Cold-start Management in POI Recommendation via Reinforcement Learning and Spatial Proximity Exploration

Discussion Paper

David Massimo, Francesco Ricci

Free University of Bozen-Bolzano, Italy

In this abstract we summarise the research presented in a paper appearing in the proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization (UMAP 2023) [1].

A major goal of any destination management organisation (DMO) is the exposure of an unbiased selection of the destination Points of Interests (POIs) to their visitors. However, with their online portals, destinations can hardly achieve that goal as they usually offer more visibility to popular POIs, compared to less popular ones. Recommender Systems (RSs) can help to generate personalised and unbiased POI recommendations, however, since they are usually trained on users' feedback (e.g., check-ins or reviews), they end up recommending well know items, to users that have a consolidated history of system interaction. Hence, RSs fail to suggest new POIs to new users, and to solve the combined new user and new item problem.

To address this limitation we have introduced a Reinforcement Learning RS, called QEXP (destination EXPloration), that is bootstrapped on logs of POI visits data to generate an initial tourist POI visits behaviour model [2]. Then the model is augmented with a second model component, derived from the knowledge of the spatial range of tourists' movements in a destination. That finally generates recommendations for relevant and new items, even to new users. QEXP can recommend POIs not yet evaluated or visited (in the training data log) by any user, provided that they possess the features of POIs experienced by tourists in the past, and they are located in the proximity of the target tourist position. Moreover, QEXP generates recommendations that replicate the observed users' POI visits behaviour also considering the inclination of tourists to explore and discover new POIs in their spatial proximity. In practice, QEXP attempts to predict the movements tourists are likely to perform but that can also lead to a satisfactory visit experience.

We have shown that QEXP can successfully accomplish the goal of a POI RS: it suggests novel and relevant items to new users. We have focused on a common application setting: a new user opens a destination portal or a mobile app, and, after having indicated the a starting point of the visit, e.g., an initial POI, requests recommendations for other POIs to visit next. We simulate this scenario and evaluate the ability of QEXP to tackle the new user and new item problems. To accomplish this goal, QEXP performance is compared with that of four baselines: "most popular"

POIs [3], a choice prediction Markov Model [4], QBASE [2], and LGLMF[5].

To test the ability of these RSs to solve the new user and new item problems, we have collected, from tourism "experts" of two target destinations (Rome and Florence), a reliable ground truth composed by novel POIs, i.e., not already present in the model training set. This data consists of a selection of new POIs that can be considered as relevant for new tourists to a destination. We have generated next POI recommendations for a uniform spatial distribution of tourists and analysed cross-sectional dimensions, such as, concentration, popularity, area and catalogue covered by the recommended POIs.

We have observed that QEXP outperforms the considered baselines in recommending relevant new items to new users. Moreover, QEXP produces a larger variety of POI recommendations, and it significantly reduces the recommendations of popular, possibly known items, while promoting the discovery of the destination full area. We argue that QEXP can support the construction of more effective POI RSs and that can be employed in "critical" destinations that are facing the issues brought by unregulated tourism policies, such as those connected with overtourism. QEXP, in combination with a well-designed HCI strategy, can help these destinations in mitigating overtourism by helping to promote uncovered areas of the destination while considering the general needs and wants of tourists.

References

[1] D. Massimo, F. Ricci, Combining reinforcement learning and spatial proximity exploration for new user and new poi recommendations, in: Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization, UMAP '23, ACM, 2023, pp. 422–431. doi:10.1145/3565472.3592966.

[2] D. Massimo, F. Ricci, Popularity, novelty and relevance in point of interest recommendation: an experimental analysis, J. Inf. Technol. Tour. 23 (2021) 473–508. doi:10.1007/s40558-021-00214-5.

[3] P. Sánchez, A. Bellogín, Applying reranking strategies to route recommendation using sequence-aware evaluation, User Modeling and User-Adapted Interaction 30 (2020) 659–725. URL: https://doi.org/10.1007/s11257-020-09258-4. doi:10.1007/s11257-020-09258-4.

[4] W. Zheng, X. Huang, Y. Li, Understanding the tourist mobility using gps: Where is the next place?, Tourism Management 59 (2017) 267–280. URL: https://www.sciencedirect.com/science/article/pii/S0261517716301509. doi:https://doi.org/10.1016/j.tourman.2016.08.009.

[5] H. A. Rahmani, M. Aliannejadi, S. Ahmadian, M. Baratchi, M. Afsharchi, F. Crestani, Lglmf: Local geographical based logistic matrix factorization model for poi recommendation, in: F. L. Wang, H. Xie, W. Lam, A. Sun, L.-W. Ku, T. Hao, W. Chen, T.-L. Wong, X. Tao (Eds.), Information Retrieval Technology, Springer International Publishing, Cham, 2020, pp. 66–78.