

Recommendation Systems with Unfairness and Feedback Loops

Lauren Patterson

BACKGROUND



Platforms serve a dominant majority and a distinct minority



Algorithms optimize their displayed products based on historical cumulative revenue.



Majority volume overwhelms the algorithm's state space and dictate recommendations



Minority items are structurally starved of exposure, regardless of their value

Main Objectives:

- Characterize self-reinforcing feedback loops and how they favor the majority population;
- Develop algorithms ensuring the minority population is served useful recommendations.

MODEL



Theoretical Results

Informal Result (K=1)



MLE and UCB converge to show the single most popular item; the one preferred by the majority population



Even UCB exploration is insufficient; group-aware algorithms are needed

Theorem: Stochastic Collapse (K=1)


Pop 1: Majority ($\alpha > 1/2$) | Pop 2: Minority
 $R_{1,t}, R_{2,t}$ = Cumulative revenues (utilities) of populations 1 and 2
 $\pi_{1,t}, \pi_{2,t}$ = Probability of showing the preferred item for populations 1 and 2

$$R_{1,t} - R_{2,t} = \Omega(T)$$
$$\pi_{1,t} \rightarrow 1 \quad \pi_{2,t} \rightarrow 0$$

EXPERIMENTAL RESULTS


K = 1

Top 6 Items Over Time (Rolling Exposure Share)



Item 0 provides the most overall utility by catering to the majority population. The learning algorithm explores, learns, and eventually realizes it only needs to show this one item.


Utility Over Time (Rolling Average)




While the majority and minority utilities are similar during initial exploration, the majority preferred item eventually dominates, resulting in the minority population not viewing their preferred item.

K = 3

Top 6 Items Over Time (Rolling Exposure Share)



Utility Over Time (Rolling Average)




Standard group blind algorithms optimize for the majority while structurally abandoning the minority

FUTURE WORK



Group-aware UCB with group-level confidence intervals



If population is observed, can different policies be used for different populations



Infer user groups from behavioral proxies and refine group-specific recommendations

