

Consumer Preference Exploration with Unexpected Recommender System

Summary (497 words)

Classical personalization models, such as recommender systems, often lead to the problem of over-specialization, filter bubbles and user boredom, as they focus primarily on identifying the most relevant products for consumers, without addressing their desire for product variety and serendipity. To deal with these issues, researchers have proposed multiple methods within the exploration-exploitation framework, where the exploration process is usually conducted randomly. However, personalized consumer preferences towards exploration, such as their desire for product variety, are not properly captured via random exploration process, which limits recommendation performance.

To address these concerns, we propose to adopt the state-of-the-art deep-learning techniques to model the unexpectedness objective in recommendations in [3] and [5], and manage to significantly improve consumer serendipitous experiences without sacrificing any accuracy metrics. In addition, we design a hybrid utility function to optimize the relevance and unexpectedness objectives simultaneously. To better calibrate the unexpectedness levels for each consumer, we combine our model with sequential modeling and the self-attention mechanism to capture personalized and session-based consumer preferences towards unexpectedness in [4], resulting in significant performance improvements in both the novelty and accuracy metrics. In particular, it could effectively identify the “variety-seekers“ in the population and provide them with more unexpected recommendations. For those consumers who prefer to stay within their own “comfort zones”, we would provide them with lower-level unexpected recommendations. This type of consumer differentiation leads to even better performance, as we demonstrate in [1].

Furthermore, we test the effectiveness of our methods in [1] and [4] not only on the offline industrial-grade data, but also through multiple large-scale online A/B tests on real customers at Alibaba watching short videos. These experiments demonstrate that, by adopting our proposed models, consumers would be significantly more likely to click on, finish watching and spend time on the recommended videos, leading to potential additional annual revenues of \$30,000,000 for the company. Due to the tangible economic impact demonstrated in our research, Alibaba has deployed our proposed unexpectedness and variety-seeking models in production to serve all the consumers on its video streaming platform.

Finally, to alleviate human feedback-loop biases when exploring consumer preferences, we propose a novel deep reinforcement learning-based model in [2] that optimizes the degree and direction of consumer preference exploration by modeling consumer trajectories in the latent space. As shown in the simulation, offline and online experiments, the proposed method could improve various business performance metrics by over 22%, while effectively exploring consumers’ interests and reducing human-feedback loop biases. These findings should help the industry to better balance between exploration vs. exploitation strategies, thus enriching the consumer preference exploration process.

In summary, in this project we introduced several methods listed in Figure 1 that effectively explore consumer preferences for serendipitous recommendations and demonstrate that these proposed methods address consumers' desire for product variety through a series of offline and online controlled experiments. Despite being recent, this project has already generated considerable impact in both the academia and the industry, and has been implemented at Alibaba and other platforms.




Figure 1: Summary of Research Publications in Designing Unexpected Recommender System

References

- [1] Pan Li, Alexander Tuzhilin, “When Variety-Seeking Meets Unexpectedness: Incorporating Variety-Seeking Behavior into Design of Unexpected Recommender Systems”, Conditional Accept at *Information System Research (ISR)*
- [2] Pan Li, Alexander Tuzhilin, “Exploring and Exploiting Consumer Preferences through Deep Reinforcement Learning and Latent Trajectory Modeling in Recommender Systems”, Second-Round Major Revision at *Management Information System Quarterly (MISQ)*
- [3] Pan Li, Alexander Tuzhilin, “Latent Unexpected Recommendations”, *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(6), pp.1-25 (2020)
- [4] Pan Li, Maofei Que, Zhichao Jiang, Yao Hu, Alexander Tuzhilin, “PURS: Personalized Unexpected Recommender System for Improving User Satisfaction”, *Proceedings of the 14th ACM Conference on Recommender System (RecSys 2020)* Full Paper with Oral Presentation; Acceptance Rate: 18%
- [5] Pan Li, Alexander Tuzhilin, “Latent Modeling of Unexpectedness for Recommendations”, *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys 2019)* Late-Breaking Result Track Paper with Poster Presentation; Acceptance Rate: 31%