On the Perils of Optimizing the Measurable

Omar Besbes — Columbia Business School

Yash Kanoria — Columbia Business School

Akshit Kumar — Yale University

We study a theoretical model of repeated user consumption in which, at each interaction, users select between an outside option and the best option from a recommendation set. We account for user heterogeneity, with the majority preferring "popular" content, and a minority favoring "niche" content. The system initially lacks knowledge of individual user preferences but can learn them.

Engagement optimization may lead to recommending only popular items, a.k.a, "popularity bias"

Utility optimization requires recommending a diverse set of options a.k.a "forced exploration"

The ability to recommend multiple items enables utility optimization without directly measuring it, and without incurring substantial reduction in engagement.

Model

We assume an infinite horizon setting with discount factor $\delta \in (0,1)$ and the recsys repeatedly recommends an assortment. Two types of products: Popular (P) and Niche (N). The base user utility for the popular product is $U_P > 0$ and niche product is drawn from a distribution $U_N \sim F$ with mean zero. In particular, we assume that $\mathbb{P}(U_N = -1) = 1 - \varepsilon, \quad \mathbb{P}\left(U_N = \frac{1-\varepsilon}{\varepsilon}\right) = \varepsilon$

The user selects between an outside option and the best of the recommended options (assume only two options are recommended at a time), i.e., $C = \operatorname{argmax}_{i \in \{\emptyset, 1, 2\}} U_i + \varepsilon_i$, where ε_i is zero mean Gumbel noise.

Key metrics of interest are Engagement $= \mathbb{E}[\sum_{t=0}^{\infty} \delta^t 1(C_t \neq \emptyset)]$ and Utility $= \mathbb{E}[\sum_{t=0}^{\infty} \delta^t (U_{C_t} + \varepsilon_{C_t,t})]$

Policies

APP (Always Popular Policy): Always recommends the popular content

PEAR (Posterior-based Exploration-driven Adaptive Recos): Starts with *diverse* recommendations and adapts over time

Thm 1. APP is engagement-optimal

Fix any discount factor $\delta \in (0,1)$, there exists a small enough $\varepsilon > 0$ such that **APP** is the uniquely engagement optimal policy.

Thm 2. PEAR is (near) utility-optimal

For sufficiently large discount factor $\delta \in (0,1)$, **PEAR** is near optimal both in terms of engagement as well as utility.

Paper

O. Besbes, Y. Kanoria, A. Kumar. **The Fault in Our Recommendations: On the Perils of Optimizing the Measruable**. *Proceedings of the 18th ACM Conference on Recommender Systems.*

Research Question

1. By optimizing for measurable proxies, are recommendation systems at risk of significantly under-delivering on utility?

   (possibly) Yes!

2. How can we optimize for utility despite not measuring it?

   Doing some forced exploration by showing diverse recommendations