Table 1: Performance comparison with baselines in terms of Precision@k and Recall@k for $k \in \{5, 10, 20\}$ on Gowalla and Yelp. The superscripts † and ‡ denote significant improvements compared to baselines and CATAPE-NoCat, respectively ($p < 0.05$).

Method	Yelp						Gowalla					
	P@5	P@10	P@20	R@5	R@10	R@20	P@5	P@10	P@20	R@5	R@10	R@20
USG	0.0282	0.0244	0.0197	0.0281	0.0523	0.0753	0.0502	0.0471	0.0413	0.0517	0.0568	0.0625
MGMPFM	0.0197	0.0173	0.0136	0.0211	0.0293	0.0493	0.0281	0.0215	0.0197	0.0263	0.0291	0.0319
BPRFM	0.0285	0.0221	0.0185	0.0296	0.0361	0.0599	0.0493	0.0443	0.0342	0.0497	0.0529	0.0581
RankGeoFM	0.0421	0.0362	0.0292	0.0392	0.0673	0.0838	0.0567	0.0501	0.0492	0.0591	0.0642	0.0718
HGMF	0.0532	0.0491	0.0401	0.0478	0.0702	0.0915	0.0798	0.0711	0.0683	0.0715	0.0773	0.0819
Metric Factorization	0.0593	0.0552	0.0481	0.0533	0.0782	0.0974	0.0821	0.0782	0.0717	0.0784	0.0814	0.0862
CATAPE-NoCat	$0.0641^{\dagger}$	$0.0613^{\dagger}$	$0.0568^{\dagger}$	0.0589	$0.0831^{\dagger}$	$0.1013^{\dagger}$	$0.0892^{\dagger}$	$0.0828^{\dagger}$	$0.0784^{\dagger}$	$0.0815^{\dagger}$	$0.0898^{\dagger}$	$0.0979^{\dagger}$
CATAPE	$\mathbf{0.0702}^{\dagger\ddagger}$	$\mathbf{0.0692}^{\dagger\ddagger}$	$\mathbf{0.0631}^{\dagger\ddagger}$	$\mathbf{0.0621}^{\dagger\ddagger}$	$\mathbf{0.0881}^{\dagger}$	$\mathbf{0.1121}^{\dagger\ddagger}$	$\mathbf{0.0924}^{\dagger\ddagger}$	$\mathbf{0.0894}^{\dagger\ddagger}$	$\mathbf{0.0813}^{\dagger\ddagger}$	$\mathbf{0.0872}^{\dagger}$	$\mathbf{0.0953}^{\dagger\ddagger}$	$\mathbf{0.1283}^{\dagger\ddagger}$

considering the hierarchical structure and incorporating the categorical information as one the most effective contextual signals in the hierarchical model. However, HGMF considers the dot product of users and POIs in measuring their similarity. Among the baselines, it is seen that Metric Factorization outperforms HGMF and other methods, suggesting that using Euclidean distance is a more precise measure of similarity as opposed to dot product.

As seen in Table 1, CATAPE significantly outperforms all of the baseline methods in terms of all evaluation metrics. This indicates that the check-in module is able to learn POI latent representation by modeling the context of users' visited POIs and the sequence of POIs. Furthermore, the results suggest that incorporating category information enables CATAPE to model the characteristics of the POIs more effectively. It is worth noting that our proposed POI embedding model can be pre-trained on a large dataset of check-ins to be used in various POI recommendation models. Also, it is seen that CATAPE-NoCat is able to outperform all the baselines significantly, indicating that learning POI embeddings only based on check-in information is able to capture complex sequential relations between POIs.

Impact of category module. To show the effect of category information on the performance, we compare the performance of CATAPE with CATAPE-NoCat. The results in Table 1 show that the performance of model significantly drops when we remove category information, indicating that the category information enables the model to capture the similarities between POIs more accurately. As mentioned in the literature [1], category information is crucial for capturing users regular habits. For instance, a user may stop by a drive-thru coffee shop every morning, just before going to their workplace. Despite the performance drop, it is seen that CATAPE-NoCat is able to outperform all the baseline methods significantly in terms of all evaluation metrics. More specifically, it is seen that CATAPE-NoCat outperforms Metric Factorization by 18% in terms of Pre@20 on Yelp and 13.57% in terms of Rec@20 on Gowalla. Finally, note that a similar experiment, removing the check-in module would not be possible because in the category module latent vectors of POIs, computed by the check-in module, are required.

5 CONCLUSIONS

In this paper, we introduced a novel POI embedding model and demonstrated the importance of characteristics of POIs in POI embedding. Our model captures the sequential influence of POIs from check-in sequence of users, as well as, characteristics of POIs using the category information. The experimental results showed that our model contributes to improving POI recommendation performance.

REFERENCES

[1] Mohammad Aliannejadi and Fabio Crestani. 2018. Personalized Context-Aware Point of Interest Recommendation. *ACM Trans. Inf. Syst.* 36, 4 (2018), 45:1–45:28.

[2] Mohammad Aliannejadi, Dimitrios Rafailidis, and Fabio Crestani. 2018. A Collaborative Ranking Model with Multiple Location-based Similarities for Venue Suggestion. In *ICTIR*. ACM, 19–26.

[3] Buru Chang, Yonggyu Park, Donghyeon Park, Seongsoon Kim, and Jaewoo Kang. 2018. Content-Aware Hierarchical Point-of-Interest Embedding Model for Successive POI Recommendation.. In *IJCAI*. 3301–3307.

[4] 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 *AAAI*, Vol. 12. 17–23.

[5] Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. 2010. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research* 11, Feb (2010), 625–660.

[6] Shanshan Feng, Gao Cong, Bo An, and Yeow Meng Chee. 2017. POI2Vec: Geographical Latent Representation for Predicting Future Visitors.. In *AAAI*. 102–108.

[7] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 173–182.

[8] Xutao Li, Gao Cong, Xiao-Li Li, Tuan-Anh Nguyen Pham, and Shonali Krishnaswamy. 2015. Rank-geofm: A ranking based geographical factorization method for point of interest recommendation. In *SIGIR*. 433–442.

[9] Xin Liu, Yong Liu, and Xiaoli Li. 2016. Exploring the Context of Locations for Personalized Location Recommendations.. In *IJCAI*. 1188–1194.

[10] Yiding Liu, Tuan-Anh Nguyen Pham, Gao Cong, and Quan Yuan. 2017. An experimental evaluation of point-of-interest recommendation in location-based social networks. *VLDB* 10 (2017), 1010–1021.

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

[12] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *UAI*. 452–461.

[13] Mao Ye, Peifeng Yin, Wang-Chien Lee, and Dik-Lun Lee. 2011. Exploiting geographical influence for collaborative point-of-interest recommendation. In *SIGIR*. 325–334.

[14] Quan Yuan, Gao Cong, Zongyang Ma, Aixin Sun, and Nadia Magnenat Thalmann. 2013. Time-aware point-of-interest recommendation. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. ACM, 363–372.

[15] Jia-Dong Zhang and Chi-Yin Chow. 2015. GeoSoCa: Exploiting Geographical, Social and Categorical Correlations for Point-of-Interest Recommendations. In *SIGIR*. 443–452.

[16] Shuai Zhang, Lina Yao, Yi Tay, Xiwei Xu, Xiang Zhang, and Liming Zhu. 2018. Metric Factorization: Recommendation beyond Matrix Factorization. *arXiv preprint arXiv:1802.04606* (2018).

[17] Pengpeng Zhao, Xiefeng Xu, Yanchi Liu, Ziting Zhou, Kai Zheng, Victor S Sheng, and Hui Xiong. 2017. Exploiting hierarchical structures for POI recommendation. In *2017 IEEE International Conference on Data Mining (ICDM)*. IEEE, 655–664.

[18] Shenglin Zhao, Tong Zhao, Irwin King, and Michael R Lyu. 2017. Geo-teaser: Geo-temporal sequential embedding rank for point-of-interest recommendation. In *WWW*. 153–162.