




Using Ranking Support Vector Machines for Group Recommendations: Restaurant Recommendations on Yelp Data

Sarah Tan¹, Rahmtin Rotabi², Giang Nguyen²
Cornell University Statistics¹, Computer Science²

ABSTRACT

There are two common approaches to group recommendation systems, namely **aggregating individual user profiles** vs. **aggregating individual recommendations** [1]. Following the later approach, we use **ranking support vector machines** [2] to build a restaurant recommendation system for individuals and groups. We propose a **happiness metric**, where how happy a user is about a restaurant corresponds to the signed distance from the restaurant to the user's hyperplane in feature space. We contrast results obtained from different ways of aggregating happiness across a group of users, such as **maximizing average happiness** and **maximizing minimum happiness**.

INTRODUCTION

- **Group recommendation systems** are particularly useful for communal activities, such as dining out.
- Such a system needs to consider the **different preferences and restrictions** of multiple individuals in a group and provide a recommendation that satisfies the group of individuals according to some criterion.

METHOD

Individual Recommendations

- For each of 4.9k individual Yelp users, we create a **ranking SVM** using features on the restaurants the user reviewed to learn a hyperplane that reflects his or her preferences.
- Label: **user's rating of restaurant (1-5)**
- Signed distance from the restaurant to the user's hyperplane is our **happiness metric**.

Group Recommendations

- **Maximize minimum happiness**

```

Input:
Group members  $p_1, \dots, p_n$ 
Corresponding hyperplanes:  $w_1, \dots, w_n$ 
Restaurants  $r_1, \dots, r_m$ 
Output: Sorted restaurants  $R_k$ 
for  $j = 1$  to  $m$  do
   $R_j.id \leftarrow r_j.id$ 
   $R_j.score \leftarrow \infty$ 
  for  $i = 1$  to  $n$  do
     $R_j.score \leftarrow \min\{R_j.score, \langle w_i, r_j.features \rangle\}$ 
  end for
end for
return sort( $R$ ) // Decreasing order on the score

```

- **Maximize average happiness**

EXPERIMENTAL EVALUATION

Data Yelp Dataset Challenge [3]

- 1,100k reviews
- 190k users
- 42k businesses
- 9 years: 2005 - present
- 5 cities: Phoenix, Las Vegas, Madison, Waterloo in Canada, and Edinburgh in the UK.

Data Processing


- 50% of users only reviewed one restaurant
- 90% of users have less than 10 reviews
- 0.5% of users have more than 1k reviews
- 50% of users connect with friends on Yelp

We removed non-restaurants and users with less than 20 restaurant reviews and ended up with **700k reviews of 14k restaurants written by 4.9k users**.

Feature Engineering

We have **262 features** including:

- **Average rating** across all reviews
- Number of reviews as an indication of **popularity**
- Binary features, one for each Yelp **category and attributes** encompassing cuisine type, services offered, ambience, noise level, etc. We exclude a category if less than 10 of 14k restaurants have it.



Model Evaluation

Training-Testing Split

For each user, we use the **most recent 20%** reviews for testing, and the **remaining 80%** for training, to reflect how a recommendation system might actually be used in practice.


Prediction Accuracy

We use the **number of inversions of the ranking** to evaluate the accuracy of our model and compare it to the baseline of the number of random permutations.

Synthetic Group Labels


To approximate communal dining situations, we draw random combinations of users' friends. We take the highest rated restaurant among a group to be the label.

USE CASE



Once upon a time, Alan Turing, Grace Hopper, and R. A. Fisher got together to discuss the secrets of the universe. They set out to pick a restaurant that would make everyone happy on average. Which restaurant should they pick? What if they want a restaurant that makes the least happy person happy?

RESULTS



CONCLUSION AND ONGOING WORK

- In general, **synthetically generated group labels** agree with restaurant that maximizes average happiness
- **Lack of actual group labels** is a problem
- We are refining our feature set to enhance the accuracy of our predictions. Since our data spans 9 years, we are considering weighing older reviews less than newer reviews.

REFERENCES

- [1] Ricci, F., Rokach, L., Shapira, B. and P.B. Kantor. *Recommender Systems Handbook*. Springer. (2011)
- [2] T. Joachims. *Optimizing Search Engines Using Clickthrough Data*. Proceedings of the ACM Conference on Knowledge Discovery and Data Mining (KDD). ACM. (2002)
- [3] Yelp. *Yelp Dataset Challenge*. http://www.yelp.com/dataset_challenge