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

Ercan Ezin, Hugo Alcaraz-Herrera, Iván Palomares

{ercan.ezin,h.alcarazherrera,i.palomares}@bristol.ac.uk

University of BRISTOL

Decision Support and Recommender Systems Research Group

The Alan Turing Institute

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.

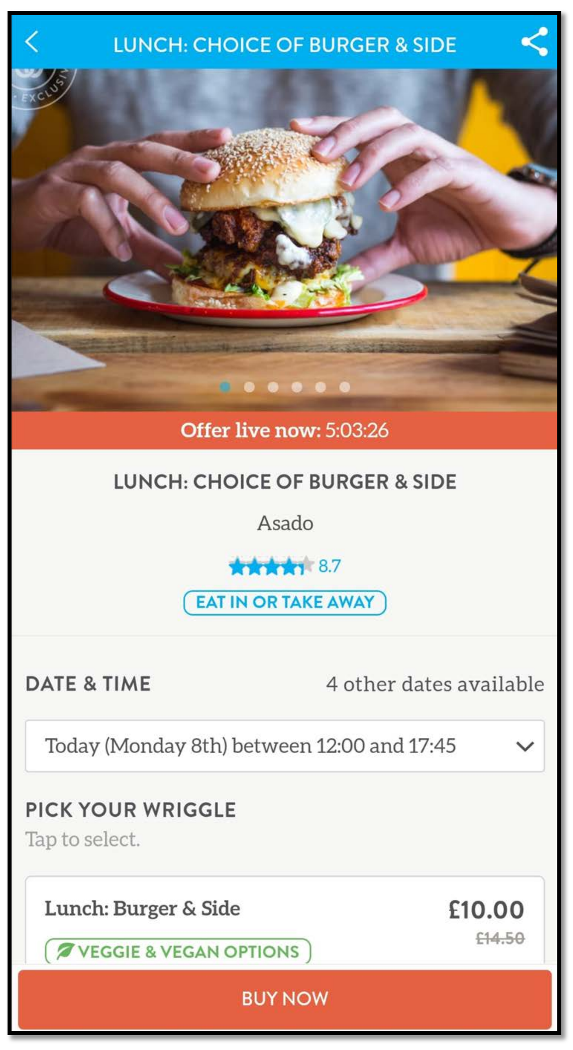

Fig:1 A deal presentation on the Wriggle App

Data Preparation

The anonymised datasets provided by Wriggle contain a history of purchased deals by every user over a period of five years, describing deals associated to food categories/cuisines. Deals are offered by over **2 thousand** restaurant and consumed by nearly **150 thousand** users in the cities of **Bristol, Cardiff and Brighton**. There is also data about every user's profile, including dietary requirements if any (**vegetarian**, **vegan**), and restaurant profiles. We filtered 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 aforesaid cities.

We consider three different time span settings for the user purchase history: **6 Months, 12 Months and entire history since 2014**.

Model

Item pre-filtering

In our model given an item set X (i.e. restaurant deals) **context information C** is firstly used to 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) Location: We calculated **average travel distance** for purchases and used this figure to filter deals within a distance; and

(iii) Dietary Requirements: although this is a user profile feature, we pre-filter suitable deals for users who are vegetarian or vegan.

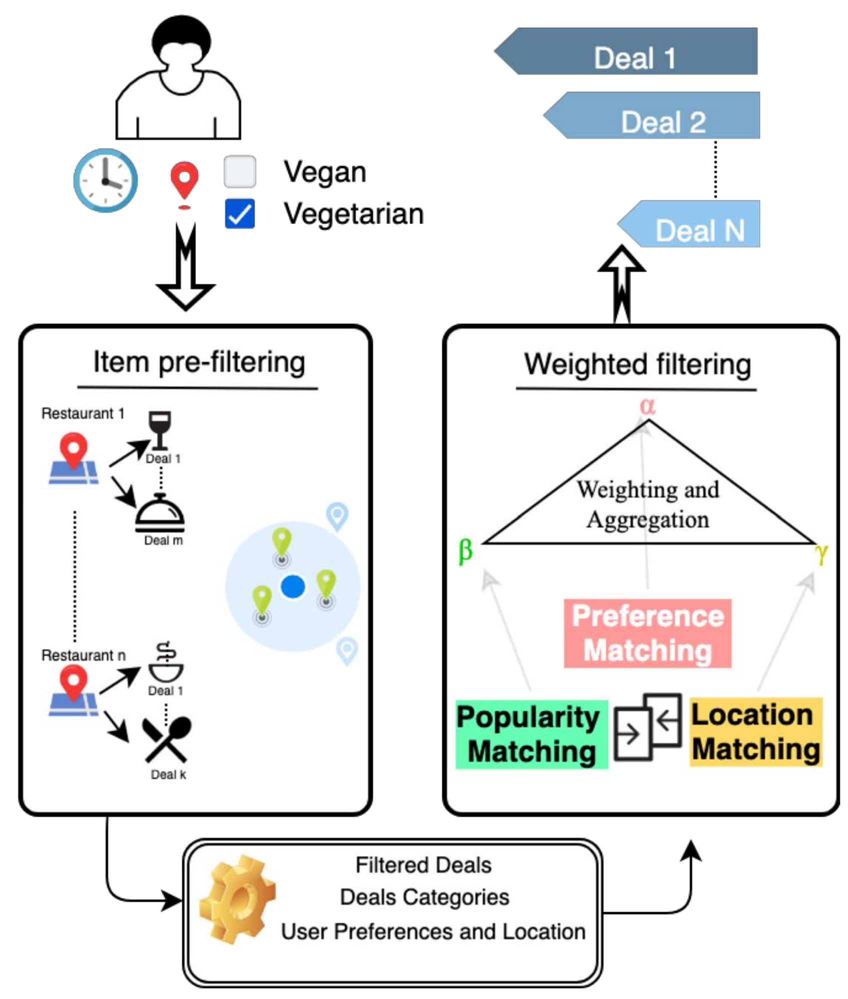

Fig:2 Data Processing Model Representation

$$\alpha+\beta+\gamma = 1$$

Weighted filtering

(i) Preference Matching(α): It calculates the similarity between preferences on food-drink categories, and the specific categorical features of a deal.

(ii) Popularity Matching(β): This process takes the restaurant popularity into account, based on the average customer rating given to the restaurant.

(iii) Location Matching(γ): It takes the distance between restaurants within a predefined radius and the current user location, thereby prioritising deals from closer restaurants.

Results

- The experiments were aiming at predicting whether **the latest deal purchased** or **restaurant visited by the user** appears in the recommendation list of size **K=10** generated by our model. We measures **Avg_Recall** and **Avg_NDCG at 10** to evaluate predictive power of the model.
- The model with optimised weight scheme tends to slightly **outperform** the version with **same weights**, in almost all cases, specially when considering **6 months time** span.
- 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**. The reason behind this might be the scarce purchase history.
- The user preference baseline **gently improves** for longer time spans which is suggesting that the more purchase history data are available, the more reliable the extracted (implicit) preference information is.

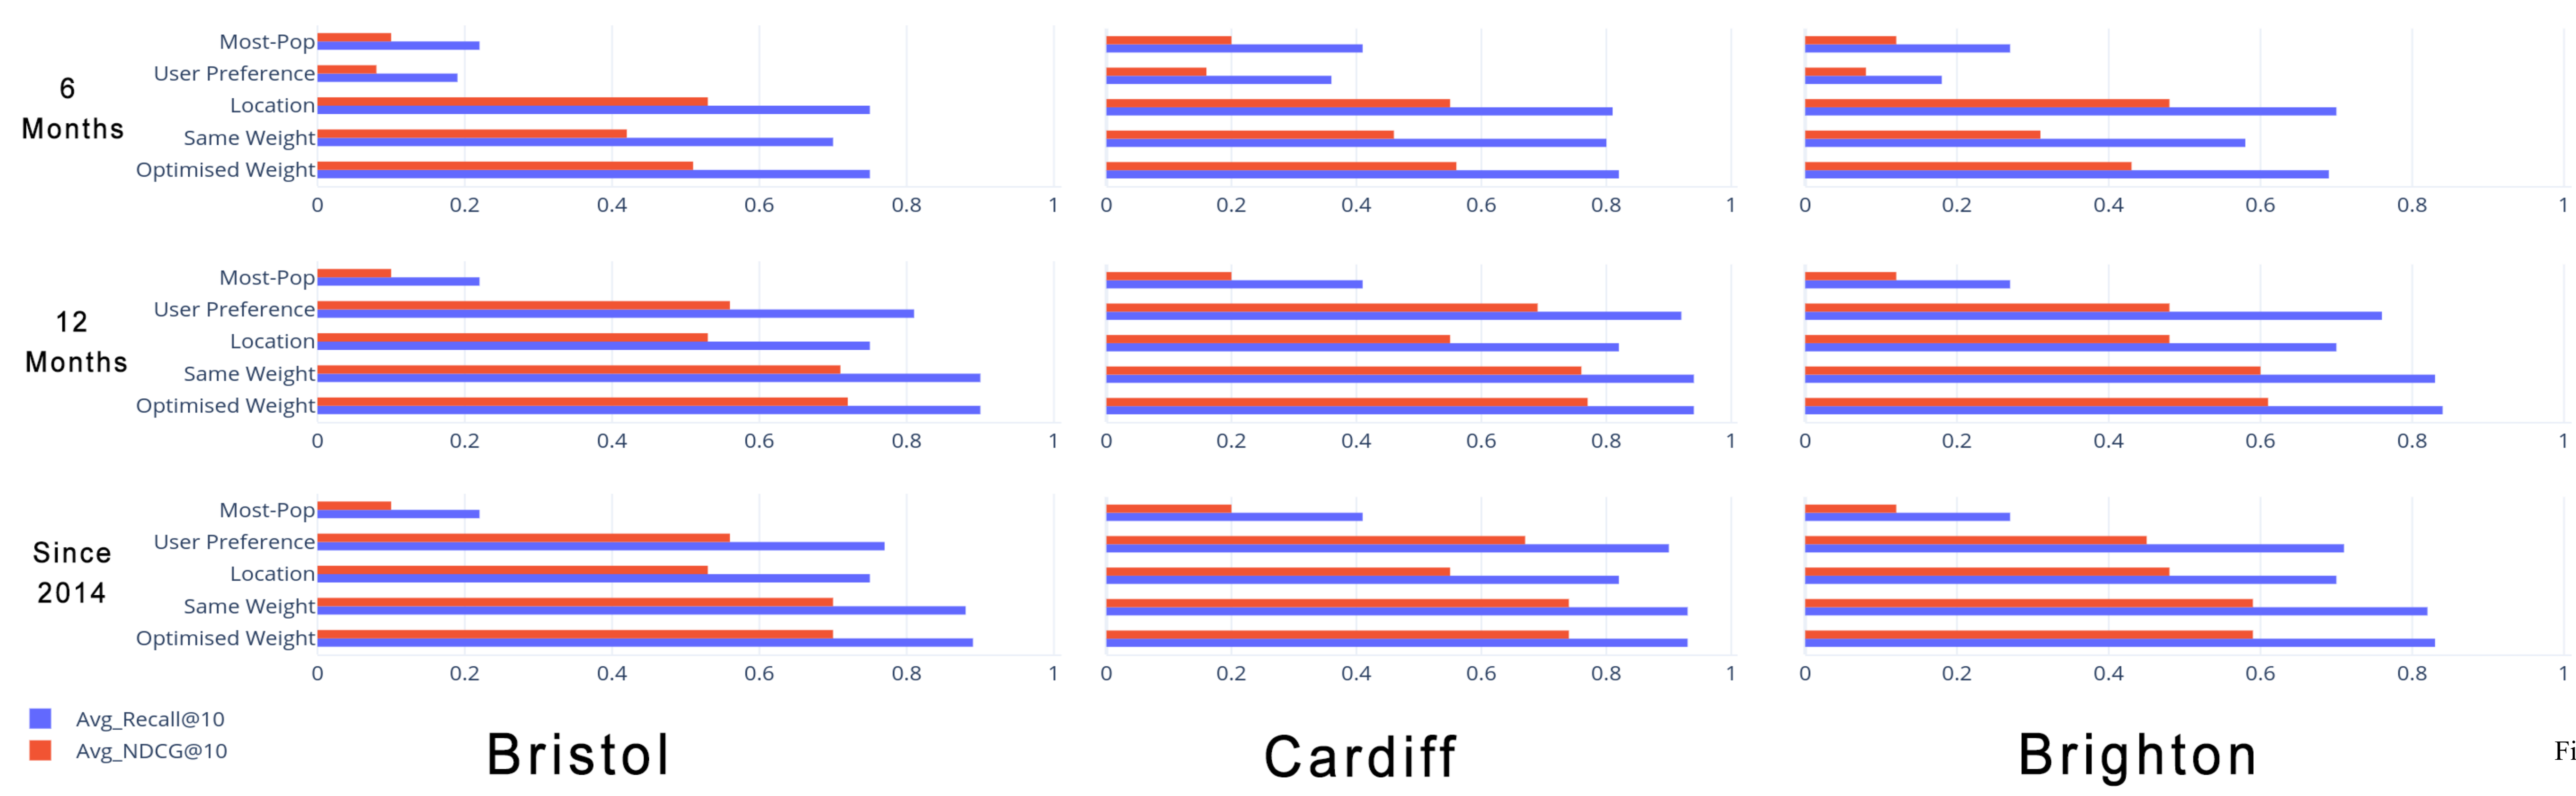

Fig:3 Experiment results for three cities.

What's Next

- Considering people often visit restaurants in **groups**, future work involves investigating preference aggregation for consensual group recommendations.
- Harnessing the capabilities of data networks in **smart cities** to enable highly **situation-aware recommendations** in real time for tourists visiting a city;
- Modeling **users' preferences** on food-drink categories more flexibly and under several decision criteria.
- Applying improved models on **open datasets** to make this research more reproducible.

Acknowledgement

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.

wriggle

dsrs.blogs.bristol.ac.uk