

Aligning Hotel Embeddings using Domain Adaptation for Next-Item Recommendation

Ioannis Partalas
ipartalas@expediagroup.com
Expedia Group
Geneva, Switzerland

ABSTRACT

In online platforms it is often the case to have multiple brands under the same group which may target different customer profiles, or have different domains. For example, in the hospitality domain, Expedia Group has multiple brands like Brand Expedia, Hotels.com and Wotif which have either different traveler profiles or are more relevant in a local context.

In this context, learning embeddings for hotels that can be leveraged in recommendation tasks in multiple brands requires to have a common embedding that can be induced using alignment approaches. In the same time, one needs to ensure that this common embedding space does not degrade the performance in any of the brands.

In this work we build upon the hotel2vec model and propose a simple regularization approach for aligning hotel embeddings of different brands via domain adaptation. We also explore alignment methods previously used in cross-lingual embeddings to align spaces of different languages. We present results on the task of next-hotel prediction using click sessions from two brands. The results show that the proposed approach can align the two embedding spaces while achieving good performance in both brands. Additionally, with respect to single-brand training we show that the proposed approach can significantly reduce training time and improve the predictive performance.

CCS CONCEPTS

• **Information Systems** → **Recommender Systems**; • **Applied Computing** → Electronic Commerce; • **Computing Methodologies** → Machine Learning.

KEYWORDS

embeddings, hotel2vec, alignment, domain adaptation, transfer learning

ACM Reference Format:

Ioannis Partalas. 2021. Aligning Hotel Embeddings using Domain Adaptation for Next-Item Recommendation. In *Proceedings of ACM SIGIR Workshop on eCommerce (SIGIR eCom'21)*. ACM, New York, NY, USA, 5 pages.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR eCom'21, July 15, 2021, Virtual Event, Montreal, Canada
© 2021 Copyright held by the owner/author(s).

1 INTRODUCTION

Online platforms often have multiple brands, for the same line of business, under the same group which may target different customer profiles. As an example, in the hospitality domain Expedia Group (EG) has brands like Brand Expedia, Hotels.com and Wotif, that can either have different profiles of travelers or be more relevant locally.

A main task in such online platforms or more generally in retail platforms, is to recommend products to customers which requires to learn an embedding space that captures their salient attributes. Hence, enable similarity comparisons that can be leveraged from recommendation systems.

In the recent years approaches that learn product embeddings from the interactions of the customers with the on-line platform have been proposed [7, 11, 20] as well as approached tailored to the hospitality domain [10, 18]. These approaches leverage the seminal word2vec model [15] in order to learn the embedding space by treating the clicked items as tokens in a sentence. While it is common to learn such embeddings in a single domain/brand, in the context of electronic commerce we would like to be able to leverage such embeddings across different domains/brands. As mentioned previously, one can learn hotel embeddings on Brand Expedia and leverage them in Hotels.com in order to bootstrap or improve the hotel embeddings in the latter case. Subsequently these hotel representations can be used in tasks like personalized recommendations.

To do so, one would need to align the embedding spaces of the different brands and use this aligned space to capture the intent of the users while searching on the on-line platform. In a very recent work, Bianchi et al. [5] study the alignment of product embeddings to enable zero-shot learning in a cross-shop scenario. Their setting is more general as in our case the multi-brands belong to the same domain, and usually we dispose a one-to-one mapping between products (hotels in our case).

In this work we propose to align embeddings from different brands using a simple regularization approach from domain adaptation. We build upon the hotel2vec model [18] for learning hotel embeddings and extend it to accommodate alignment of embedding spaces. Our approach, can also be thought of as a transfer learning one as we are able to bootstrap the learning procedure of hotel2vec in a target brand using the hotel2vec embeddings from a source brand.

The main contributions of this paper are the following:

- We propose a simple yet effective domain adaptation approach that adds a regularization term in the loss function of hotel2vec.

- We present empirical results on the task of next-hotel prediction for two brands.
- We also implement an alignment method borrowed from the task of alignment of cross-lingual embeddings. We show empirically, that to perform such a transpose to another domain one should be careful of the particularities of it.

2 RELATED WORK

In recommendation systems learning embeddings for users and items that capture the semantics is a critical part. In the domain of electronic commerce, approaches as prod2vec have been proposed in order to learn representations of products [3, 10, 11] which leverage the skip-gram model [15]. In [7] the authors propose to learn embeddings for YouTube videos by combining multiple features, which are used for candidate generation and ranking. Other approaches have also been proposed that include metadata [20] or try to capture different aspects of the product from different sources (clickstream data, text and images) [19].

Aligning embedding spaces is a topic that has been extensively studied in the NLP domain and specifically in the context of cross-lingual embeddings. A simple and efficient approach to align two embedding spaces of different languages is to learn a linear projection [14]. In [2] the authors employ an iterative approach starting from seed mappings of words (source language to target language) and do the linear projection by imposing an orthogonality constraint in the projection matrix. A survey on cross-lingual word embedding can be found in [17]. In this work we study the effectiveness of such alignment approaches in the context of hotel embeddings.

Domain adaptation is also a topic that has attracted interest in the recommendation systems space [12, 21] where the goal is to transfer knowledge from a source to a target task. In our setting we employ a domain adaptation approach in order to align different embedding spaces and we follow a straightforward regularization approach [8]. More recently, such approaches have been encompassed in a single framework for domain adaptation [13]. Other works, propose adversarial approaches for domain adaptation that learn in the same time the alignment as well as an embedding space that is invariant to the domain [6, 9, 16].

The most similar approach to our work is the one presented in [5] where the authors propose to align product embeddings to enable cross-shop recommendations. Specifically, they propose different approaches either based on content (images and text) or using the clicked products in the different shops as supervised signals in methods for alignment of cross-lingual embeddings. Our use case has a simpler setting as we dispose a mapping between products (hotels) in the two embeddings spaces we wish to align.

3 ALIGNMENT THROUGH REGULARIZED DOMAIN ADAPTATION

In order to learn representations of hotels, we use the hotel2vec model [18] which implements a neural model and is trained with the skip-gram model and negative sampling. The model learns different embeddings for click (V_c), hotel properties (V_a) and geographical information (V_g) which are fused to learn an enriched representation. Specifically, the embeddings are calculated as follows:

$$V_c = f(I_c; W_c)$$

$$V_a = f(I_a; W_a)$$

$$V_g = f(I_g; W_g)$$

where $f(x; W) = \text{ReLU}(\frac{xW}{\|xW\|_2})$ and I_c, I_a and I_g refer to the input features for the click, amenity and geographical embedding. The final hotel2vec embedding is calculated as a projection of the concatenated embeddings:

$$\text{ReLU}([V_c, V_a, V_g]^T W_e) \quad (1)$$

Let V_{h_i} be the representation for hotel h_i as calculated by equation 1 where $h_i \in H$. Then hotel2vec model minimizes the following loss function:

$$J(\theta) = - \sum_{(h_t, h_c) \in D^+} \log \sigma(V_{h_t}^T V_{h_c}) + \sum_{h_i \in N_c} \log \sigma(-V_{h_t}^T V_{h_i}) \quad (2)$$

where D^+ are the skip-gram pairs of clicked hotels in the session that are generated using a fixed length window. N_c are the negative samples which are sampled from the same market of clicked hotels in the session, as a traveler searches in a specific destination. Finally, $\sigma()$ is the sigmoid function.

In this work we propose to extend the hotel2vec model in order to accommodate embedding spaces alignment in the case of multi-brand representations.

We do so by employing a domain adaptation scheme. Denote \mathcal{D}^S the source domain where we already learned a hotel2vec model by minimizing the loss function in Equation 2, $J_s(\theta)$. Then in the target domain \mathcal{D}^T we learn hotel2vec representations by minimizing the following regularized function:

$$J_t(\theta) = J(\theta) + \lambda \|V_{h_t}^T - V_{h_t}^S\|_2$$

where $V_{h_t}^S$ refers to the corresponding embedding of hotel h_t in the source domain. Note that the embeddings of the source domain are fixed and not re-trained. Also, λ is a parameter that controls how much knowledge we would like to transfer from the source domain. As in our case we want to align the embedding spaces, we would like to constraint the model to be as close as possible in the source domain, hence set λ equal to 1.0. In the experimental section we experiment with this parameter to understand the impact in the downstream task.

Note that we define the strength of regularization globally but in a more fine version it could be defined per hotel. We leave this for future work.

The above regularizer assumes that we dispose a mapping between hotels in the two domains which is our case. Also, note that we do not suffer here from the problem of absent hotels from the training data in the source domain as the model can impute these ones.

4 EXPERIMENTAL SETTING

We evaluate the proposed approach in the task of next-item prediction where we want to predict the next clicked hotel in the session based on the previous clicked hotel.

We collect click sessions over one year of searches for both Hotels.com and Brand Expedia. It contains more than 65M user click sessions for each brand and around 700K unique hotels for Hotels.com dataset and over 1.4M for Brand Expedia respectively. We randomly split the sessions in training, validation and test with a ration of 8:1:1.

We use a system with 64GB RAM, 8 CPU cores, and a Tesla V100 GPU. We use Python 3 as the programming language and the Tensorflow [1] library for the neural network architecture and gradient calculations.

For hotel2vec we follow the experimentation methodology in [18] for tuning the hyperparameters of the model. The model in both brands is trained with $L2$ -regularization.

For the alignment approaches, we use as source domain Hotels.com hotel2vec embeddings and Brand Expedia as the target one. For the proposed approach we set the parameter λ to 1.0 in order to force the alignment of the spaces. We also tried a value of 0.5 in order to understand the behavior of the approach.

We compare the regularization approach with the linear projection alignment approach presented in [14]. In that case given the embeddings that we export from the trained models in Hotels.com and Brand Expedia domains, denoted \mathcal{V}^S and \mathcal{V}^T respectively, we solve the following optimization problem:

$$\min_W \|\mathcal{V}^S W - \mathcal{V}^T\|_2^2$$

The alignment is learned only on the common hotel embeddings for the two sets of vectors in order to avoid injecting too much noise.

Metrics: We compare the different approaches in terms of hits@k and MRR@k (Mean Reciprocal Rank). Hits@k measures the average number of times the correct selection (i.e. the hotels clicked by the users in a session) appears in the top k predicted hotels. MRR@k evaluates the average list quality of k items returned by the model.

Both metrics are calculated over the ranking induced by cosine similarity of the embedding vectors. Note that when we have a cross-brand setting this corresponds to a zero-shot learning framework.

5 RESULTS

We present in this section results for two settings: a) the zero-shot prediction of next clicked hotel where we use embeddings in a cross-brand setting and b) the supervised setting where we compare the in-brand hotel2vec model with and without domain adaptation.

5.1 Zero-shot Prediction

Table 1 presents the hits@k and MRR@k for $k \in \{10, 100\}$ to predict next-clicked hotel in the test sets for both Brand Expedia and Hotels.com.

The first two approaches, named *hotel2vec_H* and *hotel2vec_E*, correspond to single-domain trained models with no further alignment. The results, show how the corresponding embeddings behave in a cross-brand zero-shot learning approach.

The next two approaches present the corresponding metrics for the linear projection alignment approach (LP) as well as the proposed approach that employs domain adaptation (*hotel2vec_{DA}*).

Focusing only in the single-domain approaches, we can observe that the embeddings can enable zero-shot learning in a cross-brand

scenario as they can achieve good performance across all the metrics. For example, with the *hotel2vec_H* we can achieve 0.5073 of hits@100 compared to the in-brand trained model *hotel2vec_E* that achieves a 0.5335 hits@100. The same observation stands in the case of Hotels.com even though we notice that the gap between the different performances is larger.

Concerning the alignment approaches presented in the next two rows of the table, we notice that the LP method has a big drop in performance in both brands. As we mentioned earlier, even though we dispose a one-to-one mapping of the hotels between brands the method fails to keep the similarities in the projected space. This shows that extra care should be taken when aligning the embedding spaces, for example require the projection matrix to be orthogonal.

Regarding the proposed approach we can observe that is able to achieve good performance in both brands. Actually, it achieves the same performance with the in-domain model in Hotels.com, 0.7535 versus 0.7573 for hits@100, while there is a drop in performance for the Brand Expedia. We recall here that we set the λ parameter that controls the transfer of knowledge to either 0.5 or 1.0 where the latter allows us to force the alignment between brands. In the case of less regularization strength ($\lambda = 0.5$) we can observe that the metrics slightly improve. Smaller values of the parameter can balance between a full transfer of knowledge and in-brand learning. For the improvement in the Hotels.com brand we hypothesize that allowing for less regularization we are able to allow more in-domain knowledge to flow and learn a better similarity space.

5.2 Supervised Setting

In this section we present results when we compare the same in-brand model with and without domain adaptation. Note that in this case we let the neural model score the candidate hotels.

Figures 2 and 1 present the hits@10 and hits@100 respectively for hotel2vec with domain adaptation and a single-brand model. We can notice that the proposed regularization scheme improves both the metrics as well as it can reduce training time. Also, the jump-start is significant and shows that we can share knowledge between the two brands.

Table 2 presents the hits@k metric for the in-brand hotel2vec model in comparison with the proposed approach when we use the model to score the next-hotel to be clicked. We can clearly see that the domain adaptation approach outperforms the single-brand approach by 3% for hits@100 metric.



6 CONCLUSIONS

We presented in this work a simple yet effective regularization approach for aligning embedding spaces in a multi-brand scenario. For example, Expedia Group has multiple brands that operate in the same domain. The idea is to add a regularizer in the objective function of the model that learns the embeddings in order to force them to be as close as possible to the embeddings of the source brand, hence performing domain adaptation. This kind of approaches has also been explored in the past in Natural Language Processing tasks [8].

We evaluated the proposed approach in the next-hotel prediction task for two brands, that is Expedia and Hotels.com. We measured performance in terms of hits@k and MRR@k metrics. We also,

Table 1: Hits@k and MRRk for the different approaches for embeddings alignment. Note that the hotel2vec approach refers to the standard training with no alignment for the two brands.

Embeddings	Expedia				Hcom			
	hits@100	hits@10	MRR@10	MRR@100	hits@100	hits@10	MRR@10	MRR@100
<i>hotel2vec_H</i>	0.5073	0.1593	0.0393	0.0519	0.7573	0.2981	0.1102	0.1276
<i>hotel2vec_E</i>	0.5335	0.1667	0.0411	0.0545	0.7039	0.2624	0.1014	0.1178
LP	0.3593	0.1203	0.0302	0.0389	0.4715	0.1970	0.0861	0.0962
<i>hotel2vec_{DA}</i> , $\lambda = 1.0$	0.5015	0.1591	0.0392	0.0517	0.7535	0.2943	0.1094	0.1268
<i>hotel2vec_{DA}</i> , $\lambda = 0.5$	0.5042	0.1594	0.0393	0.0519	0.7553	0.2964	0.1097	0.1271

**Figure 1: Hits@100 on the test set during training. We compare the domain adaptation approach with the single hotel2vec model on Brand Expedia clicks.****Figure 2: Hits@10 on the test set during training. We compare the domain adaptation approach with the single hotel2vec model on Brand Expedia clicks.**

compared with linear projection alignment borrowed by the cross-lingual approaches for aligning embeddings of different languages [14].

Table 2: Comparison of the single-brand hotel2vec model versus the domain adaptation approach in terms of hits@k. The metrics are measured when the scores for the next-hotel prediction are produced by the model.

Approach	hits@10	hits@100
<i>hotel2vec_E</i>	17.76	56.11
<i>hotel2vec_{DA}</i>	18.10	57.81

The results, showed that the proposed approach can align the spaces of the multiple brands achieving good performance in both brands. Additionally, we showed that the regularized version of the hotel2vec model can learn faster and improve the performance. We also observed that an approach like the linear projection without taking into account some particularities of the domain leads to worse performance.

7 FUTURE WORK

For future work we would like to add a more adaptive regularization parameter that can be defined per hotel rather than being global. In that way we may wish to transfer knowledge when we are certain that a pair of hotels in the source and target brands should have the same embedding.

We would also like to explore the use of multiple source brands in order to align in the same time the embedding spaces. Multi-task approaches can be leveraged to align the different embedding spaces [13]. Also, other alignment approaches could be explored [4].

Finally, we would like to explore adversarial cross-domain adaptation for aligning the embedding spaces [9]. In this case, we want to leverage the similar features across the brands while also learning specific embeddings for each brand.

ACKNOWLEDGMENTS

We would like to thank Phong Nguyen, Cédric Houdré and Clara Jordá Lope for discussions as well as Jan Krasnodebski and Daniele Dhongi for comments on the paper. Many thanks to Petros Baltzis for engineering support.

REFERENCES

- [1] Martin Abadi, Ashish Agarwal, Paul Barham, and et. al. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/>

- Software available from tensorflow.org.
- [2] Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2017. Learning bilingual word embeddings with (almost) no bilingual data. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Vancouver, Canada, 451–462. <https://doi.org/10.18653/v1/P17-1042>
 - [3] Oren Barkan and Noam Koenigstein. 2016. Item2vec: neural item embedding for collaborative filtering. In *2016 IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 1–6.
 - [4] Federico Bianchi, Valerio Di Carlo, Paolo Nicoli, and Matteo Palmonari. 2020. Compass-aligned Distributional Embeddings for Studying Semantic Differences across Corpora. *CoRR abs/2004.06519* (2020). arXiv:2004.06519 <https://arxiv.org/abs/2004.06519>
 - [5] Federico Bianchi, Jacopo Tagliabue, Bingqing Yu, Luca Bigon, and Ciro Greco. 2020. Fantastic Embeddings and How to Align Them: Zero-Shot Inference in a Multi-Shop Scenario. *CoRR abs/2007.14906* (2020). arXiv:2007.14906 <https://arxiv.org/abs/2007.14906>
 - [6] Xilun Chen, Yu Sun, Ben Athiwaratkun, Claire Cardie, and Kilian Weinberger. 2018. Adversarial Deep Averaging Networks for Cross-Lingual Sentiment Classification. *Transactions of the Association for Computational Linguistics* 6 (2018), 557–570. https://doi.org/10.1162/tacl_a_00039
 - [7] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (Boston, Massachusetts, USA) (*RecSys '16*). ACM, New York, NY, USA, 191–198. <https://doi.org/10.1145/2959100.2959190>
 - [8] Hal Daumé III. 2007. Frustratingly Easy Domain Adaptation. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*. Association for Computational Linguistics, Prague, Czech Republic, 256–263. <https://www.aclweb.org/anthology/P07-1033>
 - [9] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-Adversarial Training of Neural Networks. *J. Mach. Learn. Res.* 17, 1 (Jan. 2016), 2096–2030.
 - [10] Cheng Grbovic. 2018. Real-time Personalization using Embeddings for Search Ranking at Airbnb. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD, 311–320.
 - [11] Mihajlo Grbovic, Vladan Radosavljevic, Nemanja Djuric, Narayan Bhamidipati, Jaikrit Savla, Varun Bhagwan, and Doug Sharp. 2015. E-commerce in your inbox: Product recommendations at scale. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 1809–1818.
 - [12] Heishiro Kanagawa, Hayato Kobayashi, Nobuyuki Shimizu, Yukihiko Tagami, and Taiji Suzuki. 2019. Cross-Domain Recommendation via Deep Domain Adaptation. In *Advances in Information Retrieval*, Leif Azzopardi, Benno Stein, Norbert Fuhr, Philipp Mayr, Claudia Hauff, and Djoerd Hiemstra (Eds.). Springer International Publishing, Cham, 20–29.
 - [13] Wei Lu, Hai Leong Chieu, and Jonathan Löfgren. 2016. A General Regularization Framework for Domain Adaptation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, 950–954. <https://doi.org/10.18653/v1/D16-1095>
 - [14] Tomáš Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. Exploiting Similarities among Languages for Machine Translation. *CoRR abs/1309.4168* (2013). arXiv:1309.4168 <http://arxiv.org/abs/1309.4168>
 - [15] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
 - [16] Saeid Motiian, Quinn Jones, Seyed Iranmanesh, and Gianfranco Doretto. 2017. Few-Shot Adversarial Domain Adaptation. In *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2017/file/21c5bba1dd6aed9ab48c2b34c1a0adde-Paper.pdf>
 - [17] Sebastian Ruder. 2017. A survey of cross-lingual embedding models. *CoRR abs/1706.04902* (2017). arXiv:1706.04902 <http://arxiv.org/abs/1706.04902>
 - [18] Ali Sadeghian, Shervin Minaee, Ioannis Partalas, Xinxin Li, Daisy Zhe Wang, and Brooke Cowan. 2019. Hotel2vec: Learning Attribute-Aware Hotel Embeddings with Self-Supervision. *CoRR abs/1910.03943* (2019). arXiv:1910.03943 <http://arxiv.org/abs/1910.03943>
 - [19] Loveperteek Singh, Shreya Singh, Sagar Arora, and Sumit Borar. 2019. One Embedding To Do Them All. *arXiv preprint arXiv:1906.12120* (2019).
 - [20] Flavian Vasile, Elena Smirnova, and Alexis Conneau. 2016. Meta-prod2vec: Product embeddings using side-information for recommendation. In *Proceedings of the 10th ACM Conference on Recommender Systems*. ACM, 225–232.
 - [21] Feng Yuan, Lina Yao, and Boualem Benatallah. 2019. DARec: Deep Domain Adaptation for Cross-Domain Recommendation via Transferring Rating Patterns. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 4227–4233. <https://doi.org/10.24963/ijcai.2019/587>