Supply-Side Equilibria in Recommender Systems

Meena Jagadeesan (UC Berkeley)

Joint work with Nikhil Garg (Cornell Tech) and Jacob Steinhardt (UC Berkeley)

Appeared at NeurIPS 2023

FORC 2024 (non-archival)

Classical View of Content Recommendation

Content

Content

Recommend video that maximizes user engagement

Classical View of Content Recommendation

Content

Content

Common perspective: a rec sys just selects content on the platform to show each user

YouTube

Recommend video that maximizes user engagement

Classical View of Content Recommendation

Content

Content

Common perspective: a rec sys just selects content on the platform to show each user

YouTube

Recommend video that maximizes user engagement

Reality: The rec sys **shapes** the content landscape.

Key Driver: Content Creator Incentives

Content creator

Content creator

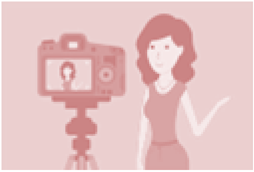

YouTube

Recommend video that maximizes user engagement

Key Driver: Content Creator Incentives

Content creator

Content creator

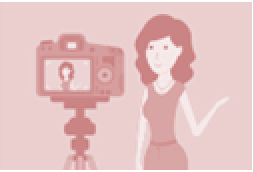

Creators **strategically design content** to be recommended.

Key Driver: Content Creator Incentives

Content creator

Content creator

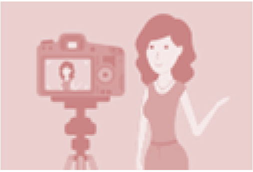

Creators **strategically design content** to be recommended.

Mainstream content

Specialized (niche) content

Key Driver: Content Creator Incentives

Content creator

Content creator

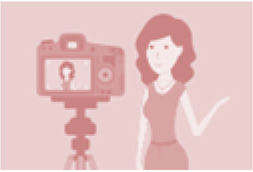

Creators **strategically design content** to be recommended.

Mainstream content

Specialized (niche) content

<u>**This work**</u>: The **design of a rec sys** implicitly influences the **content landscape** via **creator incentives.**

Our contribution

We model content creator competition as a **high-dimensional game**.

We characterize how **rec sys design** affects **content specialization at equilibrium**.

Mainstream content

Specialized (niche) content

Rec sys design ≈ embeddings learned by the rec sys

Simplified model of a recommender system

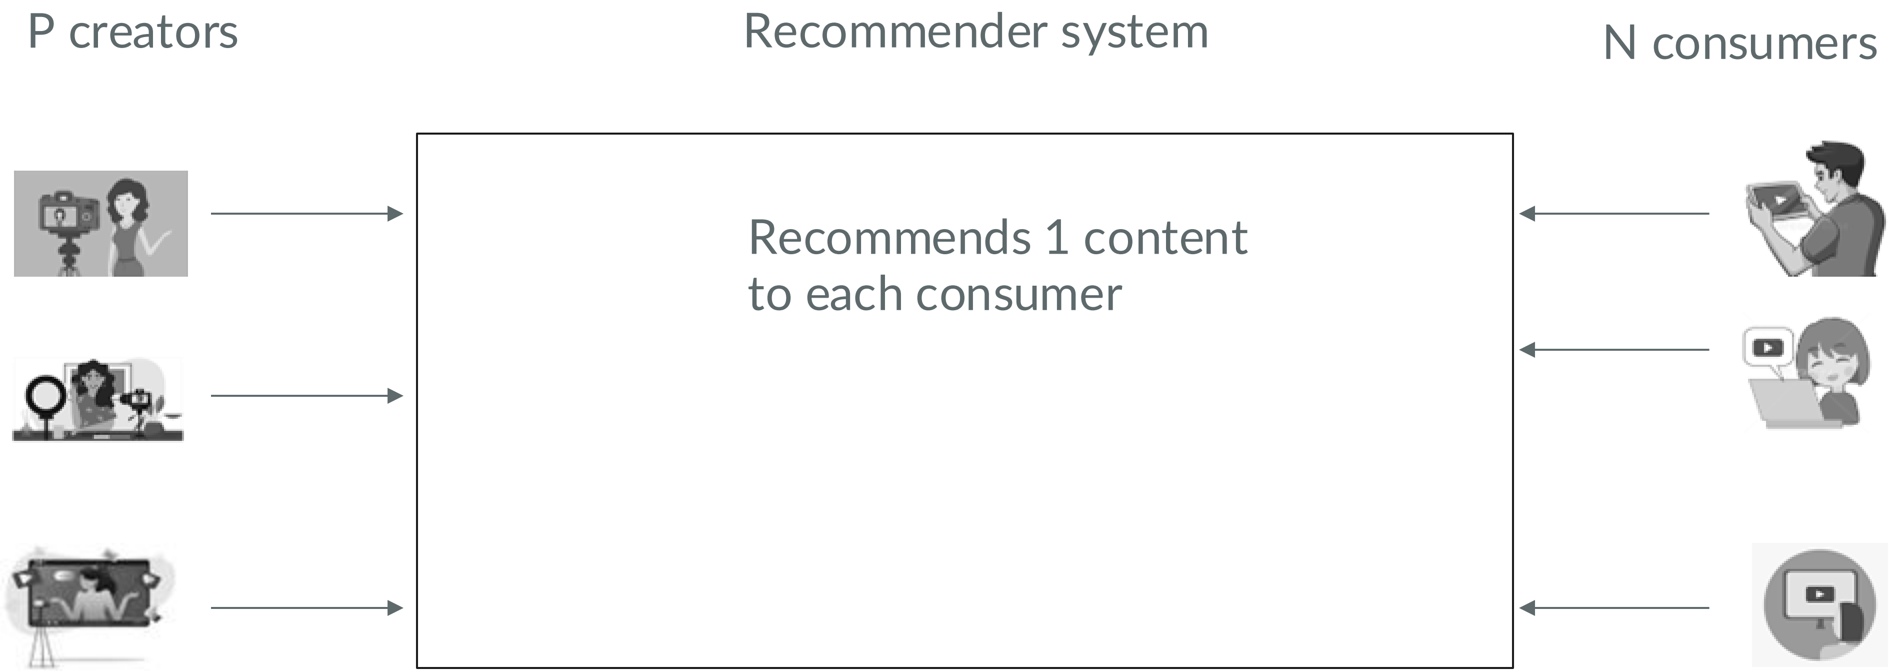

Simplified model of a recommender system

P creators

Recommender system

N consumers

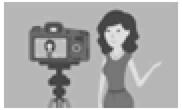

$p_1 \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

$p_2 \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

$\vdots$

$p_P \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

Recommends 1 content to each consumer

Learns consumer and content embeddings

$u_1 \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

$u_2 \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

$\vdots$

$u_N \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

Simplified model of a recommender system

P creators

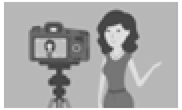

Recommender system

$p_1 \in \mathbf{R}^{\mathbf{D}}_{\geq 0}$

$p_2 \in \mathbf{R}^{\mathbf{D}}_{\geq 0}$

.
.
.

$p_P \in \mathbf{R}^{\mathbf{D}}_{\geq 0}$

Recommends 1 content to each consumer

Learns consumer and content embeddings

$$j^*(u_i; [p_1, \ldots, p_P]) = \mathrm{argmax}_{1 \leq j \leq P} < p_j, u_i >$$

$u_1 \in \mathbf{R}^{\mathbf{D}}_{\geq 0}$

$u_2 \in \mathbf{R}^{\mathbf{D}}_{\geq 0}$

.
.
.

$u_N \in \mathbf{R}^{\mathbf{D}}_{\geq 0}$

N consumers

Simplified model of a recommender system

P creators

Recommender system

N consumers

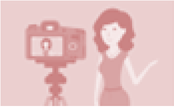

$p_1 \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

Recommends 1 content to each consumer

$u_1 \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

$p_2 \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

Learns consumer and content embeddings

$u_2 \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

$\vdots$

$$j^*(u_i; [p_1, \ldots, p_P]) = \mathrm{argmax}_{1 \leq j \leq P} < p_j, u_i >$$

$p_P \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

$u_N \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$

Each creator **strategically selects** a D-dimensional content vector.

Creator competition as a high-dimensional game

Creator action space = $\mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$ (captures both "genre" p / $||p||$ and "quality" $||p||$)

Creator competition as a high-dimensional game

Creator action space = $\mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$ (captures both "genre" p / $||p||$ and "quality" $||p||$)

Creator profit: $\mathbf{P}(p_j \mid p_{-j}, u_{1:N}) = (\sum_{1 \leq i \leq N} \mathrm{I}[\operatorname{argmax}_{1 \leq j' \leq P} < p_{j'}, u_i > = j]) - c(p_j)$

Exposure (# of recommendations won)

One-time cost of production

Creator competition as a high-dimensional game

Creator action space = $\mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$ (captures both "genre" p / $||p||$ and "quality" $||p||$)

Creator profit: $\mathbf{P}(p_j \mid p_{-j}, u_{1:N}) = (\sum_{1 \leq i \leq N} \mathrm{I}[\operatorname{argmax}_{1 \leq j' \leq P} \langle p_{j'}, u_i \rangle = j]) - c(p_j)$

Exposure (# of recommendations won)

One-time cost of production

$c(p) := ||p||^{\beta}$ ($\beta \geq 1$ captures level of nonlinearity)

Creator competition as a high-dimensional game

Creator action space = $\mathbf{R}^{\mathbf{D}}_{\geq 0}$ (captures both "genre" *p* / ||*p*|| and "quality" ||*p*||)

Creator profit: $\mathbf{P}(p_j \mid p_{-j}, u_{1:N}) = (\sum_{1 \leq i \leq N} \mathrm{I}[\operatorname{argmax}_{1 \leq j' \leq P} \langle p_{j'}, u_i \rangle = j]) - c(p_j)$

Exposure (# of recommendations won)

One-time cost of production

$c(p) := ||p||^{\beta}$ ($\beta \geq 1$ captures level of nonlinearity)

We study Nash equilibria between creators (captures content landscape).

Creator competition as a high-dimensional game

Creator action space = $\mathbf{R}^{\mathbf{D}}_{\geq 0}$ (captures both "genre" p / $||p||$ and "quality" $||p||$)

Creator profit: $\mathbf{P}(p_j \mid p_{-j}, u_{1:N}) = (\sum_{1 \leq i \leq N} \mathrm{I}[\ \mathrm{argmax}_{1 \leq j' \leq P} < p_{j'}, u_i > = j\ ]) - c(p_j)$

Exposure (# of recommendations won)

One-time cost of production

$c(p) := ||p||^{\beta}$ ($\beta \geq 1$ captures level of nonlinearity)

We study Nash equilibria between creators (captures content landscape).

Content specialization ≈ richness of set of genres p / ||p|| appearing at equilibrium

Related work

Content creator competition in rec sys (1-dimension or finite action space):

e.g., Ghosh and McAfee (2011), Ben Basat, Tennenholtz, Kurland (2017), Ben-Porat, Tenneholtz (2018, 2020), etc.

Subsequent work: *Hron, Krauth, Jordan, Kilbertus, Dean (concurrent), Yao, Li, Nekipelov, Wang, Xu (2023), Milli, Pierson, Garg (2023), Eilat, Rosenfeld (2023), Prasad, Mladeov, Boutilier (2023), Buening, Saha, Dimitrakasis, Xu (2024), Immorlica,* ***J.****, Lucier (2024), etc.*

Economic models of product selection and strategic facility location:

e.g., ***Hotelling (1929),*** *d'Aspremont, Jaskold, Thisse (1979), Salop (1979), Anderson (1992), Berry (1994), ReVelle, Eiselt (2005)*

Strategic adaptation in classification and regression:

e.g., Brückner, Scheffer (2011), Hardt, Megiddo, Papadimitriou, Wootters (2015), Perdomo, Zrnic, Mendler-Dünner, Hardt (2020)

Societal impacts of recommender systems:

e.g. Adomavicius, Bockstedt, Curley, Zhang (2013), Flaxman, Goel, Rao (2016), Kleinberg, Raghavan (2021), Milli, Belli, Hardt (2021), etc.

Our model: high-dimensional content, multi-sided competition, nonlinear costs, etc.

Our main results

We show how rec sys design shapes whether content specialization occurs.

Result: we characterize when the **content landscape consists of $>$ 1 content genre.**

- Shows role of **learned consumer embeddings**
- Shows role of **creator costs**
- Shows role of **matrix factorization parameters (which determine embeddings)**

Note*: specialization* ***improves content diversity*** *but can* ***reduce content quality*** *(see paper).*

Formalizing the set of genres in the content landscape

Let μ be a **symmetric mixed Nash equilibrium** (μ is a distribution over $\mathbf{R}_{\geq 0}^{\mathbf{D}}$) .

Formalizing the set of genres in the content landscape

Let μ be a **symmetric mixed Nash equilibrium** (μ is a distribution over $\mathbf{R}_{\geq 0}^{\mathbf{D}}$) .

$|\text{Genre}(\mu)| := \{\ p\ /\ ||p||\ \text{s.t.}\ p \in \text{supp}(\mu)\ \}$ (set of genres that arise with some probability)

Formalizing the set of genres in the content landscape

Let μ be a **symmetric mixed Nash equilibrium** (μ is a distribution over $\mathbf{R}^{\mathbf{D}}_{\geq 0}$) .

$|\mathrm{Genre}(\mu)| := \{ p / \|p\| \text{ s.t. } p \in \mathrm{supp}(\mu) \}$ (set of genres that arise with some probability)

Example: $u_1 = [1, 0]$ and $u_2 = [0, 1]$ with $P = 2$ producers and cost function $c(p) = \|p\|_2^{\beta}$.

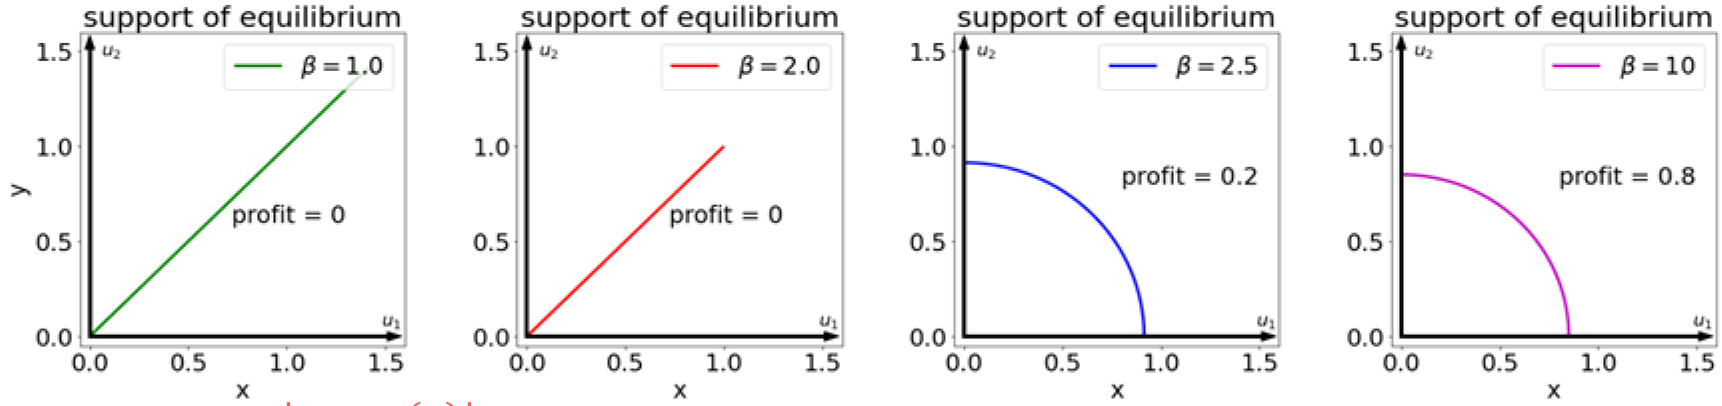

$|\mathrm{Genre}(\mu)| = 1$

(no content specialization)

$|\mathrm{Genre}(\mu)| > 1$

(content specialization)

Main Theorem: Characterization of $|\mathrm{Genre}(\mu)| = 1$ regime

Recall: $|\mathrm{Genre}(\mu)| := \{\, p \,/\, \|p\| \;\, s.t. \;\, p \in \mathrm{supp}(\mu) \,\}$ and $c(p) = \|p\|^{\beta}$.

Let $S = \{ [< u_1, p >, \ldots < u_N, p >] \;\; s.t. \;\; p \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}, \|p\| \leq 1 \}$

Let $S^{\beta} = \{ [< u_1, p >^{\beta}, \ldots < u_N, p >^{\beta}] \;\; s.t. \;\; p \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}, \|p\| \leq 1 \}$

Theorem: There exists an equilibrium μ with $|\mathrm{Genre}(\mu)| = 1$ if and only if:

$$\max \left\{ \prod_{1 \leq i \leq N} y_i \quad s.t. \quad y \in S^{\beta} \right\} = \max \left\{ \prod_{1 \leq i \leq N} y_i \quad s.t. \quad y \in conv(S^{\beta}) \right\}$$

Whether $|\mathrm{Genre}(\mu)| = 1$ relates to **convexity of S^{β}**.

Example: Sets S^{β} for different embeddings and costs

Theorem: There exists an equilibrium μ with $|\text{Genre}(\mu)| = 1$ if and only if:

$$\max\left\{\prod_{1\leq i\leq N} y_i \quad s.t. \quad y \in S^{\beta}\right\} = \max\left\{\prod_{1\leq i\leq N} y_i \quad s.t. \quad y \in conv(S^{\beta})\right\}$$

angle $\angle(u_1, u_2) = 90^o$

angle $\angle(u_1, u_2) < 90^o$

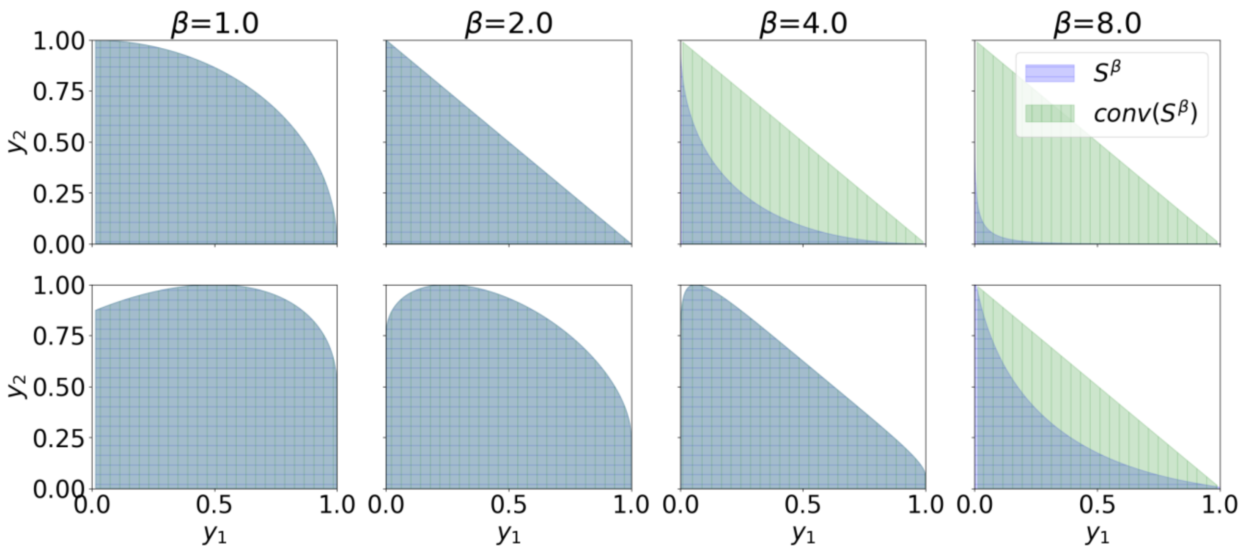

Application: Corollaries of the theorem

Specialization occurs when creator costs are sufficiently nonlinear:

Corollary: There exists ß* such that $|\text{Genre}(\mu)| = 1 \Leftrightarrow ß \leq ß^*$.

Application: Corollaries of the theorem

Specialization occurs when creator costs are sufficiently nonlinear:

Corollary: There exists ß* such that $|\mathrm{Genre}(\mu)| = 1 \Leftrightarrow ß \leq ß^*$.

Specialization occurs whenever consumer embeddings are sufficiently heterogeneous.

Corollary: Consider 2 viewers u_1 , u_2 with angle θ between them. Then:

$$ß^* = \frac{2}{1-\cos(\theta)}.$$

Corollary: Consider any set of viewers. Then:

$$ß^* \leq \frac{\log(N)}{\log(N)-\log\left(||Z||_*\right)} \text{ where } \boldsymbol{Z} := \sum_{i=1}^{N} \frac{\boldsymbol{u_i}}{||\boldsymbol{u_i}||_*}.$$

Application: Matrix factorization hyperparameters

MovieLens-100K dataset: 943 consumers and 100,000 ratings

We run nonnegative matrix factorization with D factors for varying values of D.
We interpret the learned vectors $u_1, \ldots, u_N \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$ as the consumer embeddings.

Application: Matrix factorization hyperparameters

MovieLens-100K dataset: 943 consumers and 100,000 ratings

We run nonnegative matrix factorization with D factors for varying values of D. We interpret the learned vectors $u_1, \ldots, u_N \in \mathbf{R}^{\mathbf{D}}_{\geq \mathbf{0}}$ as the consumer embeddings.

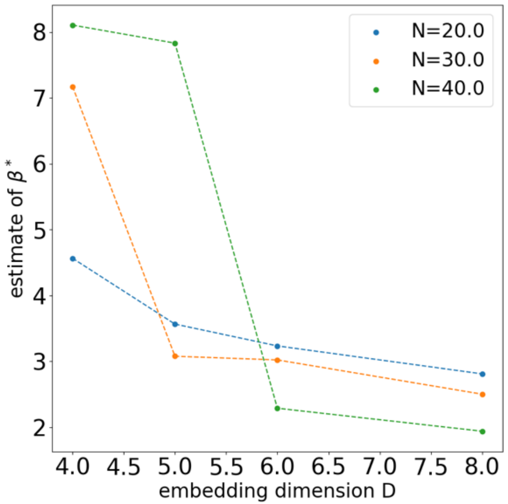

Estimate of β^* for smaller samples (N users)

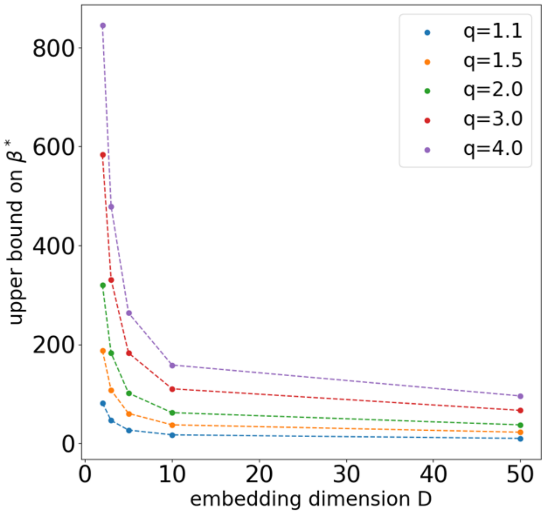

Upper bound on β^* for full dataset

Platform can induce specialization by increasing # of factors D.

Intuition: higher D -> greater personalization -> more content specialization

Conclusion

The design of a rec sys influences the **content landscape** available on the platform.

Our contribution:

We model creator competition as a high-dimensional game, and we characterized when **content specialization** occurs.

Main insight: the embeddings learned by the recommender system shape whether content specialization occurs

- Our main theorem related **content specialization** to a **convexity condition**.

*Specialization **improves content diversity** but can **reduce content quality** (see paper).*

***Future work**: Characterize how other aspects of rec sys design (e.g., ranking across slots, use of LLMs for retrieval, summarization vs. recommendation) affect the content landscape.*