

Personalised Recommendations for Context Aware Suggestions

Fabio Crestani

Faculty of Informatics

Università della Svizzera Italiana (USI)

Lugano, Switzerland

fabio.crestani@usi.ch

Abstract

Modern Information Retrieval has moved from standard text retrieval to novel applications of the same technology. Contextual suggestion is an example of this type of applications. The TREC Contextual Suggestion track addresses the problem of suggesting contextually relevant attractions to a user visiting a new city based on his/her recorded preferences from past visits to other cities. In this invited talk I will reframe the problem of representing and using context and briefly report our two past approaches to capturing the user profile to enable a system to provide more accurate and relevant recommendations. The results of our participation in the 2013 and 2015 TREC tracks, reporting how we can use such contextual information as geographical location, time, and friends' interests, show that our system not only significantly outperforms the baseline method, but also performs better than most other participants to that track, managing to achieve the best results in nearly all test contexts.

1 Introduction

The research area of Information Retrieval (IR), historically concerned with retrieving information from large archives in response to a user query, has been evolving rapidly in recent years. This evolution has brought IR researchers to deal with problems that are very different from standard IR, like for example Topic Detection and Tracking, Blog and Tweet retrieval, Knowledge Base Acceleration, Temporal Summarisation, Novelty Detection, etc. IR provides a large number of techniques

that, appropriately modified, can help provide solutions to these tasks.

Recent years have witnessed an increasing use of location-based social networks (LBSNs) such as Yelp, TripAdvisor, and Foursquare. These social networks collect valuable information about users' mobility records, which often consist of their check-in data and may also include users' ratings and reviews. A service that could be of interest to users of such networks could be related to providing them recommendation of location to visits. In fact, being able to recommend personalised venues to users plays a key role in satisfying the user needs on such social networks.

Recent research on recommending systems has focused on using collaborative-filtering technique, where the system recommends venues based on users' data whose preferences are similar to those of the target user. Collaborative-filtering approaches are very effective, but they suffer from the cold-start (i.e., they need to collect enough information about a user for making recommendations) and the data-sparseness problems. Furthermore, these approaches rely mostly on check-in data to learn the preferences of users and such information is often insufficient to get a complete picture of what the user likes or dislikes of a specific venue (e.g., the food, the view, the music). In order to overcome this limitation, recent approaches try to model the users by applying a deeper analysis on users' past ratings as well as their reviews. In addition, following the principle of collaborative filtering, they exploit the reviews of different users with similar preferences.

2 Contextual Suggestion

The TREC Contextual Suggestion Track started in 2012 and continued to 2016, the current year. It investigates search techniques for complex infor-



Figure 1: The TREC Contextual Suggestion Scenario.

mation needs that are highly dependent on context and user interests. The task was to take the representation of these user interests (profiles) and contexts and to produce a list of ranked suggestions for each profile-context pair. The scenario used consistently by the track was that of a user visiting a new city and receiving suggestions of places (e.g. bars, restaurants, museums, etc.) to visit based on what the new city made available and his preferences as extracted from the user profile (see figure 1). A full description of the task can be found (Dean-Hall et al., 2013). The similarities with collaborative filtering are obvious, the only difference is that we know too little about each individual user to be possible to use any good collaborative filtering algorithm. Obviously the track evolved over the years, slightly changing the geographical context and providing richer users' profiles, but still making it impossible to use well established collaborative filtering algorithms.

In the following we report on the approaches we followed for our 2013 and 2015 participations to this track¹ and on the use of external information to enlarge the user profile to make it possible to provide more effective contextual suggestions.

3 On the Use of External Information for Contextual Suggestion

Context has a very loose definition in the area of IR. It is related to all aspects that influence the user perception of an information need or of the relevance of a document to such information need. This includes the time, the location, the prefer-

¹In 2014 we did not take part as the author was on sabbatical. We are also taking part in the 2016 track, but the results have yet to be released, so we will not comment on the approach taken.

ences, the physical environment, and the social situation affecting the user. Capturing it, enables to differentiate from moment to moment in the life of a user, providing better suggestions.

Our approach to the TREC Contextual Suggestion task involved using external information to enrich the available information about the user and the user's context. In 2013 we enrich the user geographical context (i.e. his location in time and space), while in 2015 we enrich the contextual information about the different venues and the opinion of different users (i.e. his social context), to make it possible to provide more valuable suggestions.

3.1 TREC 2013

In (Rikitianskiy et al., 2014) we described our approach for TREC 2013, aimed at making context-sensitive recommendations to tourists visiting a new city. We presented a new approach to recommending places to users incorporating geographical information as context and exploiting data from multiple sources. Our method is based on quite a simple strategy of using the descriptions of previously rated places in closed geographical proximity to build user profiles. We also introduced a number of novel additions which have clearly lead to improved performance. In fact, the analysis of the results from the TREC evaluations performed by a large group of users, demonstrated the high level of performance delivered by our method, showing that it is able to significantly outperform the two track baselines and all other track entrants in the majority of cases. In fact, when compared to the 34 other competing systems in the track, it delivered results which were well above the median. In nearly half of all contexts, our approach was able to deliver the best set of results, confirming that the choices made during the development of the system were sensible and beneficial. More details can be found in the above cited paper.

3.2 TREC 2015

The TREC 2015 Contextual Suggestion Track changed little compared to previous years, but we experimented a quite different approach. In (Aliannejadi et al., 2016) we presented a novel method for suggesting venues to users, where the users are modelled based on venues' content as well as other users' reviews of the same venues. For the former we use the categories of the venues

enriched by keywords extracted from users' online reviews, which provide a more detailed description of the venue itself. Although the venue information is valuable for inferring "what type" of places a user may like or dislike, it does not give any clue on the reasons "why" a user rated as positive or negative a particular venue. We needed to exploit the user's opinions in order to understand what the user may have appreciated of a place and to get better recommendations for future venues. One way to obtain these opinions is mining the users' reviews and see how much they liked the venue and, more importantly, for which reasons: was it for the quality of food, for the good service, for the cozy environment, or for the location? In cases where we lacked reviews from some of the users (e.g., they rated a venue but omitted to review it) and therefore could not extract opinions, we applied the collaborative-filtering principle and used reviews from other users with similar interests and tastes. Our intuition was that a user's opinion regarding an attraction could be learned based on the opinions of others who expressed the same or similar ratings for the same venue. To do this we exploited information from multiple sources (e.g. Yelp and Foursquare) and combine them to gain better performance. In the cited paper we showed how our model outperforms all the other runs by a significant margin and was placed as the first run in the track. See the paper for details on the technique used.

4 Conclusions and Future Work

The importance of context in IR has long been recognised and context has been used in many different applications of IR and related fields. Contextual suggestion is a difficult problem because of the many and different factors that make up the context and that have an influence on the effectiveness of the suggestion. Considering all available factors and, of course, finding an effective combination of them is the best approach, but it needs to be personalised and efficiently computed to be effective. This is the current direction of research of my research group in the context of a couple of projects we are involved in. This invited talk reported on the successful results of our participation in TREC 2013 and 2015 and on how we used context as an effective mean to provide better suggestions.

Acknowledgments

The work reported in this paper was done in collaboration with my PhD students at the Faculty of Informatics of the Università della Svizzera Italiana and, in particular, with Mohammad Nejadi. The work is financially supported by the Swiss National Science Foundation (SNSF) under the grant "Relevance Criteria Combination for Mobile Information Retrieval (RelMobIR)".

References

- M. Aliannejadi, I. Mele, and F. Crestani. 2016. User model enrichment for venue recommendation. In *Proceedings of the Asian Information Retrieval Symposium (AIRS)*, Beijing, China, December.
- A. Dean-Hall, C.L.A. Clarke, N. Simone, J. Kamps, P. Thomas, and E.M. Voorhees. 2013. Overview of the TREC 2013 contextual suggestion track. In *Proceedings of The Twenty-Second Text REtrieval Conference, TREC 2013*, Gaithersburg, Maryland, USA, November.
- A. Rikitianskiy, M. Harvey, and F. Crestani. 2014. A personalised recommendation system for context-aware suggestions. In *Proceedings of the European Conference in Information Retrieval Research (ECIR)*, pages 63–74, Amsterdam, The Netherlands, March.