Short Paper

Point-of-Interest Recommendation Using Heterogeneous Link Prediction

Alireza Pourali
Laboratory for Systems, Software and Semantics (LS^3)
Ryerson University
alireza.pourali@ryerson.ca

Fattane Zarrinkalam
Laboratory for Systems, Software and Semantics (LS^3)
Ryerson University
fzarrinkalam@ryerson.ca

Ebrahim Bagheri
Laboratory for Systems, Software and Semantics (LS^3)
Ryerson University
bagheri@ryerson.ca

ABSTRACT

Venue recommendation in location-based social networks is among the more important tasks that enhances user participation on the social network. Despite its importance, earlier research have shown that the accurate recommendation of appropriate venues for users is a difficult task specially given the highly sparse nature of user check-in information. In this paper, we show how a comprehensive set of user and venue related information can be methodically incorporated into a heterogeneous graph representation based on which the problem of venue recommendation can be efficiently formulated as an instance of the heterogeneous link prediction problem on the graph. We systematically compare our proposed approach with several strong baselines and show that our work, which is computationally less-intensive compared to the baselines, is able to shows improved performance in terms of precision and f-measure.

1 INTRODUCTION

The recent interest in location-based social networks (LBSNs) such as Foursquare and Gowalla has generated a high volume of user location information on the Web [17]. This has attracted researchers to analyze social data related to users' point of interests (POI) based on their preferences and personalities, which has the potential to improve the quality of higher-level applications such as civic planning, healthcare, advertising/marketing, and crime prediction [10], just to name a few. The recommendation of a point-of-interest to a given user based on her past activity is one of the applications that has already been explored by several authors [3]. These earlier works primarily use features such as spatial and temporal activity of users where both spatial features (location of venues) and temporal features (time of visit) are taken into account. The spatial features are extracted using the geographical information of the user check-ins, which are longitude and latitude of the venues. The collection of these features are then used to train classifiers to determine and recommend a point-of-interest for a given user.

In our work, we take a different perspective on the same problem of point-of-interest recommendation by formalizing user LBSN information in the form of a heterogeneous graph which consists of different node types including users, venues, venue categories, and geographical regions, among others. We propose that the problem of point-of-interest recommendation can be viewed as an instance of the link prediction problem on heterogeneous graphs. In this paper, we systematically show (1) how a collection of LBSN information including user relationships, past user-venue interaction history, venue category information and geographical coordinates can be unified into and represented as a heterogeneous graph; (2) how meta-paths can be extracted from such a heterogeneous graph representation in order to identify potential links between users and venues (points-of-interest); and (3) that compared to much more complex baselines that incorporate spatio-temporal information into strong recommender techniques such as matrix factorization models, our proposed approach shows improved performance. We extensively compare our work with several state of the art techniques based on the gold standard dataset from [7] and show that while our heterogeneous link prediction model is quite computationally lightweight, it can offer improved performance over more computationally demanding techniques.

2 RELATED WORK

There has already been strong work on Point-of-Interest (venue) recommendation in the literature. Gambs et al. [4] have proposed an extended version of regular Markov Chains which incorporates the *n* previous visited venues of users. This model was then used for next location prediction using the users historical location visits. In their work, the next location of the users was predicted using the mobility Markov chain that was built for each individual user. Most of the location recommendation markov based models only use the geographical information of the locations without considering the context of the user footprints.

Matrix factorization has been widely used in location recommendation. In [2], matrix factorization is fused with geographical influence using Multi-center Gaussian Model (MGM) and social influence. In their work, location recommendation is based on the probability of a Gaussian distribution model, which is applied to the checked-in location centers along with the fusion framework with user preferences. However, the rich information of the check-in footprints such as the geographical context are not taken into account. In contrast, Yang et al. [18] have proposed a fusion framework, which exploits both spatial and temporal activity preferences (using tensor factorization) of users to predict their next point-of-interest. For each user, the spatial features are captured by building *Personal Functional Regions*, which are built based on frequented regions that the user visits. This model uses the spatial and temporal features separately for recommendation to users. In their work, each region is assigned with 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. Also, in this work, the lifestyle behavior of users is observed by using the temporal nature of the check-ins. It should be noted that Yang's model does not recommend venues to users and only recommends location categories.

Supervised learning models were also investigated in [9] for location recommendation. In this work, two supervised methods, linear regression and M5 trees were compared and it was

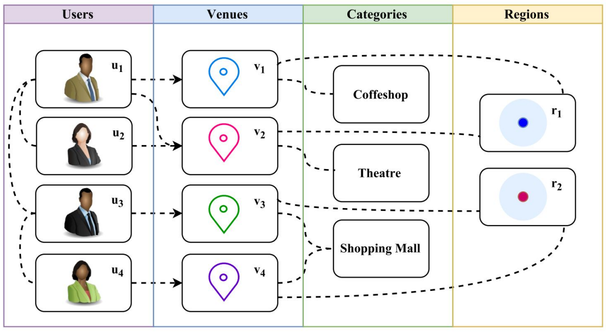

Figure 1: Sample heterogeneous graph representation of LBSN data.

found that M5 trees achieve higher prediction accuracy. Similar to other work, the set of features that were used for learning the classifiers were the user transitions between locations and spatio-temporal features of the check-in footprints. In our work, we have trained our model using the linear regression classifier and have achieved better recommendation accuracy due to using the features collaboratively and not individually compared to the work by Noulas et al. [9]. In [18] and [7], the importance of venue context information such as categories for location recommendation was also investigated. The accuracy of point-of-interest recommendations was improved by using the users' preferences that were based on the location categories they had visited. While in most existing work the current location of the user is needed for location recommendation; in our work, we show improved performance compared to the state of the art without requiring information about the last location of the user.

3 PROPOSED APPROACH

The objective of our work is to recommend a point-of-interest (venue) to a user based on her historical check-in data. Formally, given the check-in profile of a user in time interval t (Definition 1), we aim at recommending a list of venues that the user may be interested in at time interval $t + 1$.

Definition 1. ***(User Check-in Profile).*** *The check-in profile of user $u \in \mathbb{U}$ at time interval t, with respect to a set of venues $\mathbb{V}$, denoted by $CP^t(u)$, is represented by a vector of weights over the K venues, i.e., $(f_u^t(v_1), ..., f_u^t(v_K))$, where $f_u^t(v_k)$ is equal to one if user u checked in at venue v in time interval t and 0 otherwise.*

We propose to turn the problem of venue recommendation into a link prediction problem that operates over a heterogeneous graph. In addition to historical check-in data of users, there are other types of data that can be considered while recommending venues, namely venue categories, venue regions and user relationships. In this paper, we combine these data points into a unified heterogeneous graph representation model to consider them simultaneously. An illustration of our underlying representation model G can be found in Figure 1.

As illustrated in Figure 1, our representation model contains four types of nodes and four types of relations. Besides User and Venue nodes and User-Venue relations that represent historical check-in data of users, other types of data that are included consist of the following:

- *Category nodes*: In location-based social networks, venues are organized by a hierarchical category tree which provides a semantic classification of the various venues. For example, Foursquare contains a 3-level category hierarchy where categories are grouped into 10 top-level categories, such as Event, Food, Nightlife Spot and Residence. Each top-level category is classified into different subcategories. In our approach, we have infused the categories at the lowest level (Level 3) as our category nodes.
- *Region nodes*: Each region indicates a geographical area. Given the longitude and latitude of existing venues, we use X-means clustering [11] to cluster geographical coordinates and extract different regions.
- *Venue-Category relations*: Based on the hierarchical category tree defined in LBSNs for organizing venues, we assign each venue to its corresponding category in the lowest level of the hierarchy.
- *Venue-Region relation*: To identify the region of a venue, we calculate the Euclidean distance between the venue geographical coordinate (longitude and latitude) and the center of the identified regions and connect each venue to its nearest region.
- *User-User relation*: Users are connected to each other based on the friendship relation among them on the LBSN. By using this relation, potential interaction between users is also taken into account for point-of-interest recommendation.

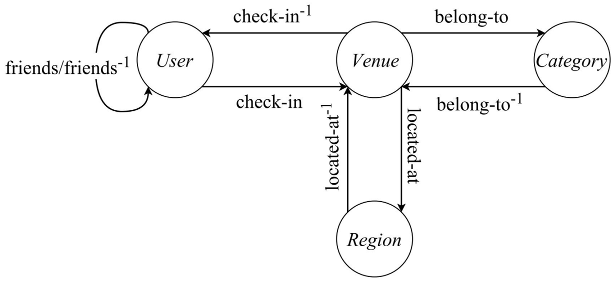

Figure 2: Network schema of the representation model.

Table 1: Meta-Paths between users and venues.

Meta-Path	Meaning of the Meta-Path
$\mathbb{U}-\mathbb{U}-\mathbb{V}$	A user visits a venue where her friends have visited
$\mathbb{U}-\mathbb{V}-\mathbb{C}-\mathbb{V}$	A user visits venues that belong to the same category
$\mathbb{U}-\mathbb{V}-\mathbb{R}-\mathbb{V}$	A user visits venues located in the same region

Having built the representation model G, in order to recommend a point-of-interest to a user $u \in \mathbb{U}$, we formulate a graph-based link prediction problem that operates over G. As our representation model is a heterogeneous graph, the neighbors of an object could belong to multiple types and the paths between two objects could have different meanings. Therefore, it is not possible to apply link prediction strategies such as Adamic/Adar and Common Neighbor, which treat all types of nodes and relations as the same in the form of a homogeneous graph [6].

Sun et al. [16] proposed the concept of heterogeneous information networks and the meta-path concept for heterogeneous information network analysis, which are now widely known and used in different data mining tasks such as ranking [8], clustering

Table 2: Performance comparison of different approaches.

City	Austin			Chicago			Houston			Los Angeles			San Francisco		
	R	P	F1	R	P	F1	R	P	F1	R	P	F1	R	P	F1
Our Approach	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**
CPOIR [7]	**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
GeoCF [19]	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 [2]	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
Markov [4]	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
ML [9]	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

[13], link prediction [1], and influence analysis [12]. In order to solve the problem of link prediction in heterogeneous graphs, Sun et al. [14] proposed the PathPredict method, i.e., meta path-based relationship prediction model to predict links between dissimilar node types. Therefore, to distinguish different types of objects and relations, following the works in [14, 15], we use the PathPredict method to determine the relevance of a point-of-interest $v \in \mathbb{V}$ for a user u. The core of the PathPredict method rests on the idea of *meta-paths*. A meta-path is a path defined over the heterogeneous network schema which can be used to define topological features with different semantic meanings. Figure 2 summarizes our representation model using a meta-structure known as the network schema.

Based on PathPredict, for the target relation $< \mathbb{U}, \mathbb{V} >$, we define a set of meta-paths starting with type $\mathbb{U}$ and ending with type $\mathbb{V}$ other than the target relation itself. We extract all such meta-paths by traversing the network schema using Breadth-First Search (BFS) within a fixed length constraint (max = 3). The extracted meta-paths and their semantic meaning are shown in Table 1. For example, The meta-path $\mathbb{U} - \mathbb{V} - \mathbb{C} - \mathbb{V}$, i.e., user-venue-category-venue considers those venues which belong to the same category of the historical check-in venues of a user as her next check-in venue.

Once the meta-paths are retrieved from the network schema, for each user-venue pair in the representation model G, we use the Degree-Weighted Path Count metric [5] to quantify each meta-path as a topological feature in the training step. For the given meta-path, Degree-Weighted Path Count penalizes paths, which pass through high-degree nodes. Then, given all user-venue pairs in the representation graph G and the extracted topological features for them, a logistic regression classifier is trained as the learning model to recommend a ranked list of points-of-interest for a given user.

4 EXPERIMENTS

In this section, we describe our experiments in terms of the dataset, setup and the details of the baselines used in the paper. The performance of our approach is then compared to the state of the art baselines and our findings are discussed.

4.1 Dataset and Experimental Setup

Our experiments were conducted on a dataset collected from the popular location-based social network of Gowalla introduced in [7]. It includes check-in data (longitude, latitude, timestamp, categories, among others.) of more than 600,000 users from Austin, Chicago, Houston, Los Angeles and San Francisco. For each user, we randomly select 70% of her check-ins to construct the training data and the remaining 30% of her check-ins for testing data as suggested in [7].

We compare our approach with several state of the art point-of-interest recommendation methods that are briefly described in the following:

(1) **BasicMF** is a classical matrix factorization techniques, which only considers users' past venue check-ins and, hence their preferences, for Point-of-Interest recommendation.
(2) **GeoCF** [19] is based on both user preference and geographical influence which are integrated into a collaborative filtering model.
(3) **MGMMF** [2] is a framework based on Multi-Center Gaussian Model, which combines both the user preference and MGM based check-in probability for Point of interest recommendation.
(4) **Markov** [4] applies Mobility Markov Chain model for predicting next venue of a user based on her mobility behavior over different time intervals and the recent venues that she has visited.
(5) **ML** [9] considers user mobility, global mobility and temporal features to describe users' check-in behavior and applies M5 decision tree to predict the next POIs of a user.
(6) **CPOIR** is one of the most related work in the literature by Liu et al [7], which proposes a Category-aware Point-Of-Interest Recommendation model that exploits the transition behavior of users between venue categories. They employ a matrix factorization model to predict the transition patterns of users' interests over categories and consequently her interests in different venues.

For evaluation purposes, we measure the performance of the methods based on Precision@K, Recall@K and F1-score as suggested in [7].

4.2 Experimental Results

In this section, inline with [7], we compare the performance of our proposed approach with other state of the art baselines when Top-6 venues are recommended by each method. The results are reported in Table 2 in terms of Recall, Precision and F1-score.

It can be observed that BasicMF model, which is solely based on user interests performs worse than others for most of the cities and in terms of all three metrics. This means that incorporating other auxiliary information such as geographical, social, and temporal features leads to improved quality of venue recommendation. Markov models that incorporate temporal features outperform BasicMF; however, they perform much less accurately than the other baselines. This is because Markov models assume that a user's mobility data is dense, as a result they may not perform so well on users' venue data in LBSNs which is very sparse.

As another observation, MGMMF and GeoCF that fuse geographical influence and user interest into their proposed approach and take into account the correlation between these features offer more accurate recommendations compared to the ML model, which exploits geographical influence as a single feature. As mentioned before, the CPOIR model incorporates two novel features, i.e., the transition patterns of a user's interests among venues and venue categories, for the purpose of point-of-interest recommendation. It can be seen that CPOIR offers superior results compared to both MGMMF and GeoCF in all cities and in terms of all metrics. This demonstrates the benefits of these two factors to improve the quality of venue recommendations.

Inspired from these insights, in order to be able to utilize the benefits of useful features in a unified model for point-of-interest recommendation, we formalize user LBSN information in the form of a heterogeneous graph. As highlighted in Table 2, it is evident that our proposed approach outperforms all the comparison methods in all cities in terms of Precision and F1-score. This means that our approach can successfully take advantage of different features, i.e., venue categories, geographical influence, the relationship between users, and the correlation between these features to produce more accurate recommendations, i.e., less false positive. However, it can be observed that our proposed approach results in lower Recall values, i.e., more false negatives compared to others. In other words, our method is not able to identify those venues through the current limited set of defined meta-paths in this paper. To cover more meta-paths, as our future work, we intend to increase the length constraint of meta-paths, which is currently set to 3, and take into account more features in our representation model such as transition of user's interests and sentiments of user-venue related comments.

5 CONCLUDING REMARKS

In this paper, we have proposed that venue recommendation in location-based social networks can be viewed as an instance of the link prediction problem on heterogeneous graphs. As such, we have systematically shown how various types of information can be incorporated into a heterogeneous graph based on which distance metrics between nodes can be employed as features to learn ranking classifiers. We find that even a logistic regression method can effectively show competitive performance compared to the state of the art despite its simplicity and light-weight computational requirements. We further found that while our proposed approach shows improved performance over the baselines in terms of precision and f-measure, it does not show competitive performance in recall. This can be attributed to the fact that we have only employed three meta-paths and a depth of three in the BFS search.

Our future work will explore two synergistic directions: 1) We will explore whether a more extensive set of meta-paths defined over the network schema can lead to improved recall or not. We will also study whether a higher search depth is able to identify more relevant information that can be included in the link prediction process. 2) Some users provide written textual feedback about their experience at venues they visit. These include textual reviews or recommendations. We are interested in the possibility of incorporating such unstructured user feedback into the network schema to see whether textual feedback, while quite sparse, can improve the venue recommendation task.

REFERENCES

[1] Bokai Cao, Xiangnan Kong, and Philip S. Yu. 2014. Collective Prediction of Multiple Types of Links in Heterogeneous Information Networks. In *2014 IEEE International Conference on Data Mining, ICDM 2014, Shenzhen, China, December 14-17, 2014.* 50–59. https://doi.org/10.1109/ICDM.2014.25

[2] Chen Cheng, Haiqin Yang, Irwin King, and Michael R. Lyu. 2012. Fused Matrix Factorization with Geographical and Social Influence in Location-Based Social Networks. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada.* http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/4748

[3] Chen Cheng, Haiqin Yang, Michael R. Lyu, and Irwin King. 2013. Where You Like to Go Next: Successive Point-of-Interest Recommendation. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013.* 2605–2611. http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6633

[4] S. Gambs, M.O. Killijian, and M.N. Prado Cortez. 2012. Next Place Prediction using Mobility Markov Chains. In *Eurosys' MPM.*

[5] Daniel S. Himmelstein and Sergio E. Baranzini. 2015. Heterogeneous Network Edge Prediction: A Data Integration Approach to Prioritize Disease-Associated Genes. *PLoS Computational Biology* 11, 7 (2015). https://doi.org/10.1371/journal.pcbi.1004259

[6] David Liben-Nowell and Jon M. Kleinberg. 2003. The link prediction problem for social networks. In *Proceedings of the 2003 ACM CIKM International Conference on Information and Knowledge Management, New Orleans, Louisiana, USA, November 2-8, 2003.* 556–559. https://doi.org/10.1145/956863.956972

[7] Xin Liu, Yong Liu, Karl Aberer, and Chunyan Miao. 2013. Personalized point-of-interest recommendation by mining users' preference transition. In *22nd ACM International Conference on Information and Knowledge Management, CIKM'13, San Francisco, CA, USA, October 27 - November 1, 2013.* 733–738. https://doi.org/10.1145/2505515.2505639

[8] Xiaozhong Liu, Yingying Yu, Chun Guo, and Yizhou Sun. 2014. Meta-Path-Based Ranking with Pseudo Relevance Feedback on Heterogeneous Graph for Citation Recommendation. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management, CIKM 2014, Shanghai, China, November 3-7, 2014.* 121–130. https://doi.org/10.1145/2661829.2661965

[9] Anastasios Noulas, Salvatore Scellato, Neal Lathia, and Cecilia Mascolo. 2012. Mining User Mobility Features for Next Place Prediction in Location-Based Services. In *12th IEEE International Conference on Data Mining, ICDM 2012, Brussels, Belgium, December 10-13, 2012.* 1038–1043. https://doi.org/10.1109/ICDM.2012.113

[10] Jon Parker, Yifang Wei, Andrew Yates, Ophir Frieder, and Nazli Goharian. 2013. A framework for detecting public health trends with Twitter. In *Advances in Social Networks Analysis and Mining 2013, ASONAM '13, Niagara, ON, Canada - August 25 - 29, 2013.* 556–563. https://doi.org/10.1145/2492517.2492544

[11] Dan Pelleg and Andrew W. Moore. 2000. X-means: Extending K-means with Efficient Estimation of the Number of Clusters. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000), Stanford University, Stanford, CA, USA, June 29 - July 2, 2000.* 727–734.

[12] Chuan Shi, Yitong Li, Jiawei Zhang, Yizhou Sun, and Philip S. Yu. 2017. A Survey of Heterogeneous Information Network Analysis. *IEEE Trans. Knowl. Data Eng.* 29, 1 (2017), 17–37. https://doi.org/10.1109/TKDE.2016.2598561

[13] Yizhou Sun, Charu C. Aggarwal, and Jiawei Han. 2012. Relation Strength-Aware Clustering of Heterogeneous Information Networks with Incomplete Attributes. *PVLDB* 5, 5 (2012), 394–405. http://vldb.org/pvldb/vol5/p394_yizhousun_vldb2012.pdf

[14] Yizhou Sun, Rick Barber, Manish Gupta, Charu C. Aggarwal, and Jiawei Han. 2011. Co-author Relationship Prediction in Heterogeneous Bibliographic Networks. In *International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2011, Kaohsiung, Taiwan, 25-27 July 2011.* 121–128. https://doi.org/10.1109/ASONAM.2011.112

[15] Yizhou Sun, Jiawei Han, Charu C. Aggarwal, and Nitesh V. Chawla. 2012. When will it happen?: relationship prediction in heterogeneous information networks. In *Proceedings of the Fifth International Conference on Web Search and Web Data Mining, WSDM 2012, Seattle, WA, USA, February 8-12, 2012.* 663–672. https://doi.org/10.1145/2124295.2124373

[16] Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S. Yu, and Tianyi Wu. 2011. PathSim: Meta Path-Based Top-K Similarity Search in Heterogeneous Information Networks. *PVLDB* 4, 11 (2011), 992–1003. http://www.vldb.org/pvldb/vol4/p992-sun.pdf

[17] Dingqi Yang, Daqing Zhang, Zhiyong Yu, Zhiwen Yu, and Djamal Zeghlache. 2014. SESAME: Mining User Digital Footprints for Fine-Grained Preference-Aware Social Media Search. *ACM Trans. Internet Techn.* 14, 4 (2014), 28:1–28:24. https://doi.org/10.1145/2677209

[18] Dingqi Yang, Daqing Zhang, Vincent W. Zheng, and Zhiyong Yu. 2015. Modeling User Activity Preference by Leveraging User Spatial Temporal Characteristics in LBSNs. *IEEE Trans. Systems, Man, and Cybernetics: Systems* 45, 1 (2015), 129–142. https://doi.org/10.1109/TSMC.2014.2327053

[19] Mao Ye, Peifeng Yin, Wang-Chien Lee, and Dik Lun Lee. 2011. Exploiting geographical influence for collaborative point-of-interest recommendation. In *Proceeding of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2011, Beijing, China, July 25-29, 2011.* 325–334. https://doi.org/10.1145/2009916.2009962