An Artificial Animated Boxer

Alexander Kolliopoulos

May 3, 2004

1 Introduction

A tempting goal of artificial life is to model independently acting human characters that appear to behave realistically to a viewer. By modeling artificial human characters engaged in one activity, such as boxing, we may gain insights into the larger problem of creating convincing humans. In the case of boxing, we are specifically concerned with reactive motion. A boxer must react quickly to incoming punches from an opponent, while attempting to catch the opponent off guard with a well placed attack. In order to make the reactive boxer interactive for a human user, we would like the system to run in real time.

There are strategies to boxing that should be kept in mind when designing a system to animate synthetic boxers. A right handed boxer will typically approach an opponent with his left side facing forward. This makes it easier to protect himself and gives his left hand the shortest distance to travel to the opponent. There are four punches used in boxing: the jab, cross (straight right), hook, and uppercut. The jab is thrown with the left arm, being used as a defensive tool or to get the opponent off guard to set up a more powerful punch. The other three punches are thrown with the right arm. The cross is the most simple of these, while the hook requires twisting and arching of the body. This makes the hook more difficult to execute and leaves the body partially open to attack. Similarly, an uppercut comes from below, but is a risky punch to throw. The cross and uppercut both have the advantage of being powerful and difficult to see since they don't come in the line of the opponent's vision. Feints are faked punches used to make it difficult for an opponent to read a boxer's body language and are also an important part of a boxer's repertoire. A boxer might have a defensive style, making him known as a grinder, or an offensive style, a puncher. A good defense will make use of blocks, dodges, feints, and jabs to wear down an opponent. With basic boxing strategy in mind, a behavioral system to make use of it may be designed.

2 The System Design

There are two main problems that need to be addressed for this project: animation and control. It is difficult to model good human motion with physically based systems [2], and they do not run in real time. Motion capture can be used to quickly generate a library of motion, but playback of independent clips of motion capture will not generate very smooth animation. The details of the animation system were implemented by Kevin Forbes for his portion of a Character Animation project. Our solution is to use quaternions to blend between clips of motion capture, creating smooth but perhaps not completely physically realistic animation. Nearly 350 different motions performed by myself were captured and manually segmented. Each action was performed multiple times in a variety of styles: amateurish, skilled, tired, energetic, defensive, and offensive. Using this library of motion captured clips, two boxers can be animated with different styles.

3 The Artificial Boxer

The behavioral decision system is based on that of artificial fishes [4]. Hence, the boxer has some internal state, perception of the external world, and an intention generator and behavior routines to select actions to be carried out by the motor controllers, in this case, the animation system.

The boxer has a fairly simple state model; the only dynamic variable that changes through the fight is the boxer's physical energy. Other parameters that affect the style of motion and decision processes are skill, boxing style, and strength. These are stored in XML files and are loaded at runtime. As a boxer throws punches or takes hits, his energy decreases. By taking a moment to clear his head, a boxer may recover a bit of energy. Also, hitting the other boxer can artificially raise the boxer's energy level as he becomes motivated by the apparent weakness in his opponent.

The boxer has a perception system capable of giving him the relative angle and distance to the other boxer, if that boxer is within view. A boxer will be visible if he is within a cone with angle 60 degrees from the view vector. To see more detailed motion, such as an incoming punch, the joint in interest must be within 45 degrees of the view vector. If the opponent disappears behind the boxer, he only knows on which side the opponent disappeared on and cannot get detailed information until he tracks the opponent again.

The intention generator is depicted in Figure 1. It functions similarly to that of artificial fishes, but the conditions return a confidence level between 0 and 1 rather than absolute true and false. Thus, the confidence acts as a probability that the boxer will consider the condition to be true. A random number $r \in [0, 1]$ is generated, and if r is less than the confidence, the condition is considered to be true. This adds a stochastic element to the intention generator, but it also presents an opportunity to perform simple learning.

Each condition is associated with an *adjustment*, which is a positive number. When the adjustment is less than 1, it means that the boxer has a lower confidence in the condition than what the condition function actually returns. For adjustments greater than 1, the boxer has a higher confidence than what is returned by the condition function. To apply an adjustment less than 1, the confidence is simply modified by $\hat{c} = \alpha c$, where $\hat{c}$ is the adjusted confidence, c is the value returned by the condition function, and α is the adjustment value for that condition. This has the effect of lowering the confidence, but there can be problems if α becomes too small, in which case the condition nearly always is assessed to be false. The designer of an artificial boxer will often need to specify states in which a condition must evaluate to true, regardless of the adjustment for that condition. To allow for this, when a condition function returns 1, no adjustment is applied. Adjustments greater than 1 are handled in a similar manner, with $\hat{c} = 1 - \frac{(1-c)}{\alpha}$ for $c \neq 0$.

Adjustments are learned by scoring paths taken in the intention generator tree. When

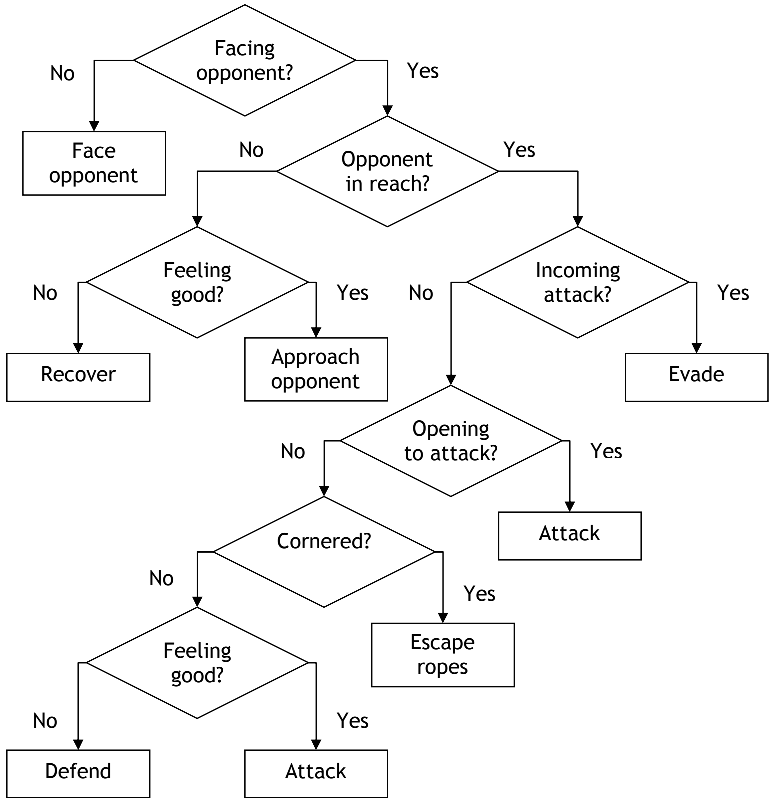

Figure 1: The boxer intention generator

an event occurs that is beneficial to a boxer, such as landing a punch, the adjustments for each condition in the path taken to the intention that resulted in the positive behavior are modified. Specifically, conditions that evaluated to true have their adjustments raised, and those that were false have their adjustments lowered. Hence, in the future, the same successful path will be more likely to be taken. Negative events, such as a blocked punch or taking an opponent's hit are scored similarly. However, some nodes in the path will be scored more often than others, so they should receive a smaller modification to their adjustment. For example, the facing condition will always receive a modification since it is at the root of the tree. To address this, each node up from the last condition before the selected intention will receive half the modification of the node further down the tree in the path.

For our purposes, the behavior routines are fairly simple. One example is the face opponent behavior routine, which queries the boxer's perception for the approximate angle to the opponent and then activates either a large or small step to the left or right, depending on what is perceived to be the quickest way to line up the enemy. The attack behavior routine will select a punch based on the distance to the opponent and the state of the boxer. No learning is applied at the behavior routine level, but it is not difficult to imagine a system of informing the behavior routines of their success or failure and allowing them to modify themselves in a routine dependent fashion. The main barrier to this would be the amount of coding required. Since each behavior routine functions in a different way, the reinforcement learning would have to be custom designed for each routine.

4 Results

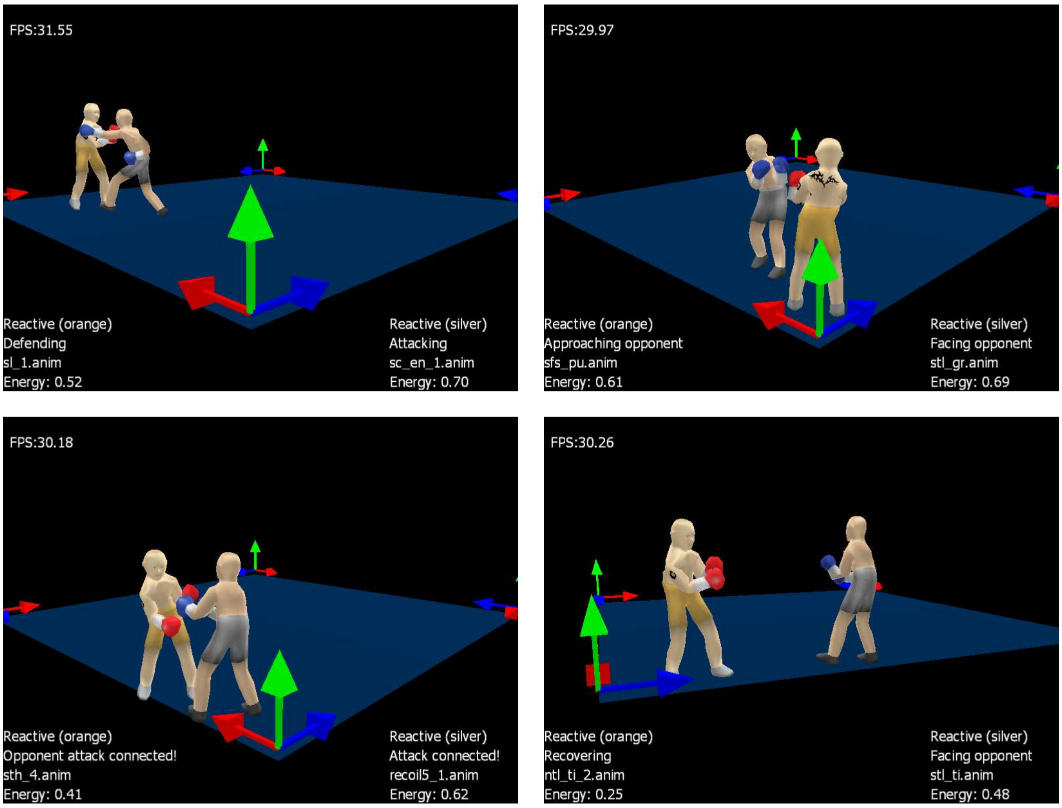

Figure 2: Two artificial boxers fighting

The learning system of the intention generator turns out to be quite sensitive to the reward and punishment scheme used. For a while during development there was trouble with the boxers discovering that they could stay far enough from each other to be safe and just stand in place recovering energy without fighting. It is perhaps interesting that the boxers chose to be pacifists despite the fact that they were intended to be fighters, but this behavior is unacceptable if the goal is to win a fight. This problem was corrected by punishing the recover intention when the boxer is not tired.

In testing with a user controlled boxer, it was found that a fairly aggressive style is effective; and it is somewhat reassuring that more than once, one of the boxers has discovered this and raised the adjustment on the opening condition significantly. Often this may result in the other boxer getting beaten so badly that he is fairly helpless. In normal conditions, one can observe some characteristics of a real boxing match, with fighters dodging punches and catching each other off guard.

Limitations can be seen in the system of adjusting the intention generator since conditions may only be favored or disfavored, making them more or less likely to be true independent of the state of the environment beyond the condition function. This is a strength of the approach for its simplicity and speed, but it also prevents any deeper cognitive learning.

5 Conclusions

The results of this project are a method of animating and controlling an artificial boxer using motion capture and a dynamic intention generator. Much of the way boxers behave is reliant on the choices of weights for rewards and punishments. However, despite the apparent simplicity of the approach with the intention generator consisting of only six conditions and seven behavior routines, a wide range of behaviors can be observed. Some interesting directions for future work in this area might be to make modifications to the intention generator's connectivity as in [3], running hundreds of simulations to automatically generate effective intention generators from a selection of conditions. A more advanced reinforcement learning system that takes into account the state of the environment and the actor as in [1] would likely lead to improved action selection if the problem of the massive possible state space in boxing could be overcome. Reactive human motion is an area with a wide range of possibilities that have yet to be explored. This project has found some degree of success in extending methods designed for much simpler creatures to automatically animating artificial human boxers.

References

[1] Bruce Blumberg, Marc Downie, Yuri Ivanov, Matt Berlin, Michael Patrick Johnson, and Bill Tomlinson. Integrated learning for interactive synthetic characters. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 417–426. ACM Press, 2002.

[2] Petros Faloutsos, Michiel van de Panne, and Demetri Terzopoulos. Composable controllers for physics-based character animation. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 251–260. ACM Press, 2001.

[3] Karl Sims. Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 15–22. ACM Press, 1994.

[4] Xiaoyuan Tu and Demetri Terzopoulos. Artificial fishes: physics, locomotion, perception, behavior. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 43–50. ACM Press, 1994.