

CSE 258, Fall 2018: Assignment 1

Instructions

In this assignment you will build **recommender systems** to make predictions related to reviews of Clothing, Shoes, and Jewelry on *Amazon*.

Solutions will be graded on Kaggle (see below), with the competition closing at **5pm, Monday November 19** (note that the time reported on the competition webpage is in UTC!).

You will also be graded on a brief report, to be submitted electronically on gradescope by the following day. Your grades will be determined by your performance on the predictive tasks as well as your written report about the approaches you took.

This assignment should be completed **individually**.

To begin, download the files for this assignment from:
<http://jmcauley.ucsd.edu/data/assignment1.tar.gz>

Files

train.json.gz 200,000 reviews to be used for training. It is not necessary to use *all* reviews for training, for example if doing so proves too computationally intensive. While these files are one-json-per-line (much as we have seen so far in class), you may find it useful to represent them more concisely in order to produce a more efficient solution. The fields in this file are:

itemID The ID of the item. This is a hashed product identifier from Amazon.

reviewerID The ID of the reviewer. This is a hashed user identifier from Amazon.

helpful Helpfulness votes for the review. This has two subfields, 'nHelpful' and 'outOf'. The latter is the total number of votes this review received, the former is the number of those that considered the review to be helpful.

reviewText The text of the review. It should be possible to successfully complete this assignment *without* making use of the review data, though an effective solution to the category prediction task will presumably make use of it.

summary Summary of the review.

price Price of the item.

reviewHash Hash of the review (essentially a unique identifier for the review).

unixReviewTime Time of the review in seconds since 1970.

reviewTime Plain-text representation of the review time.

category Category labels of the product being reviewed.

pairs_Purchase.txt Pairs on which you are to predict whether a product was purchased.

pairs_Category.txt Pairs (reviewerID and reviewHash) on which you are to predict the category of an item. (Not relevant for CSE258)

pairs_Rating.txt Pairs (reviewerIDs and itemIDs) on which you are to predict ratings. (Not relevant for CSE158)

test_Category.json.gz The review data associated with the category prediction *test* set. Again, the field that you are trying to predict has been removed.

baselines.py A simple baseline for each task, described below.

Please do not try to crawl these products from Amazon, or to reverse-engineer the hashing function I used to anonymize the data. I assure you that doing so will not be easier than successfully completing the assignment.

Tasks

You are expected to complete the following tasks:

Purchase prediction Predict given a (user,item) pair from 'pairs_Purchase.txt' whether the user purchased the item (really, whether it was one of the products they reviewed). Accuracy will be measured in terms of the *categorization accuracy* (fraction of correct predictions). The test set has been constructed such that exactly 50% of the pairs correspond to purchased items and the other 50% do not.

Rating prediction Predict people's star ratings as accurately as possible, for those (user,item) pairs in 'pairs_Rating.txt'. Accuracy will be measured in terms of the *mean-squared error* (MSE).

A competition page has been set up on Kaggle to keep track of your results compared to those of other members of the class. The leaderboard will show your results on *half* of the test data, but your ultimate score will depend on your predictions across the *whole* dataset.

Grading and Evaluation

This assignment is worth 25% of your grade. You will be graded on the following aspects. Each of the two tasks is worth 10 marks (i.e., 10% of your grade), plus 5 marks for the written report.

- Your ability to obtain a solution which outperforms the baselines on *the unseen portion of* the test data (5 marks for each task). Obtaining full marks requires a solution which is substantially better (i.e., at least several percent) than baseline performance.
- Your ranking for each of the tasks compared to other students in the class (3 marks for each task).
- Obtain a solution which outperforms the baselines on *the seen portion of* the test data (i.e., the leaderboard). This is a consolation prize in case you overfit to the leaderboard. (2 mark for each task).

Finally, your written report should describe the approaches you took to each of the tasks. To obtain good performance, you should not need to invent new approaches (though you are more than welcome to!) but rather you will be graded based on your decision to apply reasonable approaches to each of the given tasks (5 marks total).

Baselines

Simple baselines have been provided for each of the tasks. These are included in 'baselines.py' among the files above. These baselines operate as follows:

Purchase prediction Find the most popular items that account for 50% of purchases in the training data. Return '1' whenever such an item is seen at test time, '0' otherwise.

Rating prediction Return the global average rating, or the user's average if we have seen them before in the training data.

Running 'baselines.py' produces files containing predicted outputs. Your submission files should have the same format.

Kaggle

We have set up a Kaggle page to help you evaluate your solution. You should be able to access the competition via your UCSD address. The Kaggle pages for each of the tasks are:

<https://inclass.kaggle.com/c/cse158258-fa18-purchase-prediction>

<https://inclass.kaggle.com/c/cse158-fa18-category-prediction>

You may need to follow the private links to access the competition pages:

purchase prediction: <https://www.kaggle.com/t/e5845e78758440038a2d4d7faaf2f2be>

rating prediction: <https://www.kaggle.com/t/ffe2bb5970974a53860068c7190b6cce>