Swing or Take: Decision Making under Uncertainty Based on Visual Information

Dibya Ghosh **Maxwell J. Weinstein**

Abstract

We try to better understand decision making in baseball, specifically the interaction between visual stimuli from an incoming pitch and a hitter decision to swing. We employ a psychophysical model of human perception of moving 3d projectiles based on the angular momentum of a pitch to model and classify visual stimuli. Using this as a better characterization for "tunneling", we use this as our model for hitter perception of partially observed pitch paths and determine the efficacy of this model to explain hitter decision making.

1 Introduction

The human ability to track moving objects and predict its final location is a primitive skill, but is not well understood or studied. In baseball, a hitter has approximately 0.125 seconds to decide whether or not to swing. In those 0.125 seconds a hitter relies on what observers of the sport may call, instinct, or innate knowledge of how to react given the visual stimulus of a moving trajectory. Ultimately, a batter is limited by reaction time to the trajectory itself; and, a practiced hitter relies on muscle memory that allows them to slow down the time it takes to make a decision. Practice makes perfect, as they say, and in the mind of an experienced and well-played batter is some cache of thousands upon thousands of fleeting trajectories, digested partially, that they have personally taken or attacked during the course of their career.

What is unclear to a skeptical outsider is whether this decision to swing or take is arbitrarily deliberated. Given so little time to decide, are hitters generally relying on prior sentiment before the pitch is released? Or does the experienced batter update his strategy within the 0.125 seconds he has to take in the physical trajectory and make his decision on the go? Understanding that these are professionals, the answer is probably somewhere in between; a hitter will have some idea of what he is comfortable doing given the context but will jump on his opportunity if presented and realized.

Though, understanding that a batter, being human, is limited by the irreducible cost of his own reaction time; similarly, the batter is limited by his cognitive abilities to project the final location of a pitch. There is evidence that this is an ability realized early in babies, as young as 8 months [1]. How much this limits each batter in their ability to make the right decision is not well studied or understood.

Realizing that there is some hard limit to how quickly a human can react, there also is a limit to how well a batter can project, discern, and characterize a moving baseball. This ability ultimately manifests itself in the outcome of the decision to swing or not, and given that a successful major league hitter fails 2/3rds of the time to hit the ball, it is safe to say they are often fooled if not simply late.

We hypothesize herein that what truly informs a hitters decision in the 0.125 seconds he has to decide is not the trajectory itself, but the perceived trajectory by the hitter. This can be justified to

Submitted for the *Probability in the Real World* seminar.

say that a hitter has no access to the physical parameters of the pitch and can only rely on what he sees.

2 "Tunneling" and the Hitter Decision Making Pipeline

This interplay between perceptual and actual trajectory should be manifested statistically: pitches that are deemed perceptually equivalent by the batter, but are actually physically very different should yield miserable results. This leads us to the concept of tunneling:

Tunneling is best described as the art of fooling a hitter into thinking one pitch is another, by allowing the pitched ball to stay as close to the trajectory of the hypothesized pitch for as long as possible, only deviating once its too late for the hitter to change his decision.

Fooling the hitter into thinking one pitch is another entirely is not so important as to influence his decision as it is to make the outcome of said decision poor. That is if a hitter begins to swing at a pitch he projects to be even just 5 millimeters lower, he will swing and miss. In other cases, just tiniest errors in projecting the pitch will lead to weak contact and good outcomes for the pitcher.

It is our hypothesis that the hitter will project the location of the pitch based on pitches that he deems perceptually equivalent up to the point where he makes the decision. If this is true then we should be able to predict the outcome of a swing, i.e. hit or miss, with reasonable certainty knowing the residual between the actual and projected locations.

Now that we have motivated the randomness in hitter decision making, we will equip ourselves with a physical model of the trajectory, then we will create a similarity metric for perceptual trajectories, and finally test our model for predicting the decisions of batters on perceptually equivalent pitches.

3 Pitchfx Data

We scrapped 6-million pitches from the 2008-2014 Major League Baseball seasons from the publicly available pitch database, PitchFX. This dataset was collected by high resolution tracking cameras, that take high resolution pictures of each pitch thrown in the majors, and learns the physical parameters of the trajectory. That dataset also includes meta-data on every pitch, specifically the batter, pitcher, umpire involved and the corresponding outcome. We specifically are interested at looking at pitch pairs in the same plate appearance, i.e. to the same batter via the same pitcher consecutively in the same AB (at bat).

4 Physics of a Pitched Ball

We use a 9 parameter model suggested by Baseball Physicist Alan Nathan. This approximates the final location of the ball within 0.5 inches on average, as found via quality assurance testing in labs.

The parameters of initial velocity, acceleration, and position given by the PitchFx data are interpolated from the initial and final positions using a Linear SVD approach to solve for the trajectory model "...using measurement consisting of time, and screen position." [5]

So we have this in hand as our best approximation of the actual trajectory for every time t during the flight of the pitch. Using the 9-parameter parametric model as specified [5]:

$$x(t) = a_{x0}t^2 + v_{x0}t + p_{x0}$$
$$y(t) = a_{y0}t^2 + v_{y0}t + p_{y0}$$
$$z(t) = a_{z0}t^2 + v_{z0}t + p_{z0}$$

Where a_{x0} denotes the initial acceleration in the x direction, v_{y0} the initial velocity in the y direction, p_{x0} and the initial x-position of the pitch at 55 feet. The same definitions hold for the y, z coordinates, where y is depth to the plate and z is height from the ground, in inches.

5 Physics of Decision Making

In studying the cognitive limits of the human eye to perceive head on trajectories, it has been shown we use the visual direction of a moving projectile to estimate the 3d motion of an object. Moving forward we will use this as our model for how a hitter judges the trajectory of the ball and where it will end up. We better describe visual direction in the upcoming section, first with Figure 1 [2] in hand:

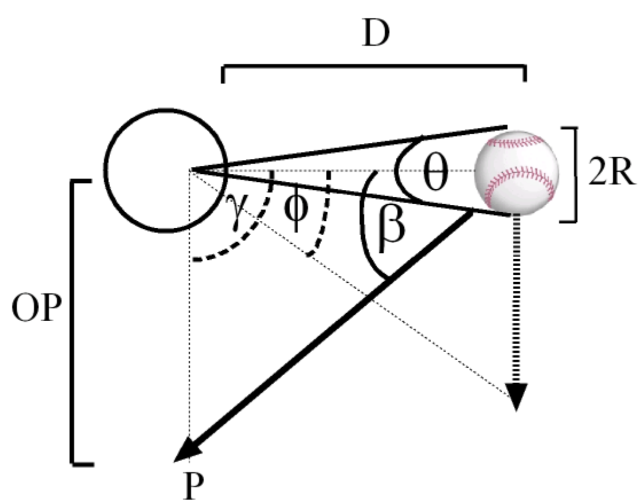

Figure 1: A depiction of the visual direction of a pitched ball

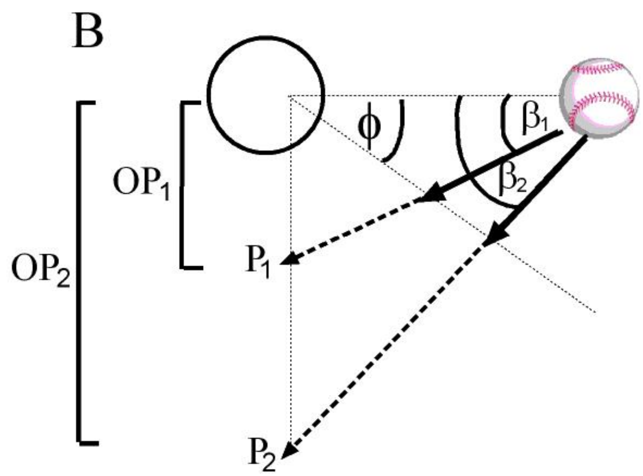

Figure 2: A depiction of the visual direction of two pitches

We define the visual direction of a moving projectile as β. That is the angle between the middle of the object in motion and the vector from that object to the observer, as depicted above.

It has been shown in multiple lab studies that the human eye cannot discriminate velocity to a high enough degree; but instead the human eye uses visual direction as the main feature in discriminating trajectories.

Rob Gray of perceptionaction.com determined in lab experiments that two pitches were deemed perceptually equivalent by the human eye unless proportional difference – $\frac{\beta_1-\beta_2}{\beta_1}$ – between the visual direction of consecutive pitches was between 0.6-0.8 [2].

This gives us a model for perceptually equivalent pitches. If two consecutive pitches have a proportional difference in visual direction below this detection threshold, we deem the pitches to be (perceptually) equivalent to the batter. With this model in hand, we uncover pitch pairs of interest where we can study hitter decision making as a function of our modeled visual information.

We calculate β given the following formula:

$$\beta = \tan^{-1}\frac{2R(\partial\phi/\partial t)}{D(\partial\theta/\partial t)}$$

Where R is the radius of a baseball. $\partial\phi/\partial t$ is the rate of change of ϕ with respect to time – the angle between the observers eye and the height of the ball to the ground; D is the distance to the observer, and finally, $\partial\phi/\partial t$ is the rate of change of the angle from the vectors from top to the bottom of the projectile.

Figure 2 shows the visual directions for two consecutive pitches. We deem these to be "perceptually equivalent" if they maintain a proportional difference below a certain threshold value up to the point where the hitter makes a decision. We will solve for our thresholds later, and describe our methodology under "Model Performance". Lab experiments, as mentioned before, suggest this threshold value to be 0.6-0.8 [2].

6 Model Comparison of Tunneling

Testing the hypothesis that hitters can be fooled into thinking a pitch looks one way when it is truly another, previous researchers have scrutinized consecutive pitches that have a small euclidean

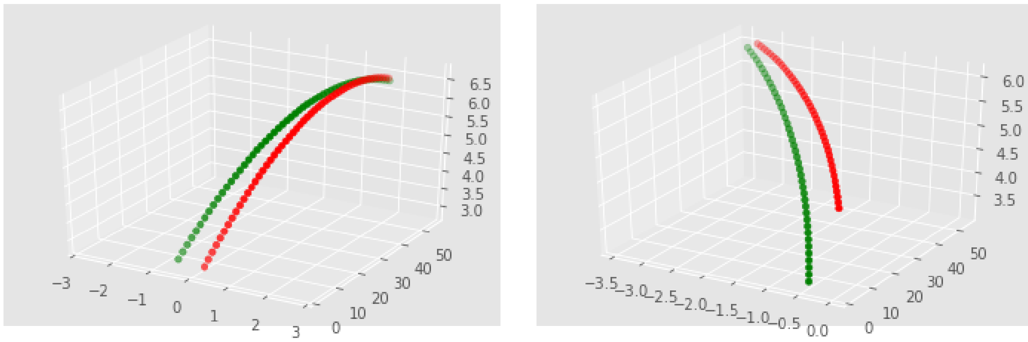

Figure 3: An example of a "tunneled" pitch pair under the old definition (on left) and under the new definition (on right).

distance at the decision point for some evidence of this effect. To do this required only looking at the pitches 23.5 feet from the plate – an arbitrarily defined "decision point".

To show how our model for perceptual equivalence differs lets consider two examples of pitch pairs under these two differing models of what the hitter deems similar trajectories.

Figure 3, shows a what this old definition would consider perceptually equivalent. In fact these pitches are perceptually equivalent under the new model, as well.

Figure 4, however is a pitch pair that is "tunneled" in our new model, but not in the old. The two pitches generally follow the same trajectory up to the decision point, despite the fact that they are released at slightly different positions. This varying release point is what makes the distance at 23.5 feet too large for it to be considered equivalent in the old model. This however is unnecessarily strict when considering that a hitter may not be able to pick up on the varying release point in the first place. As mentioned before $\frac{\beta_1 - \beta_2}{\beta_1}$ is under the detection level up until 23.5 feet, so we deem these to be equivalent to the batter.

In other words, what we have done is created a new similarity and distance measurement on the space of all pitches. As before we looked at only the distance at a point in 3d space, now we consider the entire trajectory up to the decision point, effectively using more information on the entire path.

7 Model Performance

In this section, we show how to fit this model to the pitch data given, compare it to the experimental evaluations by the baseball labs like perceptionaction.com, and then examine the predictive accuracy of our model. We top it off with an examination of our model on various contexts, to see how external factors may also affect the predictive power and accuracy of the model and of batter decision making.

We posit that there is some threshold visual direction β^* between two pitches that will determine if the batter thinks they are the same pitch. Now, although this information is not present in the dataset, we will use latent features to extract the similar information. In particular, we believe that if a batter sees two pitches with the same visual information, or that look identical, then their decision to swing on the first pitch is directly correlated with their decision to swing on the second. Now, in order to properly solve an optimization problem for this threshold, one would likely have to use complicated probabilistic techniques, and so we make the following simplifying assumption .Let $\beta(P_i, P_{i+1})$ be

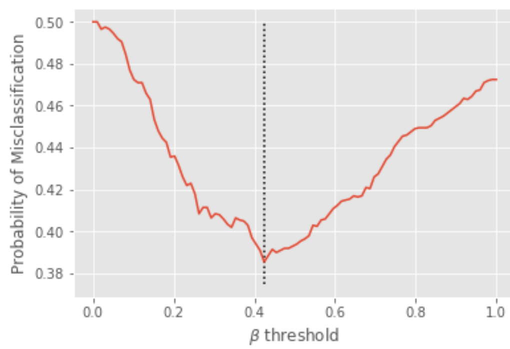

Figure 4: Determined threshold for swing based on Beta, and Gini Index, all contexts

the visual direction between two consecutive pitches, and S_i, S_{i+1} be their decisions to swing on these pitches respectively, then

$$\beta(P_i, P_{i+1}) < \beta^* \implies S_i = S_{i+1}$$

Solving for the optimal β^* in this scenario is very easy; we can reduce it to finding the best split in a 1 dimensional covariate with 1 dimensional outputs. For our loss, we chose between the Gini impurity and the information gain metrics, and ultimately decided on the gini impurity metric because the gini impurity is immediately interpretable for our problem domain. To recall, if we have multiple classes $C_1 \dots C_n$, and we observe each class for frequencies $f_1 \dots f_n$, then letting $p_i = \frac{f_i}{\sum_j f_j}$, then the *Gini Impurity* metric is given by

$$I_G(f_1 \dots f_n) = \sum_i p_i(1 - p_i)$$

Figure 4 shows us the optimal β^* for all contexts, i.e. for any count is around 0.4. We classify the hitters decision optimally at a .38 misclassification rate overall when we predict $\hat{S_{i+1}} = S_i$. Figure 5 shows us the optimal threshold for a hitter's count, i.e. where a batter has more balls than strikes. The varying β^* in these counts shows evidence for this being a key latent feature in decision making. There should be no difference in the distribution of pitch trajectories in these counts, yet the optimal β for predicting decision making is changing across these contexts. Since context obviously changes the hitters approach and their propensity to swing, we update our model to include context:

Let us formally define the sets $C_{s,b} = \{(P_i, P_{i+1}) \mid \text{Strikes on Pitch i + 1} = s, \text{Balls on Pitch i + 1} = b\}$ for all $s \in \{0, 1, 2\}$ and $b \in \{0, 1, 2, 3\}$. Then for each $C_{s,b}$ we find the optimal threshold $\beta^*_{s,b}$ to classify decision making based on the gini index described above.

With $\beta^*_{s,b}$ in hand we have a model for predicting decision making for each context.

Figure 6 compares our performance to predict swing-or-no-swing with the standard industry model – euclidean distance between consecutive pitches at 23.5 feet. The optimal misclassification rate is slightly higher at 0.41 percent. Overall our model is only correlated 0.10 with the industry standard classifier, meaning not only are we capturing different information we are doing slightly better because of it.

Lastly, Figure 7 is a classic ROC plot that shows us doing much better than predicting randomly what the hitter will do next, simply by using our optimal $\beta^*_{s,b}$ split for the appropriate context of P_{i+1}.

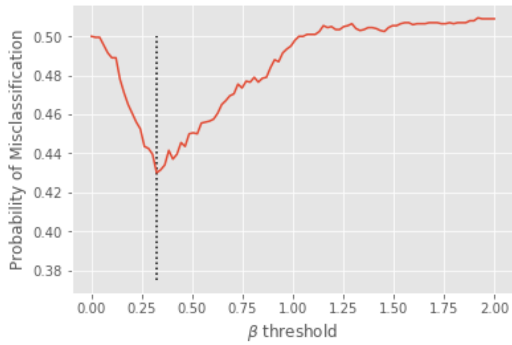

Figure 5: Determined threshold for swing based on Beta, and Gini Index, Hitter's Counts

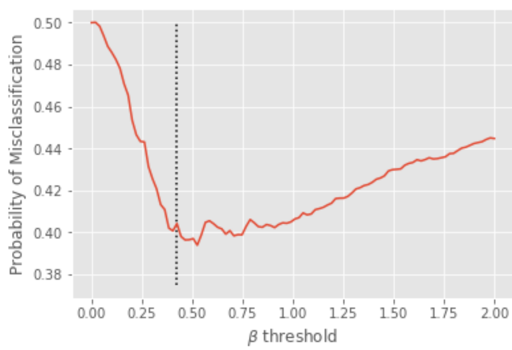

Figure 6: Determined threshold for swing based on previous definition of "Tunneling" using Gini Index

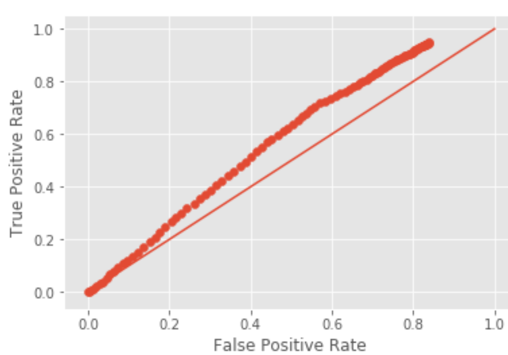

Figure 7: ROC Plot for Model accuracy in predicting swing-or-not-swing

8 Ranking Hitters through Simulation

In previous sections, we compare our approach to previous state-of-the-art models, and show that in general, our method tends to work equally, if not better than a method which fails to take psychophysics into account. However, this doesn't give us much information as to how applicable the model is in the real world. In general, there are two ways we can tackle this issue: by testing in the real world or in simulation. Clearly, if one was affiliated with the data division of a MLB team, then a large-scale study could be implemented to compare batter tendencies to the testing batter efficacy would be simple. Unfortunately, we didn't have such resources or time available to us for this survey, so we sought to see how to compare our model to a real batter in simulation.

In the following paragraphs, we shall lay out the general gist of how one *may* use simulation to compare batter prediction models to the reality. Unfortunately, our simulations were hampered by three factors: unrealistic swinging motions by the robot, difficulty in learning options for the robot (an issue with learning continuous control tasks in reinforcement learning and robotics), and time constraints. We do hope to pursue this avenue further as a more informal study in the future. Below, we lay the simulation and strategies that we attempted for this work, as well as some preliminary results.

The goal of the simulation is to have a scenario where a batter model can face pitches of various styles; this batter model is generated by combining decisions on pitches that he has seen in the last 8 years, and tested on pitches that he did not experience. On these holdout pitches, we can test how our perceived batter will do, which will allow us to rate batters and pitchers, which is a serious concern in the baseball trading and scouting industry.

We perform our simulation in Mujoco, a full-featured physics engine. We picked it for speed of execution, realistic dynamics models, and because of its excellent contact mechanics. Thanks to the PitchFX data, we have access to the trajectories for every single ball thrown between 2008 and 2014, and we design a method that allows us to replicate this trajectory in a Mujoco simulation. This allows us to visually reconstruct every trajectory for a batter to "witness". We design a robot with a bat-shaped object as it's arm: this will be our batter. It is unnecessary to simulate a human, because realistic movement of the bat is all that is necessary to evaluate performance. Herein lay our issue: it was difficult to synthesize bat movement similar to that of a professsional player.

We design an agent that allows the batter to swing "high" (in the top third of the box), "center" (in the middle third), and "low" (in the lower third). To allow the robot to perform these three actions, policies need to be learned, either via supervised or imitation learning setups, or via reinforcement learning. We didn't have the motion capture tools necessary to do imitation learning, and what we found is that learning an option via reinforcement learning is very difficult. This was the second issue with our goal.

Now given this simulation, where the batter observes visual information at every timestep (which we denote to be 0.01 seconds), the batter must make a choice of whether or not to hit the ball at the time. To construct such a model, what we do is a *nearest-neighbors* style of approach: the batter compares the current trajectory to all pitches that he has seen recently, and if the pitch is visually similar to another pitch, then take the same action he did on the old pitch (including when he swung). This defines an easy-to-construct model of batter hitting strategies. Evaluating the success of a batter simply reduces into checking how many times the ball actually makes contact with the bat. Rigorously, this model can be defined as a POMDP with partially observable agent states and binary rewards at the end of an episode. We omit a full description of this rigorous treatment for brevity.

Although we were unable to successfully create this simulation model and get experimental results with it, it remains an interesting avenue of work regardless. We think that if properly created, we can use such a model to aid scouts, and to help hitters figure out where their critical points are without having to actually interact with the batters in person. [**?**]

9 Discussion

In this work, we present a model which uses batter visual information to predict whether or not a batter will swing. As compared to previous approaches, we accurately model the perception system of the batter, and as a result, have more predictive power in our models. We extend the limited experiments run by other baseball labs, and show that the psycho-physical models do correspond better with batter data from the past decade than current models do. Although this visual information system is still very noisy, future improvements to the visual model and a more extensive investigation into pitch data will allow us to understand the batter decision process better.

References

[1] Biological Baseball
http://www.exploratorium.edu/baseball/articles/biological-baseball.html

[2] Rob Gray.
Pitch Tunneling And Perceptually Equivalent Pitches
http://perceptionaction.com/pitchtunnels/

[3] Rob Gray.
Spin Perception
http://perceptionaction.com/spinperception/

[4] Julie M Harris, Vit F Drga *Using visual direction in three-dimensional motion perception*. Nature Neuroscience, 229–233 (2005), 2005.

[5] R. Pendelton
Pitchf/x Nine Parameter Trajectory Model Fitting
http://baseball.physics.illinois.edu/PitchFX_9P_Model-4.pdf