Task-Based User Modelling for Personalization via Probabilistic Matrix Factorization

Rishabh Mehrotra
University College London
r.mehrotra@cs.ucl.ac.uk

Emine Yilmaz
University College London
emine.yilmaz@ucl.ac.uk

Manisha Verma
University College London
m.verma@cs.ucl.ac.uk

ABSTRACT

We introduce a novel approach to user modelling for behavioral targeting: *task-based* user representation and present an approach based on search task extraction from search logs wherein users are represented by their actions over a task-space. Given a web search log, we extract search tasks performed by users and find user representations based on these tasks. More specifically, we construct a user-task association matrix and borrow insights from Collaborative Filtering to learn low-dimensional factor model wherein the interests/preferences of a user are determined by a small number of latent factors. We compare the performance of the proposed approach on the task of collaborative query recommendation on publicly available AOL search log with a standard term-similarity baseline and discuss potential future research directions.

Categories and Subject Descriptors

H.3.3 [**Information Storage And Retrieval**]: Information Search and Retrieval—*User Modelling*

Keywords

Search tasks; User modelling; Personalization

1. INTRODUCTION

As a consumer of the informational content, different users have distinct preferences of information for decision making; thus accurately understanding their respective information needs and decision preferences is crucial for providing effective decision support. While human behaviours are largely determined by their own goals and preferences, the mined knowledge reveals users' underlying intentions and behaviour patterns, which provide unique signals for human centric optimization and personalization. Web search personalization has recently received a lot of attention by the research community. Personalized search leverages information about an individual to identify the most relevant recommendations for that person. A challenge for personalization is in collecting user profiles that are sufficiently rich to be useful in settings such as result ranking, query recommendations, etc, while balancing privacy concerns.

A prominent line of prior research uses long term histories to directly improve retrieval effectiveness. Various authors have considered topic based representations for personalization [1] making use of hand picked Open Directory Project (ODP) topical categories. While such topics are easily specified, significant human effort is required in labelling queries for each topic. Additionally, topical category based methods restrict user's profile coverage in a major way as different users might share the same topical profile yet perform different search tasks for different informational needs. Another line of research for personalization has focused on using term based representations wherein user interests profiles are built using terms extracted from user's browsing history following which term weights are generated using different weighing schemes. While query terms are representative of user interests, they often limit the scope of personalization as different users inherently follow different distributions over words and queries belonging to the same topic might not contain any overlapping terms which makes finding similar users difficult in such settings.

In this work, we focus on learning user profiles based on the *search tasks* users are involved with. Users interact with search engines to accomplish some task such as *arrange a trip*, *plan a wedding* etc. Such broad requirements prompts the use of multiple queries, sometimes spanning multiple sessions. We define search tasks as the group of queries a user issues to accomplish such overall intended task and advocate the use of such search tasks to build individual user models. We postulate that in a web search setting, a user representation based on the search tasks users' perform would better capture user actions, interests and preferences. Given a search log, we extract search tasks performed by users and find user representations based on these tasks. More specifically, we construct a user-task association matrix and borrow insights from *Collaborative Filtering* to learn a low-dimensional factor model wherein the actions/interests/preferences of a user are determined by a small number of latent factors. By applying probabilistic matrix factorization to the user-task association matrix, we learn *task-based user representations* for each user and evaluate the quality of the learnt user representations by making use of these representations for the task of Collaborative Query Recommendation wherein we suggest queries to a particular user based on queries issued by other similar users. We compare the performance of the proposed approach against a term similarity based baseline on publicly available AOL search logs.

2. TASK BASED USER MODELLING

Our objective is to build succinct user profiles from the search task information embedded in search logs. Existing user modelling methods for web search rely heavily on per user topical interests and hence, fail to differentiate between

RecSys 2014 Poster Proceedings, October 6-10, 2014, Foster City, Silicon Valley, USA.

users which share similar topical interests. We postulate that in web search setting, search logs contain information about various actions that users perform and profiling users based on search tasks would better capture the heterogeneity in user information.

Task Discovery in Search Logs: Our goal here is to use search log data to create a list of global search tasks. Following the approach of task discovery as proposed in [2], a task is defined as the maximal subsequence of possibly nonconsecutive queries in referring to the same latent user need which makes the set of all user tasks a partitioning of the set of all user queries. We formulate the task discovery problem as follows: given a query log QL and a user u, let T_u be the set of user tasks discovered by a query partitioning scheme π; the user task discovery problem can then be described as finding the best query partitioning strategy π^* that approximates the actual set of user tasks Θ such that:

$$\pi^* = argmax_{\pi} \xi(\Theta, T, \pi) \tag{1}$$

where function $\xi(.)$ is an accuracy measure which evaluates how well the query partitioning strategy π approximates the actual user tasks Θ. We use cosine similarity to measure this accuracy. This step is followed by clustering the user tasks identified to obtain universal tasks across all users. The final set of user tasks obtained are represented by a set of query terms and henceforth define the set of tasks used for experiments. For details, please refer Lucchese *et al*[2].

User-Task Association Matrix: Based on the extracted search tasks, we construct a user-task association matrix which represents the search tasks users have been involved with. For each user u_i, we create a bag-of-queries representation from the list of queries issued by the user and compare each user with each of these search tasks t_j obtained above. For each user-task $<u_i, t_j>$ pair, we populate the corresponding value in the user-task association matrix (R) with the cosine similarity score (r_{ij}) we obtain for the pair. For tasks in which users do not have any matching queries, we assign a score of 0 to the corresponding pair.

Probabilistic Matrix Factorization for User Representations: We wish to extract task-based user vector representations by jointly mapping users and tasks to a joint latent factor space. Following Salakhutdinov *et al* [3], we model the user-task association in terms of probabiltic matrix factorization problem and learn latent vector representation for each user from the user-task association matrix by fitting a probabilistic model. Given the user-task association matrix $\mathbf{R}$, we find the user feature matrix $\mathbf{U} = [u_i]$ and task feature matrix $\mathbf{T} = [t_j]$. The conditional distribution over the observed user-task associations $R \in \Re^{N \times M}$ is given by:

$$P(R|U,T,\alpha) = \prod_{i=1}^{N} \prod_{j=1}^{M} \left[\mathcal{N}(R_{ij}|U_i^T T_j, \sigma^2 \right]^{I_{ij}} \tag{2}$$

where $\mathcal{N}$ denotes the Gaussian distribution and I_{ij} is the indicator function which is 1 if the user i was involved in search task j. The user matrix $\mathbf{U}$ obtained as a result, contains vector representations of each of the users which is used in further experiments.

3. EXPERIMENTAL EVALUATION

A good user profile for query recommendation should capture a user's specific interests & informational needs. Based on this intuition, we evaluate performance of the proposed approach on *Collaborative Query Recommendation* where

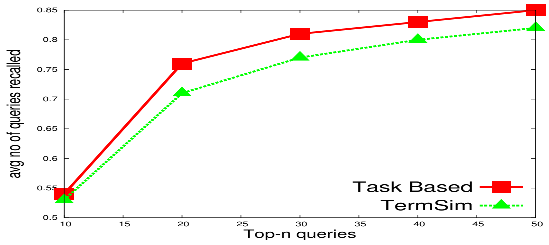

Figure 1: Performance on Collaborative Query Recommendation

the goal is to recommend queries to a user based on queries issued by similar users. We calculate the weighted frequency of a candidate query for 10 most similar users of the target user u, and selected the top n queries as recommendation. We make use of the AOL log dataset which consists of ~20M web queries collected over three months and use data for about ~1200 users who have issued more than 550 queries. We run our Task Discovery algorithm on the set of queries for each of these users which results in a total of ~0.12M tasks which we cluster using cosine similarity score to obtain a set of 1529 search tasks using which we create the user-task association matrix. Our baseline(*TermSim*) is a method that only uses bag-of-words based representation for each user where the terms are extracted from user queries & similar users found using cosine similarity between each user's bag-of-word based representations. We consider the test-set of queries in the target user as relevant, and computed average number of relevant queries matched in the recommendation query set as the performance metric.

We plot the average number of query matches between the recommended set of queries and user's own test set of queries against n where n refers to the top-n query suggestions from 10 most similar users. Our initial results (Figure 1) show that the proposed Task-Based user modelling approach(*Task-Based*) performs better than *TermSim* which demonstrates that search tasks can serve as potent user modelling tools. Since *TermSim* relies strictly on term matching for measuring user similarities, its coverage is limited: it might not capture insights for the users with too few queries or those who shared the same search interest but issued different queries or performed different tasks. Task based user modelling can help in better differentiating between users which have similar topical interests but perform different tasks. To better leverage the topical user profiles, it would be interesting to combine user topical-interest information with user task-associations to come up with a unified user model. We leave this as potential future work.

4. ACKNOWLEDGEMENTS

This work was supported in part by a Google Faculty Research Award.

5. REFERENCES

[1] P. N. Bennett, R. W. White, W. Chu, S. T. Dumais, P. Bailey, F. Borisyuk, and X. Cui. Modeling the impact of short-and long-term behavior on search personalization. In *ACM SIGIR*, 2012.

[2] C. Lucchese, S. Orlando, R. Perego, F. Silvestri, and G. Tolomei. Identifying task-based sessions in search engine query logs. In *WSDM*, 2011.

[3] A. Mnih and R. Salakhutdinov. Probabilistic matrix factorization. In *NIPS*, 2007.