Balancing Preferences, Popularity and Location in Context-Aware Restaurant Deal Recommendation: A Bristol, Cardiff and Brighton Case Study

Ercan Ezin, Hugo Alcaraz-Herrera
{ercan.ezin,h.alcarazherrera}@bristol.ac.uk
University of Bristol
Bristol, United Kingdom

Iván Palomares
i.palomares@bristol.ac.uk
University of Bristol. Bristol, United Kingdom
The Alan Turing Institute. London, United Kingdom

ABSTRACT

We propose a personalisation solution to recommend tailored restaurant deals for residents or visitors in a city. Unlike previous work on recommendations in the restaurant sector where actual venues are recommended, we focus on suggesting specific products in the form of deals offered by such restaurants. This is done by jointly filtering relevant information for the end-user based on their food-drink preferences, the popularity of the restaurant, its proximity to the user's location and temporal constraints on the availability of deals. A real case study has been conducted upon datasets provided by Wriggle, a platform for discovering local deals in various cities across England.

KEYWORDS

Personalised Tourism, Restaurant Recommendation, Preference Modeling, Context-Aware Recommendation, Weighting

1 INTRODUCTION

Personalisation services for tourism, leisure and entertainment have been investigated for recommending Points-of-Interest (PoIs) or sequences of them [1, 2], selecting suitable cities for a group itinerary [3], or recommendations in the hotel sector [4, 6], to name a few. This study focuses on recommendations in the restaurant sector, which has also attained significant attention within the tourism landscape: in medium to large cities where both residents and visitors alike search for new restaurants, cafes or bars amid hundreds or thousands of available options [7], eating or drinking out is a cornerstone activity where personalisation turns indispensable to help them finding venues that meet their taste.

Various research efforts have been made on recommending suitable restaurants based on different forms of user preferences and contextual factors [8–10]. However, these works typically focus on recommending venues, by analysing characteristics associated to the restaurant itself, without looking at specific products (e.g. dishes, drinks, deals, etc.) offered by that restaurant or analysing how they meet the specific user needs or preferences. Despite this is an important decision-making step for for customers, many of them also seek specific dishes or suitable offers/deals that meet their preferences to a deeper level of granularity. To our knowledge, this is the first study to jointly consider both (i) general aspects of restaurants (location, opening times and popularity) and (ii) specific item features (through users' preferences on specific types of food-drink deals), for recommending restaurant deals for residents and visitors in a city. Some services and apps, such as Wriggle [1], have recently arisen in which users in Bristol, Cardiff and Brighton can search for available restaurant deals in their area.

We present a model for recommending temporary deals offered by restaurants, taking account of (i) users' preferences on food-drink categories, (ii) contextual information and (iii) restaurant popularity. In our approach, the recommendable items are deals offered by restaurants, rather than restaurants "as a whole". We investigate the problem of weighting (balancing) and aggregating similarity information for the three aforesaid aspects. In addition, we conduct a case study and a preliminary evaluation with real user and restaurant deal data provided by Wriggle on three UK cities. The results hint that by setting the weighting parameters for balancing the aforesaid sources from user to user, our proposed scheme has the potential for addressing the cold start problem (e.g. first-time visitors to a city with no purchase history), hence becoming adaptable to both local residents and tourists.

2 MODEL

Let $u_i \in U$ be the ith user and U the set of all users. Denote by $C = \{c_1, \ldots, c_M\}$ the set of existing food-drink categories in the system, e.g. 'cocktails', 'tapas', 'Indian', 'Chinese', etc. Given M categories, every user u_i has associated a preference vector $P_i = (p_{i1}\, p_{i2} \ldots p_{iM})$ where $p_{ik} \in \{0, 1\}$ is a preference indicator towards category c_k by u_i. In our current version of the model, the value of p_{ik} is binary and determined depending on whether the user consumed deals under c_k or not. A restaurant deal $x_j \in X$, with X the set of all restaurant deals (item set), can have associated one or more categories $c_k \in C$. Thus, we formally define a deal as a tuple $x_j = \langle r_{x_j}, C_j, V_j \rangle$,

[1] Wriggle website: https://www.getawriggleon.com/

where r_{x_j} is the restaurant that offers the deal. C_j is the temporal context of the deal, namely start and end date at which the deal is available, and whether it is a lunchtime and/or dinner time deal. $V_j = (v_{j1}\ v_{j2} \ldots v_{jM})$ is a binary feature vector associated with the offer, in which $v_{jk} = 1$ if the deal x_j is labeled with category c_k, and $v_{jk} = 0$ otherwise. Our solution consists of two major stages: a context-based item pre-filtering stage, and a weighted filtering stage.

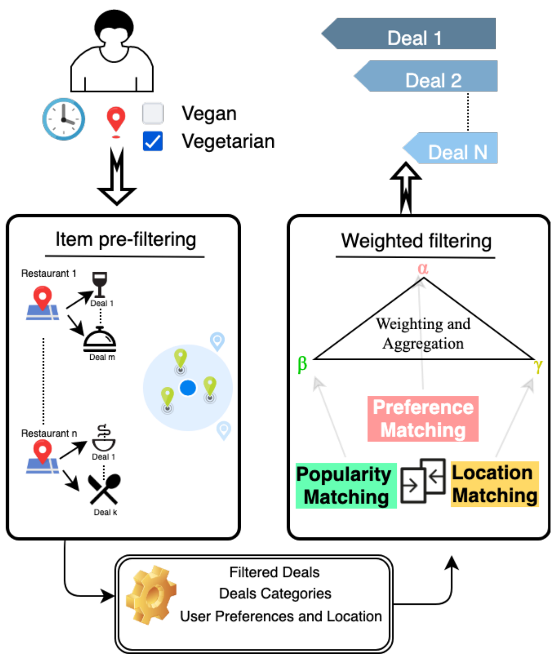

Figure 1: Architecture of the model for restaurant deal recommendation

Item pre-filtering. Unlike rating-based context pre-filtering approaches in the literature [5], in our model given an item set X (i.e. restaurant deals) context information C is firstly used to extract a subset of data related to the items relevant to that context. This is fundamental in domains where contextual limitations imply that not all existing items may be relevant or accessible by the end user at a certain place or time. For the scope of this study focused on the *Wriggle* data, we extract a subset of relevant deals to the current user and their context, accomplishing: (i) *Start-End Time*: most deals are periodical or limited and have a start-end time, therefore the currently available deals must be filtered; (ii) *Lunch or Dinner Time*: some deals are only active at lunchtime or dinner time, hence unavailable deals at a given time of the day are filtered out; and (iii) *Dietary Requirements*: although this is a user profile feature, we pre-filter suitable deals for users who are vegetarian or vegan.

Weighted filtering. This stage applies *three* matching processes and then weighs and aggregates resulting similarities:

(1) *Preference Matching*: It calculates the similarity between u_i preferences on food-drink categories, given by P_i, and the specific categorical features of a deal x_j, given by V_j. The cosine similarity is determined between both one-dimensional vectors, $m_\alpha(u_i, x_j) = sim(P_i, V_j)$. In essence, this filtering process entails a content-based approach relying on user preferences and item features of deals, hence it can easily integrate other content-based models in extant literature.
(2) *Popularity Matching*: This process takes the restaurant popularity into account, based on the average customer rating given to the restaurant. The popularity matching is calculated as the average customer rating of the restaurant r_{x_j}, thus $m_\beta(u_i, x_j) = pop(r_{x_j})$. Despite its simplicity, this solution is not personalised for the end user in question, because it is only dependent on r_{x_j}. An alternative personalised solution would be to apply a Collaborative Filtering (CF) algorithm to identify the K most similar users to u_i who rated r_{x_j}, based on their preference vectors P_i, and predicting how popular the restaurant might be for u_i.
(3) *Location Matching*: It takes the distance between restaurants within a predefined radius and the current user location, thereby prioritising deals from closer restaurants:

$$m_\gamma(u_i, x_j) = 1 - \frac{dist(u_i, r_{x_j})}{radius} \quad (1)$$

One of the contributions in this study is an adaptive weighting scheme for balancing preferences, popularity and location. Let α, β and γ be the weighting parameters or degrees of influence played by the preference, popularity and location matching, respectively. Without loss of generality, $\alpha, \beta, \gamma \in [0, 1]$ and $\alpha + \beta + \gamma = 1$. The overall matching used for selecting and recommending the top-N deals for u_i, is:

$$m(u_i, x_j) = \alpha \cdot m_\alpha(u_i, x_j) + \beta \cdot m_\beta(u_i, x_j) + \gamma \cdot m_\gamma(u_i, x_j) \quad (2)$$

We now describe a preliminary solution for adaptively setting α, β and γ for every user. It is worth noting that deeper investigation of applying more advanced optimisation or machine learning techniques to optimally set these weights, constitutes our immediate future work.

The influence of α, which refers to the user preferences on food-drink types, should rely on the size of the user's purchase history, i.e. the number of deals previously consumed. Users with a longer history have more accurately built preferences P_i than (cold) users with a short history, hence α should be higher in the former case. For users with no purchase history, e.g. first-time visitors to a city, for instance), preference information in P_i should be disregarded by setting $\alpha = 0$. Inspired by fuzzy set theory, we achieve this by setting $\alpha \in [0, \alpha_{max}]$, $0 < \alpha_{max} < 1$, such that α increases as the user history grows.

The influence of β relies on the amount of ratings received by the restaurant associated to x_j. If r_{x_j} has more customer ratings, β should be higher under the premise that frequently rated restaurants have more reliable (less biased) popularity information, and vice versa. Likewise, for a new restaurant with no ratings, we set $\beta = 0$. Using a similar principle as the one for α, we set $\beta \in [0, \beta_{max}]$, $0 < \beta_{max} < 1 - \alpha_{max}$.

The influence of γ, which refers to the proximity between user and venue, is (without losing generality) determined upon the other two parameters, as $\gamma = 1 - (\alpha + \beta)$. In other words, distance becomes more relevant if u_i has a smaller purchase history or r_{x_j} has less customer ratings. If both α and $\beta = 0$, the filtering process between a cold user and a deal offered by an unrated restaurant becomes purely location-based, $\gamma = 1$.

3 EXPERIMENTAL CASE STUDY

This section presents a case study conducted in collaboration with Wriggle, on a real dataset describing restaurants, deals and purchases made by users who used Wriggle in Bristol, Cardiff and Brighton. By using the purchase history and user profile, a sensitivity analysis is conducted on our proposed model parameters.

Dataset Description. The anonymised datasets provided by Wriggle contain a history of purchased deals by every user over a period of five years, between 2014 and 2019. Around 305K purchases are logged by 141K users. Also, a total of approximately 11K deals offered by 2153 restaurants are included in the dataset, with each deal being associated to one or multiple categories, out of a total of 63 categories describing food or drink characteristics/cuisines. There is also data about every user's profile, including dietary requirements if any (vegetarian, vegan), and restaurant profiles that contain the restaurant's average popularity based on users' rating on deals offered by that restaurant.

Experimental Setting. We filter users who have at least one purchase in the last 5 months of purchase dataset because real location data exists only for that particular period. Then, we split the user history dataset into a training and test set for three major cities, Bristol, Cardiff and Brighton, that Wriggle operates currently. We consider three different time span settings for the user purchase history: 6 Months, 12 Months and entire history since 2014. We then separate the latest deal with location information purchased by each user into the test set. Users with three or less items in their purchased history have been removed for the purpose of this experiment, leaving a consolidated purchase history of 2043 Users for Bristol, 249 for Cardiff and 643 for Brighton. Category information retrieved from deals in the purchase history is used to built preference vector of user P_i for the preference matching. Likewise, the information about restaurant popularity, opening times and location are retrieved from the restaurant-related data. For the contextual information, location data and time are inferred by retrieving the temporal information associated to the last purchased deal (test data). Finally, we consider $k = 10$ for the size of the recommendation list.

Evaluation Metrics. We recommend the top-k matching offers to the target user and investigate the predictive power exhibited by the model in recommending the (removed) latest deal purchased by each user, or the restaurant which offered it. For this end, the performance evaluation metrics employed are adapted versions of average recall@k and average NDCG@k on all users, thereby predicting the appearance of each user's latest deal or visited restaurant in her history in the recommendation list. The average recall is:

$$avg\_recall@k = \frac{\sum_{u_i \in U} y_i}{|U|} \tag{3}$$

$$y_i = \begin{cases} 1 & \text{if last deal consumer by } u_i \text{ is among top-}k, \\ \frac{1}{2} & \text{if last restaurant visited by } u_i \text{ is among top-}k, \\ 0 & \text{otherwise.} \end{cases}$$

Average Normalised Discounted Cumulative Gain at k:

$$avg\_NDCG@k = \frac{\sum_{u_i \in U} NDCG@k_i}{|U|}$$

$$NDCG@k_i = \sum_{j=1}^{k} \frac{2^{z_{i,j}} - 1}{log_2(j+1)} \tag{4}$$

where $z_{i,j} = 1$ if the last deal consumer by u_i is the jth recommended item, $z_{i,j} = 0.5$ if the restaurant last visited by u_i is at the jth recommended item, and $z_{i,j} = 0$ otherwise.

Results and Discussion. Three baseline approaches, and two versions of the proposed model with non-null weights, are considered:

Most Popular: Recommend deals based on venue popularity.
User-Preference: Recommend deals predicted on preferences over categories in deals.
Location: Recommend deals based on restaurant proximity.
Same Weight: Popularity, preferences and context are equally important for every user and restaurant, i.e. $\alpha = \beta = \gamma$.
Optimised Weight: It adaptively sets weights as explained in Section 3, with $\alpha_{max} = \beta_{max} = 0.3$. Both α (resp. β) become maximum when the user history length (resp. restaurant rating count) is greater than five.

Figure 2 summarises the average results obtained by the five models, for users in the three cities considered and the three time span settings considered. Despite a more exhaustive validation is needed, the results provide some interesting insights.

The proposed model with optimised weight scheme tends to slightly outperform the version with same weights, in almost all cases, specially when considering a shorter time span (6 months). Whilst this improvement is not significant,

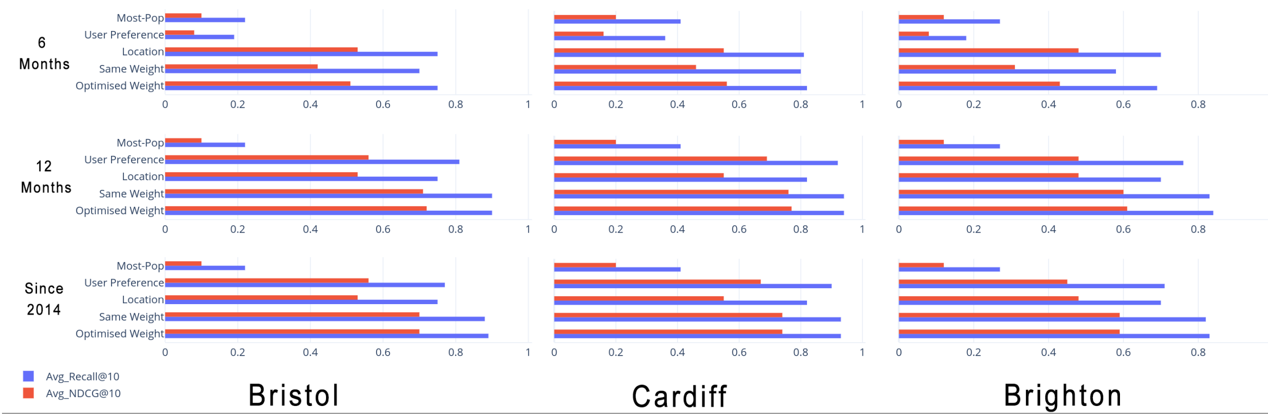

Figure 2: Comparison results in terms of average recall and average NCDG for $k = 10$

it motivates us to investigate how to further improve it by devising more user-adaptive weight optimisation methods in future work. Both two versions of our model generally outperform the three baseline approaches, however a location based recommendation has better predictive power in two of the three cities for the 6-month case. This suggests that most users may have a scarce purchase history in such a short time span, in which case prioritising restaurant proximity might increase the chances for better predictions.

Finally, the fact that the user preference baseline gently improves for longer time spans, suggests that the more purchase history data are available, the more reliable the extracted (implicit) preference information is.

4 CONCLUSION

This contribution proposes a recommendation model for suggesting restaurant deals to local and visiting users to a city by balancing their food-drink preferences, the popularity of the restaurant, and the context surrounding the user, such as his/her location. A case study has been conducted with real data provided by Wriggle, with insightful results motivating the need for follow-up research on how to optimally balance multiple information sources.

People often visit restaurants in groups whose members have diverse preferences. Accordingly, future work involves investigating preference aggregation for consensual group recommendations [12, 13]. We are also interested in (i) harnessing the capabilities of data networks in smart cities to enable highly situation-aware recommendations in real time, specially for tourists visiting a city; (ii) modeling users' preferences on food-drink categories more flexibly and under several decision criteria; and (iii) applying improved models on open datasets to make this research more reproducible.

5 ACKNOWLEDGMENTS

The authors would like to thank Rob Hall (CEO, Wriggle) and Clement Debiaune (CTO, Wriggle) for incentivising and fostering the collaboration that made this research possible.

REFERENCES

[1] D. Herzog, C. Laβ, Wolfgang Wörndl. Tourrec: a tourist trip recommender system for individuals and groups. Proceedings 12th ACM Conference on Recommender Systems (Recsys'18), pp. 496-497, 2018.

[2] A. Moreno, A. Valls, D. Isern, L. Marin, J. Borrás. SigTur/E-Destination: Ontology-based personalized recommendation of Tourism and Leisure Activities. Engineering Applications of Artificial Intelligence,

[3] E. Ezin, I. Palomares, J. Neve. Group Decision Making with Collaborative-Filtering ín the looṕ: interaction-based preference and trust elicitation. Accepted, IEEE SMC 2019 Conference. In press.

[4] M. Al-Ghossein, T. Abdessalem, A. Barré. Cross-Domain Recommendation in the Hotel Sector. Proceedings RecTour'18, in 11th Conf. ACM Recsys'18, pp. 1-6, 2018.

[5] G. Adomavicius, A. Tuzhilin. Context-Aware Recommender Systems. In F. Ricci et al. (Eds.) Recommender Systems Handbook, pp. 217-253, Springer, 2011.

[6] A. Ebadi, A. KrzyÅijak. A Hybrid Multi-Criteria Hotel Recommender System Using Explicit and Implicit Feedbacks. International Journal of Computer and Information Engineering, 10(8), 1450-1458, 2016.

[7] P. Longart. Consumer Decision Making in Restaurant Selection. PhD Thesis, Buckinghamshire New University, 2015.

[8] L. Li, Y. Zhou, H. Xiong, C. Hu, X. Wei. Collaborative Filtering based on User Attributes and User Ratings for Restaurant Recommendation. Proceedings 2nd IEEE IAEAC Conference, pp. 2592-2597, 2017. 26(1), pp. 633-651, 2013.

[9] E. Palumbo, G. Rizzo, R. Troncy, E. Baralis. Predicting Your Next Stopover from Location-based Social

[10] J. Zeng, F. Li, H. Liu, J. Wen, S. Hirowaka. A Restaurant Recommender System based on User and Location in Mobile Environment. Proceedings 5th IIAI International Congress, pp 55-60, 2016. Network Data with Recurrent Neural Networks. Proceedings RecTour'17, 11th Conf. ACM Recsys'17, pp. 1-8, 2017.

[11] M. Bressan, S. Leucci, A. Panconesi, P. Raghavan, E. Terolli. The Limits of Popularity-Based Recommendations, and the Role of Social Ties. Proceedings 22nd ACM SIGKDD International Conference, pp. 745-754, 2016.

[12] A. Delic, J. Neidhardt, H. Werthner. Group Decision Making and Group Recommendations. Proceedings 20th IEEE CBI, pp. 79-88, 2018.

[13] I. Palomares. Large Group Decision Making: creating Decision Support Systems at Scale. Springerbriefs in Computer Science, Springer, 2018.