Exploring Deep Space: Learning Personalized Ranking in a Semantic Space

Jeroen B. P. Vuurens
HHS & TU Delft
j.b.p.vuurens@tudelft.nl

Martha Larson
TU Delft & Radboud University
m.a.larson@tudelft.nl

Arjen P. de Vries
Radboud University
arjen@acm.org

1. MOTIVATION

The use of word embeddings is currently a strong trend in Natural Language Processing, which is successfully applied for tasks like analogous reasoning, question answering, and translation [2]. In [4], by learning embeddings for movies we capture factors that describe the differences between the movies. For recommendation, we rank the movie embeddings by their distance to a hyperplane, for which coefficients are learned using pairwise learning to rank to optimally rank a user's past preferences. Our experiments on Movielens1M show that using embeddings that are learned from user ratings significantly and greatly outperform state-of-the-art collaborative filtering algorithms. We show that the same architecture can also be used for content-based recommendations.

2. THEORY

When learning word embeddings, semantic differences between words are consistently encoded, e.g. gender encoding [3]. Similarly, for learned movie embeddings we observe that patterns that describe the differences between movies are also consistently encoded, e.g. movie genres, suspense (Figure 1). However, for movies there are many factors that are useful to describe why a user prefers some movies over others, e.g. favorite actor, scary elements, humor. Figure 1 illustrates that in a low-dimensional embedding space factors are not encoded independently and therefore not all items can be positioned ideally. Therefore we propose to learn high-dimensional item representations, which makes it possible to encode a greater number of factors independently.

When recommending items to a specific user, we seek to represent the user's interest over the factors that are encoded in the semantic space. An important consideration is that a user may only care about a limited number of factors. For instance, some users may like or dislike scary elements in movies, while other users are indifferent to whether a movie contains scary elements. Thus, when recommending items, we have two considerations. Firstly, in the embeddings space the 'best' recommendation candidates are positioned closely to the items the user rated highly. Secondly, the vector between two movies in semantic space should be reflected in their ranking by the extent to which the corresponding encoded factors are relevant to the user.

3. IMPLEMENTATION

In this study, we learned item embeddings using the DBOW variant of Paragraph2Vec [1]. To recommend items, we learn the coefficients to a hyperplane and use the distance to the hyperplane to rank the items. An advantage of using a hyperplane is the difference between movies that are preferred equally by the user can be ignored by choosing a parallel hyperplane. We implemented a deep learning variant to pairwise learning to rank to learn hyperplane coefficients that optimally rank a user's past preferences.

Figure 1: Encoded semantics for the most popular movies.

Table 1: Comparison of the effectiveness on MovieLens 1M. The subscripts in the column "sig. over" correspond to a significant improvement over the corresponding system, tested using McNemar test, 1-tailed, *p-value* < 0.001.

System	Recall@10	sig. over
BPRMF [1]	0.079	
UserKNN [2]	0.087	
WMRF [3]	0.089	
DS-CF-500	0.144	1,2,3
DS-CF-1k	0.151	1,2,3

4. RESULTS

To evaluate the effectiveness of our approach, we compare the results of our approach to that the MyMediaLite implementation of BPRMF, WRMF and UserKNN on Movielens1M in Table 1.

References

[1] Q. V. Le and T. Mikolov. Distributed representations of sentences and documents. In *Proceedings of ICML*, 2014.

[2] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[3] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.

[4] J. B. P. Vuurens, M. Larson, and A. P. de Vries. Exploring deep space: Learning personalized ranking in a semantic space. In *Proceedings of RecSys Deep Learning workshop*, 2016.