2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining

Neural Embedding Features for Point-of-Interest Recommendation

Alireza Pourali
Ryerson University
alireza.pourali@ryerson.ca

Fattane Zarrinkalam
Ryerson University
fzarrinkalam@ryerson.ca

Ebrahim Bagheri
Ryerson University
bagheri@ryesron.ca

***Abstract*—The focus of point-of-interest recommendation techniques is to suggest a venue to a given user that would match the users' interests and is likely to be adopted by the user. Given the multitude of venues and the sparsity of user check-ins, the problem of recommending venues has shown to be a difficult task. Existing literature has already explored various types of features such as geographical distribution, social structure and temporal behavioral patterns to make a recommendation. In this paper, we propose a new set of features derived based on the neural embeddings of venues and users. We show how the neural embeddings for users and venues can be jointly learnt based on the prior check-in sequence of users and then be used to define three types of features, namely user, venue, and user-venue interaction features. These features are integrated into a feature-based matrix factorization model. Our experiments show that the features defined over the user and venue embeddings are effective for venue recommendation.**

***Index Terms*—Point-Of-Interest recommendation, Neural Embedding, Feature-based matrix factorization**

I. Introduction

Due to the growing popularity of location-based social networks (LBSNs), point-of-interest recommendation systems have received growing attention by researchers and many service providers are now focused on developing recommendation systems for their end-users [1]–[3]. The current point-of-interest recommendation models primarily take various features, such as spatio-temporal characteristics of user check-ins, transitions between venues, human mobility patterns and category of venues into account [4]–[6]. Given the collection of these features, different machine learning methods such as deep neural network, matrix factorization and collaborative filtering models are utilized to recommend a point-of-interest to a user [7]–[9].

In this paper, we propose a new set of features based on the neural embedding of users and venues. Our features rely only on neural embeddings trained based on users' check-ins, which place users and venues within the same embedding space. We show how the shared user and venue embedding space can give way to different types of features, which are then incorporated into a feature-based matrix factorization technique for venue recommendation. The contributions of this paper include:

1) We systematically show how the check-in sequence of users can be utilized to learn neural embeddings for both users and venues within the same embedding space;
2) We introduce three classes of features based on neural embeddings of users and venues, which capture the characteristics of users, venues and their interactions;
3) We formalize how the defined features are encoded into feature-based matrix factorization for venue recommendation and extensively benchmark our work.

This paper extends state of the art by showing how user check-in sequences can be used to learn neural user and venue embeddings and how these neural embeddings can be used to define strong indicative features for venue recommendation. In our experiments, we show that features extracted from the user and venue embeddings are able to effectively outperform the state of the art methods.

II. Proposed Approach

We formally define the problem of point-of-interest recommendation as follows:

(Venue Recommendation). Let the *check-in sequence* of user $u \in \mathbb{U}$, denoted by $CS(u) = (v_u^t)$, be a sequence of venues ordered on t, where v_u^t shows a venue $v \in \mathbb{V}$ that user u checked in at time t. Based on $CS(u)$ from T consecutive time intervals, we aim at recommending a ranked list of venues for user u, each denoted by v_u^x such that $x \in T+1$. The recommendations are ranked descendingly, based on the predicted degree of interest of user u in v at time x. We first learn user and venue embeddings within the same space, then use these embeddings to derive features that can be used in feature-based matrix factorization. One of our key contributions is to learn user and venue representations in the same embedding space based on the check-in sequence of users with the expectation that a user would be closer to the venues that she has checked-in at compared to those she has not. We show that by embedding users and venues in the same space, it will be possible to construct different types of features based on these embeddings, which can be directly incorporated into a feature-based recommendation model. In the following, we introduce the embedding model to learn user and venue representations. Then, we introduce our features constructed

ASONAM'19, August 27-30, 2019, Vancouver, Canada

ACM ISBN 978-1-4503-6868-1/19/08.../$15.00
https://doi.org/10.1145/3341161.3343672

from the embeddings and finally describe how these features are exploited in a feature-based recommendation model.

A. User and Venue Embeddings

Our goal is to learn user and venue embeddings in the same space so that a user is closer to the venues that she has checked-in compared to those that she has not. To this end, given the check-in sequence of all the users in T consecutive historical time intervals, i.e., $\{CS(u) = (v_u^t)\,|u \in \mathbb{U}, 1 \leq t \leq T\}$, we map each user and venue into an L-dimensional embedding vector, denoted by e_u and e_v within the same embedding space, in the same feature space by applying a neural embedding model based on the framework shown in Figure 1.

We consider the sequence of checked-in venues of each user u as a sentence and place the ID of user u at the beginning of each sentence which results in $|\mathbb{U}|$ sentences. In the learning process, both the ID of users $|\mathbb{U}|$ and venues $|\mathbb{V}|$ are treated as word tokens and a user ID is essentially always associated with a set of check-ins. In other words, during training, for each sentence, the sliding context window will always include the first word in the sentence (i.e., user ID). Formally, given the check-in sequence of each user u, $CS(u) = (v_u^t)$, and a context window size of k, the objective is to maximize the following log probability:

$$\frac{1}{T}\sum_{t=k}^{T-k} log\, pr(v_u^t|v_u^{t-k},...,v_u^{t+k},u) \tag{1}$$

We adopt a softmax function to model $pr(v_u^t|v_u^{t-k},...,v_u^{t+k},u)$ as follows:

$$pr(v_u^t|v_u^{t-k},...,v_u^{t+k},u) = \frac{exp(e_{context}^{\top}.e_{v_u^t})}{\sum\limits_{v\,\in\,\mathbb{V}} exp(e_{context}^{\top}.e_{v_u^t})} \tag{2}$$

$$e_{context} = h(e_{v_u^{t-k}},...,e_{v_u^{t+k}},e_u) \tag{3}$$

where $e_{v_u^t}$ denotes the embedding vector of venue $v_u^t \in CS(u)$ and $e_{context}$ denotes the context vector, which is obtained by averaging over the embedding vectors of the context information (Equation 3). Further, $h(.)$ is a function that averages the embedding vectors. By sampling a target venue and its context window at each iteration, our network is trained using stochastic gradient descent and updates the parameters via backpropagation. As such, the final result of the learning process is an L-dimensional embedding for each user (e_u) and each venue (e_v).

B. Embedding-based Features

Given user and venue embeddings, i.e., e_u and e_v, which are embedded in the same space, we define three classes of features:

User features: for each user $u \in \mathbb{U}$, we propose two types of features: 1) *user embedding* feature, which is equivalent to the L-dimensional embedding vector for user u, i.e. e_u; and 2) *top-N similar users* feature that is inspired from the main idea of collaborative filtering based on which similar users share similar check-in behaviour in the future. Therefore, for a given user u, we consider her top-N similar users, denoted by S_u. This provides us with N features, which correspond to the most similar users to a target user. To compute the similarity of two users u_1 and u_2, we apply cosine similarity between e_{u_1} and e_{u_2}.

Venue features: for each venue $v \in \mathbb{V}$, we define two types of features: 1) *venue embedding* feature, which is the learnt L-dimensional embedding vector for venue v, i.e., e_v; and 2) *top-N similar venues* feature where for a venue v, we calculate its similarity with other venues through cosine similarity of their embedding vectors and then select the top-N most similar venues to v, denoted as S_v.

Global user-venue interaction feature: Given both user and venue embeddings in the same feature space, we can incorporate a global feature to denote the similarity between any given user and venue pair. The global feature of a user-venue pair such as (u, v) is calculated by the cosine similarity between e_u and e_v.

C. Venue Recommendation

These features are encoded into a feature-based matrix factorization as follows:

$$\widehat{r}_{u,v}(\alpha,\beta,\gamma) = \mu + \textstyle\sum_j b_j^{(g)}\gamma_j + \sum_j b_j^{(u)}\alpha_j + \sum_j b_j^{(v)}\beta_j + (\sum_j \alpha_j x_j)^{\top}(\sum_j \beta_j y_j) \tag{4}$$

where $\alpha \in R^{N_\alpha}$, $\beta \in R^{N_\beta}$ and $\gamma \in R^{N_\gamma}$ are the input vectors consisting of the features of user u, the features of venue v and the global feature for the pair (u, v) with the lengths of N_α, N_β and N_γ, respectively. Further, $b_j^{(g)}$, $b_j^{(u)}$ and $b_j^{(v)}$ are the global, user and venue bias parameters. The latent vectors x_j and y_j capture the j^{th} user feature and the j^{th} venue feature, respectively. The global features and bias values do not have any corresponding latent vectors. The response value, i.e., $\widehat{r}_{u,v}(\alpha,\beta,\gamma)$, predicts whether user u will check-in at venue v in the future. In the following, we show how to encode our features into Equation 4.

Encoding user features: In the user input vector $\alpha \in R^{N_\alpha}$, consisting of the user features of user u, we reserve the first $|\mathbb{U}|$ dimensions to one-hot encode the ID of user u and its top-N most similar users, i.e. S_u. Therefore, given $Q = \{u\} \cup S_u$, each user $q \in Q$ is encoded by inserting a "1" in the q^{th} dimension of the $|\mathbb{U}|$-dimensional vector. The rest of dimensions are set to "0". Then, given a user u, we use the dimensions from $(|\mathbb{U}+1|)$ to $(|\mathbb{U}+L|)$ to encode her user embeddings, e_u. Therefore, the total number of user features is $N_\alpha = |\mathbb{U}| + L$.

Encoding venue features: We reserve the first $|\mathbb{V}|$ dimensions in the venue input vector $\beta \in R^{N_\beta}$, to one-hot encode a venue v and its top-N similar venues, i.e. S_v. Then, given a venue v, we use the dimensions from $(|\mathbb{V}+1|)$ to $(|\mathbb{V}+L|)$ to encode the venue embeddings, e_v. Total number of venue features is $N_\beta = |\mathbb{V}| + L$.

Encoding global user-venue interaction feature: The idea behind the user-venue interaction feature is that a user is

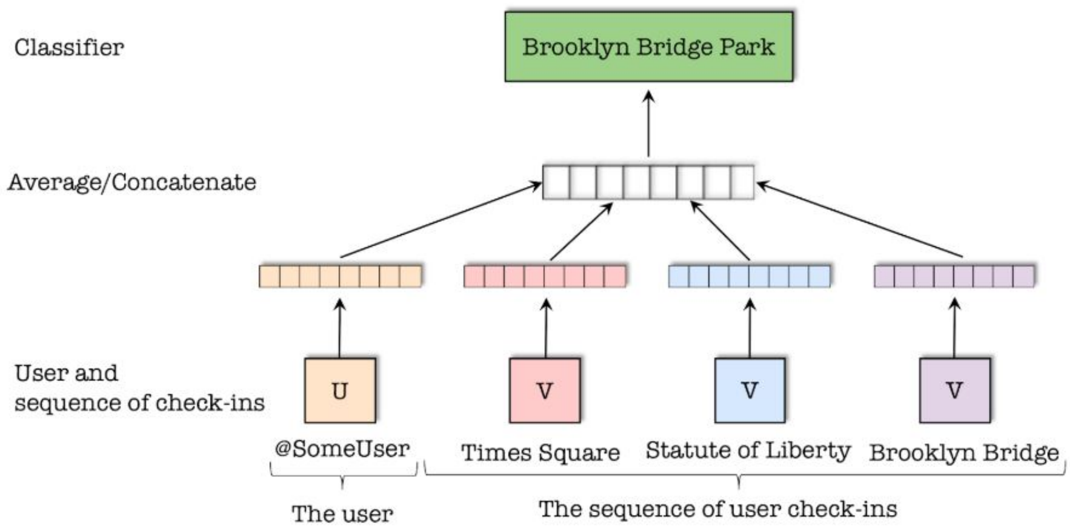

Fig. 1. A framework for learning user and venue embeddings in the same feature space.

more likely to check-in at a venue closer in the latent feature space to those she has checked in at before; therefore, the corresponding venue should receive a larger global bias value. Therefore, we define the global user-venue interaction feature as $\gamma_1 = cosineSim\,(e_u, e_v)$ and the number of global features is $N_\gamma = 1$.

With these coded features, for a user-venue pair (u, v), based on Equation 4, we have the following factorization formula:

$$\widehat{r}_{u,v}(\alpha,\beta,\gamma) = \mu + b_1^{(g)}\gamma_1 + \sum_j b_j^{(u)}\alpha_j + \sum_j b_j^{(v)}\beta_j + \left(\sum_{q\in\{u\}\cup S_u} x_q + \sum_{l=1}^{L} e_{u,l}\, x_{|\mathbb{U}|+l}\right)^{\top}\left(\sum_{p\in\{v\}\cup S_v} y_p + \sum_{l=1}^{L} e_{v,l}\, y_{|\mathbb{V}|+l}\right) \quad (5)$$

The parameters are trained by minimizing the log-likelihood loss function using stochastic gradient descent. Our goal is to construct a ranked list of venues that u may be interested in. Therefore, for a user u, we first calculate the value $\widehat{r}_{u,v}$ for all the pairs $\{(u,v)\,|\,v\in\mathbb{V}\}$ and then construct a descendingly ranked list on $\widehat{r}_{u,v}$.

III. EXPERIMENTS

A. Dataset and Evaluation Methodology

Our experiments were conducted on a dataset collected from the Gowalla LBSN introduced in [1]. The check-in data in this dataset are collected from 600,000 users from November 2010 to December 2011 that captured their check-ins in various cities of the United States, i.e., Austin, Chicago, Houston, Los Angeles and San Francisco. Table 1 shows some statistics about the number of users and venues in each city.

Our evaluation strategy and metrics are based on [1], which suggests to randomly select 70% of the check-ins of each user as the training data and leave the rest of the check-ins for testing. Further, the quality of the recommendations is measured based on Precision@6, Recall@6 and F1-Score. In terms of the parameters of our model, we train our embeddings based on a dimension size of $L = 100$. Further, we perform 10-fold cross-validation and select $N = 5$, which is the best performing value for N in top-N similar user and venue features. The hyperparameters of the matrix factorization model were set as suggested in [10], namely, the learning rate, the regularization parameter for the user factor, the regularization parameter for the venue factor and the number of latent factors were set to 0.005, 0.004, 0.004 and 64.

TABLE I
NUMBER OF USERS AND VENUES PER CITY

City	Users	Venues	$\langle \frac{\text{checkins}}{\text{user}} \rangle$	$\langle \frac{\text{checkins}}{\text{venue}} \rangle$
Austin	339	7936	138.226	5.905
Chicago	257	1704	18.732	2.825
Houston	163	6812	162.135	3.880
Los Angeles	280	3607	43.179	3.352
San Francisco	370	5447	62.405	4.239

We evaluate the performance of the three proposed features (Embedding (E), Interaction (I) and Top-N (T) features) through a combination of seven variants. For example, the variant "EI" means that embedding vectors feature and the global user-venue interaction feature are included in the model but the top-N feature is not included. The results of comparing the variants of our model are reported in terms of Precision@6 and Recall@6 in Figure 1 based on the 70-30 split that was mentioned earlier. As depicted in this figure, it can be observed that embedding features show the lowest performance. We believe the reason for this is that the number of embedding features is quite high (200 for user and venue features together) and each of the features does not convey any meaning in isolation. Therefore, the high number of highly dependent features leads to poor prediction. However, when merged with the top-N similar features the quality of the results improve. This observation shows that highly similar users and venues, regardless of their social connection to the user of interest, can serve as strong indicators of relevant venues for a user. It can also be seen in the results that when these two features are combined with the global feature, the results do not show any noticeable improvement. On the other hand, when the global feature is used alone or in combination with the top-N similar features, the accuracy of venue recommendation increases. We conclude that the top-N similar features are the most effective for point of interest recommendation due to finding similar users and venues that are located close to each other in the embedding space. Given the way the embeddings are learnt, this finding basically means that considering information from (1) similar users that have shown similar check-in behavior and (2) venues that have had similar user check-ins have the highest predictive power for point-of-interest prediction. So in summary, we find:

1) Among the features, top-N similar users and venues serve as the best performing features compared to the other features.
2) The best variant that benefits from more than one feature is the "TI" model that uses both Top-N similar features and the interaction between them, but it still has a weaker performance when compared to top-N similar features.

B. Features Analysis

These findings show that embedding users and venues in the same space has been able to identify users with similar check-in patterns and venues with similar user check-ins behavior. Given such effective embedding, the consideration of top-N features and the global user-venue interaction feature based on these embeddings will point to the users that have similar behavioral patterns and so can be used for recommending the best venue. Based on this observation, we select model "T" as the venue recommendation model to be compared to the other baselines. [vs]

C. Comparison with Baselines

For comparison, we use those baselines that have already used the same benchmark dataset and the same evaluation strategy:

- **CPOIR** [1] is a Category-aware venue recommendation model based on the transitioning behavior of each user between the categories of locations.
- **BasicMF** is a basic matrix factorization technique that only considers the users' historical check-ins and their preferences for venue recommendation.
- **POILP** [11] formulates the problem as heterogeneous link prediction over venue categories, venue regions and user relationships in addition to historical check-in data of users.
- **GeoCF** [12] proposes an integrated collaborative filtering model by using the user geographical influence and user preference for venue recommendation.
- **MGMMF** [13] is based on the captured geographical influence of users' check-ins using a Multi-center Gaussian Model fused with matrix factorization over users' social influence.
- **Markovian** [14] is based on users' mobility behavior where the locations are recommended by observing n previously visited venues.
- **UMFL** [15] treats venue recommendation as a supervised learning problem where global mobility, user mobility, and temporal features are incorporated in the model.

We perform our comparative analysis based on the strategy proposed in [1] and so only the top-6 venues are taken into consideration and reported in Table 2. We observe that the BasicMF method shows the worst performance compared to other models. This can be partly attributed to the fact that this model only employs the user's similar venue visit patterns to make recommendations. Other models such as the Markovian model which is solely based on temporal features of user footprints and recommends locations to users based on the previously visited locations outperform BasicMF. However, this model performs less accurately due to the fact that users' sequential mobility data is quite sparse. The other baseline in our work is the UMFL model, which uses a geographical feature for venue recommendation and outperforms BasicMF and the Markovian model. Nonetheless, MGMMF and GeoCF that fuse the geographical influence and user interest offer more accurate recommendations compared to the UMFL model. Also, the CPOIR model is a category-aware venue recommendation model, which is based on the transitional performance of the users between the locations and their categories. This model outperforms all of the previous models. The POILP model, which uses link prediction for venue recommendation achieves better results compared to CPOIR model in terms of precision. However, it shows lower Recall compared to the others. As reported in Table 2, our proposed model outperforms all the baselines in all three metrics. This observation indicates that by embedding user and venue vectors, a new set of features can be developed that are suitable for venue recommendation. By presenting the users and venues into the same space, not only closer users and locations are discovered, but also similar user-location pairs are found. We found that direct social connection between users is not necessarily an accurate indicator for venue recommendation. This is because users' social connections do not directly translate into visiting the same venues. In contrast, we developed an alternative feature for finding top-N similar users and venues for each user and venue, which might not even be directly connected to the user or the venue. These top users and venues have similar temporal check-in behavior and as such have a similar embedding vector. The incorporation of these top-N users and venues substantially increases performance.

IV. Related Work

There is a rich line of research on Point-of-Interest (venue) recommendation that utilizes users' check-in data. Existing point-of-interest recommendation techniques can be classified

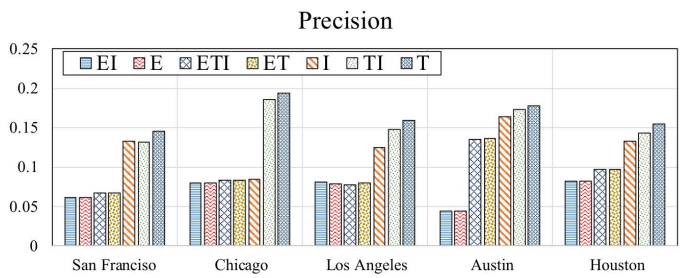

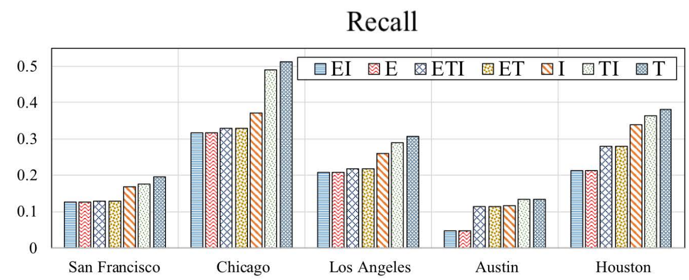

Fig. 2. The Precision@6 and Recall@6 of the seven variants of our proposed approach.

TABLE II
COMPARATIVE ANALYSIS WITH THE BASELINES DIVIDED ACROSS DIFFERENT CITIES.

City	Austin			Chicago			Houston			Los Angeles			San Francisco		
	R	P	F1	R	P	F1	R	P	F1	R	P	F1	R	P	F1
CPOIR	**0.157**	0.026	0.045	0.292	0.049	0.083	0.279	0.046	0.080	0.203	0.034	0.058	0.159	0.027	0.045
BasicMF	0.064	0.011	0.018	0.086	0.014	0.025	0.082	0.014	0.024	0.072	0.012	0.021	0.066	0.011	0.019
POILP	0.056	0.096	0.071	0.096	0.130	0.111	0.207	0.122	0.153	0.120	0.112	0.116	0.072	0.116	0.089
GeoCF	0.122	0.020	0.035	0.227	0.038	0.065	0.165	0.027	0.047	0.164	0.027	0.047	0.126	0.021	0.036
MGMMF	0.117	0.020	0.034	0.186	0.031	0.053	0.159	0.027	0.045	0.152	0.025	0.046	0.112	0.019	0.032
Markovian	0.086	0.014	0.025	0.116	0.019	0.033	0.102	0.017	0.029	0.096	0.016	0.027	0.088	0.015	0.025
UMFL	0.116	0.019	0.033	0.170	0.028	0.049	0.152	0.025	0.044	0.132	0.022	0.038	0.111	0.018	0.032
Our Approach	0.134	**0.178**	**0.153**	**0.513**	**0.194**	**0.282**	**0.380**	**0.155**	**0.220**	**0.306**	**0.159**	**0.153**	**0.195**	**0.146**	**0.167**

into different categories. Matrix Factorization [16] which is a method used for rating prediction has been widely used and integrated with various approaches in this field. Particularly, MF seeks to find the latent elements of venues and users to precisely predict the rating of the person to unvisited venues. Based on collaborative filtering, Matrix Factorization has been used to recommend venues to the users based on their similarities with others. For instance, Manotumruksa et al. [5] have addressed the problem of venues rating prediction by using Matrix Factorization on word embeddings. Liu et al. [1] has proposed a category-aware venue recommendation model that exploits the transition of users' preference among venue categories. They apply matrix factorization to predict the interest of users over venues in different categories. Manotumruksa et al. [2] have proposed a Matrix Factorization regularization technique for point-of-interest recommendation which fuses users social information and their comments using word embeddings. In their work, venues are recommended to the users by looking at their similar friends in terms of visiting locations. This has been done by using the comments left by the users along their check-in information and embedding the words they have used.

Recurrent Neural Networks (RNN) based point-of-interest recommendation models have become popular to their state-of-art performance. Manotumruksa et al. [9] have proposed a Deep Recurrent Collaborative Filtering Framework (DRCF) for venue suggestions by considering the complex user-venue interactions, users feedback and the geographical information of the locations. Spatial and temporal information have been taken into account in Liu et al. [17] extended RNN model where spatio-temporal context is modeled in each layer. Lately, Yao et al. [7] have proposed a Semantics-enriched Recurrent Model (SERM) which jointly learns the embeddings of various features and the transition parameters. The main purpose of these approaches is to enable a semantic-aware point-of-interest recommendation model to improve recommendation accuracies.

Collaborative recommendation approaches have been extensively explored in traditional recommender systems due to its positive impact on the quality of recommendations. Inspired from the presumption that friends of LBSNs share more common pursuits than non friends, several POI recommendation strategies enhance the quality of their recommendations by taking community impact straight into account such as [12]. For instance, in their approach, POI recommendations is based on a friend-based collaborative filtering model. Their model is trained based on the friendships similar preferences instead of the similarity between all the users in the social network.

Personalized point-of-interest recommendation models have been also investigated by using various techniques. Gambs et al. [14] have proposed a model for next venue prediction using the mobility Markov chain that is built for each individual user. Their model is solely based on the temporal behaviour of the user check-ins and does not consider the correlation between users' check-ins. Yang et al. [18] have proposed a fusion framework, which exploits both spatial and temporal activity preferences of users to predict their next venue. For each user, the spatial features are captured by building Personal Functional Regions, which are built based on frequented regions that the user visits. In their work, each region is assigned to a category that the user is more interested in based on her historical visits. Therefore, when the user is near each region, the category assigned to that region is used for venue recommendation.

Lately, Aliannejadi et al. [4] have proposed a personalized context-aware point-of-interest recommendation model where the venues are suggested based on a set of similarity scores. In their work, the venues are recommended contextually by using a combination of context and user based similarity scores such as frequency-based scores, venue tags score and review-based score. Finally, the similarity scores are combined by using linear interpolation and a ranked list of venues are used for point-of-interest recommendation. Also, in this work, the lifestyle behavior of users is observed by using the temporal nature of the check-ins. Pourali et al. [11] have viewed the problem of venue recommendation as an instance of the link prediction problem. They have incorporated various types of information into a heterogeneous graph such as the users, locations, categories of the locations, and the association between the users.

In this paper, we represent users and venues within the same feature space solely based on check-in data of users by utilizing neural embeddings that have already been widely used in related areas [19], [20]. Then, We utilize the user and venue embeddings to define three types of features which are incorporated in a feature-based matrix factorization model for venue recommendation.

V. Concluding Remarks

We have shown how a sequence of user check-ins can be used to learn user and venue representations in order to introduce new types of features. Results from earlier methods had shown that explicit social connections of each user does not necessarily lead to effective recommendations and as such, we built user relationships based on finding the top-N similar users and venues based on the embeddings. This feature showed to be the strongest for predicting points of interest. Additionally, we used the similarity of the embeddings of a user-venue pair to serve as a feature, which was a strong feature but not as strong as the top-N similar user and venue features. Finally, we considered the raw embeddings of users and venues to serve as features, which did not show a good performance. We also systematically explored the possibility of building models that use a combination of these features and found that the top-N similar features are the strongest features both in isolation and when combined with other features. We showed that compared to state of the art, our variant that only relies on top-N features shows a noticeably better performance.

References

[1] X. Liu, Y. Liu, K. Aberer, and C. Miao, "Personalized point-of-interest recommendation by mining users' preference transition," in *Proceedings of the 22Nd ACM International Conference on Information & Knowledge Management*, ser. CIKM '13. New York, NY, USA: ACM, 2013, pp. 733–738. [Online]. Available: http://doi.acm.org/10.1145/2505515.2505639

[2] J. Manotumruksa, C. Macdonald, and I. Ounis, "Regularising factorised models for venue recommendation using friends and their comments," in *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, ser. CIKM '16. New York, NY, USA: ACM, 2016, pp. 1981–1984. [Online]. Available: http://doi.acm.org/10.1145/2983323.2983889

[3] S. A. M. Falavarjani, F. Zarrinkalam, J. Jovanovic, E. Bagheri, and A. A. Ghorbani, "The reflection of offline activities on users online social behavior: An observational study," *Information Processing and Management*, 2019. [Online]. Available: https://www.journals.elsevier.com/information-processing-and-management

[4] M. Aliannejadi and F. Crestani, "Venue appropriateness prediction for personalized context-aware venue suggestion," in *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, Shinjuku, Tokyo, Japan, August 7-11, 2017*, 2017, pp. 1177–1180.

[5] J. Manotumruksa, C. Macdonald, and I. Ounis, "Matrix factorisation with word embeddings for rating prediction on location-based social networks," in *Advances in Information Retrieval - 39th European Conference on IR Research, ECIR 2017, Aberdeen, UK, April 8-13, 2017, Proceedings*, 2017, pp. 647–654. [Online]. Available: https://doi.org/10.1007/978-3-319-56608-5_61

[6] Z. Yao, "Exploiting human mobility patterns for point-of-interest recommendation," in *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM 2018, Marina Del Rey, CA, USA, February 5-9, 2018*, 2018, pp. 757–758. [Online]. Available: https://doi.org/10.1145/3159652.3170459

[7] D. Yao, C. Zhang, J. Huang, and J. Bi, "SERM: A recurrent model for next location prediction in semantic trajectories," in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 - 10, 2017*, 2017, pp. 2411–2414. [Online]. Available: https://doi.org/10.1145/3132847.3133056

[8] C. Yang, L. Bai, C. Zhang, Q. Yuan, and J. Han, "Bridging collaborative filtering and semi-supervised learning: A neural approach for POI recommendation," in *SIGKDD*, 2017, pp. 1245–1254.

[9] J. Manotumruksa, C. Macdonald, and I. Ounis, "A deep recurrent collaborative filtering framework for venue recommendation," in *CIKM*, 2017, pp. 1429–1438.

[10] T. Ge, Z. Sui, and B. Chang, "Exploiting collaborative filtering techniques for automatic assessment of student free-text responses," in *CIKM*, 2013, pp. 1493–1496.

[11] A. Pourali, F. Zarrinkalam, and E. Bagheri, "Poi recommendation using heterogeneous link prediction," in *EDBT*, 2018, pp. 481–484.

[12] M. Ye, P. Yin, W. Lee, and D. L. Lee, "Exploiting geographical influence for collaborative poi recommendation," in *SIGIR*, 2011, pp. 325–334.

[13] C. Cheng, H. Yang, I. King, and M. R. Lyu, "Fused matrix factorization with geographical and social influence in location-based social networks," in *AAAI*, 2012.

[14] S. Gambs, M. Killijian, and M. P. Cortez, "Next place prediction using mobility markov chains," in *Eurosys' MPM*, 2012.

[15] A. Noulas, S. Scellato, N. Lathia, and C. Mascolo, "Mining user mobility features for next place prediction in location-based services," in *ICDM*, 2012, pp. 1038–1043.

[16] Y. Koren, R. M. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *IEEE Computer*, vol. 42, no. 8, pp. 30–37, 2009.

[17] Q. Liu, S. Wu, L. Wang, and T. Tan, "Predicting the next location: A recurrent model with spatial and temporal contexts," in *AAAI*, 2016, pp. 194–200.

[18] D. Yang, D. Zhang, V. W. Zheng, and Z. Yu, "Modeling user activity preference by leveraging user spatial temporal characteristics in lbsns," *IEEE Trans. SMC: Systems*, vol. 45, no. 1, pp. 129–142, 2015.

[19] H. Fani, E. Bagheri, and W. Du, "Temporally like-minded user community identification through neural embeddings," in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 - 10, 2017*, 2017, pp. 577–586. [Online]. Available: https://doi.org/10.1145/3132847.3132955

[20] H. Fani, E. Jiang, E. Bagheri, F. Al-Obeidat, W. Du, and M. Kargar, "User community detection via embedding of social network structure and temporal content," *Information Processing and Management*, 2019. [Online]. Available: https://www.journals.elsevier.com/information-processing-and-management