The Simultaneous Boxing Game

Blaise Durkin

Abstract

The purpose of this paper is to lay out a framework for a model of boxing. Boxing is a form of combat that dates back to antiquity. Because modern boxing is a restrictive combat sport, it allows for a simpler model of hand to hand combat, compared to a more complex sport such as MMA. Nevertheless, the boxing model herein is a simplified model of real boxing. In reality, boxing is more dynamic and nuanced, which would require more complex physical models. For simplicity, this paper observes a model of boxing with a small and finite set of actions in which the two players simultaneously choose their best response given a probability distribution for the actions their opponent could take. This paper describes the rules for the game mechanics and the results that arise.

Introduction

Given the long history of boxing, there is an abundance of research on the topic from various academic fields. Previous literature has examined the many fine details of boxing, such as the energy, aggressiveness, and style of fighters, as well as the physiology of boxing in determining outcomes (Thomson 2015). However, the model herein ignores all of these features and focuses only on the outcomes that arise when players are maximizing their payoff in a repeated one shot simultaneous game. The most similar system to this one, uses a markov chain to model MMA (Holmes et al. 2023). The system herein is composed of a deterministic process where players choose their best response each round (subgame) in the repeated simultaneous game. This paper formalizes the strategic game of boxing in a system that can be translated analogously to other applications, such as: mixed martial arts, as well as cooperative games like dancing. The paper begins with a benchmark model, and continues to describe the improvements which would construct a more realistic game. The 2-player paired-sequences which the model generates create believable fights which can be utilized practically, such as: fight choreography, and training data for learning.

The game consists of 2 players (boxers), who have the same action space: a set of 8 actions: 4 attacks and 4 defenses. Assume both players have their left leg as their leading leg and their feet are not being displaced; they are simply moving their upper bodies. Like real boxing, in this game players are trying to successfully punch their opponent while successfully avoiding being punched by their opponent.

In reality humans have a lot more degrees of freedom which creates a near infinite set of actions. To simplify the model we only have 8 actions: Jab, Cross, Hook, UpperCut, SlipLeft, SlipRight, Duck, Back.

Player Action Definitions

Jab	Straight left
Cross	Cross right
Hook	Hook with left hand
UpperCut	Upper-cut with right hand
SlipL	Lean left
SlipR	Lean right
Duck	Lean forward
Back	Lean back

As each player has the same action space, there are 64 possible outcomes. Some of these outcomes are obvious. If both players act defensively then nothing happens and they both receive a payoff of 0. However some of the outcomes are not obvious and need to be clarified.

Payoff Matrix

	Jab	Cross	Hook	UpperCut	SlipL	SlipR	Duck	Back
Jab	[-0.5, -0.5]	[1, -1]	[1, -1]	[-1, 1]	[-0.5,0.5]	[-0.5, 0.5]	[0.5, -0.5]	[1, -1]
Cross	[-1, 1]	[-2, -2]	[-2, 2]	[2, -2]	[2, -2]	[2,-2]	[-0.5, 0.5]	[-0.5, 0.5]
Hook	[-1, 1]	[2, -2]	[-2, -2]	[-2, -2]	[-0.5,0.5]	[2,-2]	[-0.5, 0.5]	[2, -2]
UpperCut	[1, -1]	[-2, 2]	[-2, -2]	[-2, -2]	[1, -1]	[-0.5, 0.5]	[1, -1]	[-0.5, 0.5]
SlipL	[0.5, -0.5]	[-2, 2]	[0.5, -0.5]	[-1, 1]	[0, 0]	[0, 0]	[0, 0]	[0, 0]
SlipR	[0.5, -0.5]	[-2, 2]	[-2,2]	[0.5, -0.5]	[0, 0]	[0, 0]	[0, 0]	[0, 0]
Duck	[-0.5, 0.5]	[0.5, -0.5]	[0.5, -0.5]	[-1, 1]	[0, 0]	[0, 0]	[0, 0]	[0, 0]
Back	[-1, 1]	[0.5, -0.5]	[-2, 2]	[0.5, -0.5]	[0, 0]	[0, 0]	[0, 0]	[0, 0]

Outcome Payoff: [row player payoff, column player payoff]

Observe the bottom right quadrant (in purple), is where both players act defensively and both receive a payoff of 0. Green cells represent a successful defense; the attacker receives a payoff -0.5 and the defender with a payoff of 0.5. Red cells represent both players punching each other in the face very hard and thus both receive a payoff of -2. Yellow/Orange cells represent players getting hit, either; -0.5 if not hard, -1 if medium hard, -2 if very hard. The successful attacker has equal but opposite payoffs of 0.5, 1, 2 depending on how hard they hit. The magnitude of the payoff, whether 1 or 2, depends on the rotation around the spinal axis. Attacks with more rotation generate more force, and thus have a payoff of magnitude equal to 2.

Creating a simplified model of boxing means determining the outcome of ambiguous situations. Here is a list of the outcomes I found to be ambiguous and my rationale for determining the outcomes: Jab vs. Cross: Jab is faster so jab wins; Jab vs. Hook: Jab is faster so jab wins; Jab vs. UpperCut: UpperCut can be combined with a right slip so it wins; Cross vs. Hook: Hook is faster and its trajectory collides with cross so it wins. Other milder cases of ambiguity were resolved by balancing the payoff matrix so that a single action does not obviously dominate the other actions.

Game Model 1 - Fixed Beliefs

Players act to maximize their expected payoff at each round. Thus at each round the player chooses the action with the greatest expected payoff. From the perspective of the player choosing action x_i and with an opponent's action x_j:

For an action x in set X, at time t, we have the expected payoff $E[U(x)]_t$.

$$E[U(x_i)]_t \ = \ \sum_{\forall x_j \in X} [\ U(x_i,\ x_j) \ * \ P(x_j \mid x_{j,\,t-1}) \ ]$$

Player i, having calculated the expected utility for their action space, then chooses the action $x_{i,\,t}$ that will maximize their expected utility, given their set of next possible actions X_{it}.

$$x_{i,\,t} \ = \ argmax_{x_i \in X_{it}} \ E[U(x_i)]_t$$

There exists a probability distribution, $P(X \ = \ x_t \mid x_{t-1})$ for each $x_{t-1} \in X$.

Probability Distributions

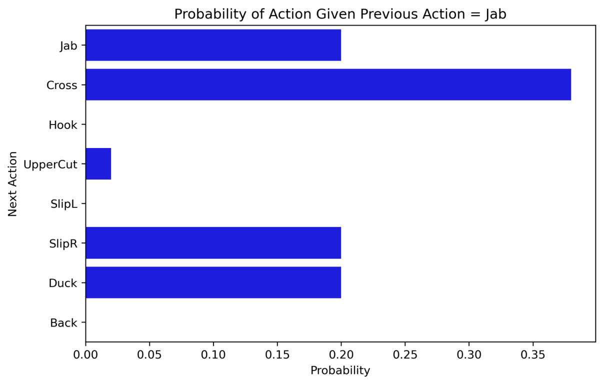

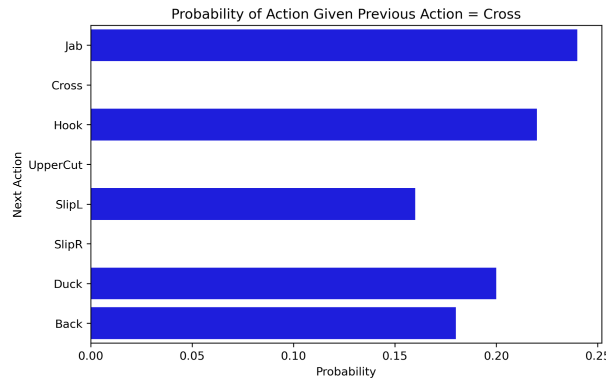

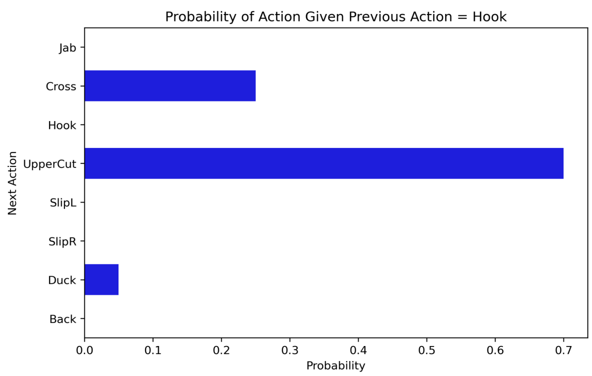

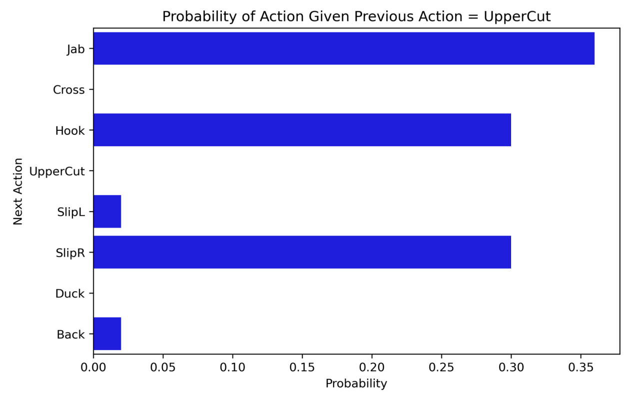

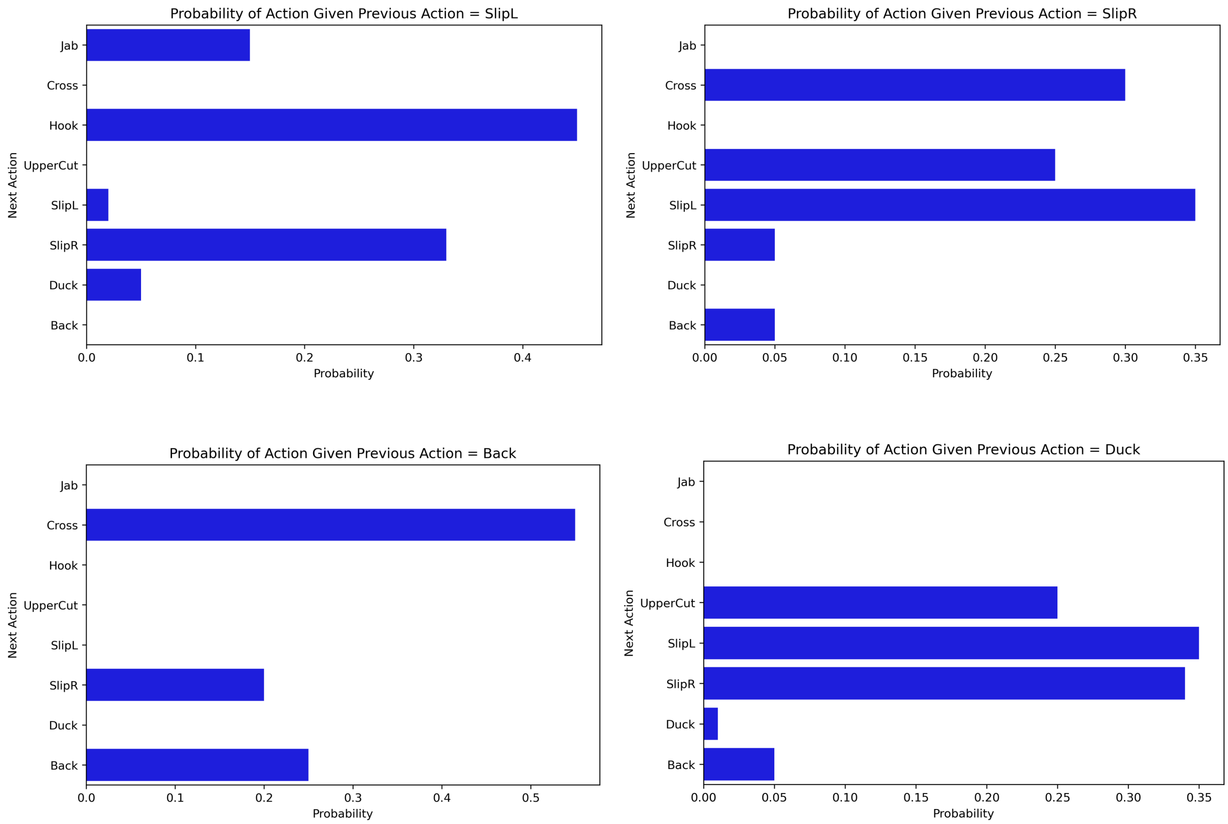

The probability distribution, $P(X = x_t \mid x_{t-1})$ represents the probability of choosing an action given the last action chosen. This serves as a set of prior beliefs about the opponent's action behavior and it is fixed over time. The probability distributions generalize human mechanics, including spine torsion, and power generation. Thus the probability distributions reflect a natural transition between actions in a sequence.

Because each action depends on the previous action, the game is initialized with player i at action x_i and player j at action x_j, where x_i, $x_j \in X$. Thus there are 64 possible initial conditions, resulting in 64 possible game outcomes. Yet, many games converge to the same equilibrium sequences.

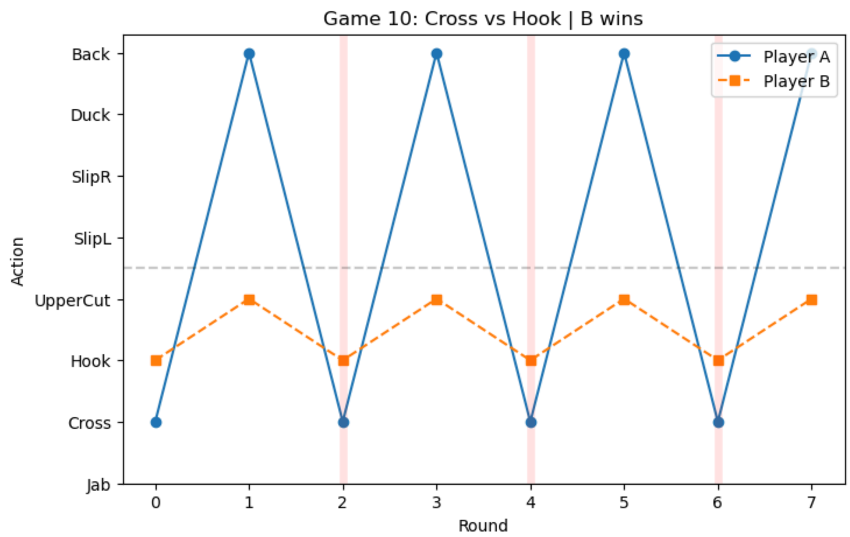

Game 10 is the most common sequence, occurring 32/64 of the games. Next, Game 1 occurs 25/64, Game 27 occurs 4/64, Game 8 occurs 2/64, and Game 36 occurs 1/64. All 64 initial conditions eventually converge to only 5 different sequences: Game 10, Game 1, Game 27, Game 8, and Game 36. For games that converge to Game 10, the winner will be the player that switches between Hook and Uppercut.

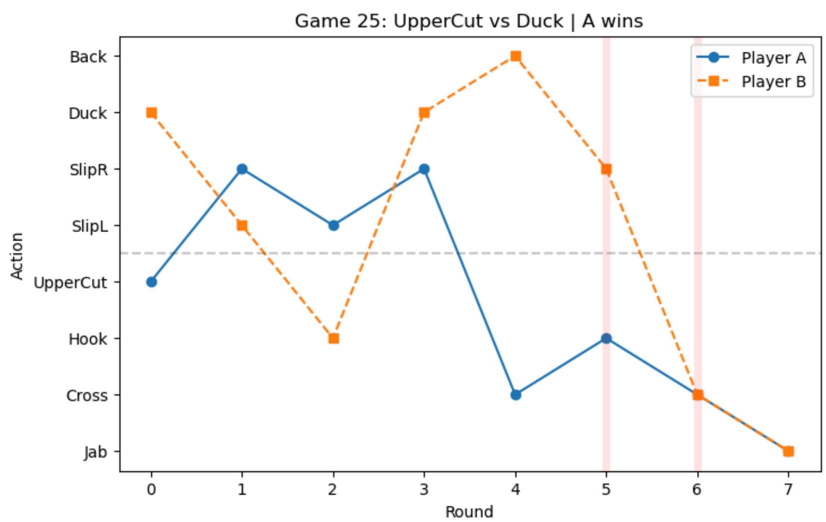

Game 25 is an example of a game that takes longer to converge, but after round 6 the game is fixed at Game 1 (Jab vs. Jab).

Game 8 is when one player does not deviate from Back while the opponent does not deviate from Jab. Game 27 is when both players do not deviate from SlipL. Finally, Game 36 is when both players do not deviate from Back. These 3 game states: Game 8, Game 27, and Game 36 could be avoided by including a stalling penalty into the game mechanics; players receive a negative payoff if they consecutively avoid attacking the opponent.

Game Model 2 - Updated Beliefs

The 2nd version of this model has updated beliefs. That is, the probability distributions are updated after each round to reflect the actual frequency of actions taken given the previous action. The magnitude of change depends on the payoff of the outcome; large negative payoffs have a larger effect on the probability distribution to reflect the fact that bad outcomes have a more significant effect on beliefs. The magnitude of the change will determine how fast players diverge from an equilibrium. Therefore, because of the updates in beliefs, players no longer converge to a fixed state for infinite time.

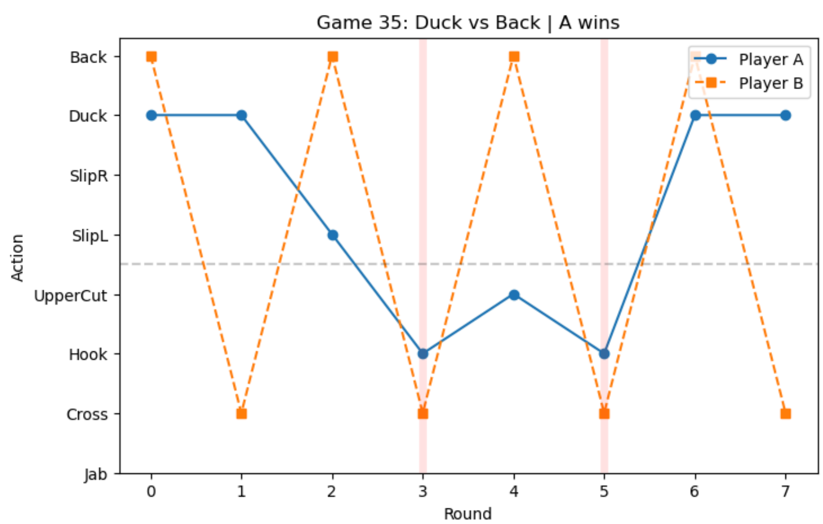

This is an example of the original Game 10; however, Player A does not indefinitely repeat the same sequence: Hook, UpperCut, Hook,…

Discussion

This model succeeds in generating realistic sequences of moves, but due to its simplicity, it fails to characterize a realistic boxing match. In reality actions should be described with continuous variables; punches can vary in speed, force, angle, etc. Nevertheless, a realistic model of boxing can arise from a discrete system. Further improvements in realism can be achieved by factoring in more dependencies into the model. This model consists of 64 possible outcomes; however, the outcomes in a more realistic model depend not only on both players' simultaneous actions but also on their last action. This would account for the trajectories and possible collisions of the players, and reduce the ambiguity of a 64 outcome model. Furthermore, this model does not account for the effects of being punched hard, such as: loss of consciousness, and also does not account for depreciating levels of energy over time. Factoring in player attributes such as: endurance, strength, skill, and style, as well as variables that change throughout the game such as: energy and speed, will generate more realistic fights. Endurance determines the decay rate of energy; less energy decreases speed and makes certain actions more desirable. Strength would affect the payoff magnitudes; more strength means more robustness to damage accrued by the opponent. Speed would reduce the ambiguity of outcomes, and would also increase the magnitude of payoffs; faster collisions do more damage. Increasing skill would increase the set of actions, and style would change the transition states (probability distributions). Pairing these additional features with a more complex intelligence system, that factors in a longer sequence of actions rather than only the last action taken, will produce AI boxing agents which could be utilized in video games. Research conducted by Facebook AI uses deep reinforcement learning to model boxing from motion capture video data (Won et al. 2021). This model includes the player to player distance, player orientations: whether they are facing each other, the players energy, and cumulative penalties. Including penalties into the game can prevent undesirable outcomes such as stalling. Combining all of the aspects of previous research will enable the creation of a realistically boxing simulation.

Conclusion

This paper provides a novel approach to simulating boxing. The simple model of boxing herein provides a benchmark model, which can be improved by the addition of more features, and can also be readapted analogously to similar strategic games such as MMA, or even cooperative games like dancing. The boxing simulation was successful in producing believable fight sequences despite the simple mechanics.

References

Thomson, E. (2015). (thesis). *The development of an amateur boxing simulation protocol*. University of Chester, Great Britain.

J. Won, D. Gopinath, J. Hodgins (2021). *Control strategies for physically simulated characters performing two-player competitive sports*. Facebook AI Research, USA

Holmes, B., McHale, I. G., & Żychaluk, K. (2023). A markov chain model for forecasting results of mixed martial arts contests. *International Journal of Forecasting*, *39*(2), 623–640. https://doi.org/10.1016/j.ijforecast.2022.01.007