
HM Recommendation System

Author(s)

csci567_id05: Zhijie wu, Fanyou Meng, Honghu Luo

Abstract

1 In this project, we need to recommend top 12 personalized articles to each customer
2 shown in the H&M dataset. In the process, we used some data preprocessing
3 techniques to reduce the data size and applied several baseline models. Since the
4 results of a single baseline model did not meet our expectations, we employed
5 several ensemble techniques in the final step to further optimize the results.

6 1 Introduction

7 **Team ID:** csci567_id05

8 **Team Member:** Zhijie Wu, Fanyou Meng, Honghu Luo;

9 **Team Strategy:** Baseline + Ensemble.

10 1.1 Overview

11 To get our scheme closer to the given dataset, we first took a deep look at the importance of features
12 (via sklearn.ensemble) and the data distribution of the dataset (provided by Kaggle). We then
13 discovered that the dataset was so large that most matrix-based recommendation algorithms could not
14 handle it. Therefore, we adopted various data preprocessing techniques in different baseline models
15 to reduce the size of the H&M dataset, such as time constraints and clustering methods, etc. As
16 mentioned in the abstract section, we experimented different baseline models, the best ones were:
17 **SVD_Sort (0.0221), K-means + Predictions on transaction dates Strategy (0.0226), K-means**
18 **+ SVD_Sort (0.0218)**. According to the characteristics and data requirements of different models,
19 we performed different data processing and filtering techniques for each model, and fine-tuned the
20 hyperparameters as much as possible. In the second section, we will introduce the specific details and
21 corresponding codes of each baseline. After this, we chose to apply the ensemble techniques with the
22 predictions generated by these better performing models, which will be described in detail in Section
23 III. Meanwhile, in Section IV, we also briefly introduce other attempted schemes and their results.

24 2 Baseline

25 In this section, we will introduce some baseline models that we used and the results we obtained.

26 2.1 SVD_Sort Strategy

27 2.1.1 General Ideas

28 In this strategy, I employed a package called **reco [3]**, and we considered popularity to be one of the
29 most valuable tags to help us make recommendations. Then, to combine popularity and purchase
30 date, I added another column called **pop_score**, which means that if an item was purchased recently,
31 it will have a higher **pop_score** than if it was purchased a long time ago; According to such tuple:

32 { 'user ': " customer_id ", 'item ': " article_id ", 'value ': ' pop_score ' }

33 I trained a singular value decomposition (SVD) matrix which was helpful for us to extract features
34 and correlations from the user-item-pop_score matrix.

35 In the next step, I got four datasets with the highest volume of transactions and iterated through all the
36 customers to find if they were in these four training sets, if so, this algorithm would work based on
37 the results of the SVD and the list of most popular articles (via pop_score) to recommend. Otherwise,
38 the algorithm would only make recommendations based on the list of most popular articles.

39 2.1.2 Data Processing and Cleaning

40 **Data Processing:** In this part, I paid more attention on recent data, so for the SVD matrix, I only
41 selected 16 weeks of transaction data to build this model. I used pop_factor to show the popularity of
42 one article, (which equals 1 divided by the number of days between the date of purchase and the last
43 day of record) then summed up pop_factor of all articles as a new attribute named pop_score for each
44 article.

45 Also, I took the transaction data of every four weeks as the four training sets, in which we could find
46 the dataset of our customer group, and the smaller the sequence number of the dataset (from train1 to
47 train4), the higher its priority.

48 2.1.3 Result

49 As a result, I ran this model for about 5 times with minor changes on parameters. For instance, I may
50 revise the learning_rate of Funk_SVD to 0.002 and reset iterations to 150; then I increased the size
51 of the training set. However, the results were not much different, the score always hovered around
52 0.0220, and the highest score via this algorithm was 0.0221 which was used as one of our baselines.

53 2.1.4 How to Run the Code

54 As you can see, there is a file in the zip file named H&M svd that is the code for this algorithm, and
55 after run the script:

```
56 pip install -r requirements.txt
```

57 you can get all the dependencies needed, then revise the path of those source files like transaction.csv,
58 article.csv and so on, to make sure it can read these files successfully, then run H-M-svd.ipynb file,
59 you can finally get a submission.csv file.

60 2.2 K-means + SVD_Sort Strategy


61 2.2.1 General Ideas

62 In this case, we decided to use K-means algorithm firstly to divide all customers to three clusters
63 based on their attributes, and then divided transaction.csv to corresponding clusters. Then we used the
64 separated dataset to train our SVD_Sort model (which has been introduced in the previous section)
65 so that we could make more precise recommendation since different clusters might have different
66 popular articles. The difference between this algorithm with the previous one is that we can find the
67 most popular articles among a cluster of customers and make recommendation, which we thought
68 was much more rational than the previous one.


69 2.2.2 Data Processing and Cleaning

70 **K-means:** After trying several baseline models, we started thinking about clustering methods, which
71 can be used to train individual recommenders for each group of customers and provide complementary
72 information to the models. For clustering, we used the K-means algorithm which is supported by
73 well-developed open source packages. The random forest regressor imported from the sklearn public

74 package can rank the importance of features from a dataset. It accepts the entire dataset as input and
 75 returns feature importance rankings. We decided to cluster based on customer information instead of
 76 article information, and let a random forest regressor rank the features of the customers. As the graph
 77 (b) shown, we used the top five user features as the input to the K-means algorithm. In K-means,
 78 it iterated from K= 1 to 10 and generated the graph (a), and finally we use the Elbow Method to
 79 determine the final value of K to be 3 (although 4 seems to be fine too). As the result, we clustered
 80 customers to 3 groups and produced a new file customers_clustered.csv.



(a) Elbow Method For Optimal K



(b) Feature Importance Ranking

81 **SVD_Sort:** After reading customers_clustered.csv, we divided the transaction data into 3 groups
 82 with respect to the corresponding customer cluster, and then used such (customer, transaction) pairs
 83 to make recommendations for customers in each group.

84 2.2.3 Result

85 **SVD_Sort:** Compared to the previous algorithm, this time I chose to use about 40 weeks' data
 86 to build SVD matrix since the data granularity of the previous version was not satisfied, and even
 87 revised volumes of the training set to cover more customers in our transaction data. (If we choose all
 88 transaction data, it will exceed the limit of RAM in the **google collab**). As a result, our final score
 89 was 0.0218. (If we didn't change the size of the training sets, the result didn't change much. Because
 90 when we only use 16 weeks' data for SVD, leading to lots of our prediction were just the same as the
 91 12 most popular articles recently). Though our final score was not improved, we used this as another
 92 important baseline.

93 2.2.4 How to Run the Code

94 **K-means:** Simply run k-means-and-feature-importance-for-customers.ipynb with the raw H&M
 95 dataset as input.

96 **SVD_Sort:** After K-means, we can get a new csv file which has one more column named clusters to
 97 show which group this customer is in, then run h-m-Kmeans-svd.ipynb to run this algorithm, and
 98 finally will get a submission.csv file.

99 2.3 K-means + Predictions on transaction dates Strategy

100 2.3.1 General Ideas

101 This strategy uses the K-means algorithm to cluster customers into 12 categories with digital rep-
 102 resentation of all attributes of customers. Then we shaped the transaction into weekly purchases,
 103 and use the data in the last week to train our model. For customers that we don't have many data to
 104 analyze, we can use popular articles in the last week. Lastly, we use an equation with respect to x,
 105 which is number of day before the week that we are predicting on

$$(25000/\sqrt{x}) + 150000 * e^{-0.2x} - 1000$$

106 to compute the values of each transaction to make predictions. After combining processed prediction
107 and popular articles, it can produce recommendations for each customer.

108 **2.3.2 Data Processing and Cleaning**

109 **Processing:** It changes all customer attributes from word representation to digit representation. And
110 it uses the standardScaler normalization method to normalize customer data.

111 **Cleaning:** During data processing, it might involving many steps for copying data. Once we found
112 out that it will not be used again, we will delete the data to prevent excessive RAM usage.

113 **2.3.3 Result**

114 This strategy makes predictions on each cluster group of customers at a time, and combines them at
115 the end. It gives 0.0226 in comparison to our previous models, they are around the same. But this
116 strategy is very time efficient and it only takes 150s to execute and that includes the reading process
117 for huge .csv files.

118 **2.3.4 How to Run the Code**

119 After installing the dependencies and changing the path for input files, you can simply run the .ipynb
120 file on kaggle with an accelerator GPU, and it will produce the recommendation for customers.

121 **3 Ensemble**

122 Generated better recommendation results by aggregating the results of each baseline model.

123 **3.1 General Ideas**

124 When our baseline models have reached a stalemate, we learned an ensemble approach from an open
125 source material on Kaggle, which allowed us to reuse the predictions generated by the less-than-ideal
126 baseline models and synthesized a more robust result - which demonstrated a much higher accuracy
127 during the test.

128 **3.2 Results and Methods**

129 At very beginning, I selected the top 3 best performing models among all the baseline models:
130 **SVD_Sort** (0.0221), **K-means + Predictions on transaction dates Strategy** (0.0226), **K-means +**
131 **SVD_Sort** (0.0218). To reduce any possible bias, I first trained an **NCF** recommendation model with
132 a small amount of data - this model computes each (customer, article) pair and then makes predictions.
133 For this reason, it is difficult for this algorithm to make recommendations on such a large amount of
134 data. However, here I limited its computations to the articles that have been recommended by the first
135 three baseline models, which greatly reduced its workload and had the ability to re-examine all the
136 previously recommended products from a more objective perspective. The final result was 0.0231
137 that was a huge improvement from the previous baseline models.

138 In the next step, we found an open source prediction which has a high test score (0.0238) and I
139 applied the ensemble tricks again with it and all the predictions we generated before. The methods I
140 tried was the following:

141 **Majority Voting:** This is a very simple algorithm that assumes all recommendation models have
142 equal weights and selects the top 12 articles with the most mentions. The test result of this method
143 finally reached 0.0239.

144 **Average Weights:** I used each model's test score divided by the sum of all test scores as the weight
145 for each model and selected the top 12 highest weighted articles for each customer. The test result is
146 0.0239 that is indistinguishable from the majority voting method.

147 **Random Weights Fine-tuned on the Dev Set:** This time I used random weights to reduce the
148 inherent bias that can be present with classical methods. The random weights were then slightly
149 fine-tuned on a small development set. The final result reached 0.0240 after several adjustments.

150 3.3 How to Run the Code

151 Run recommenders-ncf.ipynb with one extra input res_blend.csv that aggregated the results from the
152 top 3 baselines. The input file res_blend.csv is an output of h-m-ensemble-magic.ipynb. Run h-m-
153 ensemble-magic.ipynb with input files submissions_{i}.csv, which are the outputs from baselines.

154 4 Conclusion

155 4.1 Challenge

156 Looking back on the entire project, we encountered many challenges such as large data volume and
157 high time cost.

158 **Large data volume:** At the early stage of this project, we checked a lot of information online to give
159 us inspiration to deal with this problem, and the first tough challenge we met was that the amount
160 of our data was so huge that it always exceeded the limit of RAM in the **kaggle** platform, then we
161 migrated our code to **Google Colab**. However, things didn't go well, then we had to choose partial
162 data to run our code on some particular models.

163 **High Time Cost:** When we wanted to build a recommendation system, we needed to train cor-
164 responding model, which was really time consuming. For instance, when we used NCF to make
165 predictions, it took more than 20 hours to train such a model, but it eventually reported errors because
166 of the limit of RAM, which wasted a lot of time.

167 4.2 Achievement

168 Before the start of this project, the members of our group had not been exposed to such a large amount
169 of data, so we did not conduct data screening but started building the model directly. However, due to
170 the large amount of data, many models collapsed during the operation, and some of them were very
171 expensive to generating predictions. In fact, we have tried more algorithms than mentioned in the
172 report. Due to the difficulties, we had to re-observe the data and realized that data filtering can greatly
173 improve the efficiency after reviewing various materials. For example, the clustering algorithm does
174 a good job. Meanwhile, many features in a sparse dataset are probably not important at all, we found
175 we can abandon them boldly.

176 In addition, through our actual practice, we observed that the amount of data used for model training
177 is not the more the better. When we trained with the most recent 8 weeks of data, the results were
178 not much different from 16 weeks of data, and the results of one year of data were more similar to
179 the results of training with 16 weeks of data. What this actually tells us is that in such a dataset that
180 is closely tied to the timeline, the distant data is likely to be meaningless, and this makes sense in
181 reality. The type of product that customers loved 2 years ago may have fallen out of favor lately.

182 For this project, we slightly adopted customer features in training, and did not have a good method
183 and not enough time to apply article features, which included important image information. We guess
184 that's why our scores got stuck in a certain stalemate and struggled to improve. In the future, we will
185 consider adding article and image information to the scheme design.

186 4.3 Failed Version

187 In this part, we will briefly describe some algorithms that have been tried but discarded due to
188 unsatisfactory test scores.

189 4.3.1 Turicreate Model [1]

190 At the beginning, I thought this turicreate model was really helpful for our problem, and this article[1]
191 was mainly talking about how to build a recommendation system based on **History Purchase Data**.
192 Then I came to **Turi Api Documentation [2]** to learn about this model, and I found that this model
193 had many strategies to make recommendation, for instance: Popular-based recommenders, item
194 content recommenders and item similarity models (I mainly use these three strategies to build our
195 recommendation system).

196 However, when I tried to use popular-based recommenders to plug in our training set, I found that our
197 transaction.csv file was so large that it was impossible to use pandas.pivot_table to build a matrix for
198 the whole data. (That always exceeded the limit of Ram in Kaggle Platform). Then I had to migrate
199 my code to **Google Collab** and gave up normalized matrix. Then I tried to use collaborative filtering
200 model to build this recommendation system via two different strategies: Cosine Similarity and
201 Pearson Similarity to give recommendation based on user-item matrix, and even use some strategies
202 like RMSE to evaluate our model. As a result, I took about one week to learn this model and built a
203 recommendation system with three different strategies, but the highest score of all was only 0.00450.
204 From partial to all data, I tested 7 times in total, each time combining different models and strategies
205 to make recommendations, but didn't end up with a satisfactory score (much lower than expected).
206 Therefore, I decided to give up this scheme and use other models as our baseline.

207 4.3.2 Two Towers

208 The Two Towers model is a deep neural network architecture that first analyzes the characteristics of
209 customers and articles separately, and then aggregates the information together in the final layers to
210 make recommendations. It was developed by tensorflow-recommenders and it already had a decent
211 test score when I only entered 8 weeks of transaction data. However, after including all 16 weeks of
212 data, this model took over 8 hours to train and 5 minutes per 1000 customers to make predictions.
213 After finishing training the model, I didn't have enough time and equipment to support it to complete
214 all the predictions, so I had to give up this method.

215 Reference

216 [1] Tjokro, Moorissa. "How to Build a Recommendation System for Purchase Data (Step-by-Step)." Medium,
217 DataDrivenInvestor, 15 Oct. 2018, [https://medium.datadriveninvestor.com/how-to-build-a-recommendation-](https://medium.datadriveninvestor.com/how-to-build-a-recommendation-system-for-purchase-data-step-by-step-d6d7a78800b6)
218 [system-for-purchase-data-step-by-step-d6d7a78800b6](https://medium.datadriveninvestor.com/how-to-build-a-recommendation-system-for-purchase-data-step-by-step-d6d7a78800b6).

219 [2] "Recommender." Recommender - Turi Create API 6.4.1 Documentation, 9 Oct. 2020,
220 <https://apple.github.io/turicreate/docs/api/turicreate.toolkits.recommender.htm>.

221 [3] mayukh18, Mayukh18. "Reco: A Simple Yet Versatile Recommendation Systems Library in Python."
222 GitHub, 30 Mar. 2020, <https://github.com/mayukh18/reco>.

223 [4] "Introducing Tensorflow Recommenders." The TensorFlow Blog, 23 Sept. 2020,
224 <https://blog.tensorflow.org/2020/09/introducing-tensorflow-recommenders.html>.

225 [5] Philipe, Hervind. "H&M: Faster Trending Products Weekly." Kaggle, Kaggle, 22 Mar. 2022,
226 <https://www.kaggle.com/code/hervind/h-m-faster-trending-products-weekly/notebook>.

227 [6] Seeda, Pathairush. "A Complete Guide to Recommender System-Tutorial with Sklearn, Surprise, Keras,
228 Recommender." Medium, Towards Data Science, 13 Oct. 2021, [https://towardsdatascience.com/a-complete-](https://towardsdatascience.com/a-complete-guide-to-recommender-system-tutorial-with-sklearn-surprise-keras-recommender-5e52e8ceace1)
229 [guide-to-recommender-system-tutorial-with-sklearn-surprise-keras-recommender-5e52e8ceace1](https://towardsdatascience.com/a-complete-guide-to-recommender-system-tutorial-with-sklearn-surprise-keras-recommender-5e52e8ceace1).

230 [7] Morty, Tarick. "[Lb 0.0240] H&M Ensemble Magic - Multi Blend." Kaggle, Kaggle, 29 Apr. 2022,
231 <https://www.kaggle.com/code/tarique7/lb-0-0240-h-m-ensemble-magic-multi-blend>.

232 [8] Lebovitz, Ben. "K-Means and Feature Importance for Articles." Kaggle, Kaggle, 29 Apr. 2022,
233 <https://www.kaggle.com/code/beezeus666/k-means-and-feature-importance-for-articles>.