Interpolating Item and User Fairness in Multi-Sided Recommendations

Qinyi Chen[1], Jason Cheuk Nam Liang[1], Negin Golrezaei[1], Djallel Bouneffouf[2]

[1]Massachusetts Institute of Technology, [2]IBM Research

Online recommendation systems, from Amazon's product suggestions to Netflix's movie recommendations, have become fundamental in bolstering user engagement and driving revenue on digital platforms. Nevertheless, concerns regarding fairness in algorithmic recommendations arise as online platforms become increasingly integral to social and economic activities. Regulatory efforts, such as the Digital Markets Act and the Algorithmic Accountability Act, have also started to call for fairness in these domains.

Most online recommendations are executed by multi-sided platforms that recommend items (services, products, contents) to users (consumers of the services or goods). Central to them are two stakeholder groups: (1) the **items** (e.g., Amazon products, Netflix movies, LinkedIn job postings) and (2) the **users**, who purchase or engage with these services. The **platform** itself, while moderating the exchange between items and users, may have its unique goals, making it another independent stakeholder. Maintaining fairness on such multi-sided platforms, however, can be particularly challenging due to the intricate trade-offs between the interests of multiple stakeholders, who each have distinct objectives that do not necessarily align.

In face of the challenges, we aim to answer the following two questions: (1) *What constitutes a fair recommendation within a multi-sided platform?* and (2) *How would a platform implement a fair recommendation in a practical setting?* In our work, we address these two questions via the following contributions.

A novel fair recommendation framework. We formulate a fair constrained optimization problem, called Problem (FAIR), that takes a multi-sided point of view. Problem (FAIR) allows the platform to optimize its revenue while concurrently interpolating item and user fairness, resulting in an appropriate middleground for all stakeholders. The framework is flexible in the following aspects:

- Our framework provides the platform with quantifiable parameters to decide the extent of trade-offs it is prepared to accept to achieve fairness for other stakeholders, also known as the "price of fairness".
- Our framework allows the platform to specify any desired outcomes and any fairness notions it deems appropriate for the items/users, rather than adhering to a single predefined outcome or notion.
- Given that our framework is structured as a constrained optimization problem, it is versatile enough to encompass other operational considerations or include additional stakeholder groups.

A Fair Online Recommendation algorithm for Multi-sided platforms (`FORM`). We present `FORM`, a low-regret algorithm tailored for fair online recommendations where user data is not known in advance and learning co-occurs with fair recommendation. `FORM` learns the solution to Problem (FAIR) under data uncertainty and partial feedback by employing (i) relaxation of the fairness constriants and (ii) randomized exploration with perturbations, which balances learning and fairness. Our theoretical analysis demonstrates that our algorithm well approximates the optimal solution to Problem (FAIR), and is guaranteed to ensure the ideal outcome for all stakeholders (platform, item, user) even amidst data uncertainty. Our real-world case studies on an Amazon review data and the Movie Lens data again showcase the efficacy of our algorithm.

It is worth noting that the design of `FORM` contributes both methodologically and technically. Methodologically, contrary to many prior works that treats learning and enforcing fairness as two separate tasks, we highlight that data uncertainty can in fact impact fair recommendation quality while enforcing fairness can hinder effective learning. `FORM` well navigates the delicate balance between learning and fairness. From a more technical point of view, the design of `FORM` also overcomes the technical hurdle of managing *uncertainty in fairness constraints*. While most existing works on online learning with constraints require certain access to constraint feedback, our setup makes even verifying the satisfaction of item/user fairness impossible. This demands novel technical tools in the design of `FORM`.

The full version of this paper is available at https://arxiv.org/abs/2306.10050.