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 Havard UCSC

(Work done at ByteDance)

NeurIPS 2024

Recommender systems are everywhere

But sometimes, they are not so good

relevant but ***monotonous*** content

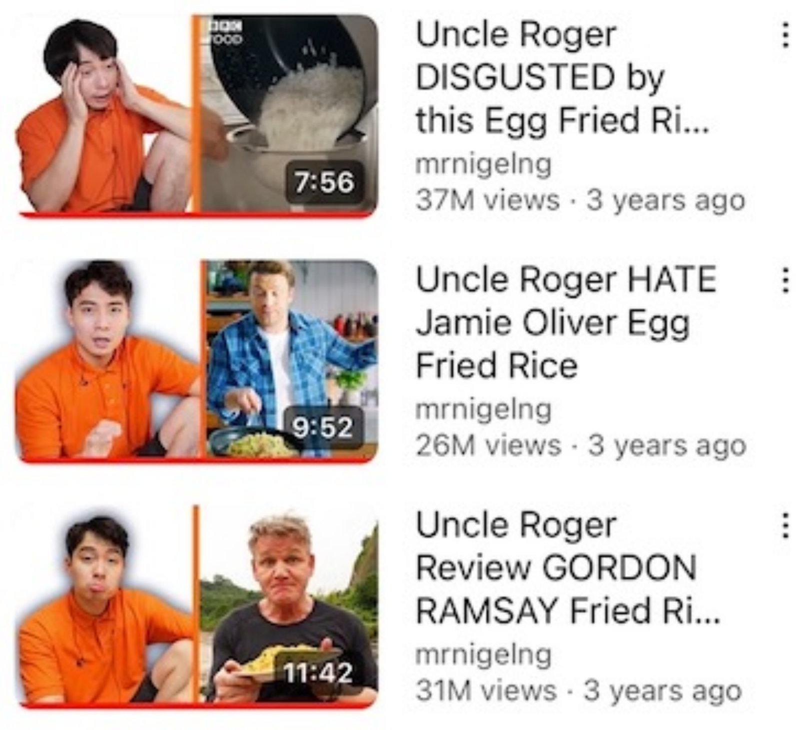

Filter Bubble

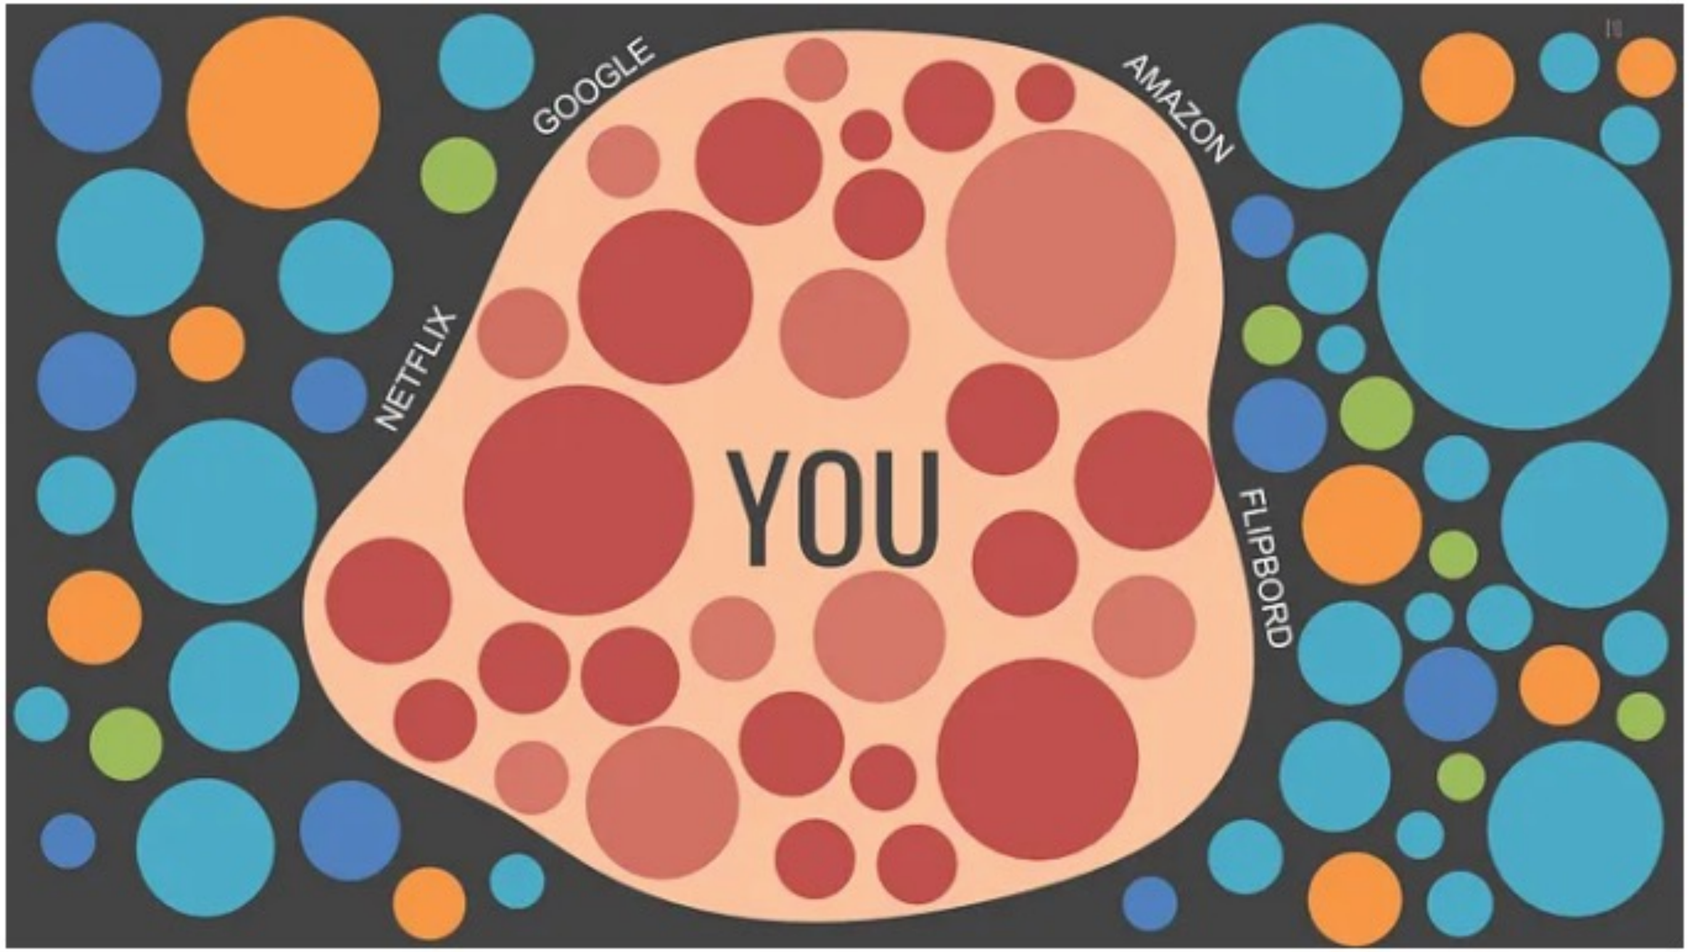

"The Internet is showing us what it thinks we want to see, but not necessarily what we need to see. Your filter bubble is your own personal, unique universe of information that you live in online." (Pariser, 2011)

Polarization

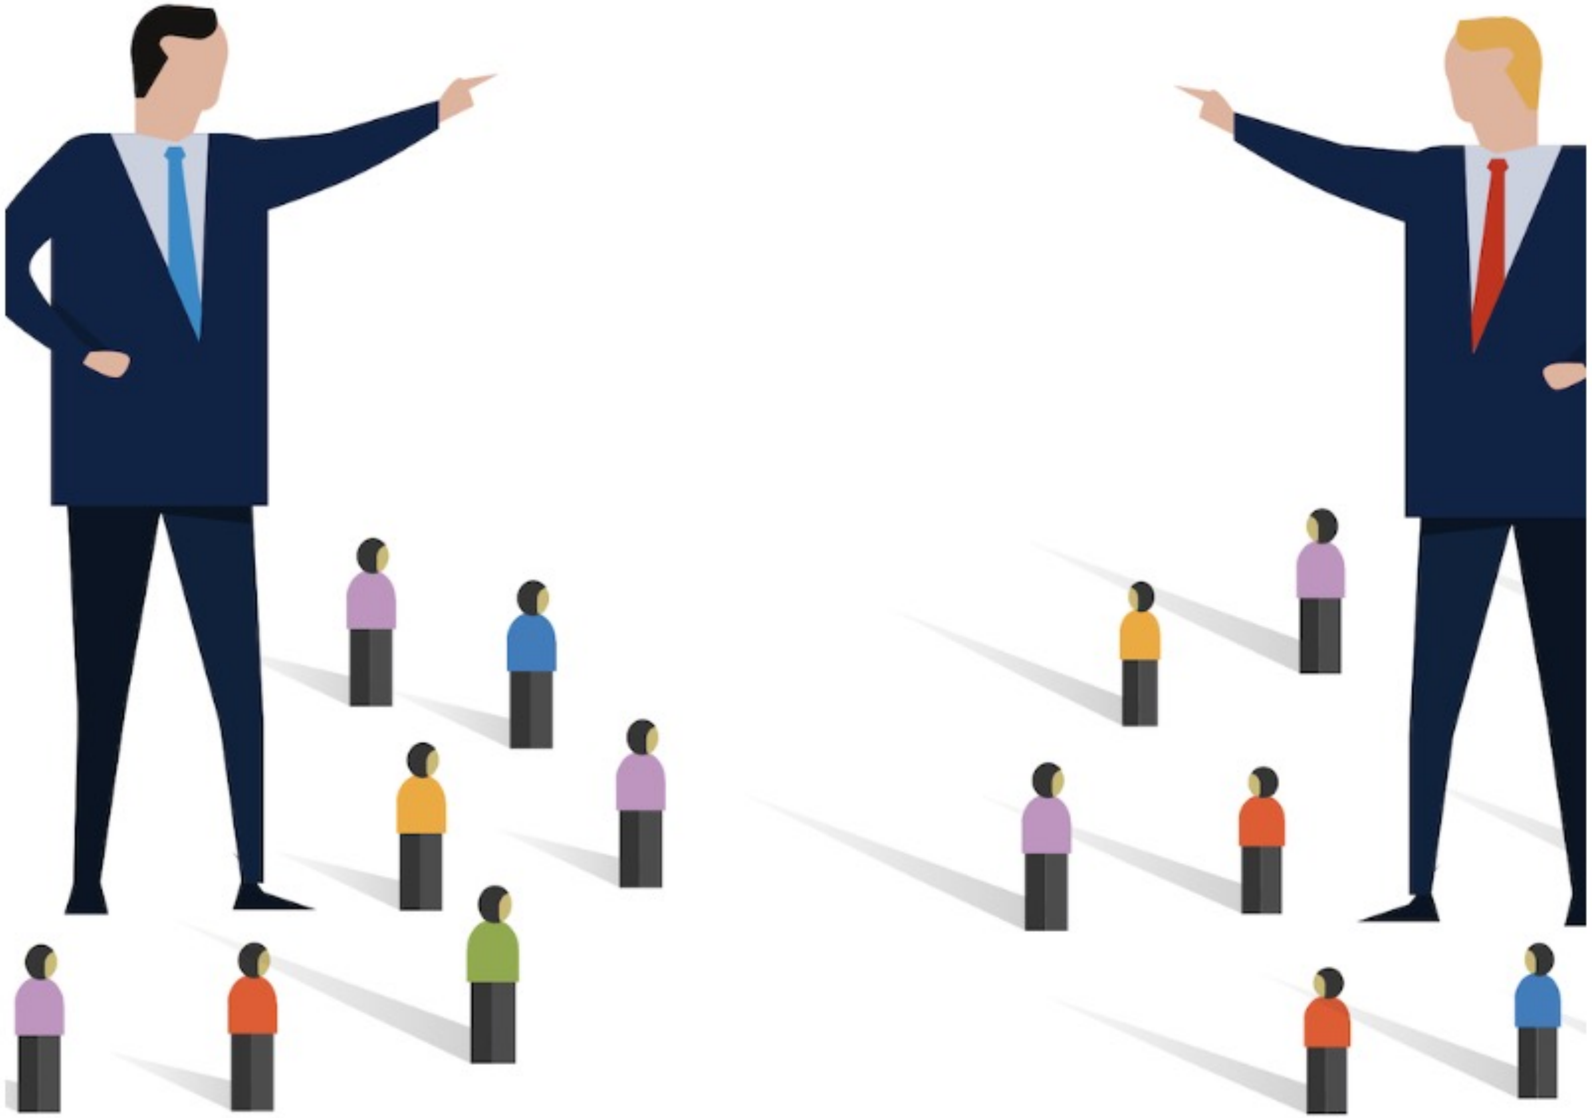

Besides recommendation relevancy, **diversity** matters!

relevant but ***monotonous*** content

less relevant but more ***diverse*** content

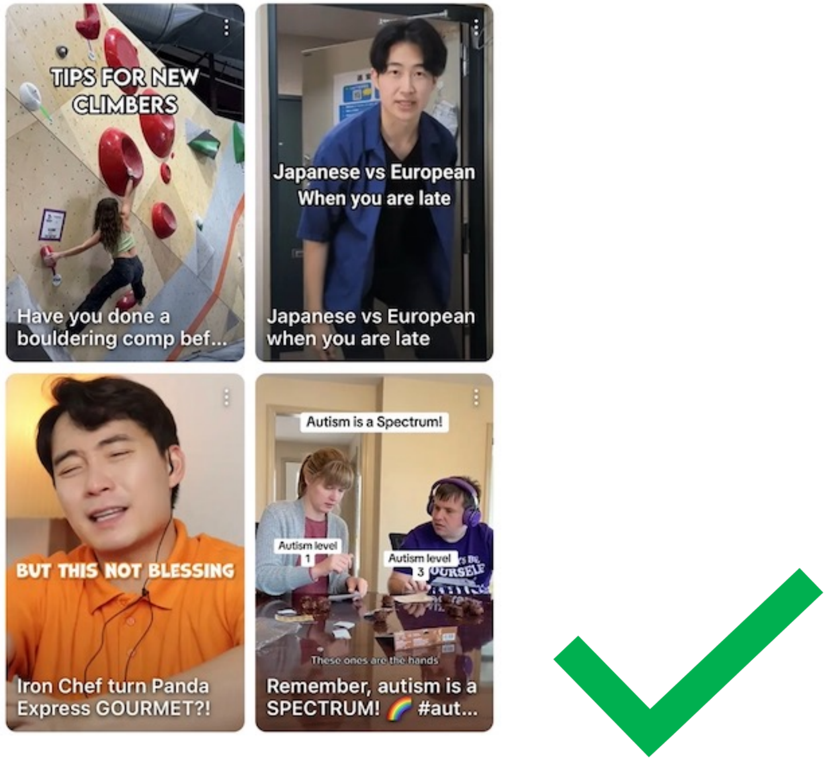

Previous methods to improve diversity:

- Re-ranking:

  [1] Carbonell & Goldstein. The use of mmr, diversity-based reranking for reordering documents and producing summaries. SIGIR 1998

  [2] Ziegler, McNee, Konstan, & Lausen. Improving recommendation lists through topic diversification. WWW 2005

  …

- Setting diversity-boosting objectives:

  [3] Zhang & Hurley. Avoiding monotony: improving the diversity of recommendation lists. RecSys 2008

  [4] Su, Yin, Chen, & Yu. Set-oriented personalized ranking for diversified top-n recommendation. RecSys 2013.

  [5] Wilhelm, Ramanathan, Bonomo, Jain, Chi, & Gillenwater. Practical diversified recommendations on YouTube with determinantal point processes. CIKM 2018.

  …

Although those methods are effective in a *static* system,
A real-world recommender system has *dynamic influences* on both content users and content creators.

Our 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.

Outline

- Introduction
- Model: User-Creator Feature Dynamics
- Main Results: Diversified Recommendation Leads to Polarization
- Ways to Mitigate Polarization

Model: User-Creator Feature Dynamics

Model: **User-Creator Feature** Dynamics

- m users, each having a preference/feature vector $u_j^t \in \mathbb{R}^d$
  - Let $U^t = [u_1^t, \dots, u_m^t]$
- n creators, each having a feature vector $v_i^t \in \mathbb{R}^d$
  - Let $V^t = [v_1^t, \dots, v_n^t]$
- Assume that the features vectors have unit Euclidean norm: $||u_j^t|| = ||v_i^t|| = 1$
- Relevancy/similarity is captured by $\langle v_i^t,\ u_j^t \rangle = \cos\left(angle\left(v_i^t,\ u_j^t\right)\right)$

Model: User-Creator Feature **Dynamics**

- At each time step $t = 1, 2, \dots,$
  - **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":

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

Illustration for User Update

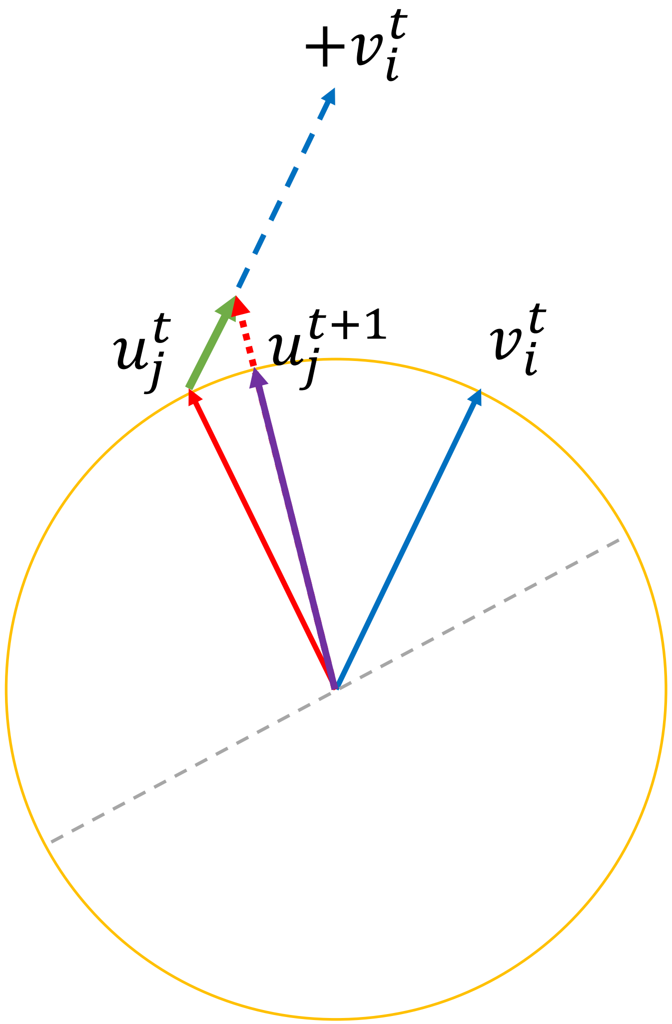

Illustration for User Update

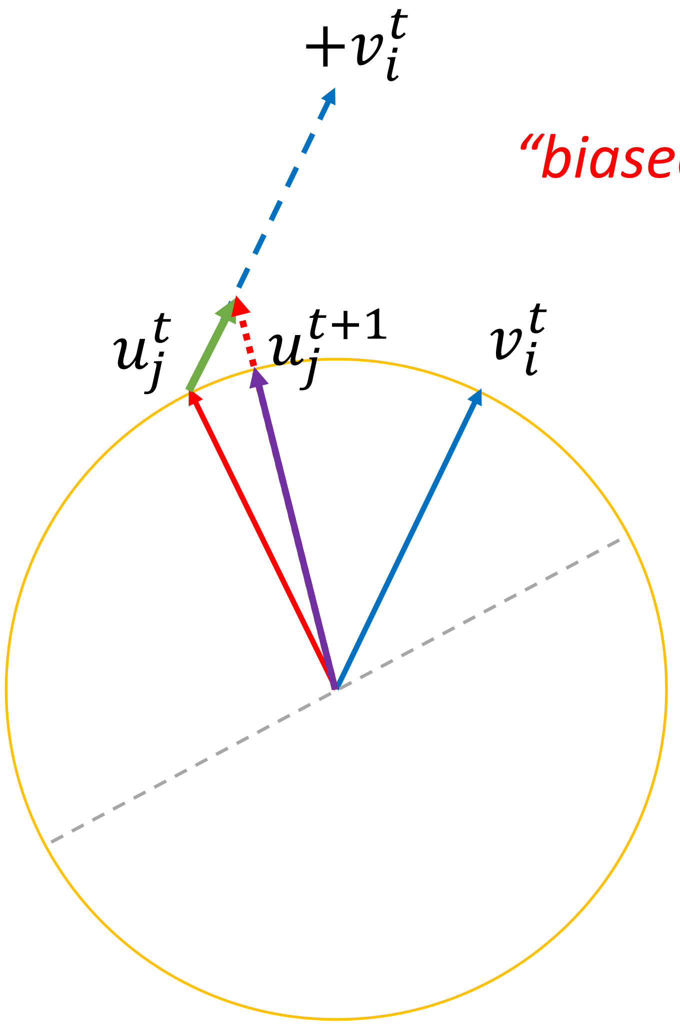

"biased assimilation" [1]

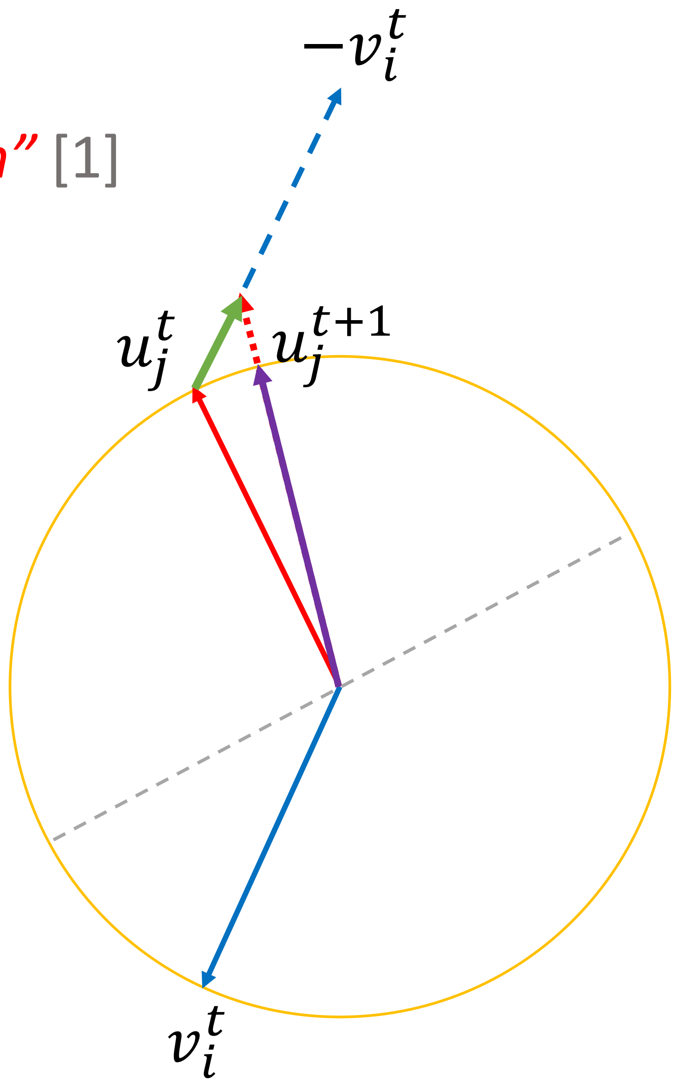

[1] Dean & Morgenstern. Preference Dynamics Under Personalized Recommendations. EC 2023.

Model: User-Creator Feature **Dynamics**

- At each time step $t = 1, 2, \dots,$
  - **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":

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

"biased assimilation"

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

creators want to attract "fans"

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

Comparison with previous works

Works	Adaptive Users?	Adaptive Creators?	Creator Reward	Dynamics or Equilibrium?	Content Adjustment Model
Ours	Yes	Yes	User engagement	Dynamics	Conditioned on previous time step; implicit cost of content adjustment
[15]	No	Yes	Exposure	Dynamics	Conditioned on previous time step; explicit cost of content adjustment
[42]	No	Yes	User engagement	Dynamics	Freely choose without cost
[35]	No	Yes	User engagement	Dynamics	Freely choose without cost
[23]	No	Yes	Exposure	Equilibrium	Freely choose with cost
[20]	No	Yes	Exposure	Equilibrium	Freely choose without cost
[7]	No	Yes	Exposure	Equilibrium	Freely choose without cost
[2]	No	Yes	User engagement	Equilibrium	Freely choose without cost
[43]	No	Yes	Designed by a welfare-maximizing platform	Dynamics	Freely choose without cost
[14]	Yes	No[1]	N/A	Dynamics	N/A
[41]	Yes	No[1]	N/A	Dynamics	N/A
[3]	Adaptive and adversarial	No[1]	N/A	Dynamics	N/A

[1]: These works study the design of recommendation algorithms for the platform with a fixed set of content, without explicitly modeling the content creators.

Outline

- Introduction
- Model: User-Creator Feature Dynamics
- Main Results: Diversified Recommendation Leads to Polarization
- Ways to Mitigate Polarization

Simulation Results: Initial State

$\eta_u = 0.1$
$\eta_c = 0.1$
Softmax with $\beta = 1$

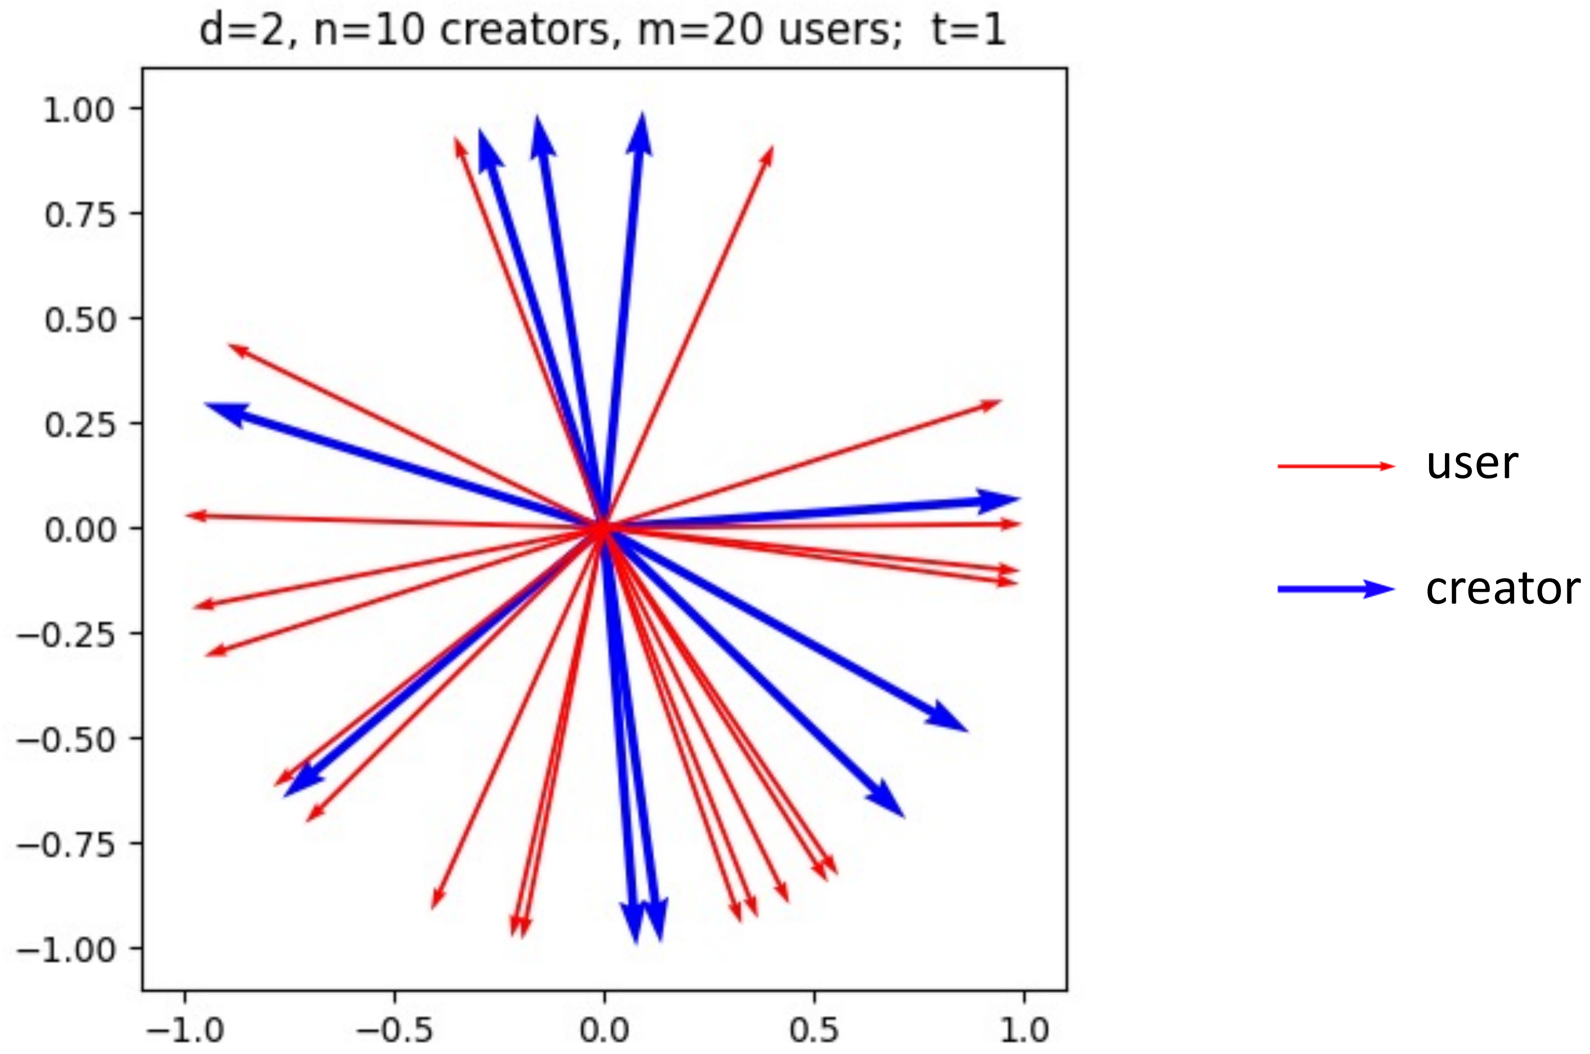

Simulation Results:

$\eta_u = 0.1$
$\eta_c = 0.1$
Softmax with $\beta = 1$

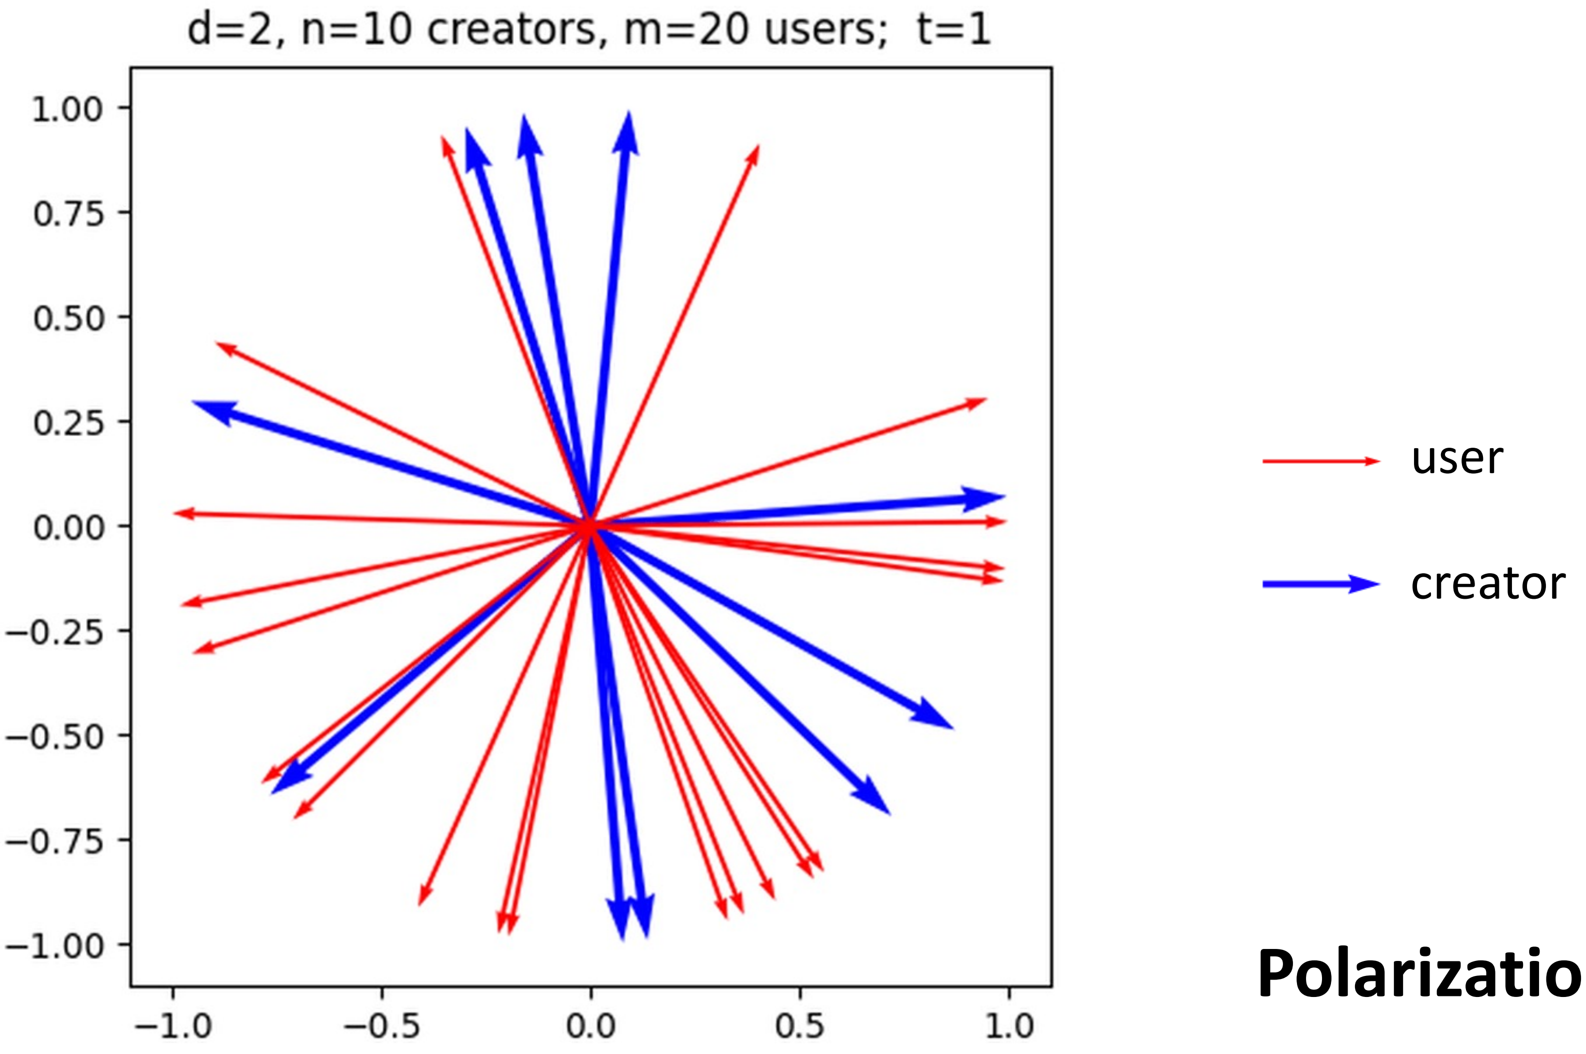

Polarization!

Simulation results for $d = 3$

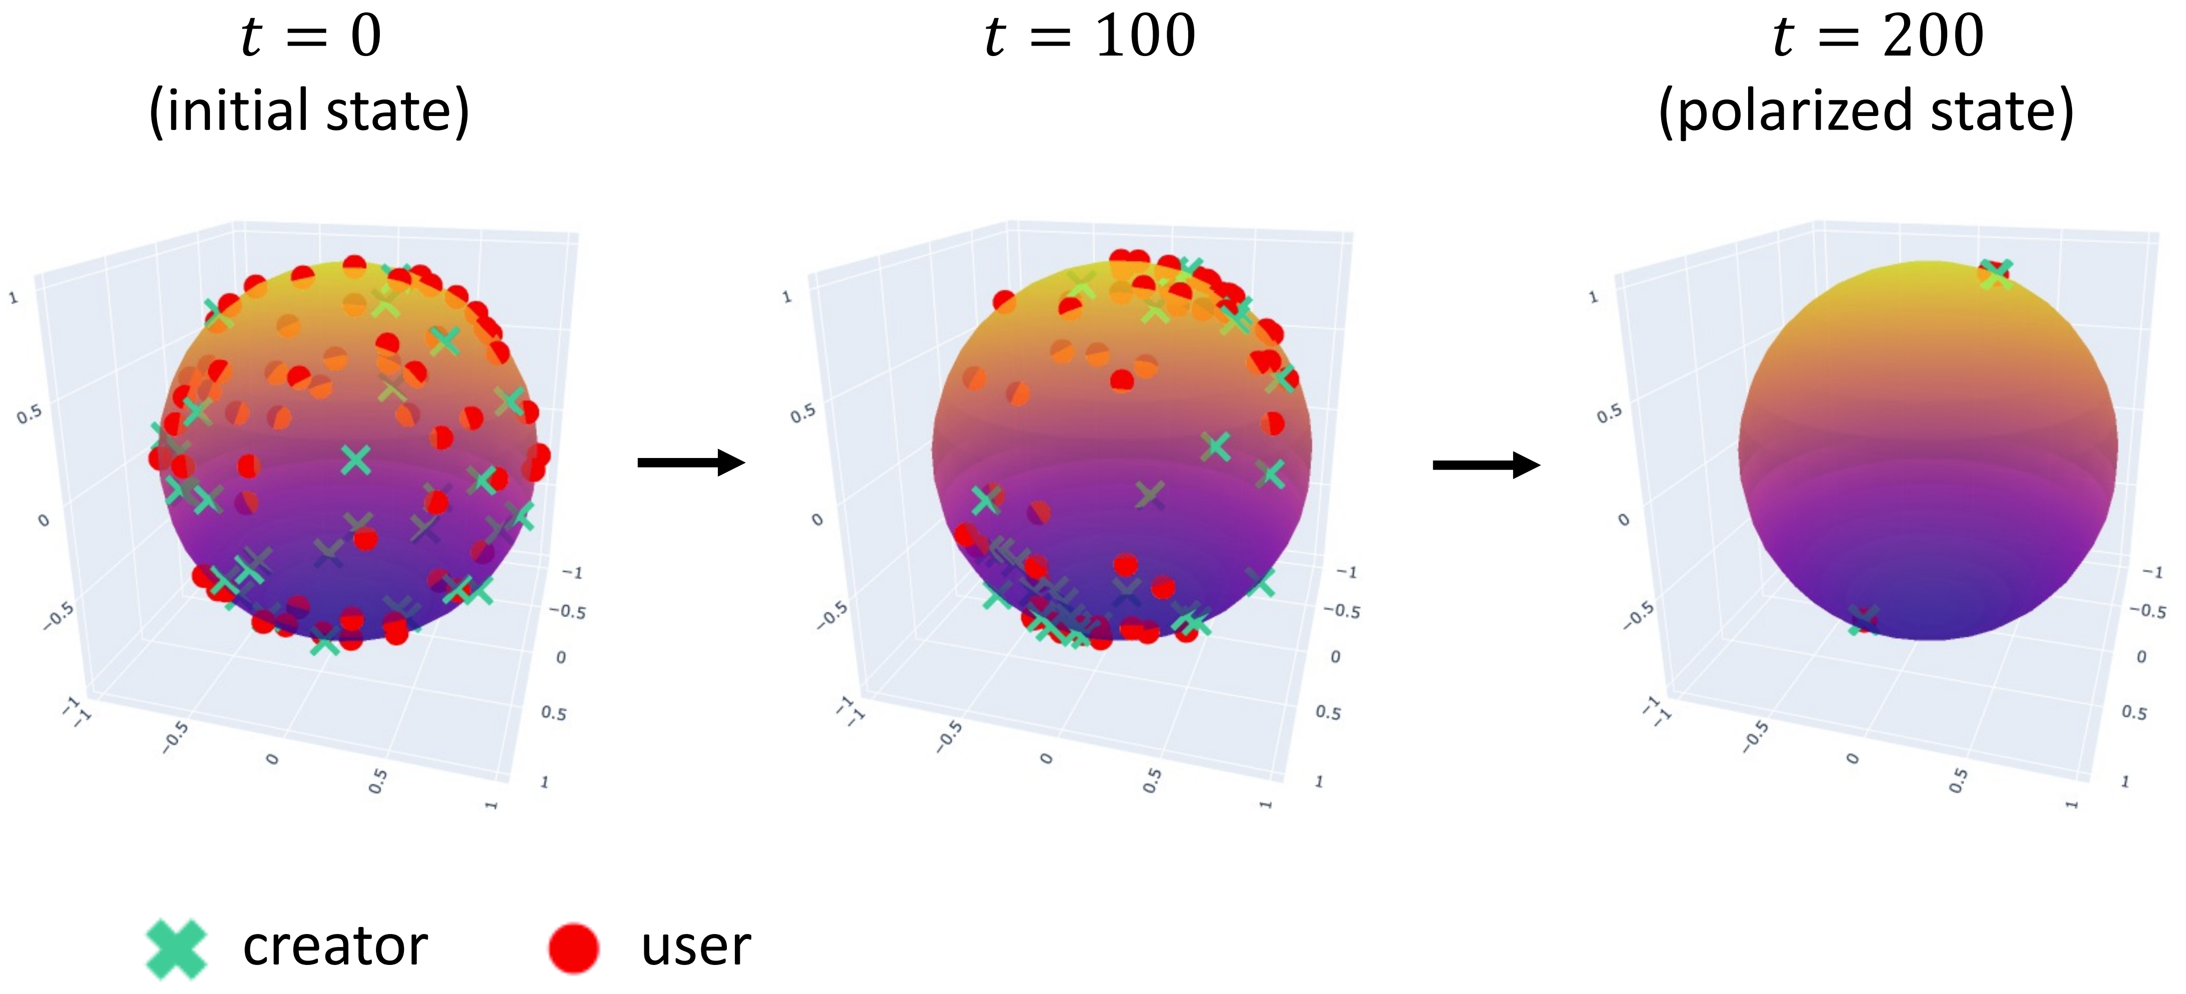

Main Theoretical Result:

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).

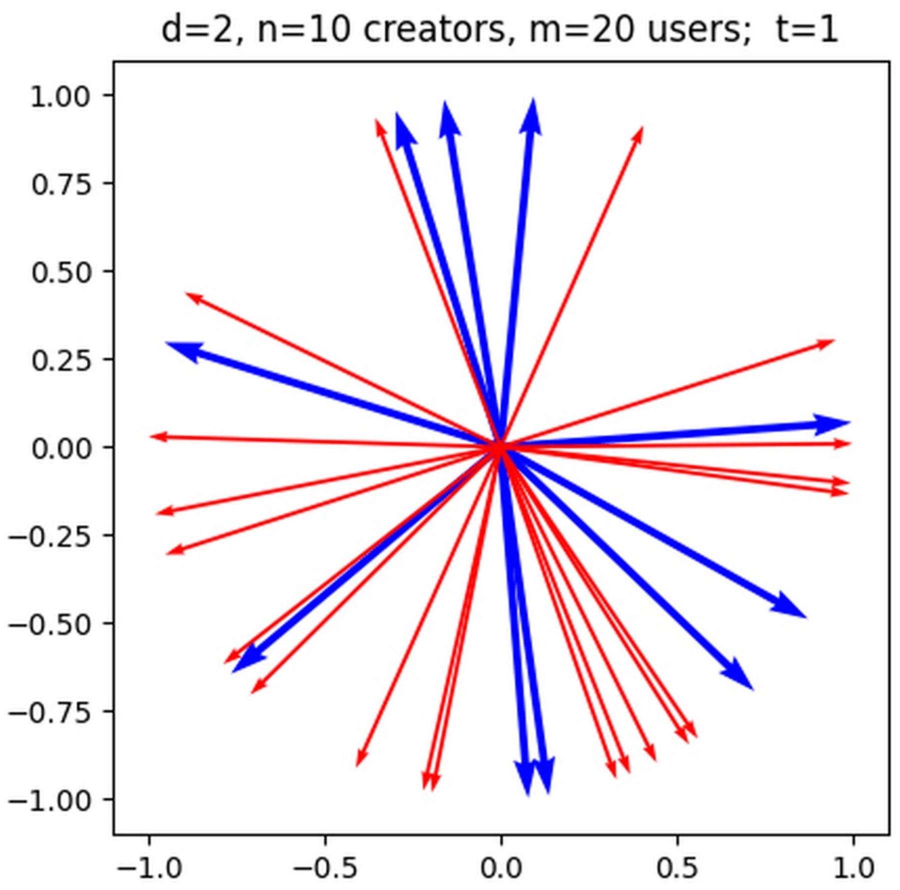

Main Theoretical Result:

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).

Implication: simple diversification techniques cannot prevent polarization in recommender systems with dual influence!

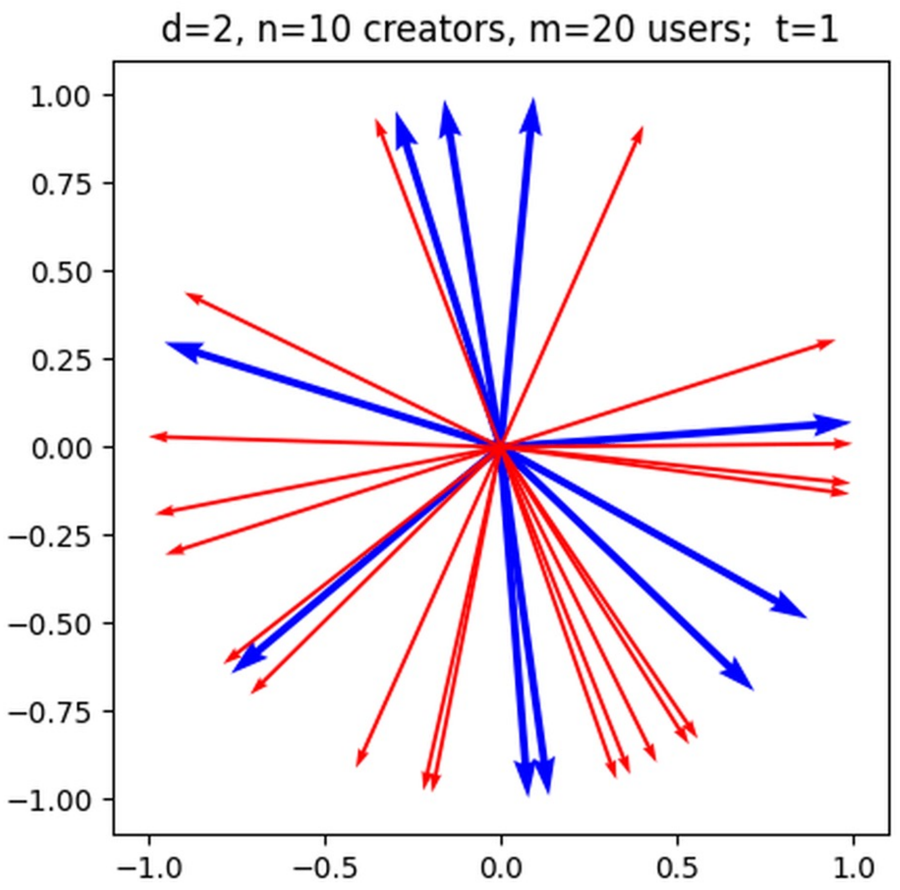

Intuition:
Why does diversified recommendation lead to polarization?

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

users creators

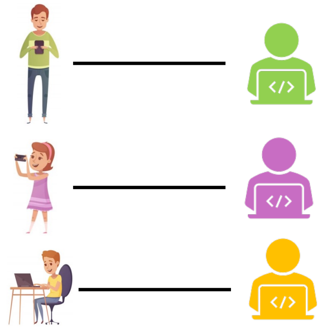

Uniform recommendation:

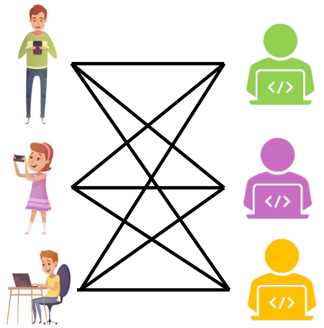

Creators update

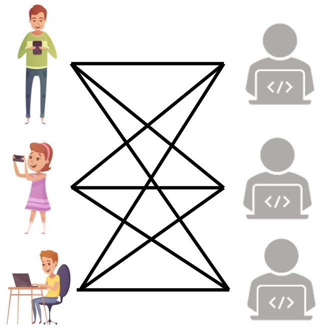

Under diversified recommendation,
different creators will be matched with more similar sets of users,
hence, they will update towards a more similar direction.

Proof of Theorem 1: Absorbing Markov Chain

- Consider $X^t = (U^t, V^t)$ as the state of a Markov chain (with infinite state space)
- Transition $X^t = (U^t, V^t) \rightarrow X^{t+1} = (U^{t+1}, V^{t+1})$ is memoryless and stochastic

Lemma 1 (absorbing):
For any $r \in [0, 1]$, the set of r-polarization states are absorbing (*once enter, never leave*)

Lemma 2 (finite path to polarization):
For any initial state X^t, for any $r \in (0, 1]$, there exists a sequence of transitions:

$$X^t \rightarrow X^{t+1} \rightarrow \cdots \rightarrow X^{t+T_r}$$

that leads to r-polarization.

r-polarization:

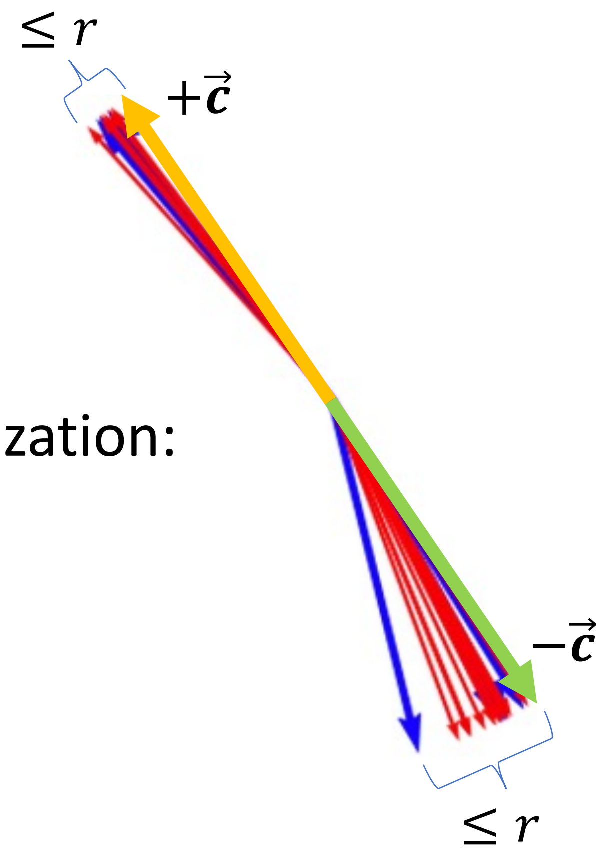

Proof of Theorem 1: Absorbing Markov Chain

- Consider $X^t = (U^t, V^t)$ as the state of a Markov chain (with infinite state space)
- Transition $X^t = (U^t, V^t) \rightarrow X^{t+1} = (U^{t+1}, V^{t+1})$ is memoryless and stochastic

Lemma 1 (absorbing):
For any $r \in [0, 1]$, the set of r-polarization states are absorbing (*once enter, never leave*)

Lemma 2 (finite path to polarization):
For any initial state X^t, for any $r \in (0, 1]$, there exists a sequence of transitions:

$$X^t \rightarrow X^{t+1} \rightarrow \cdots \rightarrow X^{t+T_r}$$

that leads to r-polarization.

Proof of Theorem 1:
Consider $\mathrm{Prob}[X^t \rightarrow X^{t+1} \rightarrow \cdots \rightarrow X^{t+T_r}]$:

- For each user, every creator can be recommended to the user with probability $\geq p_0$, so

$$\mathrm{Prob}[\, X^t \rightarrow X^{t+1} \,] \geq p_0^m.$$

- This implies

$$\mathrm{Prob}[X^t \rightarrow X^{t+1} \rightarrow \cdots \rightarrow X^{t+T_r}] \geq p_0^{mT_r} > 0.$$

So,
Prob[not enter r-polarization after KT_r steps]

$$\leq \left(1 - p_0^{mT_r}\right)^K \rightarrow 0 \quad \text{as } K \rightarrow +\infty$$

Proof of Lemma 2

Induction on the number of creators n

Lemma 2 (finite path to polarization)**:** For any initial state X^t, for any $r \in (0, 1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \cdots \to X^{t+T_r}$$

that leads to r-polarization.

Base Case $(n = 1, m \geq 1)$:

The system deterministically converges to 0-polarization:

- In particular, for any $r > 0$, the system converges to r-polarization in $T_r^1 < +\infty$ steps

A *potenial function* argument

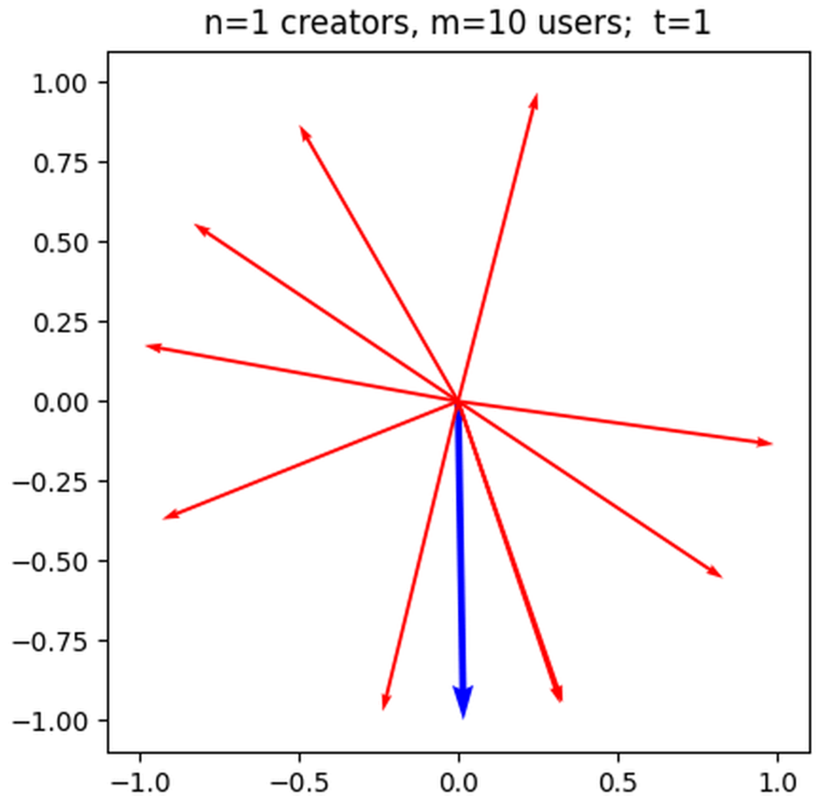

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:**
For any initial state X^t, for any $r \in (0, 1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \dots \to X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the “reflection” of one clusters:

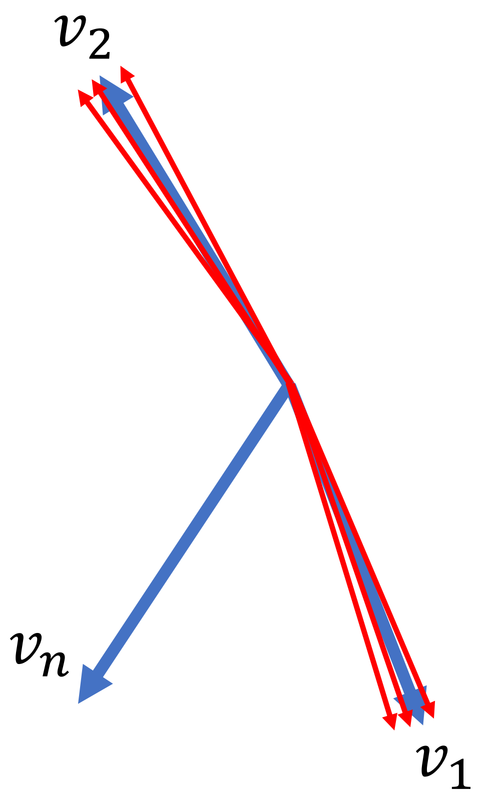

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:**
For any initial state X^t, for any $r \in (0,1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \cdots \to X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the "reflection" of one clusters:

"$2r$-consensus"

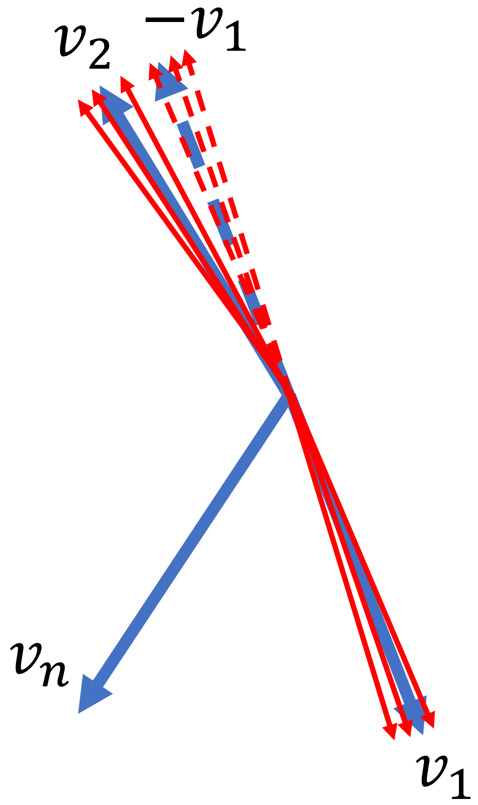

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:**
For any initial state X^t, for any $r \in (0, 1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \cdots \to X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the "reflection" of one clusters: $2r$-consensus
- Consider the position of v_n:

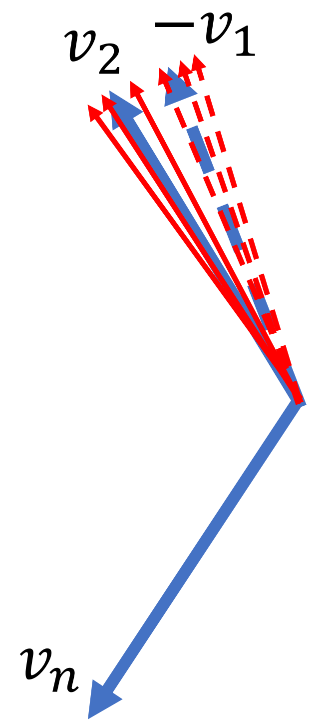

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:**
For any initial state X^t, for any $r \in (0, 1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \dots \to X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the “reflection” of one clusters: $2r$-consensus
- Consider the position of v_n:
  - (1): < 90º with the cluster
  - (2): > 90º with the cluster
  - (3): some < 90º, some > 90º

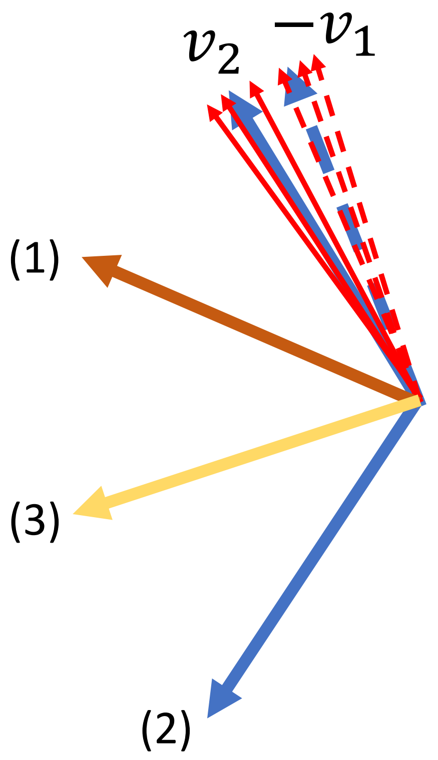

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:**
For any initial state X^t, for any $r \in (0,1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \dots \to X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the "reflection" of one clusters: $2r$-consensus
- Consider the position of v_n:
  - (1): < 90º with the cluster
  - (2): > 90º with the cluster
  - (3): some < 90º, some > 90º

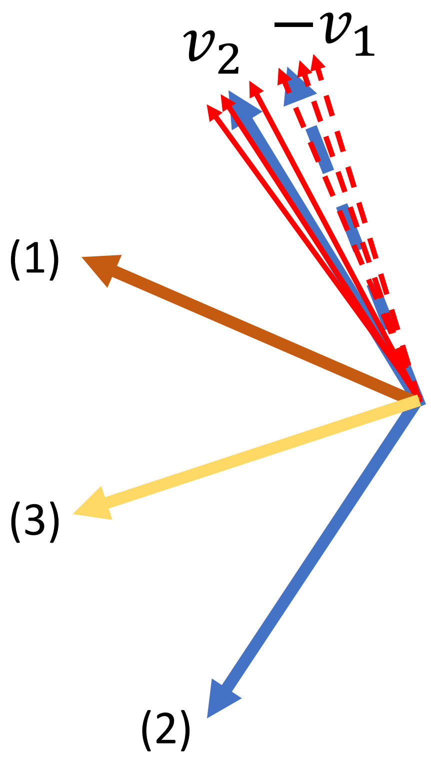

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:**
For any initial state X^t, for any $r \in (0, 1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \dots \to X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the “reflection” of one clusters: $2r$-consensus
- Consider the position of v_n:
  - **(1): < 90º with the cluster: Converge to consensus**
  - (2): > 90º with the cluster
  - (3): some < 90º, some > 90º

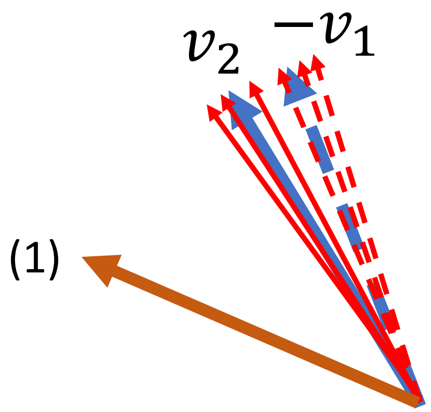

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:**
For any initial state X^t, for any $r \in (0, 1]$, there exists a sequence of transitions:

$$X^t \rightarrow X^{t+1} \rightarrow \dots \rightarrow X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the "reflection" of one clusters: $2r$-consensus
- Consider the position of v_n:
  - (1): < 90º with the cluster: Converge to consensus
  - **(2): > 90º with the cluster: Converge to polarization**
  - (3): some < 90º, some > 90º

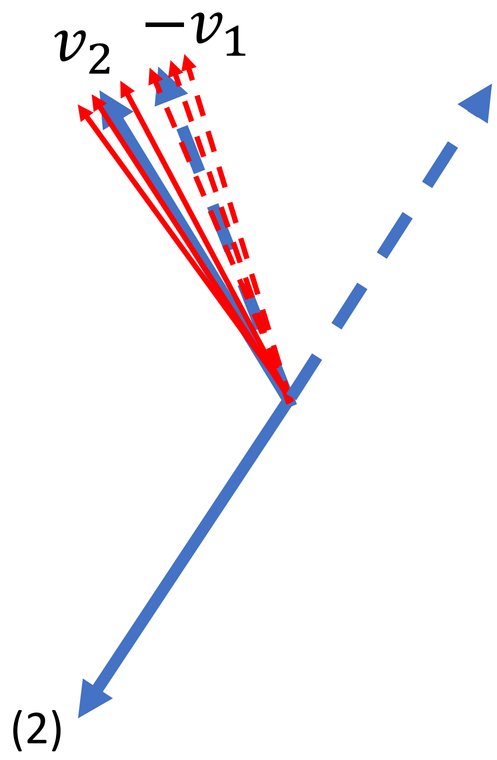

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:** For any initial state X^t, for any $r \in (0,1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \cdots \to X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the “reflection” of one clusters: $2r$-consensus
- Consider the position of v_n:
  - (1): < 90° with the cluster: Converge to consensus
  - (2): > 90° with the cluster: Converge to polarization
  - **(3): some < 90°, some > 90°**

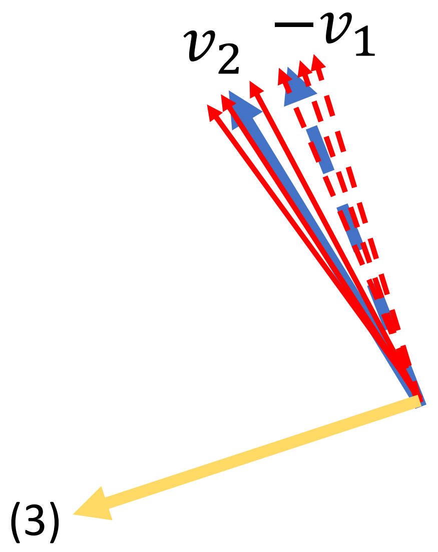

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:**
For any initial state X^t, for any $r \in (0,1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \dots \to X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the "reflection" of one clusters: $2r$-consensus
- Consider the position of v_n:
  - (1): < 90º with the cluster: Converge to consensus
  - (2): > 90º with the cluster: Converge to polarization
  - **(3): some < 90º, some > 90º**

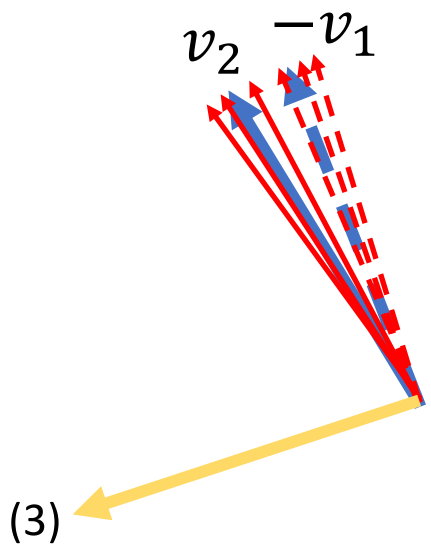

Recommend v_n *to some users with angle* < 90º

Proof of Lemma 2: Inductive Step

Lemma 2 (finite path to polarization)**:**
For any initial state X^t, for any $r \in (0,1]$, there exists a sequence of transitions:

$$X^t \to X^{t+1} \to \cdots \to X^{t+T_r}$$

that leads to r-polarization.

- Consider the *subsystem* consisting of $n-1$ creators and m users
- There exsits a path of length T_r^{n-1} leads the subsystem to r-polarization
- Consider the “reflection” of one clusters: $2r$-consensus
- Consider the position of v_n:
  - (1): < 90º with the cluster: Converge to consensus
  - (2): > 90º with the cluster: Converge to polarization
  - **(3): some < 90º, some > 90º**

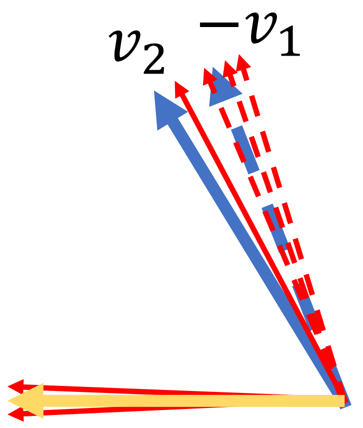

Becomes (1)!

Recommend v_n *to some users with angle* < 90º

Outline

- Introduction
- Model: User-Creator Feature Dynamics
- Main Results: Diversified Recommendation Leads to Polarization
- Ways to Mitigate Polarization

Possible ways to mitigate polarization

- **Uniform recommendation** or **setting diversity-boosting objectives**:
  - $p_{ij}^t > 0$

Possible ways to mitigate polarization

- ~~**Uniform recommendation** or **setting diversity-boosting objectives**:~~
  - $p_{ij}^t > 0$

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 $\langle u_j^t, v_i^t \rangle \geq \tau$

Proposition:

Under top-k or threshold truncation, there exist stable states with more than two clusters:

$\frac{n}{k}$ clusters for top-k $d + 1$ for threshold $\tau = 0$

Effect of top-k truncation: more than two clusters

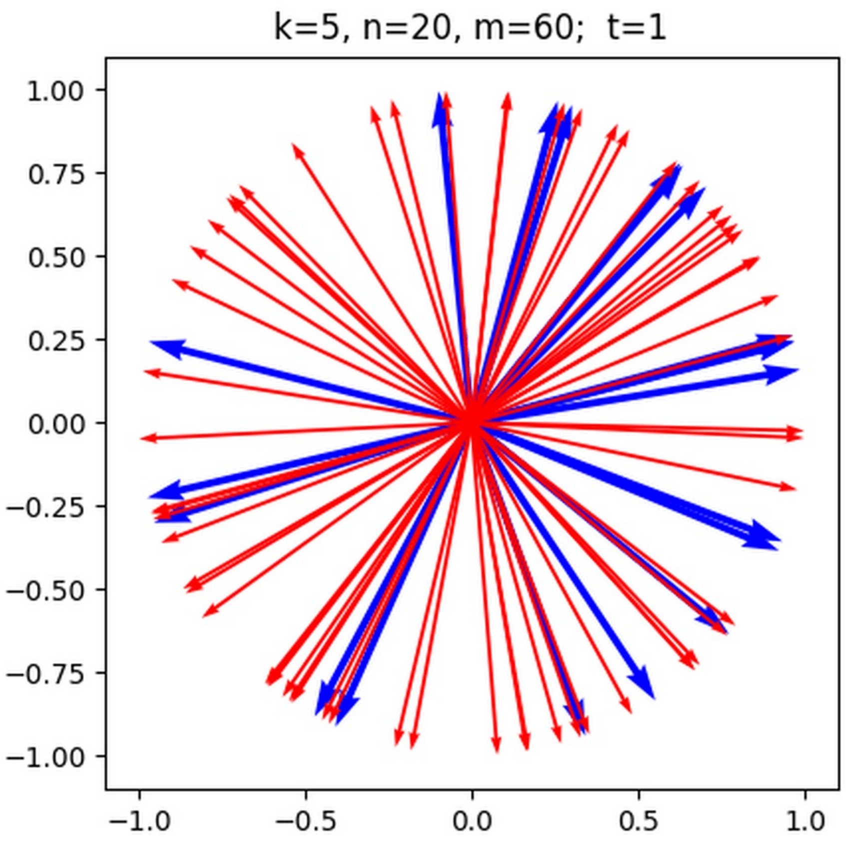

Proposition:

Under top-k truncation, there exist stable states with $\frac{n}{k}$ clusters.

Effect of top-k truncation: reduced polarization

Average pairwise distance between creators:

$$\frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{k \neq i} |v_i - v_j|$$

Average relevancy of creators wrt users:

$$\frac{1}{m} \sum_{j=1}^{m} \sum_{i=1}^{n} p_{ij} \langle v_i, u_j \rangle$$

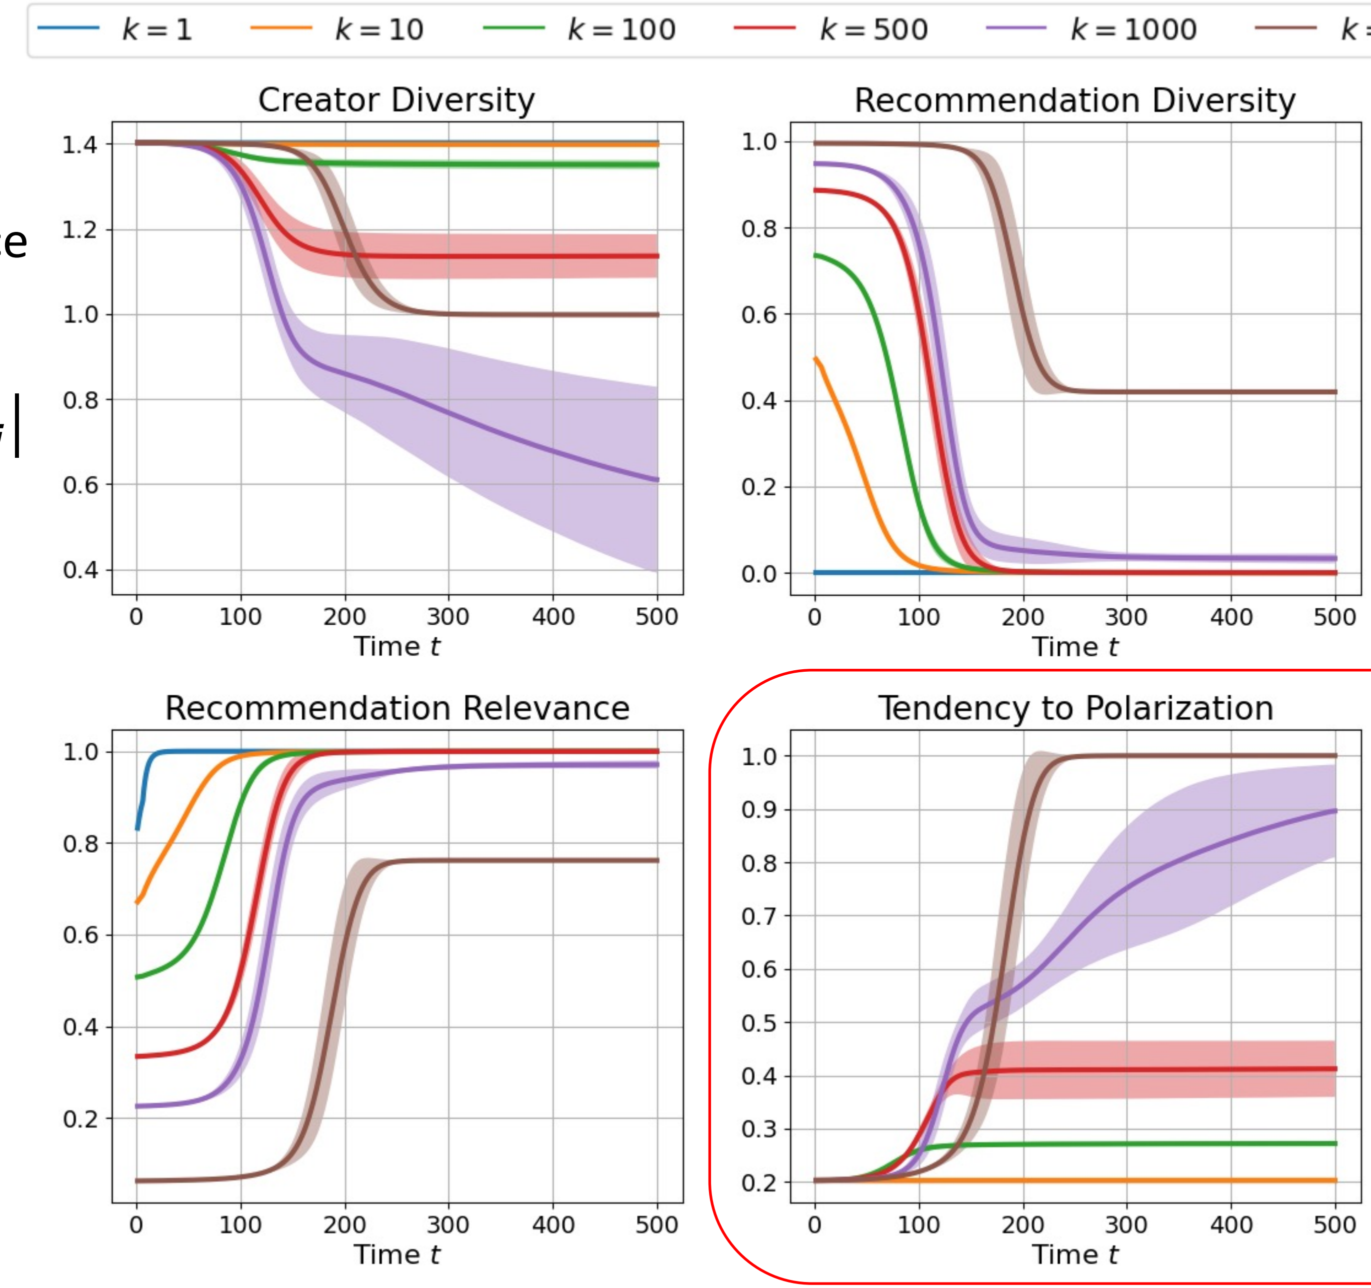

Weighted variance of creators wrt to a user:

$$\frac{1}{m} \sum_{j=1}^{m} \sum_{i=1}^{n} p_{ij} |v_i - \overline{v_j}|^2$$

Tendency to Polarization: Average absolute inner product between creators:

$$\frac{1}{n^2} \sum_{i=1}^{n} \sum_{k=1}^{n} |\langle v_i, v_k \rangle|$$

Effect of threshold truncation

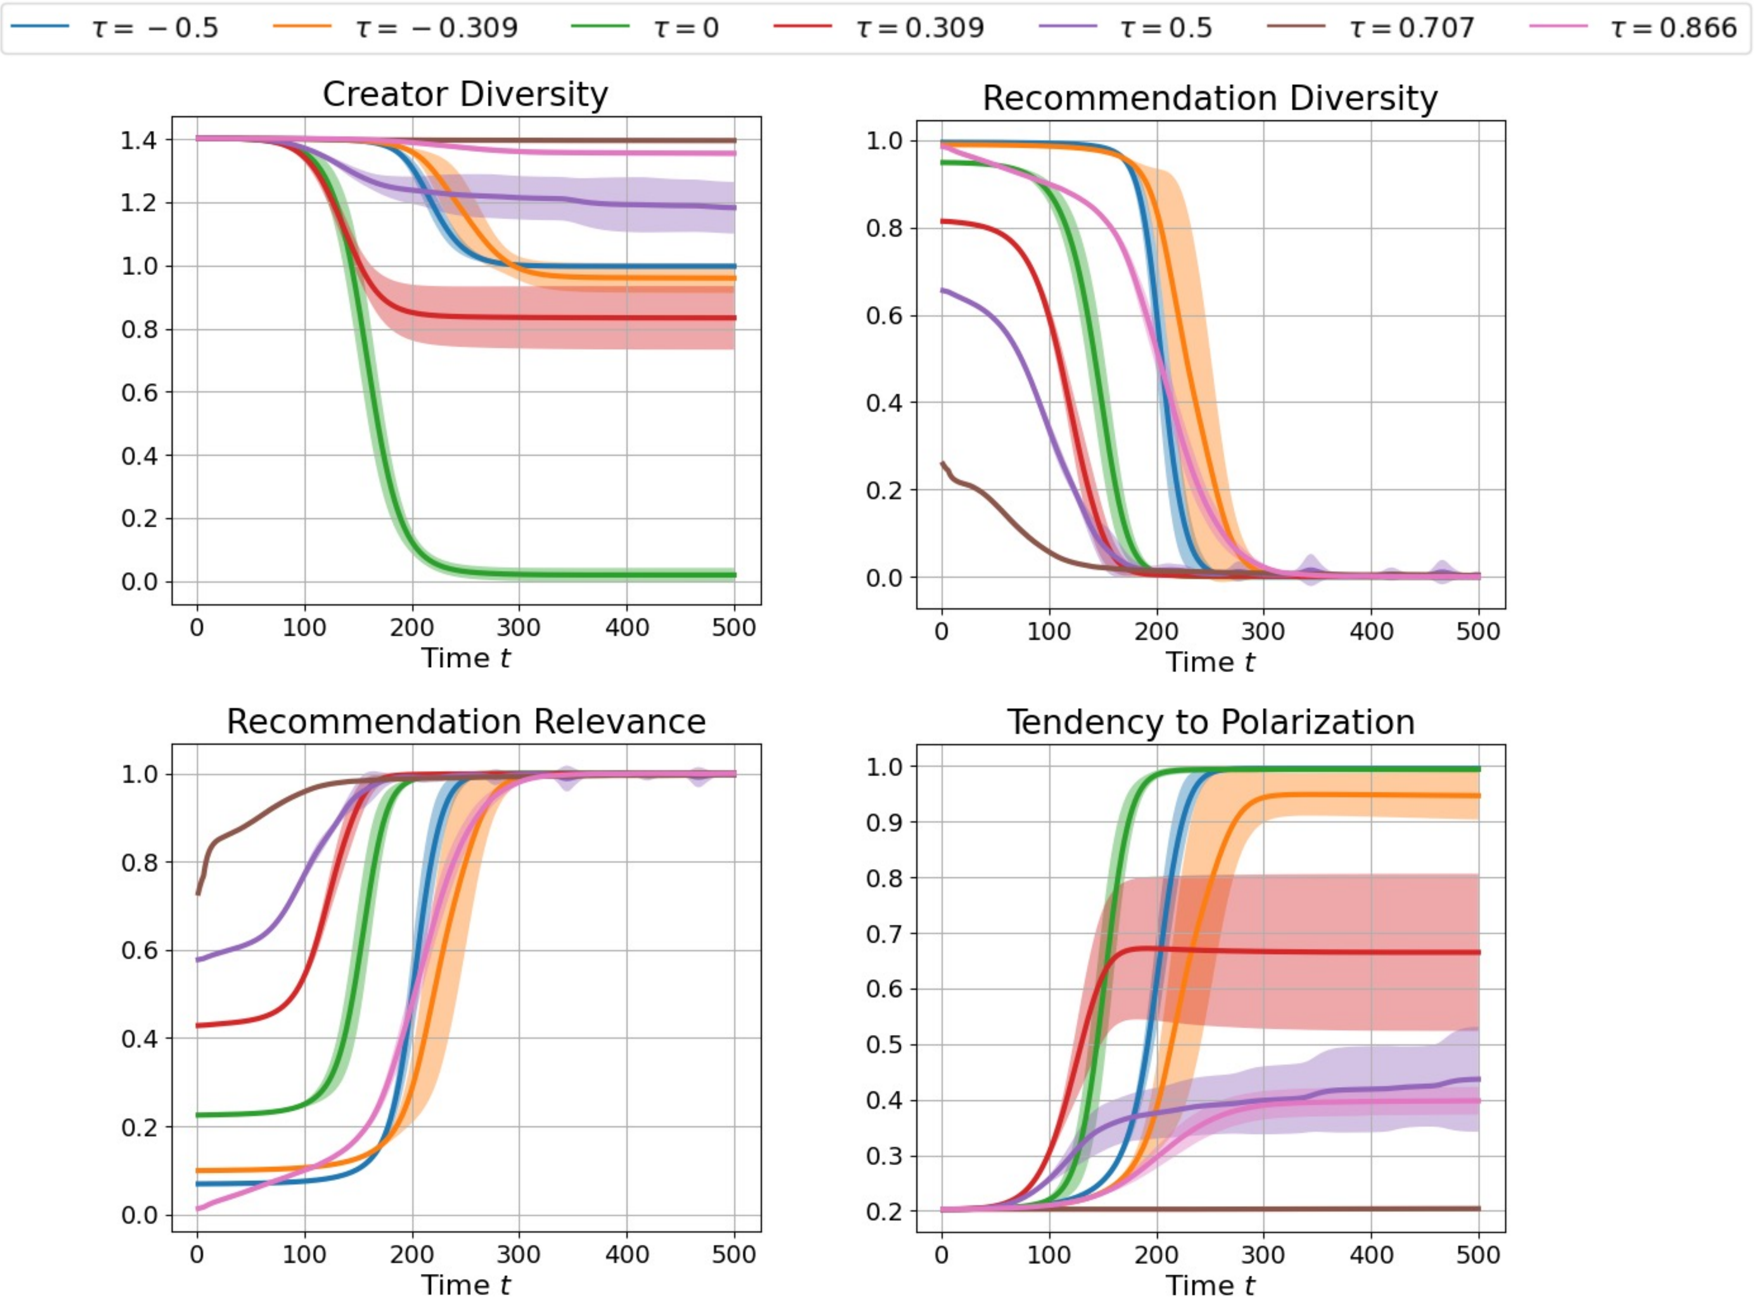

Increasing relevancy mitigates polarization

Besides **top-k truncation** and **threshold truncation**,

We can also just **increase** β in the softmax function: $\frac{\exp(\beta \cdot \langle v_i^t, u_j^t \rangle)}{\sum_{k\in[n]} \exp(\beta \cdot \langle v_k^t, u_j^t \rangle)}$

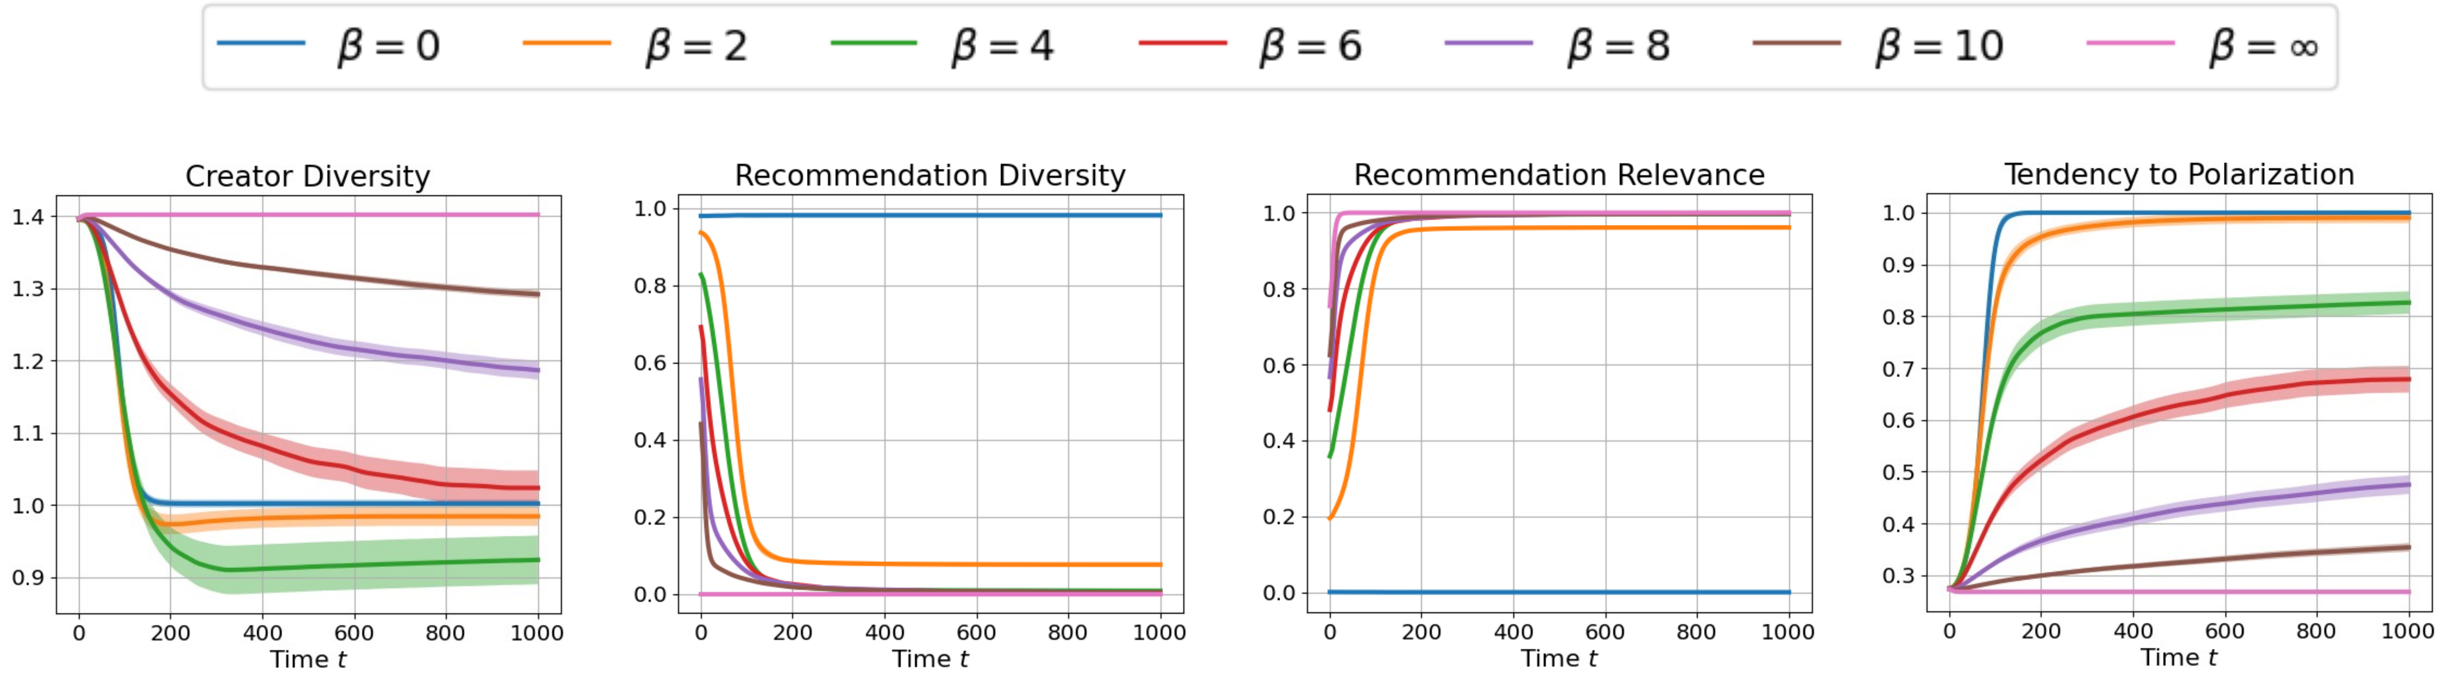

The effect of β

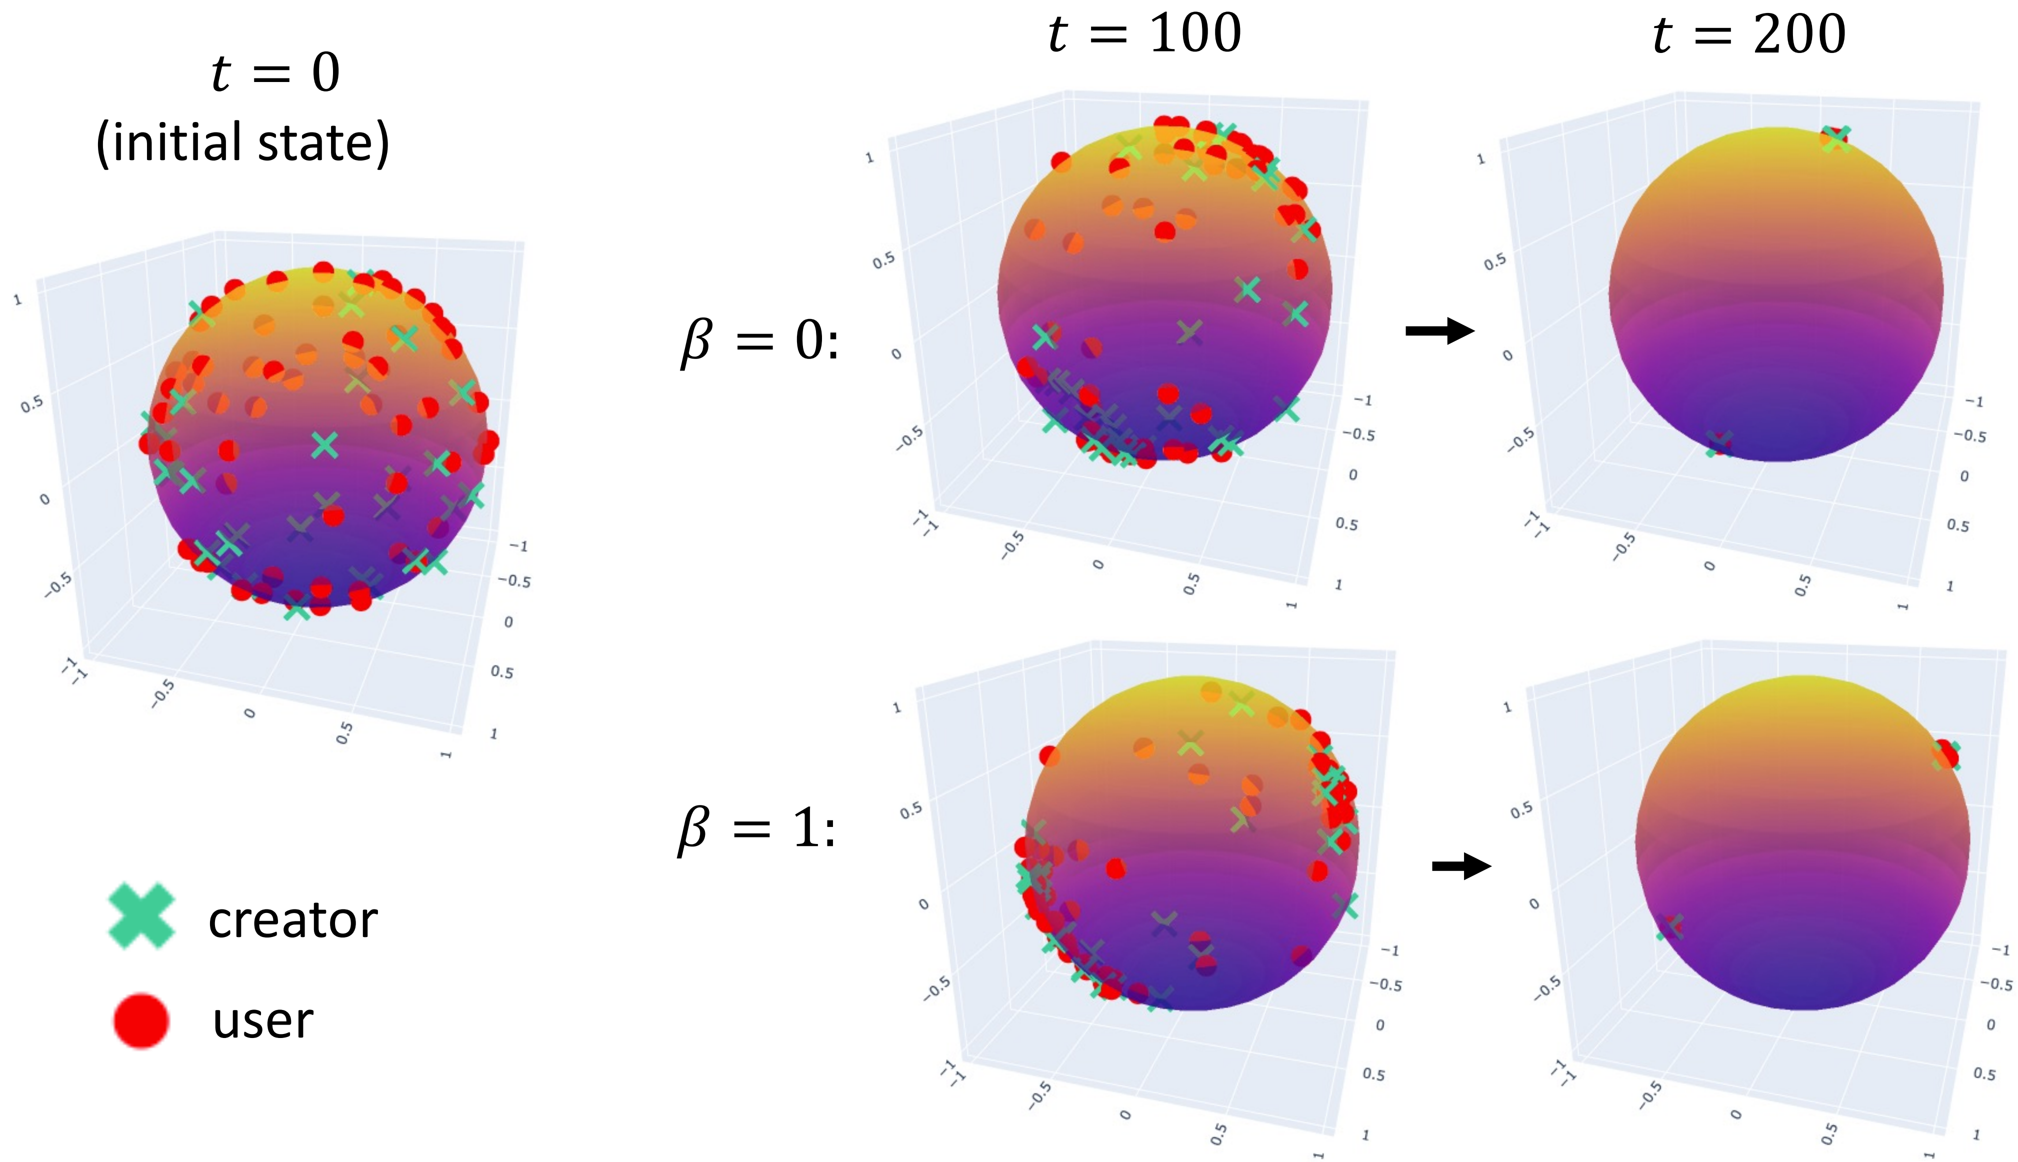

Larger β (higher relevancy) results in more clusters (higher creator diversity & less polarization)

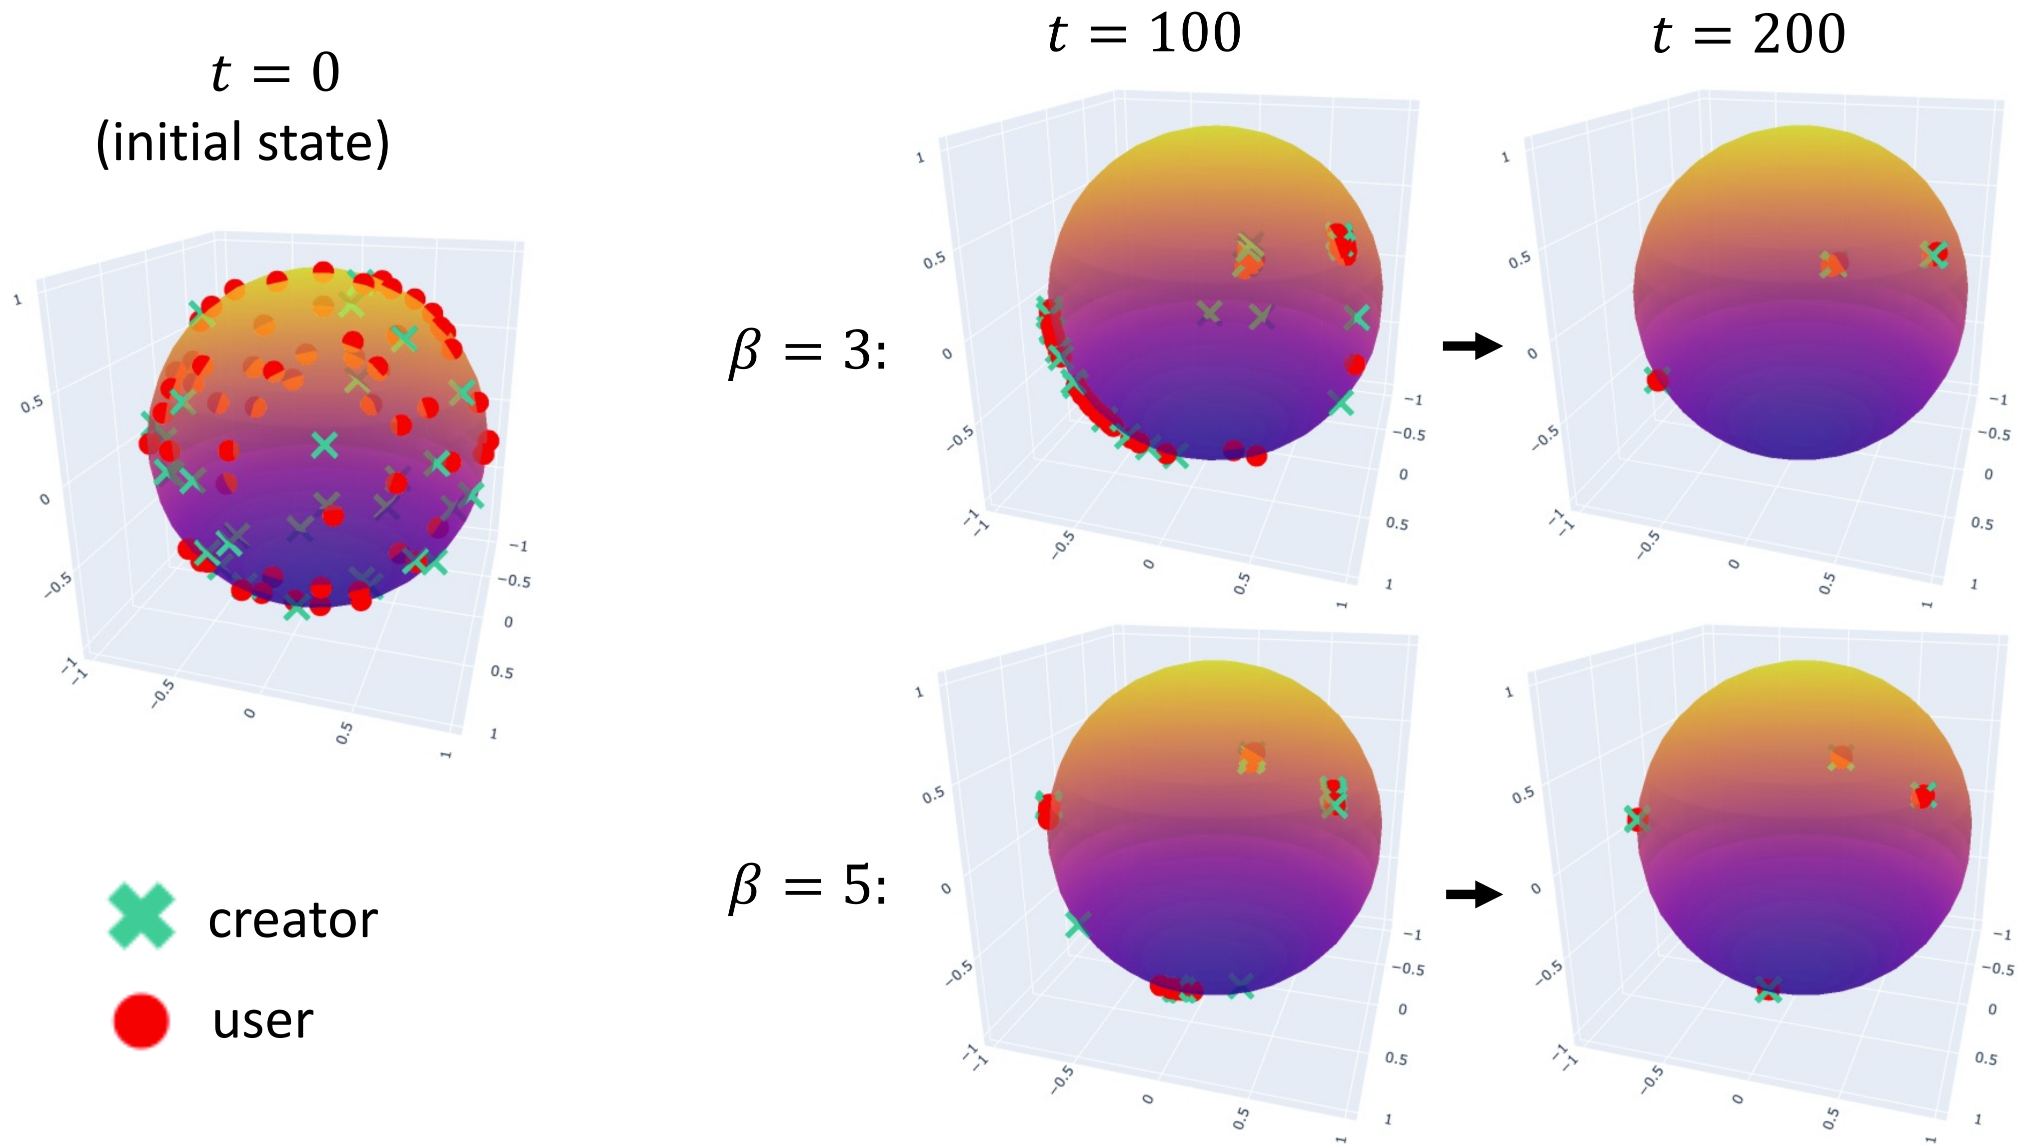

Summary

- We provide a theoretical model to capture the dual influence of recommender systems.
- Simple diversification techniques cannot improve diversity in the long run.
- Increasing relevancy reduces polarization.
- The *tradeoff* between **the diversity of recommendations to users** and **the diversity of the entire system** is worth exploring.

To design diverse and healthy recommender systems, we have to take into account the *multi-sided influences* of such systems in the real world.

See our paper for details:

Tao Lin, Kun Jin, Andrew Estornell, Xiaoying Zhang, Yiling Chen, Yang Liu
User-Creator Feature Polarization in Recommender Systems with Dual Influence. (NeurIPS 2024)