

Collaborative Bandits for News Recommendation

Supervisor: Panayiotis Danassis

June 12, 2017

1 Project Overview & Goal Description

Recommender systems have become an essential part of modern on-line businesses. Due to the vast amount of information from million of sources around the world, a key challenge of today's websites is to engage users by recommending new items according to user preferences (context). This project focuses on the problem of news recommendation, aiming to improve the quality of the recommended articles and increase the click through rate (CTR).

Contextual multi-armed bandits [1] [2] have been shown to perform well in these scenarios because of their ability to balance exploration and exploitation when making decisions under uncertainty. While effective, standard contextual bandit algorithms do not take into account collaborative effects that arise from the interaction of the users with the items. In this project, we will try to address this limitation by following the work of Li et al. [3] where the authors dynamically group both the users, based on the items under consideration, and the items, based on the similarity of the clusterings induced over the users. We will test the developed model using a news recommendation dataset (e.g. Yahoo! Front Page dataset, or similar) and compare it with alternative methods.

2 Required Skills

Good programming skills are required.

Being passionate about the topic and good English skills are a must.

References

- [1] Wikipedia, "Multi-armed bandit — wikipedia, the free encyclopedia," 2017, [Online; accessed 2-June-2017]. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Multi-armed_bandit&oldid=779388270
- [2] L. Zhou, "A survey on contextual multi-armed bandits," *CoRR*, *abs/1508.03326*, 2015.
- [3] S. Li, A. Karatzoglou, and C. Gentile, "Collaborative filtering bandits," in *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. ACM, 2016, pp. 539–548.