University of Glasgow Terrier Team at TREC 2016: Experiments in Contextual Suggestions

Jarana Manotumruksa
University of Glasgow
Glasgow, Scotland, UK
j.manotumruksa.1@research.gla.ac.uk

Craig Macdonald
University of Glasgow
Glasgow, Scotland, UK
craig.macdonald@glasgow.ac.uk

Iadh Ounis
University of Glasgow
Glasgow, Scotland, UK
iadh.ounis@glasgow.ac.uk

ABSTRACT

For TREC 2016, we focus on tackling the challenges posed by the Contextual Suggestion by investigating the use of user-generated data (e.g. textual content of comments and venue's information) in location-based social networks (LBSNs) to suggest a ranked list of venues to users. In particular, we exploit a word embedding technique to extract user-venue and context-venue preference features to train learning-to-rank models. In addition, we leverage each venue's information (e.g. number of check-in) to extract venue-dependent features. We train learning-to-rank models using these features on the TREC 2015 Contextual Suggestion dataset. We submit two runs (*uogTrCs* and *uogTrCsContext*) where *uogTrCsContexl* is a context-aware approach. The batch experimental results show that *uogTrCS* is competitive, performing above the TREC median in terms of NDCG@5 and P@5 and outperforms *uogTrCsContext*.

1 INTRODUCTION

Similar to previous years, for the batch experiment in TREC 2016, the Contextual Suggestion track asks participants to suggest a ranked list of venues to users, based upon their profiles and preferred contexts [2]. The user's context contains the user's location (e.g. a city where he is looking for venues to visit) and contextual preferences: namely the *duration* of their trip (daytime, nighttime, weekend, longer), the *season* of the year (Spring, Summer, Autumn and Winter), the *group* of people the user is intending to visit the venue with (alone, friends, family and others) and the *type* of the trip (business, holiday and other). In the following, we first describe how we exploit word embedding techniques to leverage the textual content of comments to model user's preferences and characteristic of venues (Section 2. Then, we describe a set of features used to train learning-to-rank models in Section 3. Our proposed approaches for the batch experiments are described in Section 4. Section 5 highlights our submitted runs and their achieved performances.

2 MODELLING USER'S PREFERENCES AND CHARACTERISTICS OF VENUES

In this section, we describe how we leverage user-generated data from Location-based Social Networks (LBSNs) to improve the quality of the venue suggestion. In particular, we exploit word embeddings to model user's preferences and characteristic of venues from the textual content of comments. Our approach builds upon an intuition – which was shown to be effective in the TREC 2013 and 2014 Contextual Suggestion tracks – that people who have similar opinions about venues are likely to share similar venue preferences [4, 8]. However, our approach differs from the approaches of Yang et al. [8] and Manotumruksa et al. [4], in that we model the user's preferences and the characteristics of the venues by using word embeddings to better represent the users and venues as vectors in the word embedding space, rather than using traditional retrieval models. Further details and experiments for this approach can be found in [3].

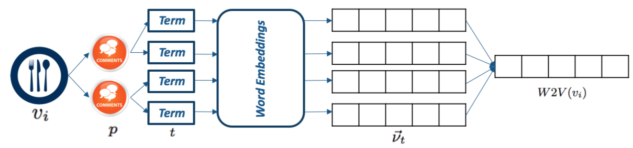

Figure 1: Modelling the characteristics of venues from comments using Word2vec.

To model the characteristics of venues, as illustrated in Figure 1, we exploit word embeddings to infer a vector-space representation of a venue from its comments. In particular, for each term t occurring in the comment posts $P_i = \{p_{i,1}, ..\}$ of venue v_i, we use a pre-trained word embedding model to represent the term as a embedding vector $\vec{v}_t$, and then aggregate those embedding vectors to model the characteristics of the venue as follows:

$$W2V(v_i) = \sum_{p \in P_i} \sum_{t \in p} \vec{v}_t \qquad (1)$$

By doing so, we obtain a vector-space representation $W2V(v_i)$ of the characteristics of venue v_i. Similar to the characteristic of venue, as illustrated in Figure 2, we model model the user's preferences, $\vec{UV}_j$, of user u_j by summing the vector of the venues rated in the user's profile U_j:

$$\vec{UV}_j = \sum_{v_i \in U_j} W2V(v_i) \times r_{i,j} \qquad (2)$$

We then estimate the similarity between venue v_i and the user preferences $\vec{UV}_j$, by calculating the cosine similarity between $\vec{UV}_j$ and $W2V(v_i)$. Typically, the user's preferences $\vec{UV}$ are modelled separately for positively (rating of 3 or 4) vs. negatively (0 or 1) rated venues in the user's profile.

In the TREC 2016 Contextual Suggestion track, the users can explicitly express their preferred context such as what time of the day they will visit the city and who they are going to visit venues with. Given a set of user's preferred contexts (e.g. day time, summer and friends), for each contextual dimension d, we exploit the word embeddings to identify a list of terms T_d that are related

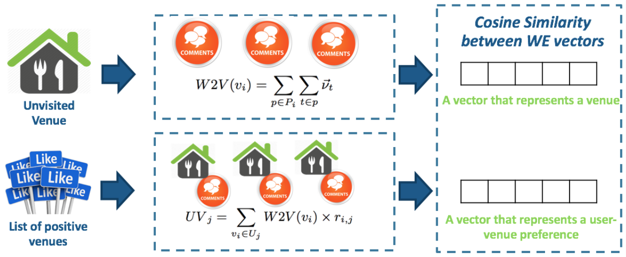

Figure 2: Modelling the characteristics of user's preferences using Word2vec.

to contextual dimension d. For example, terms that are close to the word 'Family' in the vector space are "grandparents", "supporting", "group", "mom", "kid", etc. Observing such words frequently in the comments posted for a venue on a LBSN intuitively increases our belief in that venue being relevant for recommendation scenarios involving *Family* contextual requirements. Then, we model the user's contextual preferences $\vec{CV}$ on contextual dimension d as follows (further details can be found in [3]):

$$\vec{CV}_d = \sum_{t \in T_d} \vec{v}_t \qquad (3)$$

Note that we explored various techniques to aggregate vector-space representations apart from summing (e.g. averaging, maximum as well as concatenation of different aggregations). Based on initial experiments, the summing is the most effective for this task.

3 CONTEXT-AWARE VENUE RANKING

In the section, we describe how we leverage the vector-space representation of characteristic of venue $W2V(v_i)$, user's preferences $\vec{UV}_j$ and user's contextual preference $\vec{CV}_d$ to extract useful features that can train learning-to-rank model to effectively generate ranked list of venues relevant to users' profiles. In addition, Deveaud *et al.* [1] and Mccreadie *et al.* [5] showed effectiveness of venue recommendation systems use any venue dependent sources of evidence such as number of comments and ratings in the TREC 2014-2015 Contextual Suggestion tracks. Therefore, we also extract venue dependent features from venue's information retrieved from Foursquare LBSN. The following summarises the features used in this work:

- **2 user-venue preference features (UV)**: Cosine similarity between the vector-space representation of characteristic of venue $W2V(v_i)$ (Equation (1)) and the vector-space representation of the user's preference of the user $\vec{UV}_j$ (Equation (2)) – one feature for each positive and negative venues user's preference.
- **4 context-venue preference features (CV)**: Cosine similarity between the vector-space representation of venue $W2V(v_i)$ and user's contextual preference $\vec{CV}_d$ (Equation (3)) – one feature for each contextual aspect.
- **6 venue-dependent features (V)**: Number of check-ins, number of likes, number of comments, number of photos, average Foursquare rating, and unique number of users (further details can be found in [1]).

Table 1: Results of our runs in the Contextual Suggestions track. Figures in bold represent the top performances.

	UV	V	CV	NDCG@5	P@5	MRR
TREC Median	-	-	-	0.2562	0.3931	**0.6014**
uogTrCS	√	√	×	**0.2756**	**0.4207**	0.5886
uogTrCsContext	√	√	√	0.2582	0.3828	0.5475

4 BATCH EXPERIMENT

In the batch experiment of the Contextual Suggestion track, we submitted two runs making use of learning-to-rank, but differing in terms of the sets of features deployed from those mentioned in Section 3. In particular, we submitted two runs, namely *uogTrCs* and *uogTrCsContext* – in particular:

- *uogTrCs* is a learning-to-rank approach that considers only user-venue preference features (UV) and venue-dependent features (V);
- *uogTrCsContext* is a context-aware learning-to-rank approach that consider all features mentioned in Section 3.

We trained both runs on the TREC 2015 Contextual Suggestion track test collection, by splitting user-context pairs 66% and 33% across training and validation sets. We use LambdaMART, a state-of-the-art listwise learning-to-rank technique [7], as the learner. Following Manotumruksa *et al.* [3], for word embeddings, we use the Word2Vec tool[1], training a skip-gram model [6] on the comments dataset from Foursquare (over 2.7M comments). We use the validation set to determine the word embedding parameters (e.g. vector size and window size). Before training the model, we perform stemming and remove standard stopwords from the comments.

5 BATCH RESULTS

Table 1 reports the performance of our two submitted runs (*uogTrCs* and *uogTrCsContext*) together with the TREC Median using the official measures, namely NDCG@5, P@5 & MRR. The columns UV, V and CV indicate type of features that are used by the runs. From the table, we observe that the *uogTrCs* run is the most effective of the tworuns across all three measures. Indeed, *uogTrCs* is competitive, performing above the TREC median in terms of both NDCG@5 and P@5. Overall, the results for both of our runs exhibit promising performance comparable with the TREC median, and hence merit further study in the future. By comparing the experimental results in Table 1 with the results reported by Manotumruksa *et al.* [3], consistent conclusions can be observed. In particular, the experimental results from these works conclude that modelling user's preferences $\vec{UV}_j$ (Equation (2)) and the characteristics of the venues $W2V(v_i)$ (Equation (1)) using the word embeddings is effective. However, in [3], Manotumruksa *et al.* reported the usefulness of user's contextual preferences $\vec{CV}_d$ (Equation (3)) in improving the performance of venue recommendation systems, while in TREC 2016 *uogTrCsContext* did not perform well in comparison with *uogTrCs*. Further failure analysis is required in order to answer why context-venue preference features are not effective.

[1]https://code.google.com/archive/p/word2vec/

6 CONCLUSION

For TREC 2016, we submitted two runs that exploit word embedding technique to extract useful features from user-generate data in Foursquare LBSN to train learning-to-rank models. The experimental results from the TREC 2016 Contextual Suggestion track show that modelling user's preferences and characteristics of venues using word embeddings is effective and overall our run, *uogTrCs* is competitive, performing above the TREC median in term of NDCG@5 and P@5. For future work, we plan to conduct a failure analysis to answer why context-venue preference features cannot improve the effectiveness of venue recommendation systems.

REFERENCES

[1] Romain Deveaud, M-Dyaa Albakour, Craig Macdonald, and Iadh Ounis. 2014. On the Importance of Venue-Dependent Features for Learning to Rank Contextual Suggestions. In *Proc. of CIKM*.

[2] Seyyed Hadi Hashemi, Charles LA Clarke, Jaap Kamps, Julia Kiseleva, and Ellen Voorhees. 2016. Overview of the TREC 2016 Contextual Suggestion Track. In *Proc. of TREC*.

[3] Jarana Manotumruksa, Craig Macdonald, and Iadh Ounis. 2016. Modelling User Preferences using Word Embeddings for Context-Aware Venue Recommendation. *arXiv preprint arXiv:1606.07828* (2016).

[4] Jarana Manotumruksa, Craig Macdonald, and Iadh Ounis. 2016. Predicting Contextually Appropriate Venues in Location-Based Social Networks. In *Proc. of CLEF*. Springer, 96–109.

[5] Richard McCreadie, Saul Vargas, Craig MacDonald, Iadh Ounis, Stuart Mackie, Jarana Manotumruksa, and Graham McDonald. 2015. University of Glasgow at TREC 2015: Experiments with Terrier in Contextual Suggestion, Temporal Summarisation and Dynamic Domain Tracks.. In *Procs of TREC*.

[6] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv:1301.3781* (2013).

[7] Qiang Wu, Chris J. C. Burges, Krysta M. Svore, and Jianfeng Gao. 2008. *Ranking, Boosting, and Model Adaptation.* Technical Report MSR-TR-2008-109. Microsoft.

[8] Peilin Yang and Hui Fang. 2013. Opinion-based user profile modeling for contextual suggestions. In *Proc. of ICTIR*.