User-Creator Feature Polarization in Recommender Systems with Dual Influence

Tao Lin*, Kun Jin*, Andrew Estornell, Xiaoying Zhang, Yiling Chen, Yang Liu
Harvard Google ByteDance ByteDance Harvard UCSC

Motivation 1: **Relevancy** vs **Diversity** in Recommender Systems

Purely maximizing *recommendation relevancy* might cause undesirable outcomes:

- Filter bubble
- Polarization
- …

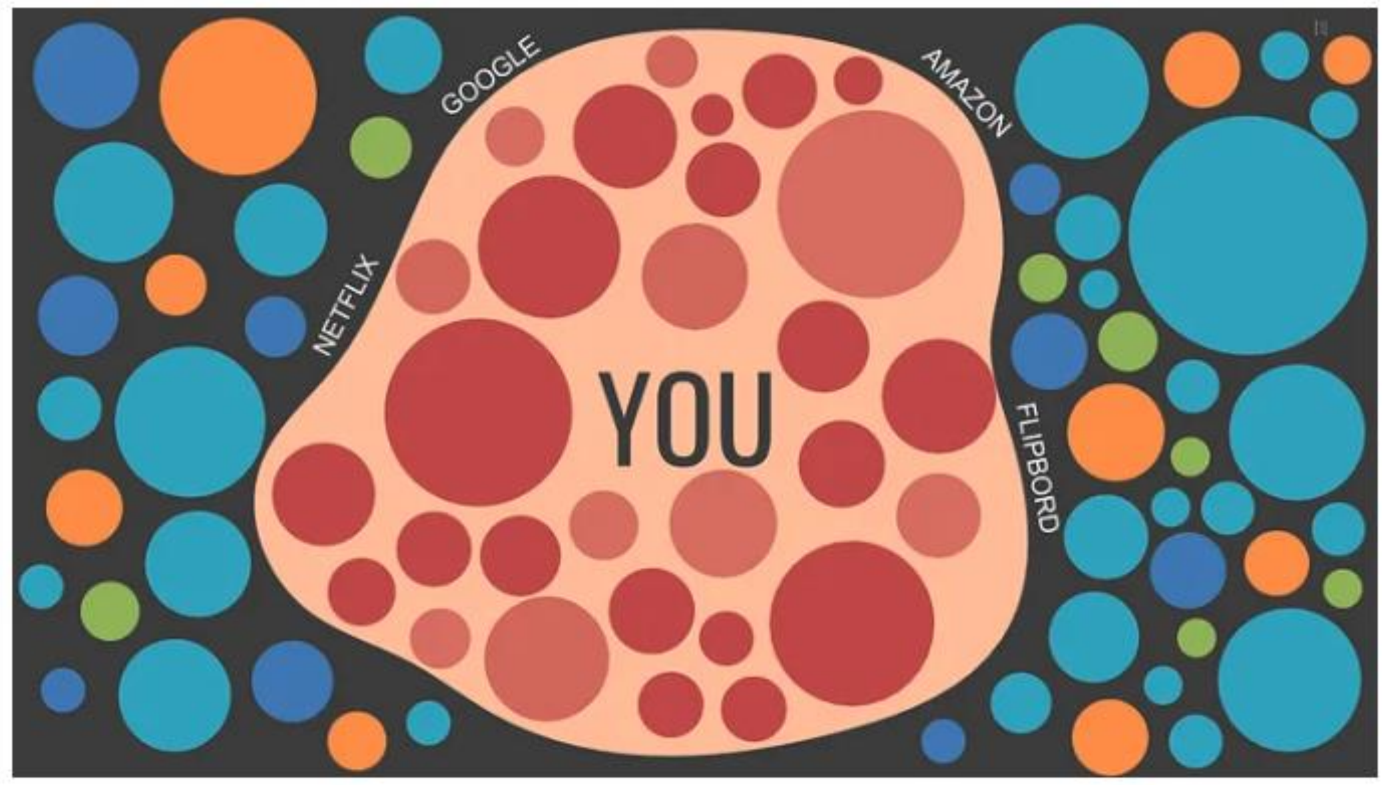

So, previous works have proposed various *diversity-improving* techniques:

- Re-ranking
- Setting diversity-boosting objectives
- …

Motivation 2: Recommender Systems **Dynamically Influence** Both Users and Creators

Although previous diversity-improving techniques are effective in a *static* system,

A real-world recommender system has *dynamic influences* on both content users and content creators.

Our Research Finding:

Due to the dynamic dual influence on users and creators,

- simple diversification techniques cannot improve the diversity of a recommender system in the long run.
- What's more, such techniques might cause polarization.

Model: User-Creator Feature Dynamics

- m users, each having a preference/feature vector $u_j^t \in \mathbb{R}^d$
- n creators, each having a feature vector $v_i^t \in \mathbb{R}^d$
- Assume that the feature vectors have unit norm: $||u_j^t|| = ||v_i^t|| = 1$
- Relevancy/similarity between creator and user is $\langle v_i^t, u_j^t \rangle = \cos\left(angle(v_i^t, u_j^t)\right)$

At each time step $t = 1, 2, \ldots$, the following events happen in order:

- **Recommendation:** For each user $j \in [m]$, a creator $i \in [n]$ is randomly sampled with probability $p_{ij}^t = p_{ij}^t(U^t, V^t)$ and recommended to that user.
  - *Example: Softmax probability function* $p_{ij}^t(U^t, V^t; \beta) = \frac{\exp\left(\beta \cdot \langle v_i^t, u_j^t \rangle\right)}{\sum_{k \in [n]} \exp\left(\beta \cdot \langle v_k^t, u_j^t \rangle\right)} > 0$
- **User Update:** The preference of each user $j \in [m]$ moves "towards" the recommended creator if the user likes the creator, otherwise moves "away":

$$u_j^{t+1} = \mathcal{P}\left(u_j^t + \eta_u \cdot \mathrm{sign}\left\langle v_{i_j^t}^t, u_j^t \right\rangle \cdot v_{i_j^t}^t\right)$$

"biased assimilation" [1]

- **Creator Update:** Each creator $i \in [n]$ is updated towards the weighted average of the matched users:

creators want to attract "fans"

$$v_i^{t+1} = \mathcal{P}\left(v_i^t + \eta_c \cdot \frac{1}{|\text{matched users}|} \sum_{j \in \text{matched users}} \mathrm{sign}\langle u_j^t, v_i^t \rangle \cdot u_j^t\right)$$

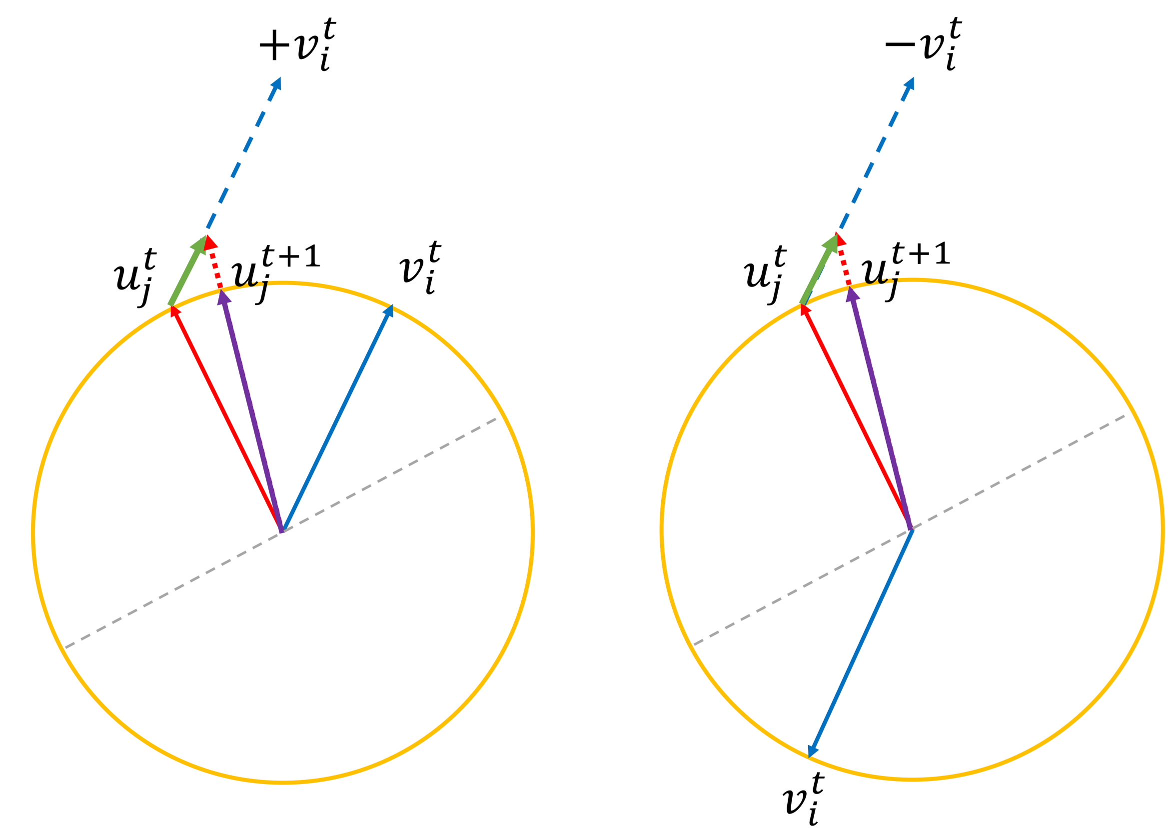

Illustration of user update

Main Theoretical Result: Diversified Recommendation Leads to Polarization

Theorem 1: For any n, m, d, and for any initial state, assuming $0 < \eta_u < \eta_c/2 < 1/4$, **as long as the recommendation probability satisfies** $p_{ij}^t > p_0 > 0$, the user-creator feature dynamics must eventually ***polarize*** (i.e., converge to two opposite directions).

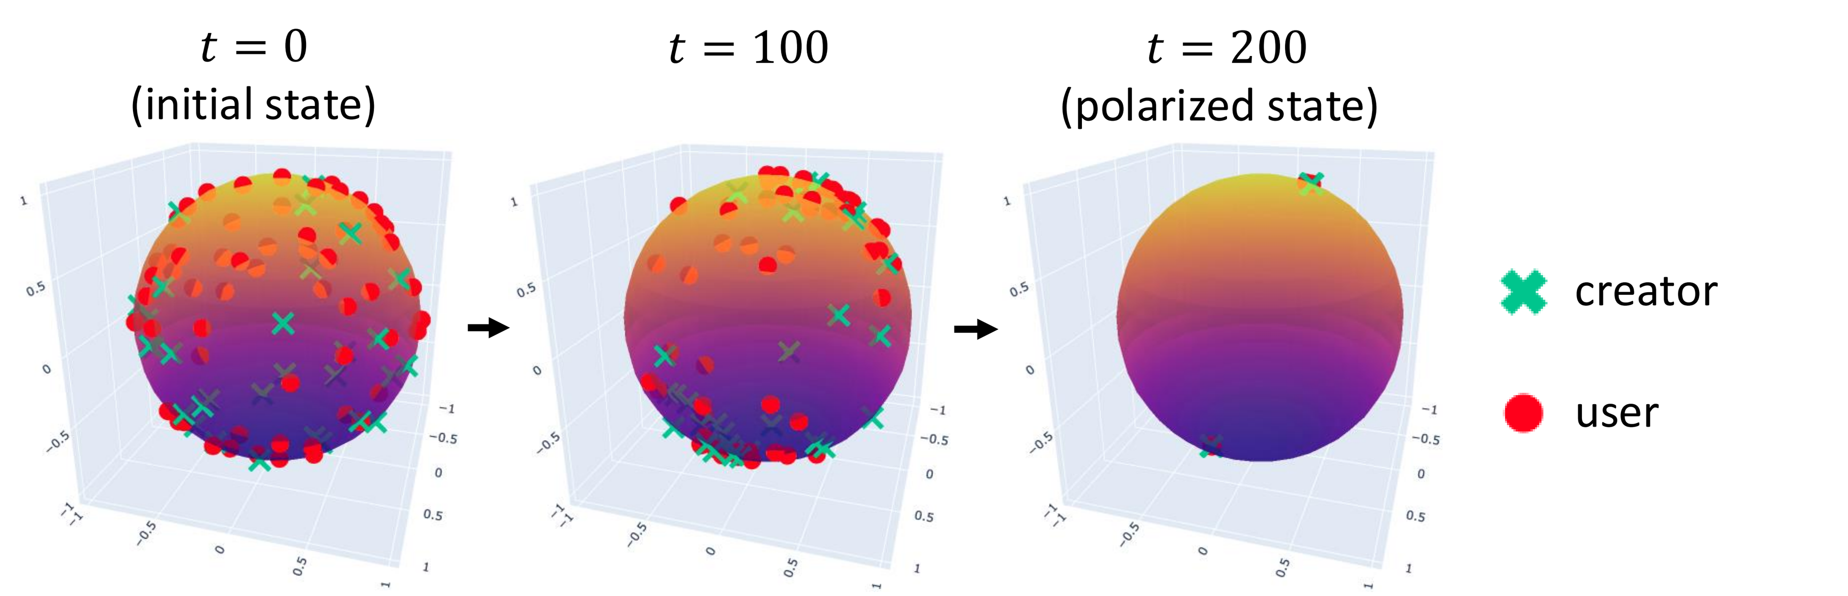

A Possible Way to Mitigate Polarization: Just Increase Recommendation Relevancy!

We tested some methods for improving *relevancy* and *efficiency*:

- **Top-k truncation**: for each user $j \in [m]$, sort the creators by the inner products $\langle u_j^t, v_{(1)}^t \rangle \geq \cdots \geq \langle u_j^t, v_{(k)}^t \rangle \geq \cdots \geq \langle u_j^t, v_{(n)}^t \rangle$. Only recommend one of the first-k creators.
- **Threshold truncation**: Only recommend creators with $\langle u_j^t, v_i^t \rangle \geq \tau$

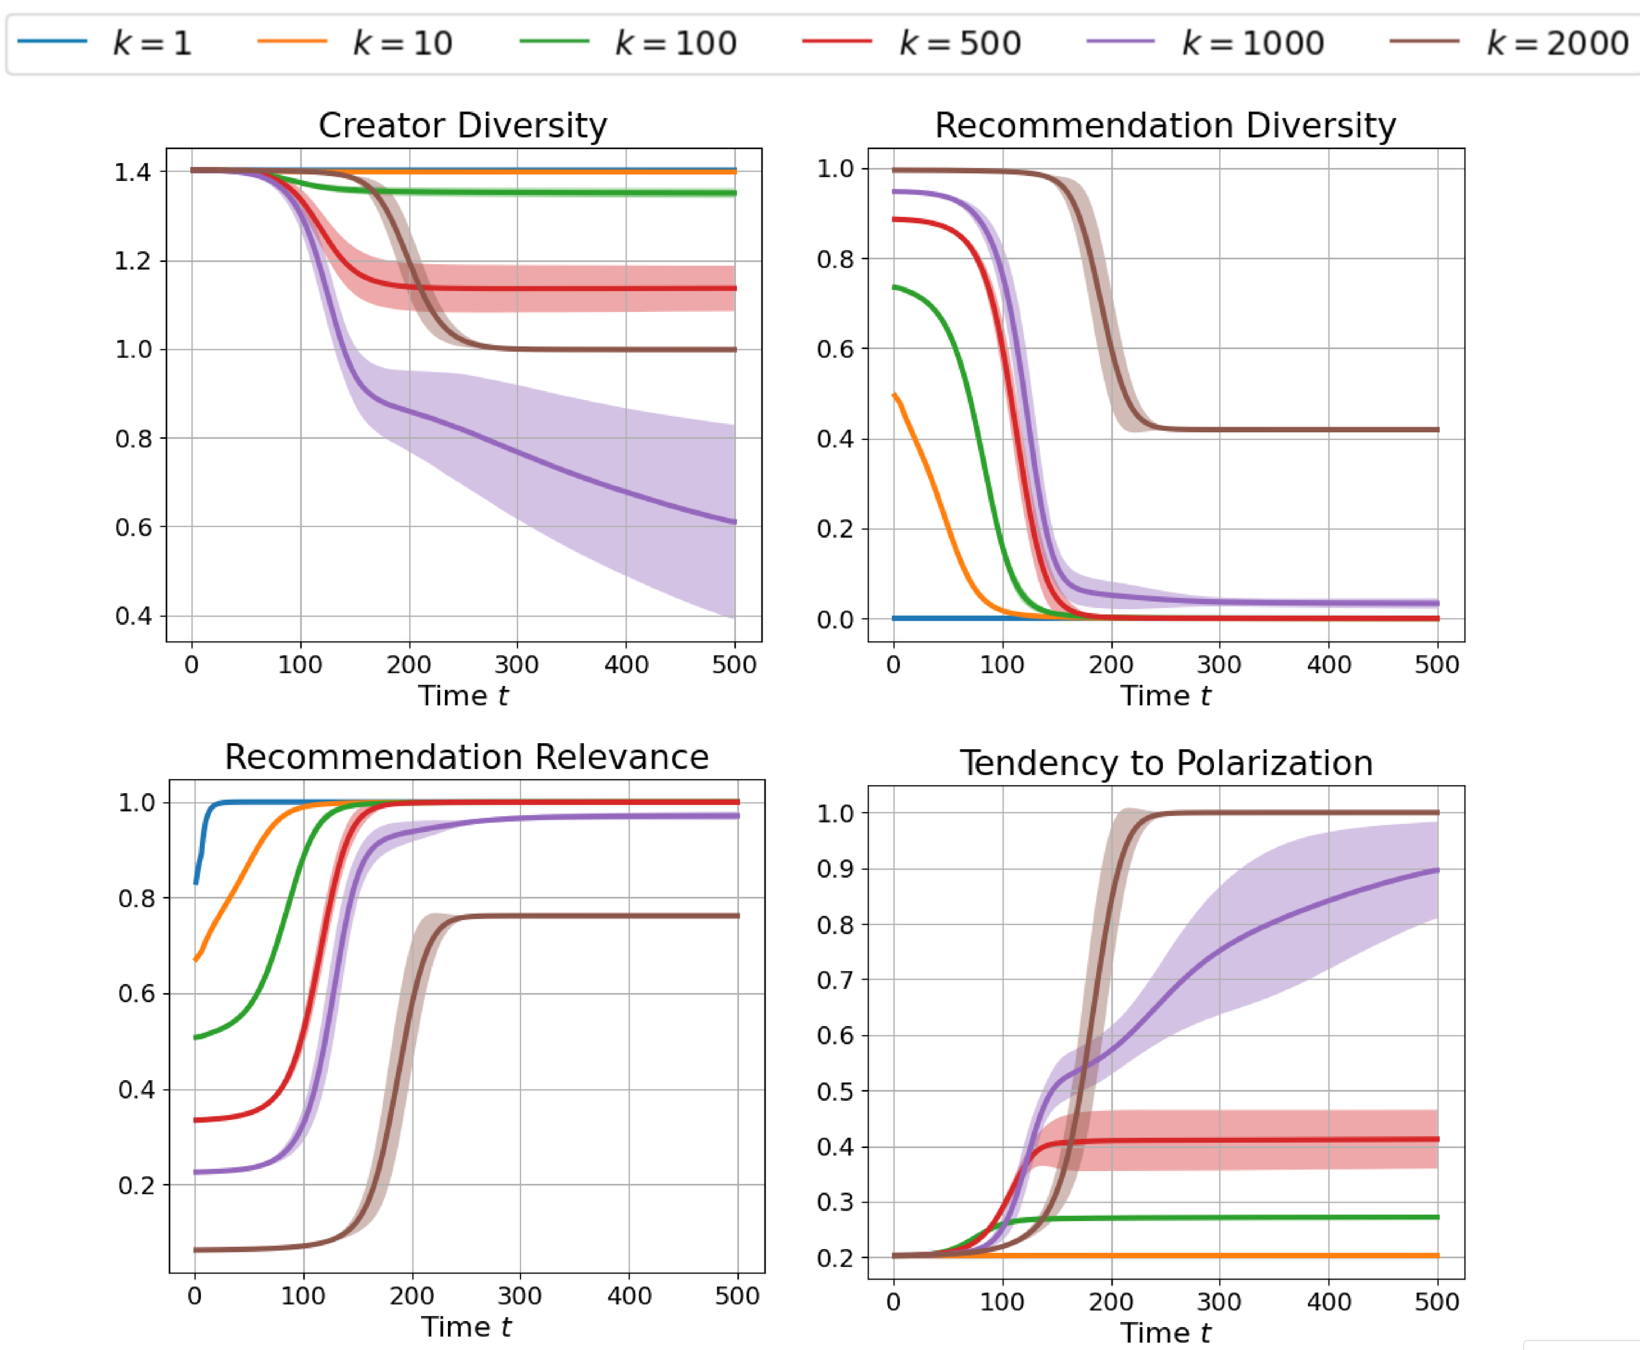

Reference:
[1] Dean & Morgenstern. Preference Dynamics Under Personalized Recommendations. EC 2023.

Our paper: