Understanding Boss Battles: A Case Study of Cuphead

Andrew Wood and Adam Summerville
California State Polytechnic University
ajwood@cpp.edu and asummerville@cpp.edu

Abstract

"Bosses" – powerful, difficult enemies – have been a part of video games for the majority of their existence. Despite their integral role in many games, they have rarely been the focus of study. Cuphead is a run-and-gun style 2D action game with a large pool of boss battles, making it an ideal game for comparing and contrasting bosses, to determine what makes a compelling (or frustrating) boss battle. In this case study, we developed an ontology of 'shoot-em-up'-genre attacks and bosses. With this system for codifying a boss battle, we cluster the attacks using a Gaussian Mixture Model – which are then used to represent a boss as a "bag-of-attacks". We then use multinomial regression to predict the player experience of a boss given the parameterized boss.

Keywords

Cuphead, Boss Battles, Video Games, Clustering, Regression, Player Experience, SHMUP Ontology

Introduction

Boss battles exist in a multitude of genres – usually representing difficult, singular encounters unlike what the player experiences in the rest of the game. Cameron Wood suggests that boss characters are usually significantly stronger than other enemies, often having some significance to the plot of the game's story[1]. Bosses often have a unique attack-move set, and Mike Stout explains that they often pose as a roadblock to test the player's developed skills up to that point in the game[2]. This paper aims to piece together what information is important in determining what makes a video game 'boss battle' good, specifically in the game "Cuphead" [3].

Cuphead is classified as a 'classic run and gun action game heavily focused on boss battles.'[4] By its nature, it also falls under the categories of 'action', 'SHoot-eM-UP' (SHMUP), and 'bullet hell' games. As Heather Alexandra explains, the game is notoriously praised for its high level of difficulty and precise mechanical gameplay[5]. The game boasts 19 unique bosses to fight against, each with their own design, music, animations, attack sets, and challenges. This array of bosses presents a perfect setup for analyzing and directly comparing boss battles.

While each boss is unique, they share many features – similar attacks, patterns, platforming challenges, etc. This work is focused on determining which of these factors makes a 2D action shoot-em-up boss battle enjoyable. To do this we:

1. Develop an ontology for SHMUP attacks and bosses
2. Cluster SHMUP attacks to find similarities across different attacks
3. Perform a regression to determine which attacks and boss features are engaging and which are frustrating

While bosses have been the focus of some other work – we are the first to try to perform a "static analysis" (i.e. an analysis that does not rely on run-time data) of a boss encounter. In the rest of the paper we first describe how other works have considered boss battles, we then describe the SHMUP ontology, we then perform the analysis of the bosses, and finally we present the results of our study.

Related Work

While bosses have been a part of video games since their introduction in *DND* [6] – a game that emulates many of the characteristics of the table top game – they have seen much less research than something like video game levels. Siu et al. [7] develop a programming model for describing bosses in 2D action games, which is the basis for work by Butler et al. [8] to generate bosses via program synthesis. These two works strive to develop a model for bosses that is flexible enough to describe many different behaviors, while still being constrained enough so that random generation is able to produce playable bosses. In contrast to the work described in this paper – their work does not strive to understand, or generate, engaging/effective boss battles. However, it must be said that the ontology we developed for describing a SHMUP boss is lossy, and there is no way to map from a parameterized boss to a playable one.

Agriogianis [9] examines the history of bosses, and discusses features that he finds to be enjoyable in boss battles but performs no quantitative analysis to try to analyze boss battles.

Summerville et al. [10] did a similar analysis – find metrics to predict qualitative human ratings – however, their work focused on game levels in the domain of *Super Mario Bros.*. Fasterholdt et al. [11] did a deep dive into understanding jumping in 2D platforming games, which has many similarities to the data collection phase of this work – frame-by-frame analysis of game video – although their work was focused on finding the different parameters for jumps, not assessing the player experience that derives from those jumps.

Player experience modeling is its own field of study [12] with many different approaches such as physiological measurement or psychological survey. However, these approaches tend to try to assess a player's holistic experience instead of trying to predict the experience based on parameters found in the game. Guckelsberger et al. [13] attempt to remove the human in the loop aspect of player experience modeling by using artificial agents, but their approach requires an artificial agent – which might not be feasible to create.

Ontologies of video game features have been created, most notably the Game Ontology Project [14] of Zagal et al. However, the Game Ontology Project is designed to cover a wide variety of games and does not cover the specifics required to perform this work.

Research Approach

In this research, we first developed an ontology of the bosses' attack-moves so that their unique attacks could each be represented as a set of values to be directly compared with one-another, instead of attempting to describe attacks in plain English and compare their characteristics on a categorical level. We represent bosses as a *bag-of-attacks* (inspired by *bag-of-words* models in Natural Language Processing [15]) – i.e. each boss is represented as a count of the attacks of each type that it has represented. To find the attack type, the attacks are clustered via a Gaussian Mixture Model (GMM) – such that all attacks are grouped with attacks of similar style. To get the parameters for the attacks, we developed a tool for extracting data on these attacks from video-recorded gameplay of Cuphead taken from YouTube.

Given the parameters for the attacks and the attack clusters, we use the bag-of-attacks combined with other boss features (health, level type, etc.) as the independent variables for 6 multinomial regressions, each with its dependent variable being Likert-style ratings of a bosses' specific play aesthetic (a la Hunicke et al. [16]).

Method

Quantitative Data Extraction

For extracting data on boss attacks, we used gameplay video from YouTube[17]. We used YouTube videos recorded at 30 frames per second, with a resolution of 1280 x 720 pixels. Measurements were mostly taken using units of frames, pixels, and seconds. We created a custom tool in Python using Jupyter Notebook for taking measurements from individual frames of the video, shown in Figure 1. Each of the four considered bosses had all of their attacks parameterized according to the ontology – discussed in the next section.

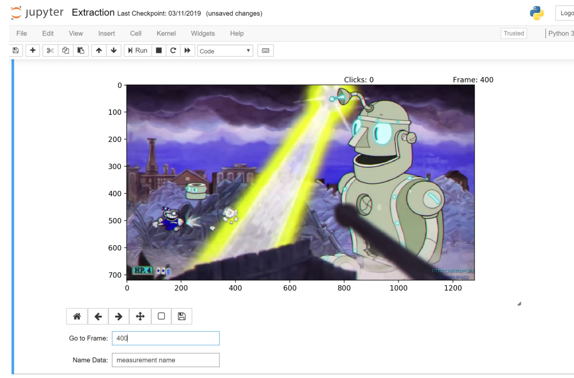

Figure 1: The frame data extraction tool. Users can step through the video frame by frame, recording measurements to pixel accuracy.

SHMUP Attack Ontology

To be able to represent the attacks in a way such that a boss can be represented as a bag-of-attacks, we require a set of primitives to be able to represent all of the attacks – thus the development of an ontology for describing the attacks in a SHMUP game. Initially, we referenced an unofficial catalog of terms that referred to common attacks in shoot-em-up games[18] and tried to categorize the attacks from Cuphead. We then used one-hot encoded values to list which category they belonged to. This quickly revealed its limits as some of the attacks did not fit into any of the catalog terms, while others fit into multiple. We also faced the issue of having to come up with a parameter set for each of the pre-defined attacks to define the variations within that attack category. Ultimately, this approach proved too complicated and would have produced extremely sparse data sets unfit for regressions.

Instead, we chose to define an attack as "any object onscreen that would damage the player if contact was made between the player character and the object". The primitive features of all attacks are:

- **Size** – Measured as the area of an attack in pixels^2. The size of an attack affects how difficult the attack is for a player to dodge, as well as how much screen space is available for the player to move around in, restricting player navigability.
- **Speed** – Measured in pixels/second. Fast-moving objects leave less reaction time to dodge, and make tracking attacks more difficult.
- **Concurrent Spawn Count** – An integer count of the number of identical objects spawned simultaneously, or near-simultaneously. Multiple objects spawned at the same time typically results in an array of attacks that move across the screen synchronously, generally creating wall-like formations that require precise movements for the player to avoid them.
- **Number of Cycles** – A count of the number of times the 'concurrent spawn wave' is launched in relatively quick

succession. An attack with 5 'concurrent spawns' and 3 'cycles' implies 3 successive waves of 5 bullets each.

- **Infinite Spawns** – A binary variable used to indicate if the number of cycles is theoretically infinite, where the number of cycles is heavily dependent on the player's rate of progress in the boss battle (e.g. if the player stopped attacking entirely, the boss would keep launching the same attack until its health reached a particular threshold).
- **Time Between Attack Waves** – Measured by the number of frames between the spawn of one wave and its subsequent wave, when an attack has 2 or more cycles. A value of 0 is given when there is only one attack cycle. Short intervals indicate a lot of attacks being launched in quick succession, making them more difficult to dodge and giving the player less time to plan movements.
- **Attack Cue Duration** – The attack cue is measured in number of frames between the start of a cue animation and the frame when the subsequent attack has fully formed on-screen. Short attack cues can make an attack very difficult to dodge, and longer cues can make attacks feel more fair.
- **Lifespan** – A measure (in frames) of an attack's maximum time on-screen. A longer lifespan indicates more 'clutter' on screen as other attacks are added over time, thereby increasing the number of objects for the player to keep track of and also increasing difficulty.
- **On-Screen Until Dead** – A binary value indicating if the 'lifespan' is theoretically infinite, assuming the player were to stop attacking. These objects would, therefore, stay on-screen until they were destroyed by the player.
- **Health** – The amount of 'hit-points' an attack can take before being destroyed. Increased health correlates with more screen time; low health indicates attacks that can be destroyed quickly as an alternative to dodging the attacks.
- **Invincibility** – A binary variable indicating whether an attack can be destroyed or not. It works in conjunction with the 'health' feature to represent 'infinite health', instead of using a large value as a proxy for infinity.
- **Parryable** – A decimal value between 0 and 1 that represents the fraction of attack objects that are parryable. Parryable attacks cannot be shot down, but their damage can be negated if the player is able to time a special move when they come into contact with the attack. More parryable attacks implies slightly easier avoidance – they are harder to avoid than those that can be shot and destroyed, but easier to avoid in that damage can be negated, even if not properly dodged.
- **Spawn Location** – This is actually 4 separate binary measurements. They are mutually exclusive, meaning that only one of the 4 values can be assigned 1, with all others being 0 for any given attack. The spawn is said to be from 'Off-screen', in 'Open Space', from the 'Boss', or from 'Another Attack'. Attacks that spawn from off-screen may be poorly telegraphed and difficult to dodge, and attacks from open space are often hard to predict and generally spawn near the player. Attacks from the boss are usually well-telegraphed by a boss animation. Attacks that spawn from other attacks may come as surprise and be relatively close to the player when spawned, making them challenging to dodge.

These features are able to describe the attacks such that they can be faithfully represented.

Attack Clusters

As mentioned above, the original idea to represent bosses as a bag-of-attacks led to a SHMUP glossary that lacked the ability to cover all attacks in Cuphead; however, the goal of finding commonalities in attacks so as to be able to group them together remained. Toward this end, we clustered the attacks using a Chinese Restaurant Process Gaussian Mixture Model [19]. A standard Gaussian Mixture Model represents the clusters as k different n-dimensional Gaussian distributions, and the data points are samples from those distributions. A Chinese Restaurant Process Gaussian Mixture Model does not fix the number of clusters a priori, and uses an infinite Dirichlet Process to determine the number of clusters as part of the inference. This was necessary for clustering the attacks as there was no predetermined number of attack types. For this work, Scikit-learn [20] was used to perform the inference. The mixture model found 15 clusters for the attacks – which post-hoc we describe as:

- **Standard Medium** – medium-sized attacks that generally move slowly and predictably, mostly spawning off-screen
- **Small Destroyable** – relatively small attacks that are mostly destroyable, and spawn from a boss or other attack, moving moderately fast
- **Bullet Hell** – standard 'bullet'-type attacks that typically aren't destroyable and move quickly in a linear fashion
- **Large Threats** – mostly large invulnerable attacks that either specifically target the player, or move very fast
- **Rumor Spawn 2 & 3** – the 2nd and 3rd-phase Rumor Boss (Large; little movement; high HP)
- **Vertical Attacks** – these 2 attacks span the vertical height of the screen
- **Jabs** – these 2 attacks are large 'jabs' by the primary boss on-screen (quickly moves forward, then back)
- **Medium, Short Cue** – medium size projectiles with short attack cues
- **Large Area** – very large-area attacks (some attacks – some bosses themselves)
- **Moving Walls** – these 2 attacks are both slow 'wall-like' attacks that force the player to move around them
- **Hilda Spawn 3** – Hilda Berg's huge final phase form
- **Frequent Beams** – these 2 attacks trigger on and off frequently, spanning the length of the screen
- **Grim Spawn 1 & 3** – the 1st and 3rd-phase Grim Matchstick spawn
- **Khal Spawn 1** – the 1st-phase Khal's Robot spawn
- **Khal Spawn 3** – the 3rd-phase Khal's Robot spawn

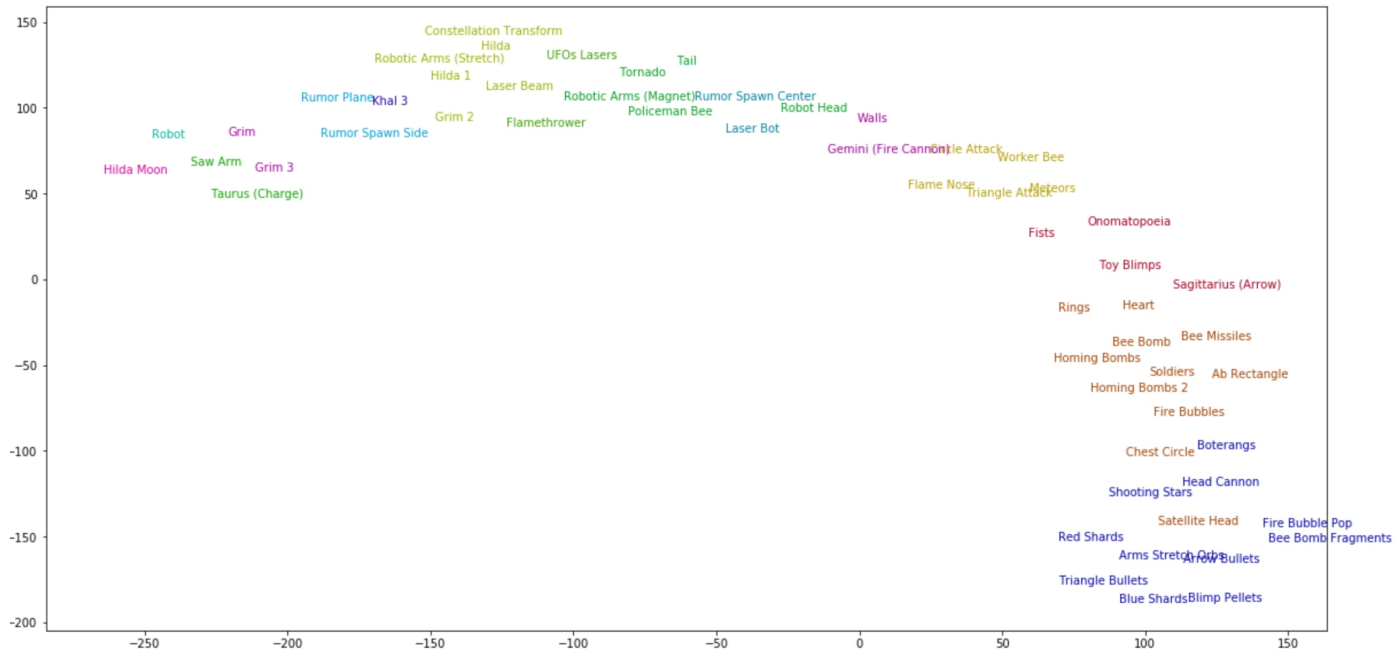

Figure 2: The t-SNE embedding of the boss attacks, showing how the boss attacks are similar (or dissimilar) from each other. Each cluster has a different color coding.

To visualize the clusters, we used t-distributed Stochastic Neighbor Embedding (t-SNE) [21] – a visualization tool that performs a dimensionality reduction by placing datapoints such that similar points are close to each other, while dissimilar points are far apart. This embedding can be seen in Figure 2. The exact positions of the attacks do not matter, but rather which attacks wind up closer to which is important. The upper-left portion of the mapping sees the unique, singular attacks for some of the bosses wind up close to each other, while the lower-right portion represents the more standard “bullet” style attacks – with the manifold representing increasing complexity to the attacks as one travels from the lower-right to the upper-left.

Player Experience Modeling

The ultimate goal of this research is to determine what aspects of boss battles correlate to positive player experiences. For the purposes of this study, the different boss aspects we chose to focus on are:

- **Platforming Difficulty**
- **Ability to Recognize Attack Patterns/Cues**
- **Avoiding Attacks**
- **Fight Duration / Boss Hitpoints (HP)**
- **Ability to Hit the Boss**
- **Number of Simultaneous Attacks**

These 6 categories will be referred to from here on as the bosses’ *aesthetics*. In order to find the correlations between a boss’ characteristics and its aesthetics, an understanding of the players’ impressions are needed. To get this information, we conducted an online survey.

The survey was formatted with three different parts.

1. **Player Background** – background information about the player, so that we could differentiate between responses from Cuphead/SHMUP fanatics and more casual video game players if needed
2. **Preferred Boss Aesthetics** – What players felt was important for boss battles in general on a 3 point Likert-like scale
3. **Favorite/Least Favorite Boss Aesthetics** – Players’ favorite and least favorite battles, out of the 19 bosses in Cuphead. With each of the responses, they are asked to provide ratings of their chosen bosses’ aesthetics using a 3 point Likert-like scale [22].

Due to the ordinal nature of the ratings, we *(1)* used mode as the distributions’ centers in the results section and *(2)* performed a categorical (as opposed to a linear) regression.

There were 41 participants, and the results of section 2 can be seen in Figure 3. We see the most stress on fairness, with the runner-ups being a challenging fight and a test of the player’s strategy. A wide range of bosses were chosen as participants favorite (#1 being Grim Matchstick) and least favorite (#1 being Dr. Khal’s Robot). Due to the time consuming nature of gathering the boss attack information, we could only consider 4 bosses – Rumor Honeybottoms, Hilda Berg, Grim Matchstick, and Dr. Khal’s Robot – which account for 27 different participant responses. In future work, we would like to expand the number of bosses considered. Figure 4 shows the boss aesthetics for all bosses.

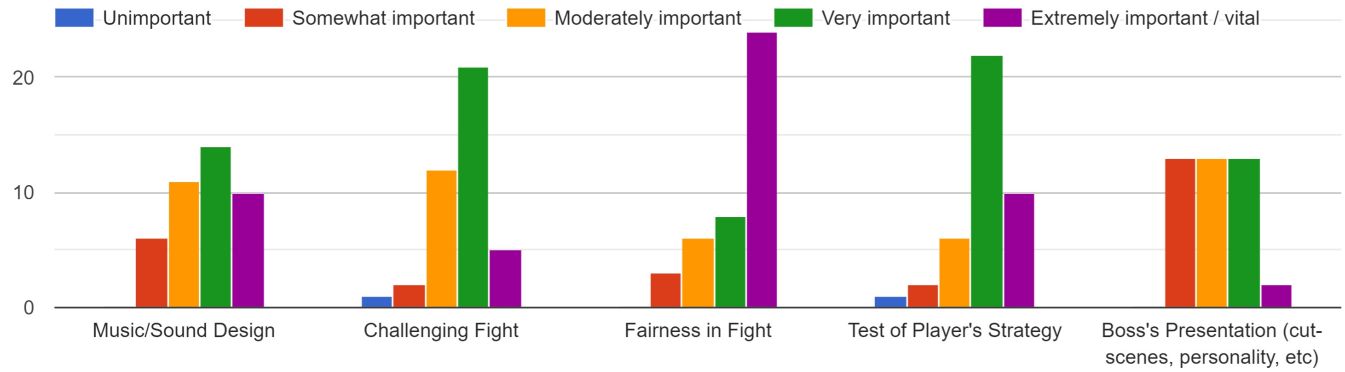

Figure 3: Bar graphs from the user study; describing the user ratings for how important each aspect is in creating an enjoyable boss battle. We see that fairness is the single most important reported factor, with challenge and test of strategy being close behind. Sound and presentation are much less important to players, although at least 5% of participants considered them vital.

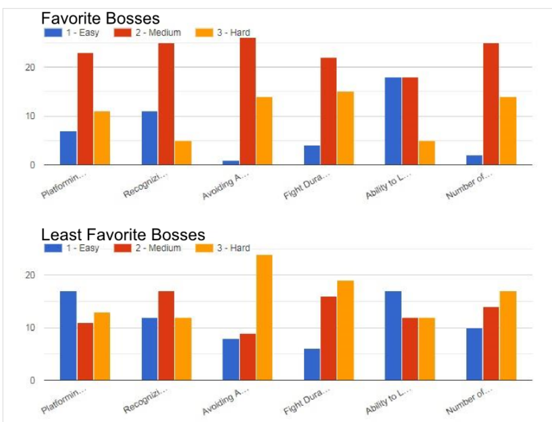

Figure 4: Bar graphs from the user study displaying the sum of players' ratings for their chosen bosses. We see that the favorite bosses tend to be rated as Medium across all fields, and the least favorite bosses are rated as harder to avoid.

Categorical Regression

While the participant ratings are ordinal (Easy, Medium, Hard), we chose to use a multinomial regression to predict the ratings. The key reason for this is that ordinal regressions assume that the effects of the independent variables are the same at all parts of the scale; an assumption that is unlikely to hold as a variable might easily have a different effect between Easy-Medium vs Medium-Hard. We chose to use a Least Absolute Shrinkage and Selection Operator (LASSO)[23] multinomial regression – as the LASSO regularization leads to sparse weights, i.e. performs variable selection. A key motivation for this work is to determine which features of a boss lead to which aesthetics – our hypothesis is that not all types of attack will be explanatory for some of the aesthetics.

Results

Table 1 shows the predictions made for each of the boss aesthetics from each of the 6 regressions. The predictions are listed, with the bold text representing that the model correctly predicted the mode of the survey results for that category and non-bolded labels represent an incorrect prediction (with the correct predictions in parentheses). As can be seen from the table, the 6 regressions correctly predicted a total of 20 of the bosses' aesthetic impressions, missing only 4 out of 24 – 83.33% accuracy.

It should also be noted that all of the incorrect predictions came from *attack avoidance* and *fight duration,* and for two of the bosses – *Hilda Berg* and *Grim Matchstick*. All of the incorrect predictions predicted Hard when Medium should have been predicted. Our hypothesis was that attack avoidance would be related to platforming difficulty, cue recognition, and the number of simultaneous attacks – however, Hilda Berg has the same ratings for those as Dr. Khal's Robot, but Dr. Khal's Robot attack avoidance is Hard as opposed to Medium for Hilda Berg. Seemingly, the ability for a player to avoid a boss's attacks are less correlated with those features than we hypothesized, meaning future work is required to determine what specifically makes a boss's attacks difficult to avoid. Similarly, we hypothesized that the health of a boss would be a good predictor for fight duration, but in actuality, this does not seem to be the case – Dr. Khal's Robot and Rumor Honeybottoms have 1250 and 1000 HP respectively, while Hilda Berg and Grim Matchstick have 2200 and 1200 respectively. Again, it would seem that there are boss features related to fight duration that are not being adequately captured by the current ontology.

Digging in to the learned weights of the regressions, there are some very clear trends in where and why certain aspects of the bosses hold so much weight in determining ratings. For example, there is a strong correlation between predicting a boss is difficult (rating it 'hard') in terms of the 'number of simultaneous attacks', and the number of attacks a boss has from cluster **Bullet Hell**. This makes sense because cluster **Bullet Hell** represents the 'classic bullet-hell' types of at-

	Platforming Difficulty	Cue Pattern Recognition	Attack Avoidance	Fight Duration	Ability to Land Hits	Simultaneous Boss Attacks
Rumor Honeybottoms	**Hard**	**Medium**	**Hard**	**Hard**	**Hard**	**Hard**
Hilda Berg	**Medium**	**Medium**	Hard (Medium)	Hard (Medium)	**Easy**	**Hard**
Grim Matchstick	**Hard**	**Medium**	Hard (Medium)	Hard (Medium)	**Medium**	**Medium**
Dr. Khal's Robot	**Medium**	**Medium**	**Hard**	**Hard**	**Hard**	**Hard**

Table 1: The results of the six multinomial regressions. Bold labels represent a correct prediction of the modal participant response and non-bolded labels represent an incorrect prediction (with the correct predictions in parentheses). We see attack avoidance and fight duration to be the most difficult to predict, with both Hilda Berg and Grim Matchstick being predicted as hard for both (when they should have been medium).

tacks, meaning that these attacks are generally very small, move very fast, and come out in waves of several projectiles. It makes sense that this type of attack should be the best indicator of a boss having a lot of simultaneous attacks.

Of course, there are instances where the results don't seem to represent what one would logically think to be strong correlations. One example of this is from the regression that predicted a player's impression of the difficulty of platforming during a specific battle. In the context of the game, 'platforming' refers to the difficulty in jumping between platforms (flooring) without falling to the bottom of the screen, where a player would take damage for failing to stay on grounds. When the prediction weights are analyzed here, an analyst might expect that the parameter representing a level as a 'platforming' level versus a 'flying' level (where the player does not have to platform because they can fly) should be the strongest determiner in predicting whether or not a boss battle is considered difficult to platform in. However, the weights actually suggest that small, destructible objects correlate better with whether or not a battle is difficult for platforming.

Conclusion and Future Work

Ultimately, a player's opinion towards a particular boss battle is just that - an opinion. It can be extremely difficult to formulate a numerical representation of a game that is accurate enough to create predictive models that can accurately represent player opinions. Selecting the right parameters to represent game information in a relatively simple way is a challenge, and the ontology developed for this research can certainly be improved (as will be described below). It is especially difficult to predict player preferences in general because it is entirely subjective, and not all people agree on what makes a boss battle good or enjoyable.

However, the results of this study were quite positive, and we believe the work can be improved even further. In this research, we were able to successfully use clustering techniques to categorize attack types in meaningful ways, and with this information run regressions to represent relationships between boss information and player satisfaction with surprising accuracy. The models are in no way perfect and certainly could not be relied on to predict satisfaction under any stakes, but the results of the study were promising.

Boss battles as a whole are an area of Game Design that have not been studied very much with formal research; however, the results of this study imply that meaningful research can and should continue to be done.

There are various improvements that could be made to this research to potentially produce stronger results in terms of more accurate/sensible attack clusters, better-fit regressions, and stronger correlations. For example, We believe more attack information could have been preserved and would strengthen correlations if measurements were recorded for 'height' and 'width', instead of by 'area'. This is because logically, in the game, a projectile's size/area doesn't matter as much as the length of the object. For example, there are some attacks that are 'walls'; they are tall and thin. These attacks are extremely difficult to dodge because the player must move a far distance to get around them. However, when only the area is observed, a regression wouldn't identify that a wall with small area is difficult to dodge because it would have similar area to smaller projectiles. Unfortunately, making this change would mean re-extracting a lot of data, which would be quite time consuming.

For future work, we would like to gather data on more bosses. We believe that with this added data, our regressions may show clearer trends. Having extra boss' data that was not processed in the regression could also serve as a test of the model's ability to effectively predict aesthetic ratings on a boss whose data did not influence the formulation of the model itself. We could potentially take the boss' attack data, fit it into the existing clusters, and use those values with the existing regression models to predict what the Likert responses would be, and compare that to the existing survey results.

Another analysis we would like to complete is: given the aesthetic ratings of a boss and a player's self-determined experience level with games matching the genre of Cuphead, could a model predict the players preference of the boss; could it predict if it is a favorite or least favorite? If we are able to successfully complete this analysis, we would then like to do an end-to-end prediction on boss preference. Given a boss' attack set and a player's preferences, could we categorize the attacks, use that information to predict the boss' aesthetic ratings, and based on those predicted ratings, predict a player's impression of the boss?

We believe that these improvements and additional analysis would provide a good foundation for beginning to determine: What makes a good boss battle?

References

[1] Cameron Wood. Designing a boss fight: Lessons learned from modern games. In *https://gamedevelopment.tutsplus.com/tutorials/designing-a-boss-fight-lessons-learned-from-modern-games–gamedev-2373*, 2012.

[2] Mike Stout. Boss battle design and structure. In *https://www.gamasutra.com/view/feature/134503/boss_battle_design_and_structure.php?page=3*, 2010.

[3] StudioMDHR. Cuphead, 2017.

[4] Cuphead wiki, 2019.

[5] Heather Alexandra. Cuphead: The kotaku review, 2017.

[6] Gary Whisenhunt and Ray Wood. dnd, 2017.

[7] Kristin Siu, Eric Butler, and Alexander Zook. A programming model for boss encounters in 2d action games. In *Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2016.

[8] Eric Butler, Kristin Siu, and Alexander Zook. Program synthesis as a generative method. In *Proceedings of the 12th International Conference on the Foundations of Digital Games*, page 6. ACM, 2017.

[9] Theodore Agriogianis. The roles, mechanics, and evolution of boss battles in video games. 2018.

[10] Adam Summerville, Julian RH Mariño, Sam Snodgrass, Santiago Ontañón, and Levi HS Lelis. Understanding mario: An evaluation of design metrics for platformers. In *Proceedings of the 12th International Conference on the Foundations of Digital Games*, page 8. ACM, 2017.

[11] Martin Fasterholdt, Martin Pichlmair, and Christoffer Holmgård. You say jump, i say how high? operationalising the game feel of jumping. 2016.

[12] Josef Wiemeyer, Lennart Nacke, Christiane Moser, et al. Player experience. In *Serious Games*, pages 243–271. Springer, 2016.

[13] Christian Guckelsberger, Christoph Salge, Jeremy Gow, and Paul Cairns. Predicting player experience without the player.: An exploratory study. In *Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, pages 305–315. ACM, 2017.

[14] José P Zagal and Amy Bruckman. The game ontology project: Supporting learning while contributing authentically to game studies. In *Proceedings of the 8th international conference on International conference for the learning sciences-Volume 2*, pages 499–506. International Society of the Learning Sciences, 2008.

[15] Zellig S Harris. Distributional structure. *Word*, 10(2-3):146–162, 1954.

[16] Robin Hunicke, Marc LeBlanc, and Robert Zubek. Mda: A formal approach to game design and game research. In *Proceedings of the AAAI Workshop on Challenges in Game AI*, 2004.

[17] Prosafia Gaming. Cuphead - all bosses & ending, 2017.

[18] Bulletmagnet. Unofficial shmups glossary. In *http://www.sega-16.com/2005/04/unofficial-shmups-glossary/*. Sega-16, 2005.

[19] Michael D Escobar and Mike West. Bayesian density estimation and inference using mixtures. *Journal of the american statistical association*, 90(430):577–588, 1995.

[20] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[21] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.

[22] Rensis Likert. A technique for the measurement of attitudes. *Archives of psychology*, 1932.

[23] Fadil Santosa and William W Symes. Linear inversion of band-limited reflection seismograms. *SIAM Journal on Scientific and Statistical Computing*, 7(4):1307–1330, 1986.