

Stephany Liu (liu48185)

William Zeng (wizeng)

Exploring Video Game Recommendation Techniques

Background

Although extensive research has been done on the front of movie recommendation systems, relatively less work has been published on the subject of video game recommendations. Video games are ubiquitous in America, with 75% of American families owning at least one video game console and 60% of Americans playing video games. However, one major problem is that there is no formal, centralized platform to help consumer decide which video games or consoles to purchase. Given that consoles are typically around \$300 and each video game is around \$60, making the right decision about purchasing a video game is very important. Given that the video game industry is worth \$22.41 billion in revenue in the US according to the ESA, which is expected to increase much more in future years with the advent of mobile gaming, offering high quality recommendations for video games could potentially result in great profits.

Objective

In order to help consumers decide which video games they would enjoy and want to purchase, we will analyze data about a gamer's trends and favorites, and explore various recommendation techniques backed by video games ratings and sales data to determine other games that the user might enjoy. We can then predict how much a user would enjoy a previously unseen game.

Approach

Data Source

Through use of GiantBomb's API, we scraped data about 100 user reviews of games for our preliminary dataset. This dataset included details about the user who wrote the review, the date of the review, the title and description of the game, and the rating given to the game (from 1 to 5 stars, with only whole numbers allowed), among other metrics. We scraped game metadata and ratings from IGN, obtaining the rating, category, and console for over 14000 games. In the future, this dataset can be expanded to include reviews from more popular gaming social networks, such as Steam, Playfire, and Internet Game Database (IGDB.com).

Classification of Scores

As a baseline, we used metadata (username and description of game review) obtained from GiantBomb to classify a review to a score (1 to 5) using a random forest classifier, which is effective in dealing with high-dimensional data such as this dataset.

Collaborative Filtering

We will then explore the merits of item-item and user-user collaborative filtering, various learning algorithms, and different methods of feature engineering in order to determine an optimal way to recommend video games.

Evaluation

In our baseline, we evaluated our results by calculating the accuracy score of the prediction using the random forest classifier. In future work, we hope to evaluate our results by computing the root-mean squared error of the predicted enjoyment for unseen video games. The predicted enjoyment can be computed as some combination of the time spent playing a game, the

number of trophies accomplished per game, and the rating given to a game, depending on the consistency of our final dataset.

Baseline Summary

For our baseline, we focused on using just word features of review descriptions to predict the rating (whole number in the range of 1 to 5). This is a basic metric for predicting a user's reception of the game, and provides a foundation for the rest of the project as it utilizes much of the framework for more advanced machine learning algorithms and sets up our pipeline. For our algorithm, we preliminarily implemented a random forest classifier.

To train our baseline, we created a training set by using numpy to vectorize username and descriptions of game reviews into counts by word and tf-idf matrix, respectively, using those as features. We then arbitrarily split this set into a train and test set, and trained our random forest classifier using the training half.

To test our baseline, we used the features of the test set to predict the review the user would have given the game. We evaluated based on the accuracy of the classifier in predicting scores, or the number of correct scores divided by the total number of reviews tested.

Preliminary Results

As previously mentioned, our baseline task was to predict the rating of a game based on the review description. We chose this since it required developing most of the infrastructure for predicting how much a user would enjoy a previously unseen game.

Upon training a random forest classifier using scikit-learn, we achieved an accuracy of 0.48 in predicting the ratings of games. Compared to 0.2, the expected value of the accuracy from random guessing one of five ratings, this score is significantly better, indicating that our baseline is on the right track.