

# Personalized Ranking for Context-Aware Venue Suggestion

Mohammad Aliannejadi  
Università della Svizzera  
italiana (USI)  
Lugano, Switzerland  
alianm@usi.ch

Ida Mele  
Università della Svizzera  
italiana (USI)  
Lugano, Switzerland  
ida.mele@usi.ch

Fabio Crestani  
Università della Svizzera  
italiana (USI)  
Lugano, Switzerland  
fabio.crestani@usi.ch

## ABSTRACT

Making personalized and context-aware suggestions of venues to the users is very crucial in venue recommendation. These suggestions are often based on matching the venues' features with the users' preferences, which can be collected from previously visited locations. In this paper we present a novel user-modeling approach which relies on a set of scoring functions for making personalized suggestions of venues based on venues content and reviews as well as users context. Our experiments, conducted on the dataset of the TREC Contextual Suggestion Track, prove that our methodology outperforms state-of-the-art approaches by a significant margin.

## CCS Concepts

•Information systems → Recommender systems; Personalization;

## Keywords

User models, contextual suggestion, LBSNs, review mining.

## 1. INTRODUCTION

Nowadays, almost all mobile devices have Internet access which allows users to search for information wherever they are and whenever they need to. Users often rely on their mobile devices when they are looking for events to participate, activities to do, and interesting nearby venues to visit. In this paper we focus on venue suggestion, which consists of proposing a list of places that can be interesting for the user. This is an important task, since the traditional manual search for the best venue among the myriad of available ones may be time consuming and not easy to do, especially when the user is visiting a new city or, for example, wants to spontaneously plan the night out with friends. Venue suggestions can be made by considering the preferences of the user, which are mined from the venues that the user has previously visited and can be further improved by the user's

context (e.g., the user is alone or with family, she is on a business trip or on a romantic weekend).

In this paper we aim at making personalized suggestions by taking into account both the users' preferences and her context. Our approach assigns a score which depends on users' preferences, opinions, and context in order to rank the candidate suggestions. Our model captures the user's preferences and understands her tastes by leveraging the venues' categories and the user's opinions. These are extracted from the online reviews often available in LBSNs. The model is then enriched by adding the contextual information of a specific user (e.g., season, group type).

The experiments on a TREC collection demonstrate that our approach performed very well compared to other state-of-the-art approaches [4].

## 2. RELATED WORK

Recently, due to the availability of Internet access on mobile devices and on the fact that contextual information can be provided by the sensors of the mobile, researchers have been focusing their interest in context-aware suggestions for venues. Compared to recommending news or products, the task of suggesting venues in a city raises further challenges, since it needs to consider not only the preferences of the users but also other constraints related to the context, such as the city, season, and people who accompany the user.

Content-based approaches make suggestions for venues by simply matching the venues' content (e.g., description and categories) with the user's preferences. Rikitianskii et al. [6] proposed to apply Part-of-Speech tagging to the venues' descriptions in order to get the most informative terms for a venue, which are then used to create positive and negative profiles. For each user, they trained a binary classifier using such profiles to rank the candidate suggestions.

Review-based approaches aim to build enhanced user profiles using their reviews. Reviews provide a wealth of information that can be extracted to enable a system to deal with the data sparsity and cold-start problems. Yang et al. [7] use reviews from Yelp to extract users' opinions. Given a pair (*user*, *venue*), they created positive and negative profiles for each pair by extracting data from all users' reviews. The list of suggestions is then ranked by using the similarity scores between all pairs of profiles.

In this paper, we propose a combination of content-based and review-based approaches. We use content to model users' interest and reviews to incorporate users' opinions in the model. We use also the context of a user since contextual information plays an important role in venue suggestion.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

SAC 2017 3-7 April, 2017, Marrakesh, Morocco

© 2016 ACM. ISBN 978-1-4503-2138-9.

DOI: 10.1145/1235

### 3. USER MODELING

#### 3.1 Frequency-based Score Component

The first component is based on the frequencies of venue categories and taste tags. We first explain how to calculate the score for venue categories. The score for tags is calculated analogously.

Given a user  $u$  and a her history of rated venues  $(v_i)_{h_u} = \{v_1, \dots, v_n\}$ , each venue is assigned with a list of categories  $C(v_i) = \{c_1, \dots, c_k\}$ . We define the category profile of a user as follows:

**DEFINITION 1.** A *Category Profile* is either positive or negative. A *Positive-category profile* is a set of all distinct categories belonging to venues that a particular user has previously rated positively. A *Negative-category profile* is defined analogously for the venues that are rated negatively.

We assign a user-level-normalized frequency value to each category in the positive/negative category profile. The user-level-normalized frequency for a positive/negative category profile is defined as follows:

**DEFINITION 2.** A *User-level-Normalized Frequency* for an item (e.g., category) in a profile (e.g., positive-category profile) is calculated as follows:  $cf^+(c_i) = \frac{\text{count}(c_i)}{\sum_{v_k \in h_u} \sum_{c_j \in C(v_k)} 1}$ . A user-level-normalized frequency for negative category profile,  $cf^-$ , is calculated analogously.

Given a user  $u$  and a candidate venue  $v$ , the category-based similarity score  $S_{cat}(u, v)$  between them is calculated as follows:

$$S_{cat}(u, v) = \sum_{c_i \in C(v)} cf^+(c_i) - cf^-(c_i). \quad (1)$$

We calculate the category similarity score from two sources of information, namely, Foursquare ( $S_{cat}^F$ ) and Yelp ( $S_{cat}^Y$ ).

**Venue Tags Score.** We further enrich the category-based model using “taste tags” which are the most salient words extracted from the users’ reviews. We can leverage them to have a crisper description of the venues and improve our suggestions. We create positive and negative tag profiles for each user following Definition 1. Similar to the category scores, we assign a user-level-normalized frequency following Definition 2 to each tag,  $tf(t)$ , in the positive and negative tag profile. The tag similarity score is then calculated similar to Equation 1.

#### 3.2 Review-Based Score Component

A further component uses the reviews to understand the motivation of the user behind a positive or negative rate. Indeed, modeling a user solely on venue’s content is very general and does not allow to understand the reasons why the user enjoyed or disliked a venue. Our intuition is that a user’s opinion regarding an attraction could be learned based on the opinions of other users who gave the same or similar rating to the same attraction.

An alternative to binary classification would be a regression model, but we believe it is inappropriate since when users read online reviews, they make their minds by taking a binary decision (like/dislike). The binary classifier is trained using the reviews from the venues a particular user has visited before. We used the positive training samples

which are extracted from the positive reviews of positive example suggestions,

Since the users’ reviews contain lots of noise and off-topic terms, we calculated TF-IDF score as our feature vectors for training the classifier. As classifier we used Support Vector Machine (SVM) [3] with linear kernel and consider the value of the SVM’s decision function as the score since it gives us an idea on how close and relevant a venue is to a user profile.

For each user we trained two SVM classifiers using reviews from Yelp and TripAdvisor. The corresponding scores are named  $S_{rev}^Y$  and  $S_{rev}^T$ , respectively.

#### 3.3 Context-based Score Component

Contextual information is very important for improving the quality of venue suggestions. In this section, we propose two scores for measuring the similarity between the context of a user and the information about a place. Note that we are able to measure the contextual appropriateness of a venue to a given user only based on those contextual signals which are available on the LBSNs (i.e., the season, the trip, and the group type). Our basic idea is to compare the current user’s context with the distribution check-ins of a particular venue over that context. We assume that the distribution of check-ins over a contextual signal reveals the level of the venues’ appropriateness to that context. In the rest of this section we explain the score used for the season, a similar method is applied for the travel score.

**Season Score.** If we know in which season a user has been visiting a candidate venue, we can leverage the distribution of check-ins over seasons for a better ranking of venue suggestions. For those reviews which do not indicate the season, we assumed that most of the people leave reviews on LBSNs soon after they visit a place, and we can compute the distribution based on the reviews’ timestamps.

Let  $S$  be the set of seasons and  $s_u$  be the season a particular user visited a place,  $SP(s, v)$  is a function returning the number of check-ins by other users for venue  $v$  in season  $s$ . Hence, we define the season score for user  $u$  visiting  $v$  as:

$$S_{ctx}^{season}(u, v) = SP(s_u, v) - \frac{\sum_{s_i \in S, s_i \neq s_u} SP(s_i, v)}{\sum_{s_i \in S, s_i \neq s_u} 1}, \quad (2)$$

where  $SP(s_u, v)$  is the number of check-ins for the venue  $v$  in the same season of the user  $u$  and  $\sum_{s_i \in S, s_i \neq s_u} 1$ , is the number of seasons other than the season when the user data was recorded. This score effectively detects if a venue is appropriate for a specific season by dividing the four seasons into two buckets: one is the user’s current season and the other one is given by the other seasons for which the average is computed.

**Travel Score.** We can assume two more dimensions in user’s context which can be leveraged to enhance the personalized ranking of venues. These two types of information are *Trip Type* and *Group Type*. Trip type indicates whether the user is visiting a venue for business or leisure, while group type defines the group that is accompanying the user in her trip (e.g., family, friends, etc.) We looked into the information available on some LBSNs to find the best possible match between such contextual information and the information about places. Some LBSNs track and report distribution of traveler types who visit a particular place. Therefore, we map these two contextual dimensions onto the available information from TripAdvisor. The score  $S_{ctx}^{travel}$  is calculated similar to Equation 2.

## 4. EXPERIMENTAL RESULTS

**Dataset.** We evaluate our model on the dataset of TREC Contextual Suggestion Track 2015. More in details, given a set of example venues as user’s preferences and some contextual signals, the task consists in returning a ranked list of candidate venues which fit the user’s profile and context.

**Evaluation Metrics.** We evaluate the performance of our proposed model by reporting P@5 (Precision at 5) and MRR (Mean Reciprocal Rank). Our model uses **Category**, **Tags**, **Reviews**, and **Context**, so we call it **CaTReCx**.

**Baselines.** We compare our method with the three top-ranked participants in the TREC Contextual Suggestion Track 2015 [4]. Our first baseline is the best performing run [1] (labeled as **BASE1**) which uses four scores (reviews from Yelp, categories from Yelp and from TripAdvisor, and keywords from Foursquare). It ranks the venues based on the linear combination of these scores. The second baseline is the second best run [5] (labeled as **BASE2**). It utilizes factorization machines for venue recommendation. The instances which are fed into the factorization machine are composed of three blocks representing user, context, and venue features. BASE2 uses Foursquare as its source of information. The third best run [7] (labeled as **BASE3**) creates positive and negative profiles for each user and adds to them all the reviews of similar users from Yelp. It also creates positive and negative profiles for venues. The venues are then ranked by linearly combining the similarity scores of all profile pairs. Finally, the ranked list of venues is modified by applying a number of contextual filters.

**Results.** We rank the venues considering all the aforementioned scores as features for LambdaMART learning-to-rank technique. We conduct our experiments using a 5-fold cross validation across the training data.

Table 1 shows the performance of our model as well as of the baselines. Experimental results demonstrate that our system outperforms all the three baselines by a significant margin when it uses all the three sources of information. Note that in order to perform a fair comparison between our work and the baselines, we also report our system’s performance using only the same source of information used by the respective baseline. More in details, CaTReCx achieves a 5.34% improvement in terms of P@5 and a 3.93% improvement in terms of MRR over BASE1. Our approach also exhibits an improvement over BASE2 using only the data from Foursquare (F) in terms of P@5. Moreover, it beats BASE3 in terms of both P@5 and MRR by a large margin. For completeness, we report the median performance of all participants of TREC. The performance of our methodology compared to the TREC median performance proves the effectiveness of our model. In particular, since our approach combines multimodal information from multiple LBSNs, it can significantly improve the precision of venue suggestions.

It is worth noting that we tried different classifier and regression algorithms for the review-based score component; however, since the SVM classifier with a linear kernel exhibited a much better performance than the other models we do not report the results of the others. In fact, as we discussed in our previous work [2] SVM is a perfect match for this classification problem since the number of positive training samples are dramatically more than the negative samples. Most classifiers tend to correctly classify the class with more training examples, while SVM is not affected by the relative size of the classes.

**Table 1: Performance of our model compared to three baselines and the median performance of TREC participants. *Y* represents Yelp as source of information, while *T* and *F* represent TripAdvisor and Foursquare, respectively.**

	Y	T	F	P@5	MRR
CaTReCx	✓	✓	✓	<b>0.6171</b>	<b>0.7695</b>
BASE1	✓	✓	✓	0.5858	0.7404
CaTReCx	✗	✗	✓	<b>0.5782</b>	0.7188
BASE2	✗	✗	✓	0.5706	<b>0.7190</b>
CaTReCx	✓	✗	✗	<b>0.6095</b>	<b>0.7453</b>
BASE3	✓	✗	✗	0.5583	0.6815
TREC Median	-	-	-	0.5090	0.6716

## 5. CONCLUSIONS AND FUTURE WORK

In this paper we described a personalized ranking model for context-aware venue suggestion. Our model aimed to capture various types of information from multiple sources which can be important to a user for visiting a venue. The experimental results on the TREC Contextual Suggestion Track dataset demonstrated our system effectiveness compared to the state of the art.

As future work, it would be interesting to define a generative probabilistic model which predicts user tags for a new venue by modeling empirically the mapping between the tags the user selected for a venue and its content.

## 6. ACKNOWLEDGMENTS

This work was partially supported by the RelMobIR project of the Swiss National Science Foundation (SNSF).

## 7. REFERENCES

- [1] M. Aliannejadi, S. A. Bahrainian, A. Giachanou, and F. Crestani. Univ. of lugano at TREC 2015: Contextual suggestion and temporal summarization tracks. In *TREC*, 2015.
- [2] M. Aliannejadi, I. Mele, and F. Crestani. User model enrichment for venue recommendation. In *AIRS*, 2016.
- [3] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 1995.
- [4] A. Dean-Hall, C. L. A. Clarke, J. Kamps, K. Julia, and E. M. Voorhees. Overview of the TREC 2015 contextual suggestion track. In *TREC*, 2015.
- [5] R. McCreadie, R. Deveaud, S. Mackie, M. Jarana, G. McDonald, V. Saul, C. Macdonald, and I. Ounis. Univ. of Glasgow at TREC 2015: Experiments with terrier in contextual suggestion, temporal summarisation and dynamic domain tracks. In *TREC*, 2015.
- [6] A. Rikitianskii, M. Harvey, and F. Crestani. A personalised recommendation system for context-aware suggestions. In *ECIR*, 2014.
- [7] P. Yang and H. Fang. Univ. of Delaware at TREC 2015: Combining opinion profile modeling with complex context filtering for contextual suggestion. In *TREC*, 2015.