’s posture based on the threat parameters. A sample anticipation pose is shown in yellow in Figure 2. Modifying a motion capture pose using the anticipation rules can leave the character in an unbalanced state, therefore in the next step, our system optimizes this out-of-balance pose to re-balance the character without “undoing” the results of the behavior rules. This optimization procedure produces the *sustained pose* shown in orange in Figure 2.

Once we have generated the pose that the character should achieve to anticipate the incoming threat, we synthesize the anticipation response movement through a unique blending process that brings the character from the original motion pose to the sustained pose. We found empirically that more visually pleasing anticipation is created when properties such as balance and *point-to-point* hand trajectories are incorporated into the motion synthesis step. We perform our blend over a time interval that puts the character into the anticipation pose ‘just in time’ for the impact. Then, at the time of impact, we create a dynamic response using the approach described by Zordan et. al [34], shown in red in Figure 2. The physically based reaction acts as a plausible response over a small time period following the impact after which the system blends back to motion capture data. Our basic technique is augmented to include the ability to generate an additional *startle* response that reveals a second form of anticipation action. This reflex-based response follows quickly after the instant in time when the threat is recognized and precedes the sustained anticipation response described. The startle is computed in a manner similar to that of the sustained response, but its timing and characteristics are consistent with the description of startle behavior described in the psychology literature.

This paper presents three primary contributions:

(1) We introduce a set of primitive behaviors that have been gleaned from the field of psychology and that are present when animals (specifically monkeys) and humans attempt to anticipate and avoid an impending impact to the front or side of the head.

(2) We propose a parameterization for these behaviors with respect to the approaching threat and describe an implementation of them as procedural rules for achieving a protective response that includes defensive behaviors such as blocking and turning away. We combine these anticipatory behaviors and a dynamic response within a practical hybrid forward kinematic, inverse kinematic, and dynamic simulation framework for creating movement leading up to, during, and after an impact.

(3) The final contribution of this paper lies in the generalization of the rules to impacts to regions other than the front and sides of the head. We discuss how the intuition for each defensive behavior can be applied to produce anticipatory responses for impacts to the rest of the body and we demonstrate by extending the rules to anticipate impacts to the back of the head and along the trunk of the body.

II. RELATED WORK

Our approach builds on findings from the psychology field as well as from techniques originating in the field of computer animation. We describe the background in computer animation first before highlighting the relevant findings in psychology that motivate our approach.

A. Computer Animation

Our goal of generating character response to the environment is shared by several researchers [7], [24], [2], [16] (among others). The approaches described for generating the response to an interaction take into account the physical components related to the impact, either in the form of a simulated collision [7], [33], [22], [34] or by directly modifying the dynamic parameters of the motion such as joint velocities [24], [2] or momentum [17], [16]. The result of creating these kinds of changes is character motion that gives the impression of responding *physically* to the impact. In addition, recently, researchers have coupled these changes with transitions to new motion sequences following the impact in order to capture more *stylized* and more *complete* response behaviors [22], [2], [16], [34]. In a similar fashion, we include both a simulated reaction and the potential for transitioning to a new motion behavior, after the impact if the conditions merit. However, in comparison to all of these previous examples, our effort is unique in that we emphasize the preparation leg of an interaction. To the best of our knowledge, ours is the first work which addresses anticipation response in a thorough and explicit manner.

In order to generate believable anticipatory movement, we edit the character's motion leading up to the impact while taking balance into account. Motion capture editing and procedural motion generation have been hot topics for animation researchers in the last decade, leading to a wide variety of novel approaches. Specifically, we classify our method as a constraint-based motion editing solution applied to a single motion capture sequence. An abundance of other work in character motion synthesis has focused on similar techniques with various goals and results (*eg.*[5], [31], [11], [19], [27], [1], [13], [21]). See the survey by Gleicher [12] for a more thorough discussion. Within this framework, we describe our technique as employing specialized motion blending between motion capture sequences and procedurally generated poses for the dedicated purpose of responding to impending impacts. What makes our technique stand out in this context is both our focus on the creation of a responsive pose and our motion blending approach for achieving that pose which takes into account specific constraints while blending, such as point-to-point hand movement.

For believable movement, a character must remain balanced. As such, we edit the character's motion leading up to the impact while taking balance into account. Similar to our approach, a small group of researchers have presented methods for maintaining or correcting balance by directly modifying the motion data [3], [4], [30], [23], [17]. The advantage of these techniques over physically based ones is that they do not require a dynamic balance controller, which is generally more complex than a kinematic balancing approach such as the one we describe. Instead they provide a means for controlling the appearance of balance by ensuring that computed parameters such as the center of mass (COM) and zero moment point (ZMP) remain plausible. The balance problem is then generally posed as a constrained optimization problem which can be solved using an inverse kinematic-like approach that takes into account the mass displaced while adjusting joint angles [3], [4], [23] or by making adjustments to the entire body to correct balance while minimizing changes to the original motion [30], [17]. These two types of solutions also reveal a clean division between controlling balance with the COM or the ZMP. We found through experiments that the ZMP, as defined in [30], and the COM produce very similar values across the examples in our database of standing responses to impact. Therefore, for ease of implementation and to maximize our control over the process, we use the COM to adjust balance. The details of our specific method are described in Section IV, but we note that, in contrast to previous techniques, *our approach includes a means to stabilize for balance by planting the non-support foot* if it is not already in contact with the ground.

An additional area in computer animation is highlighted in the procedural animation research of Neff and Fiume [23] and Perlin et al. [25], [26]. This work is similar to our own in that it focuses on procedural animation that is derived from combinations of rules. The power of creating complex character poses from simple components is appealing because many animations can be produced once the rule set is in place. For example, Neff combines a reach rule with a parameter that ranges from interest to repulsion and that is encapsulated in the stance of the character. We note that developing rules does require a (somewhat) skilled animator or valid biomechanical and psychological data to create good primitive components. However, once these rules are in place, the ways in which they can be combined make make the benefit worth the overhead. In comparison to previous work, there are specific details that make our work unique. First, we focus on a multi-layered approach specifically for response and, while this testbed is focused, we incorporate it into a complete response system for pre and post impact response. Also, we ease the generation process by building a parametric model of the threat and control the response indirectly by changing the threat's behavior. Finally, we focus on reflexive and low-level cognitive behaviors surrounding anticipatory response, and thus draw from psychological studies associated with such activities. We describe our findings from the related psychology literature next.

B. Anticipation in Psychology

Researchers in psychology have studied defensive movements for many years. Generally, psychologists are interested in understanding the brain control mechanisms associated with such motions. However, the studies involving both monkeys and humans also describe typical movements associated with defensive behavior. In recent work, Cooke et al. performed experiments on monkeys to compare the defensive reactions caused by electrical stimulation to cortical areas to those caused by an external air puff [6]. The 'threat' was simulated as a puff or burst of air directed at the monkey. The authors qualitatively describe six movements that were evoked by the air puff. Among these six motions were three that are particularly interesting when preparing the entire body for an impact: shoulder shrug, retraction of the head from the direction of impact, and arm movements which depend on the location of the threat. These movements include bringing the hand to the area near the head when the puff was directed at the head and bringing the elbow to the torso when the puff was directed at the side of the torso. In the same study Cooke and colleagues observe that responses include an initial spike of muscle activity and then a sustained muscle response. The initial spike is attributed to a startle reflex which is relatively insensitive to stimulus type [18], [32]. The sustained response, on the other hand, is comprised of a set of spatially directed movements and can involve turning the head, ducking, withdrawing from the direction of the stimulus, navigational veering during locomotion, and blocking the object with a body part like the forearm to protect another part such as the ribcage [14], [15], [18], [29].

In 1992, King et. al studied adult human defensive movements in response to looming visual stimuli [15]. In their experiment, participants were distracted with the task of playing a video game. While playing the game, the participants were presented with a visual stimulus approaching from the side of the head. The authors describe several defensive movements including head and upper body motion as well as raising of the hands, blinking and bracing of the shoulders. They also present latencies associated with the defensive movements and we utilize this data for the timing associated with our defensive rules.

In 2001, Li and Laurent reported findings of their study which involved understanding the response of individuals to an impending collision to the head from a ball with various velocities and eccentricities [20]. Of particular interest in this paper are their findings on how angle and velocity affect the speed of the participant's movement. They define a *time to collision* as the time between initial reaction movement and when the ball hits (or would hit) the subject. They found that this time gets larger with eccentricity, which is defined as the angle between the subject's gaze and the stimulus (threat). They attribute this fact to the observation that a subject underestimates the time to collision at large angles and therefore reacts sooner. This may stem from a natural survival mechanism since our vision in the periphery is not as accurate and therefore underestimating and reacting quickly is the conservative response. They also found that defensive movement speed does not vary with angle, however, it does vary with velocity of the ball. The velocity of the ball affected both the maximum movement speed of the participants as well as how quickly the participants reached that maximum speed.

For our purposes, we extract from these findings a set of rules for creating movements in anticipation of impacts. Based on our survey, our implementation involves rules for recoil (shoulder shrug and arm retraction), blocking (the stimulus with the hand), and leaning and turning away (to withdraw and protect from the direction of the stimulus). To produce a response, we bind the rules to appropriate parameters of the approaching threat. And, we include two distinct phases in our anticipation reaction: startle and a sustained response.

III. PSYCHOLOGY INSPIRED DEFENSIVE RULES

Our defensive rules are derived from the reaction behaviors described in the previous section. In particular, we develop primitives for anticipation, broken down into rules for leaning and turning away, as well as recoiling from, and blocking an impending impact. These rules modify the character's pose based on attributes of the threat. The behaviors described in the literature survey mostly refer to impacts to the front and sides of the head. Therefore, we define the defensive rules with respect to impacts to the same areas. However, in Section VI we discuss and provide examples of how the rules generalize to impacts to other areas of the body.

A. Threat Parameterization

In order to generate a reaction to an impending collision, we define a model of the threat that is about to collide with the character. Our model is comprised of a few parameters that allow an animator to efficiently choose the desired threat and elicit the appropriate response. We describe the threat model with the following parameters:

- the angle of approach
- the speed
- the relative size

The angle and speed can be interpreted as literal descriptors of the threat. To aid the animator, we parameterize relative speed using a slider over a range of 23 to 30 $\frac{m}{sec}$ which roughly corresponds to the range of speeds of a spiked volleyball [9]. The size parameter is an abstract quantity that refers to either the volume, the mass, or both. We allow the animator to choose the velocity and position of the threat *at* the time of the collision, however, an anticipation response requires knowledge of the threat some time *before* the collision. One approach would be to initialize the threat's position and velocity at the time of recognition directly, however, this would require that the animator place the threat very carefully. Rather, we ask the animator to specify the amount of time before the impact that the threat is to be recognized, within a limited range. In this fashion, we control the time until impact

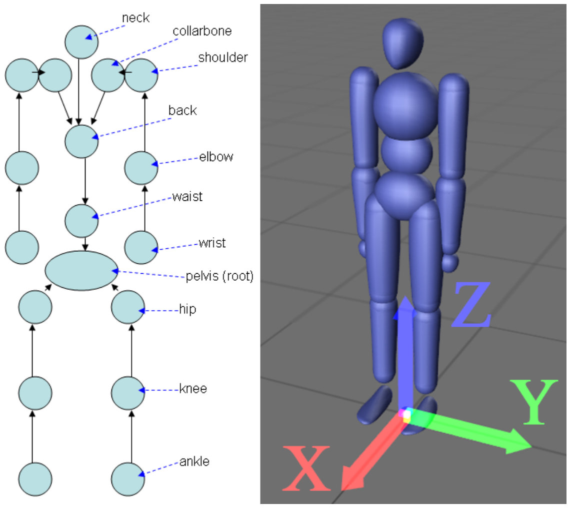

Fig. 3. *left:* The hierarchy of joints forming the character's skeleton. *right:* The character's identity pose and the global coordinate frame.

to have values consistent with the literature. From this time duration, the system uses a simple ballistic prediction to compute the starting conditions of the threat. Specific details on timing are summarized in Section V-B.

B. Character Model and Posture

Our character consists of a hierarchical skeleton rooted at the pelvis and a right-handed coordinate frame with the Z-axis up and the X-axis in the facing direction of the character. The elbows and knees are modeled as hinge joints while all other joints are ball and socket joints (see Figure 3).

The character's 'origination pose' is taken from the motion capture sequence based on the selected recognition time (described later). The orientation of the root body and the joint angles of the upper body are to be defined by the defensive rules but, first, before the rules are evaluated, the character is put into a stable posture with both feet planted on the ground. If only one foot is on the ground, the support polygon is very small and the character could easily be knocked over. If only one foot is planted on the ground in the origination pose, we look ahead some small time (10 frames, sampled at $120\ hz$) in the animation and move the root and unplanted foot according to the new data. This serves not to find a location to plant the foot, but to maintain the dynamics of the foot by moving it according to the motion capture data. The unplanted foot is then placed at the projected (X,Y) ground location using inverse kinematics and this becomes the new origination pose. This approach works well in practice when the foot is moving in a horizontal or downward direction. However, if the character was in a deliberate upwards motion on one foot , such as a martial arts kick, our technique would not produce acceptable motion. This situation would require planning to attempt to bring the foot back down to a stable location before the impending impact.

C. Rule Definitions

Given the threat model parameters and the origination pose, several anticipation rules compute joint angles for the *sustained response* pose. We present the rules in the order that they are applied. *Lean Away*, *Turn Away*, and

Lean and Turn Away			
Joint	*X*	*Y*	*Z*
neck	50	50	65
back	25	25	8
waist	10	10	12
root	15	15	20

Recoil and Startle			
Joint	*X*	*Y*	*Z*
clavicle	35	25	0
shoulder	30	0	0
elbow	0	150	0

TABLE I

MAXIMUM ROTATION VALUES (DEGREES). THE RULES CHOOSE ANGLES BETWEEN 0 AND THE MAXIMUM VALUE TO ROTATE THE JOINTS.

Recoil must be applied before *Block* because *Block* places the hands in a position with respect to the threat based on the pose generated by the other rules. If we were to evaluate *Block* first, for instance, then *Lean Away* would move the final hand positions and invalidate the block.

1) Lean Away: Cooke et al. report that during sustained response, monkeys withdrew their heads from the direction of the stimuli. The lean away rule causes the character to move the head away from the threat by bending along the spine. Intuitively, this rule serves to not only attempt to avoid the object but to increase the distance the threat must travel before impact and therefore increase the amount of time available to prepare for and block the object. Note that the lean rule could also have been implemented to lean in a direction perpendicular to the threat path in order to avoid the collision by creating a ducking or dodging effect.

To determine the magnitude for the lean, we rely on the assumption that a slow moving threat should cause the character to withdraw less than a faster moving threat. Therefore, we compute a lean control term defined as $threatSpeed/(maxThreatSpeed - minThreatSpeed)$ and scale the maximum lean angle (see Table I) by this control value to determine the amount of lean. Then, we compute a vector that is perpendicular to the threat's velocity and in the horizontal plane and use this as the lean rotation axis. The resulting lean angle is distributed among the pelvis (root) orientation and the waist and back joints at 30%, 20% and 50% respectively (see Figure 4, second row). This distribution was chosen empirically by an animator and can be controlled to suit the animator's needs. Changing these numbers will result in the same total lean, however, the distribution to the joints involved in the total lean will be modified based on the values chosen by the animator.

2) Turn Away: After applying the *Lean* rule, we employ the turn away rule to make the character rotate about the spine with the goal of orienting the back of its head toward the threat. Intuitively, it is less painful and less damaging to be struck in the back of the head than in the face. King et al.[15] describe this behavior as generally ranging from a small contralateral turn of the head to a large contralateral pivoted movement involving both the head and upper body.

To implement this rule, we begin by finding the angle necessary to meet the goal of facing the back of the head toward the threat. We then rotate the neck up to its maximum (Z) angle (see Table I) and rotate the torso joints according to the remaining angle to align the back of the head with the threat. The torso rotation, if required, is distributed among the back, waist, and root body at 50%, 30%, and 20% respectively, up to the maximum angle values for each. As in ’lean away’, these values were determined empirically. The direction of the rotation depends on the threat’s approach; if the threat approaches from the left, the character turns right and vice versa. We assume that a larger object is perceived as more dangerous and therefore should require a more drastic response to ensure maximum protection from the threat and so we define a turn control parameter and initialize it to $threatSize/(maxThreatSize - minThreatSize)$. We scale all computed turn angles by the turn control parameter to scale down the amount of rotation for small threats (see Figure 4, third row).

3) Recoil: The recoil rule “buries” the head into the shoulders by rotating the clavicle joints upward. This action helps to protect the neck, a vital and sensitive area. We associate this protective reaction with the size of the threat, assuming recoil is most severe for large or heavy threats such as a medicine ball and less severe for a smaller or lighter threat such as a beach ball. The recoil rule should also make the character “smaller” by retracting the arms to the chest therefore providing a smaller target to a predator or colliding object. In our implementation, recoiling the shoulder and elbow joints would have no effect on the final pose because these joints are also controlled by the *Block* rule that is instantiated *after* the recoil rule. Thus, as we describe in the next section, the *Block* rule is designed to incorporate this aspect of recoil. Here, we define a recoil control term and initialize it using the threat size (as above). The system uses the value of this term to scale the angles for the clavicle joints between zero and their maximum angle (see Table I). The effect of the recoil rule can be seen in Figure 4 in the third row. For a more in-depth discussion of recoil behaviors, we refer the reader to the expressive stance work of Neff and Fiume [23].

4) Blocking: Our final rule implements the blocking component of protection which serves to deflect the impending impact. Researchers report that after the startle response, the arms generally move to block the impending threat [18], [29]. Cooke et al. [6] report that movement was primarily made toward a guarding/blocking posture where the hand protected the head from the location of the air puff directed at the monkeys. Therefore, we define a blocking posture for the arms as one in which the hands are positioned on the approach vector of the threat, between the threat and the body (see Figure 4, top row).

To customize characteristics of the blocking behavior, we rely upon the assumption that humans block smaller (less threatening) objects with their arms further away from their bodies than they would for larger objects, for which they draw the hands and arms in close to the body. Intuitively, an extended arm is less protective than an arm pulled in close to the body where it can guard more of the vital areas. This behavior is consistent with the recoil properties described from the previous rule. Further, one could argue that compared to full extension, a somewhat recoiled position is also a better blocking posture for absorbing the energy of heavy impacts. In addition, with the shoulder and wrist at fixed locations, the elbow is free to sweep out a circle on the plane with normal equal to the vector from the shoulder to the wrist. Thus, we also parameterize this “swivel” angle by the threat’s size with larger threats causing the elbows to pull in closer to the body and smaller threats leaving the elbows away from the body, thereby incorporating the same set of assumptions.

For blocking, we again define control terms, as described above, for the reach distance and swivel angle and compute each from the size of the threat. Here, the control value is used to scale the reach distance from near to far and the swivel angle from in(near the body) to out(away from the body) to produce the final reach and swivel values. Coupling these two blocking parameters produces a range of blocking responses (see Figure 4, third row).

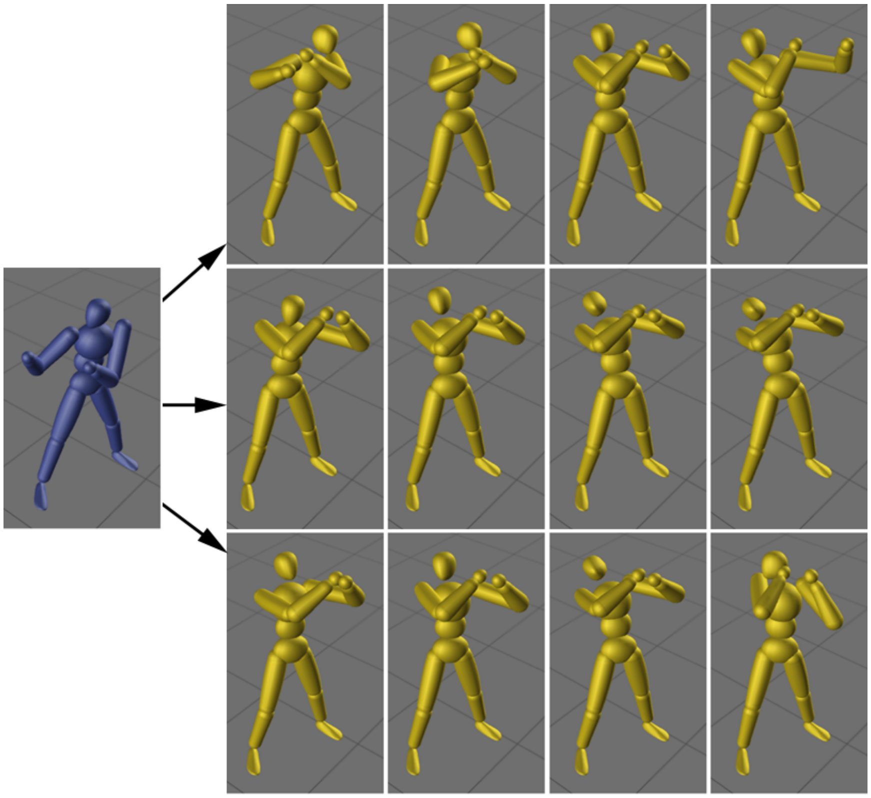

Fig. 4. Changes in anticipation as the threat parameters change. The leftmost image in the figure shows the origination pose with both feet planted on the ground. The three rows to the right represent poses generated by the rules as threat properties change, one per row. The top row shows the threat's position changing from the character's right to left. The middle row shows the threat's speed varying from slow to fast. The bottom row shows the size varying from small to large.

D. Startle

The rules as described define a protective posture referred to as the *sustained pose*. Several researchers [6], [18], [32] have also reported a startle response that occurs before the sustained response. It is described as a short latency, bilaterally symmetric response to intense stimuli that is thought to put the body in an initial protective pose. The startle response affects the shoulders, head, arms, lips, ears, and eyes in monkeys. Since we are not concerned with facial animation in this paper, we address the startle response with respect to the shoulders, arms, and head motion. In particular, during startle, the shoulders shrug, the arms retract, and the head is centered by pulling it into the shoulders. Based on this description, the startle is in effect a very brief form of the *recoil* response, based on the above description.

To create a startle response, we incorporate the described behaviors by defining a maximum recoiled pose using the angles from Table I. To make startle a much more subtle effect, rather than applying a recoil pose directly, we

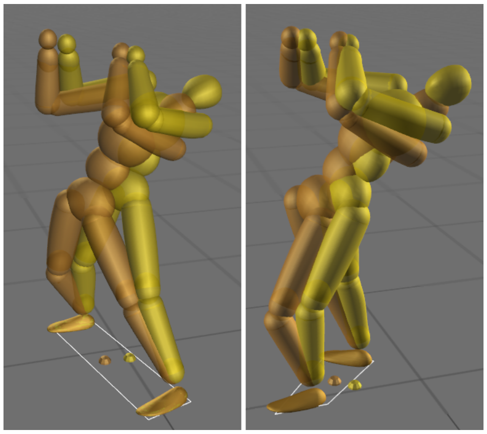

Fig. 5. Two before and after comparisons of static balance. In each image, the yellow pose represents the result of the rules before optimization while the orange pose represents the final sustained pose after optimization. The COM and support polygon are also shown. *left:* The yellow pose was already in balance and the optimization braced for impact. *right:* The yellow pose was out of balance and the optimization both balanced and braced the character for impact.

compute the actual startle pose by interpolating between the recoiled pose and an appropriate posture taken from the motion capture data following a short lookahead. (Again, we refer the reader to Section V-B for a careful discussion of timing.) Empirically, we determined that weights for this interpolation at 0.2 (recoil) and 0.8 (motion capture) lead to a subtle, compelling startle behavior. In addition to the shoulder and arm reaction, the head is turned 40% of the way *toward* the threat to give the appearance of perception. In this manner, the character initially “flinches” when the threat is perceived and then proceeds with the more deliberate sustained defensive action. Due to the short duration and reflexive nature of the startle response, we choose not to consider balance in its computation. Further, in some cases, startle is inappropriate and even distracting, therefore we present it as an option for the animator.

IV. BALANCE

The previously described rules put the character into a protective pose, however, they may also leave the character unbalanced. We would like to balance the character while changing the posture as little as possible from the original motion capture sequence. Rather than make modifications to simply move the center of mass (COM) into the support polygon, we assume that the motion capture sequence is already balanced and we let this sequence guide the COM motion. That is, the system moves the character’s projected COM based on the motion capture data. But, motivated by the need to remain balanced *after* the collision, the system also adjusts the COM to ‘brace for impact’ by adding a slight shift of the COM *toward* the threat to anticipate being knocked backwards and off balance.

We employ a quasi-Newton (BFGS) optimization [28] to move the current center of mass, COM_{cur}, to a desired position, COM_{des}, that both balances and braces the character for impact. COM_{des} is computed starting from a small lookahead to a future motion capture frame’s COM, and then applying a small, animator-controlled offset (10 cm) in the direction of the threat in order to brace the character for impact. Given the COM_{cur} and COM_{des}, the optimizer performs a search to minimize the error, err_{bal}, defined as the Euclidean distance between COM_{des} and COM_{cur}. In implementation, our approach is unique in that it treats the legs as two IK chains and runs the balance search to determine the optimal pelvis position while preventing the feet from moving. The effect of the optimization is clearly shown by the change in pose and COM in each image of Figure 5.

It is important for the balance technique to integrate well with the rules so that protective properties of the pose are not lost during optimization. We therefore introduce additional terms to measure the error between the current optimized pose and the unbalanced sustained anticipation pose. These terms measure properties of the rules: hand positions, orientation of torso body parts, and pelvis height. To avoid unnatural foot skate, we also include an error term that measures the differences in the positions of the feet between the two poses. Each term is computed using appropriate distance metrics (Euclidean or Quaternion) and the results are summed, yielding the total error:

$$err_{total} = err_{bal} + err_{handPos} + err_{torsoOrient} + err_{pelvisHt} + err_{footPos}. \tag{1}$$

The optimization minimizes the error by manipulating a limited set of ten degrees of freedom: the pelvis position and seven rule control terms (lean, turn, recoil, arm reach left and right, and swivel angle left and right). By allowing these control values to change, the optimization finds a tradeoff between balancing and protecting in a single posture.

V. MOTION GENERATION

After computing the sustained pose and (optionally) the startle pose, the anticipatory response motion is generated through a careful interpolation process described next. We highlight some details related to our findings regarding this interpolation process and then we discuss some specifics of the timing scheme and the dynamic response phase.

A. Interpolation Synthesis

Sustained anticipatory motion is generated by interpolating from the motion capture data to the target anticipation pose and back again while maintaining balance and point-to-point hand trajectories. Along with the sustained anticipation, we also optionally include the *startle* target pose.

To start, our blending algorithm determines the support foot, which we define as the one closest to the projected COM. It then treats this as the root of a kinematic chain that branches up and out to the rest of the body. Note, our algorithm assumes in this case that at least one foot is on the ground for the duration of the desired anticipation - but it makes no further assumptions with respect to foot contact (i.e. one foot on the ground is sufficient.) The initial (unbalanced) version of the synthesized motion is computed by blending this rooted kinematic chain using spherical linear interpolation (slerp). With a simple ease-in, ease-out weighting for each blend parameter, the interpolation blends from the motion capture, moving in time, to the startle pose (if used by the animator) and then to the sustained pose (see Figure 6). The motion is then blended back to the original motion capture sequence following a short lookahead. This complete motion blend can also be viewed as a displacement map with the startle and sustained poses serving as the displacement keyframes. Note that this interpolation is only a starting point as the

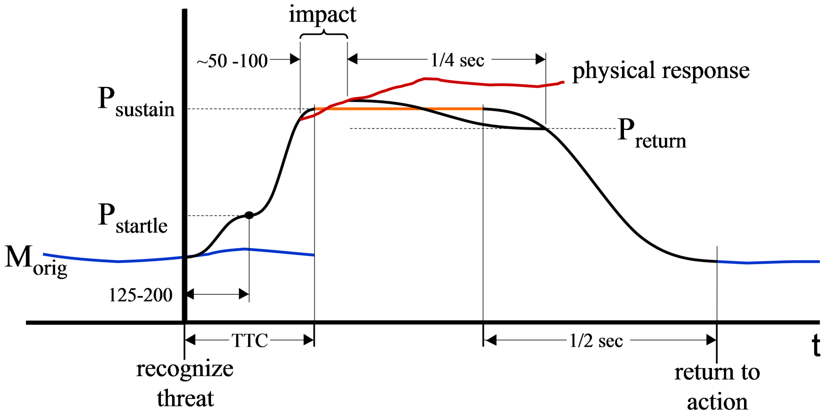

Fig. 6. Schematic of the timing sequence for our approach. All values are timing numbers, in $msec$ unless otherwise specified. Motion sequences are denoted by curves (blue for motion capture, orange for anticipatory response, red for dynamic response, and black for blended transitions).

non-support foot is free to slide and the computed poses will not (necessarily) be in balance. We take care of both of these artifacts during a subsequent re-balancing phase.

1) Arm Interpolation: We found that our animations look more visually pleasing if we include two additional factors for the arm movement in the interpolation synthesis step. First, to make the hand motion look more directed, the system moves the hand along a linear point-to-point path from starting position to ending position across each portion of the described blend (as opposed to solely interpolating the arm joint angles.) Next, the system controls the swivel angle of the arm so that the elbow does not unnaturally swing in or out during interpolation. To accomplish this goal, the system computes the initial swivel angle from the motion capture data and the target swivel angle from the anticipation pose and then interpolates between these during the blend process. To integrate these two adjustments into the described synthesis algorithm, we treat each arm as an explicit IK chain with the desired positions and swivel angles driving the output arm motion.

2) Balance: A second balance optimization phase ensures the character remains balanced throughout the synthesized motion sequence. Similar to the approach presented in Section IV (for computing a single balanced pose), the system moves the COM_{cur} to the COM_{des} across the blended motion. In this case, the COM_{des} is interpolated along the straight line from the COM of the starting frame to the COM of the ending frame. We run an optimization, as before, to find the transformation that moves the pelvis and minimizes only the err_{bal} component. Given that our balance optimization treats the legs as independent IK chains, we also lock down the non-support foot in this step.

B. Timing

Timing plays a critical role in anticipation and is key to producing realistic movement. To create our results, we make certain assumptions about when the character should respond based both on practical needs and by drawing from the known timing measurements taken from our literature survey. Figure 6 provides a detailed schematic of the timing of our approach.

Recall that the animator provides input regarding the relative time at which the threat is recognized and we call this time interval the *time to contact* (TTC) and assume this is the amount of time the character has to respond to the threat and reach the sustained anticipation pose. In doing so, we impose the assumption that the character should reach the anticipation pose “just in time”. Assuming we begin our timer, as in Figure 6, at the point the threat is recognized, the system looks ahead (TTC seconds) in the motion capture data sequence for the origination pose. Since the impact collision is computed based on the character’s posture and threat geometry, the estimated time of impact is not exact - thus, to ensure our “just in time” assumption is upheld, we build a small amount of slack into the TTC. This results in a collision being initiated some time just before the anticipation pose is fully realized and avoids the character freezing and awkwardly holding a static pose before the impact occurs.

In practice, we allow the animator to choose a recognition time, and therefore the TTC, that falls within a reasonable range roughly based on measurements described in the literature. King et. al state that defensive head movements are extremely fast [15] and they report the response latency (time between the threat entering the visual field and initiation of defensive response) of an adult human to a looming visual stimuli to be approximately $350\ msec$, on average, with a low of $240\ msec$ and a high of $600\ msec$. This response latency is also in agreement with studies of spatial trajectories for aimed movements of rhesus monkeys and for studies of human response to visual targets which both report latencies to be around $250\ msec$ on average [10], [8]. Li and Laurent [20] report that the time from initiation of movement to contact with the stimuli ranges from approximately $700\ msec$ to $1000\ msec$ on average (affected by both speed of the threat and eccentricity.) However, their study involved a (relatively) slow moving stimulus when compared to our threat speeds. Thus, we provide a range with an aggressive lower bound ($250\ msec$) and conservative upper bound ($1000\ msec$) for TTC. Startle, on the other hand, is a faster, lower latency response and we set the duration for this response to $130 - 180\ msec$ based on findings of Cooke et. al. [6].

The rest of the timing for the animation depends on what happens next. If the interaction is avoided and no collision occurs, the character will hold the sustained pose for a short time (about a third of a second) before returning to normal activity. More interestingly, if the interaction leads to an impact, the motion generation is handed over to a physics based subsystem that computes a dynamic response for the character motion. The dynamic response sequence is followed briefly (around $50 - 100\ msec$) until the collisions end to allow a short burst of pure physically based motion before returning smoothly back to motion capture controlled movement.

C. Dynamic Response

To generate a physically based reaction once an interaction has lead to a collision, we incorporate a modified version of Zordan et al.’s [34] technique for dynamic response. In brief, the original algorithm uses the motion capture to initialize a dynamic simulation of the character that responds physically to the contact by integrating impact forces and a joint-space torque controller. The system then searches for the closest match among a database of motion-captured reactions to generate a response. However, their technique does not predict the motion of the threat nor does it prepare the character for the interaction before impact.

In our version, we note a few deviations from the original algorithm. Foremost, the simulation is initialized by the synthesized anticipation sequence and then after a brief delay the motion returns back to the original motion. In addition, unlike the previous work, we found it more critical to tune the controller gains in order to generate the desired effect. We believe this is due to our goal of making the character act more purposefully. That is, while Zordan et al. used the controller to keep the character from appearing unconscious (i.e. like a ragdoll) following an impact, we expect the controller to help support the anticipatory behaviors by providing control over the perceived tension of the character, which becomes visible as the threat hits the blocking arms for example.

In our implementation, the simulation is active for a short time, derived from a fixed, user-controlled delay ($50\ msec$) up to the time required until the collisions cease. For the final animation, this newly synthesized physically based motion is used as the interpolant for the return to the original motion (see the red line in Figure 6). The result is a physically plausible but ‘un-canned’ response with anticipation that is generated based on the specific dynamic effects of the threat and followed by a smooth return to the original motion capture behavior. As an alternative to returning to the original motion, we can follow the method outlined by Zordan et al. and allow a search to return the movement to a new motion based on the simulated reaction and a repository of possible reaction examples. While the selection between these two alternatives could be made algorithmically (based on the size of the impact forces, for example), in our implementation we allow the animator to make this choice.

VI. GENERALIZING THE PROTECTIVE BEHAVIORS

The approach as presented thus far produces a wide range of dynamic interactions, however it is limited to responses for impacts to the front and sides of the head. In this section, we revisit the intuition for the rules applied to these impacts and outline how the same intuition might be applied to extend the rules to impacts to other parts of the body and therefore expand the application of the approach.

Each of the psychology inspired rules can also be viewed intuitively as a response designed to protect the head in some specific way. *Lean Away* causes the character to move its head away from the threat. Intuitively, this action serves to distance the threat from the head, effectively increasing the amount of time that the character has to move its hands into position to block the threat. Assuming that impacts to the face are more painful and damaging than impacts to the back of the head, the *Turn Away* rule makes intuitive sense as well because it aligns the back of the head with the approach vector of the threat. *Recoil* causes the character to pull its shoulders up toward its ears and arms in toward its torso, making a smaller target. This response reduces the surface area of vital parts of the body such as the neck and the ribcage, which are targets for many predators. Finally, *Block* causes the character to place its hands and arms in the path of the approaching threat. This behavior can be explained as using a less vital appendage (such as the arm) to lessen or avoid the impact of the approaching threat to a more vital area (such as the head.)

A simple extension covers the cases where the threat approaches the head from behind. Applying the intuition

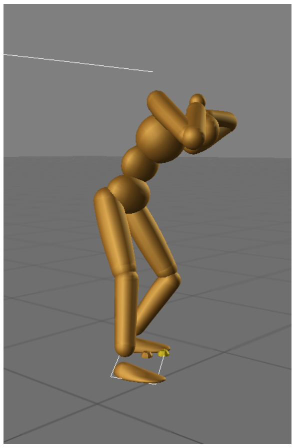

Fig. 7. Example pose for a threat approaching the character's head from behind.

described above, the character should perform all of the described protective behaviors, however, we propose an adjustment to the block rule to account for the uncertainty of the threat's approach vector as it arrives from behind. To reflect the lack of precise knowledge about the threat's location, the system places the hands in a roughly protective position as opposed to placing them on the approach vector of the threat (see Figure 7). It is also very awkward to reach away from the head when the hands are positioned behind the head, therefore, we do not control the reach distance for impacts to the back of the head. To account for the character's uncertainty about the size of the threat, we scale the size and speed control parameters to produce a conservative (more exaggerated) response. In other words, when approaching from behind, the character assumes a faster, larger threat than it would for the same threat approaching from in front. With this extension in place, we can generate protective poses in response to threats approaching the character's head from anywhere within a 360° circle.

Next, we address extensions for anticipating impacts from threats aimed along the length of the trunk. Upon careful consideration, the major difference between these interactions and those geared for the head is that the strategy for *Lean Away* becomes more challenging. In particular, we must be careful to control the appropriate joints for leaning away and for maintaining balance. Recall, the goal of the lean away rule that we described above is to distance the head from the threat. For impacts to the trunk, the character must accomplish this same goal by both pulling the body part to be hit away from the threat and by counteracting this effect with an opposing action for the remaining body parts, thereby maintaining balance. For example, to withdraw from an impact to the pelvis, the character must "lean away" at the ankles while simultaneously bending forward at the waist to counteract the previous motion, thus throwing the pelvis back and away from the threat while maintaining balance. An alternative description of the same motion is to modify the lean rule simply by negating the resulting joint angles. Thus, the upper body leans toward the threat instead of away from it, causing the the pelvis to shift *away* from the threat during the balancing phase. In both cases, the motion is in agreement with the described intuition in that it moves the precise target (the pelvis region) *away* from the threat, therefore increasing the amount of time to react and

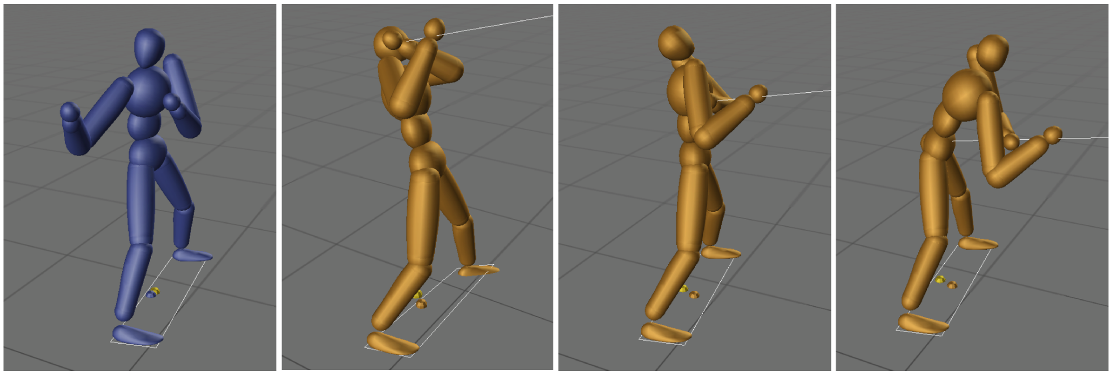

Fig. 8. The blue pose is the original motion capture pose from which the examples are generated. The three orange poses are examples of sustained anticipation responses for threat targets that range from the head (*(top-right)*) to the chest (*(bottom-left)*) and to the pelvis (*(bottom-right)*)

block. Therefore, we can parameterize this behavior along the length of the spine by blending the actions described for the pelvis with those previously described for the head. The result leads to the anticipation poses shown in Figure 8.

Finally, although we have not implemented impacts to the legs, we can still consider the implications of such an impact in terms of the intuition behind each rule. For example, the character may prepare for a collision to the upper legs by shifting its weight to the leg that is further away from the threat (increasing the time to block) and raising the closer leg while bending at the knee (essentially recoiling the leg making the target smaller). The character may turn so that the threat will strike the side of the leg instead of the front (assuming that it is less painful to be hit in the side of the leg than the shins). The arms may be used solely for balancing or they may be moved to block the threat if it deflects off the ground up toward the upper leg or pelvis area. The described response would require a more robust balancing approach and is left as future work.

VII. RESULTS AND CONCLUSIONS

The described anticipation implementation is capable of generating a wide variety of examples. In the animations that accompany this paper, we present several of these examples including cases that vary the threat's approach direction and its dynamic parameters, namely its size and speed. In each, we are able to generate a unique anticipation response for the given threat properties and then return the character to its original actions after a physical impact. To show the improvement of our system over simply applying a dynamic response, we compare our results with anticipation to results generated under the same conditions but without anticipation. We also show that our system encapsulates the capabilities of a dramatic dynamic response following a large impact by transitioning from the anticipation action to a new behavior, such as falling down (see Figure 9).

The presented approach produces compelling results for many situations, however, there are limitations as to when it should be used. Although we have described our algorithm with respect to impacts to the front of the

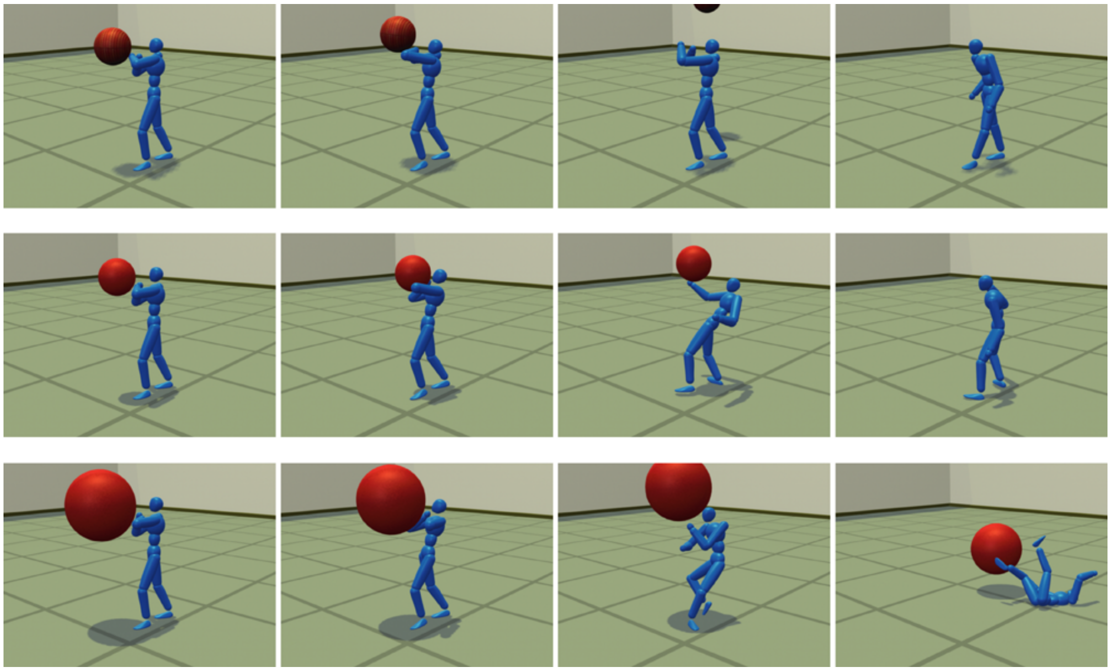

Fig. 9. Three filmstrips show animations generated by our system. Here we demonstrate the result of returning to the original motion (top) or a new motion (middle and bottom) as well as the result of mild (middle) or extreme (bottom) physical threat properties.

head, we have discussed generalizations to impacts to other parts of the body and demonstrated with anticipatory responses to the back of the head, and to areas along the trunk of the body. Although the intuition behind the responses for these areas adheres to the same rules presented for head and upper body impacts, the implementation of these rules can require more complex procedural reactions. We have demonstrated impacts to the torso, however, more work is needed to fully understand impacts to the torso and lower body and to produce more compelling reactions.

In addition, there are circumstances under which the approach will fail for impacts to the head. For example, as discussed previously, the approach for planting the foot will not work if the character was in a deliberate upward motion such as a martial arts kick. Furthermore, our approach is intended for situations in which the character must respond to an unavoidable threat. Therefore we do not provide solutions that include completely ducking or sidestepping to avoid the threat. Finally, the solutions are dependent on a motion capture library. While the dynamic response will provide a realistic response in terms of energy exchange, for a small period of time, after the dynamic response, the motion may appear unnatural in cases where we do not have motion capture data to continue the ’absorption’ of energy.

Although our initial approach to this problem has been a procedural one, we are considering alternative solutions. In particular, we are investigating a data-driven approach using motion capture data. Data will be captured as an actor responds to a harmless projectile such as a nerf ball (or even an imaginary projectile). The response will clearly be an approximation to an actual response since the actor will be aware that the threat is not dangerous, however, this

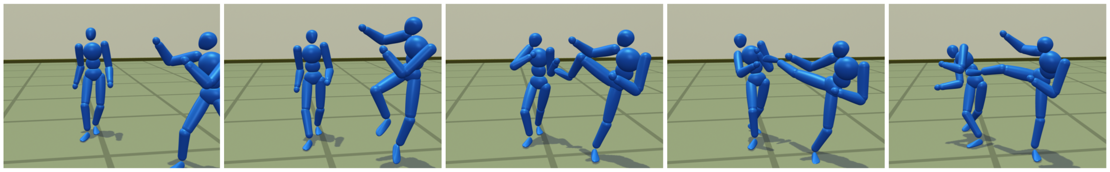

Fig. 10. In this filmstrip, the character approaching from the right is attacking the character on the left. The attacker's foot is used to initialize the threat parameters for the system. Using the attacker's foot as the threat, the system creates a defensive response for the character on the left.

approximation may be more accurate than the procedural approximation presented here. We are exploring methods for choosing the appropriate response sequences and blending them to produce realistic anticipatory responses.

Because the approach is procedural, it is relatively cheap to compute responses. The computation of the response pose requires evaluation of constant time rules and runs in real time as is demonstrated in the accompanying video. The balance computation requires an optimization that generally converges within 10 iterations. The motion blending component also requires procedural blending and an optimization that converges quickly. Therefore, computing and moving to the anticipation pose can be accomplished in real-time, however, a truly compelling response requires the physical response component which is currently the most computationally expensive aspect of the method.

Although we primarily consider medium size, spherical threats (*eg.* volleyballs), the approach generates compelling responses for threats of various sizes and speeds. As an example, we present a martial arts sequence where the attacking character's foot is treated as the impending threat and the defending character responds to the threat, striking an anticipatory pose to protect against the blow to the chest (see Figure 10). We suspect that the responses to different types of threats, such as sticks, splashes of water, and even smaller spherical threats (*eg.* baseballs), may need to vary from our results to remain realistic. Furthermore, we have only considered cases where the character was performing an upright motion (walking, fighting, etc.) with relatively low momentum before the impact. Such motions allow the freedom of modifications like leaning and planting the feet. If the character was initially sitting or running, realistic responses would likely a different set of response rules.

Finally, our solution only loosely integrates the eccentricity work of Li et. al into the response. In particular, they describe how threat speed and eccentricity affect the movement velocity and how soon a person responds to an impending collision respectively. Although we do build upon their results to determine an appropriate TTC range, we use the TTC as the total time to respond in order to uphold our desire to allow the character to achieve the anticipatory pose "just in time". A more realistic solution would separate the time to respond from the response speed, allowing the character to respond sooner for threats approaching from the periphery without dictating the speed of the response. For a threat approaching on the periphery, the character should achieve the protective pose before the collision occurs (a conservative response), while for a threat in front, the character could achieve the pose "just in time". This approach, however, would require a more realistic technique for "holding" a pose that is reached conservatively before the impact. A possible solution would be to use a physical simulation to model the overshoot involved with striking the protective pose and attempting to hold it before the impact occurs. We leave this extension as future work.

In conclusion, we have presented an approach for generating protective poses in anticipation of an impending collision. We have identified a minimal set of primitive behaviors from psychological studies of how humans and monkeys react to impending collisions. We have developed rules for these behaviors, parameterized by the characteristics of impending threats, and shown that this minimal set of behaviors is sufficient to produce compelling anticipatory responses to impending collisions. Our hybrid kinematic and dynamic solution for producing and achieving protective poses results in motion that maintains balance constraints and natural hand trajectories while not straying far from the underlying motion sequence. Finally, we have provided intuition for extending the behaviors and demonstrated that this intuition can be applied to generalize to impacts to other regions of the body.

ACKNOWLEDGMENT

The authors would like to thank Matthew Fast, Nick van der Horst, Bill Chiu, and Robert Rose for their contributions to this project. We would also like to thank Sony Computer Entertainment for their help with collecting motion capture data and Alias Wavefront for donating their software. This work was supported in part by the National Science Foundation (CCF-0237706 and CNS-0423733).

REFERENCES

[1] Y. Abe, C. K. Liu, and Z. Popović, “Momentum-based parameterization of dynamic character motion,” in *2004 ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, July 2004, pp. 173–182.

[2] O. Arikan, D. A. Forsyth, and J. F. O’Brien, “Pushing people around,” in *SCA ’05: Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*. New York, NY, USA: ACM Press, 2005, pp. 59–66.

[3] R. Boulic, R. Mas, and D. Thalmann, “Position control of the center of mass for articulated figures in multiple support.” *Proc. 6th Eurographics Workshop on Animation and Simulation*, pp. 130–143, 1995.

[4] ——, “A robust approach for the control of the center of mass with inverse kinetics.” *Computers & Graphics*, vol. 20, no. 5, pp. 693–701, 1996.

[5] A. Bruderlin and L. Williams, “Motion signal processing,” in *Proceedings of ACM SIGGRAPH 95*, ser. Computer Graphics Proceedings, Annual Conference Series, R. Cook, Ed. Addison Wesley, 1995, pp. 97–104.

[6] D. Cooke and M. Graziano, “Defensive movements evoked by air puff in monkeys,” *Journal of Neurophysiology*, vol. 90, no. 1, pp. 3317–3329, 2003.

[7] P. Faloutsos, M. van de Panne, and D. Terzopoulos, “Composable controllers for physics-based character animation,” in *SIGGRAPH ’01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM Press, 2001, pp. 251–260.

[8] M. Flanders and P. Cordo, “Kinesthetic and visual control of a bimanual task: specification of direction and amplitude,” *Journal of Neuroscience*, vol. 9, no. 2, pp. 447–453, 1989.

[9] B. Forthomme, J.-L. Croisier, G. Ciccarone, J.-M. Crielaard, and M. Cloes, “Factors correlated with volleyball spike velocity,” *American Journal of Sports Medicine*, vol. 33, no. 10, pp. 1513–1519, 2005.

[10] A. Georgopoulos, J. Kalaska, and J. Massey, “Spatial trajectories and reaction times of aimed movements: effects of practice, uncertainty, and change in target location,” *Journal of Neurophysiology*, vol. 46, no. 4, pp. 725–743, 1981.

[11] M. Gleicher, “Retargetting motion to new characters,” in *SIGGRAPH ’98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM Press, 1998, pp. 33–42.

[12] ——, “Comparing constraint-based motion editing methods,” *Graph. Models*, vol. 63, no. 2, pp. 107–134, 2001.

[13] E. Hsu, K. Pulli, and J. Popović;, “Style translation for human motion,” *ACM Trans. Graph.*, vol. 24, no. 3, pp. 1082–1089, 2005.

[14] S. King and A. Cowey, “A defensive response to looming visual stimuli in monkeys with unilateral striate cortex ablation,” *Neuropschologia*, vol. 30, no. 1, pp. 1017–1024, 1992.

[15] S. King, C. Dykeman, P. Redgrave, and P. Dean, “Use of a distracting task to obtain defensive head movements to looming visual stimuli by human adults in a laboratory setting,” *Perception*, vol. 21, no. 2, pp. 245–259, 1965.

[16] T. Komura, E. S. Ho, and R. W. Lau, “Animating reactive motion using momentum-based inverse kinematics,” *Compute Animation and Virtual Worlds*, vol. 1, no. 16, pp. 213–223, 2005.

[17] T. Komura, H. Leung, and J. Kuffner, “Animating reactive motions for biped locomotion,” in *VRST ’04: Proceedings of the ACM symposium on Virtual reality software and technology*. New York, NY, USA: ACM Press, 2004, pp. 32–40.

[18] C. Landis and W. Hunt, *The Startle Pattern*. Farrar and Rinehart, 1999.

[19] J. Lee and S. Y. Shin, “A hierarchical approach to interactive motion editing for human-likefigures,” in *Proceedings of ACM SIGGRAPH 1999*, A. Rockwood, Ed. Los Angeles: Addison Wesley Longman, 1999, pp. 39–48.

[20] F.-X. Li and M. Laurent, “Dodging a ball approaching on a collision path: Effects of eccentricity and velocity,” *Ecological Psychology*, vol. 13, no. 1, pp. 31–47, 2001.

[21] C. K. Liu, A. Hertzmann, and Z. Popović, “Learning physics-based motion style with nonlinear inverse optimization,” *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 1071–1081, Aug. 2005.

[22] M. Mandel, “Versatile and interactive virtual humans: Hybrid use of data-driven and dynamics-based motion synthesis,” 2004, *Master’s Thesis, Carnegie Mellon University*.

[23] M. Neff and E. Fiume, “Methods for exploring expressive stance,” in *SCA ’04: Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*. New York, NY, USA: ACM Press, 2004, pp. 49–58.

[24] M. Oshita and A. Makinouchi, “A dynamic motion control technique for human-like articulated figures,” *Computer Graphics Forum (Eurographics 2001)*, vol. 20, no. 3, pp. 192–202, 2001.

[25] K. Perlin, “Real time responsive animation with personality,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 1, no. 1, pp. 5–15, Mar. 1995.

[26] K. Perlin and A. Goldberg, “Improv: A system for scripting interactive actors in virtual worlds,” in *Proceedings of SIGGRAPH 96*, ser. Computer Graphics Proceedings, Annual Conference Series, Aug. 1996, pp. 205–216.

[27] Z. Popović and A. Witkin, “Physically based motion transformation,” in *Proceedings of ACM SIGGRAPH 1999*, 1999, pp. 11–20.

[28] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C*. New York: Cambridge University Press, 1994.

[29] W. Schiff, “Perception of impending collision,” *Psychological Monographs: General and Applied*, vol. 79, no. 11, pp. 1–26, 1965.

[30] S. Tak, O. young Song, and H.-S. Ko, “Motion balance filtering,” *Computer Graphics Forum*, vol. 19, no. 3, pp. 437–446, August 2000.

[31] A. Witkin and Z. Popovic, “Motion warping,” in *Proceedings of ACM SIGGRAPH 95*, ser. Computer Graphics Proceedings, Annual Conference Series, R. Cook, Ed. Addison Wesley, 1995, pp. 105–108.

[32] J. Yeomans, B. Scott, and P. Frankland, “Tactile, acoustic, and vestibular systems sum to elicit the startle reflex,” *Neuroscience Biobehavioral Reviews*, vol. 26, no. 1, pp. 1–11, 2002.

[33] V. B. Zordan and J. K. Hodgins, “Motion capture-driven simulations that hit and react,” in *ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, July 2002, pp. 89–96.

[34] V. B. Zordan, A. Majkowska, B. Chiu, and M. Fast, “Dynamic response for motion capture animation,” *ACM Trans. Graph.*, vol. 24, no. 3, pp. 697–701, 2005.