Meta AI

We propose to improve **platform governance** using a method for sampling sets of users that are **diverse**, **fair** and **relevant**. We implement this method in a **scalable** way using **determinantal point processes**.

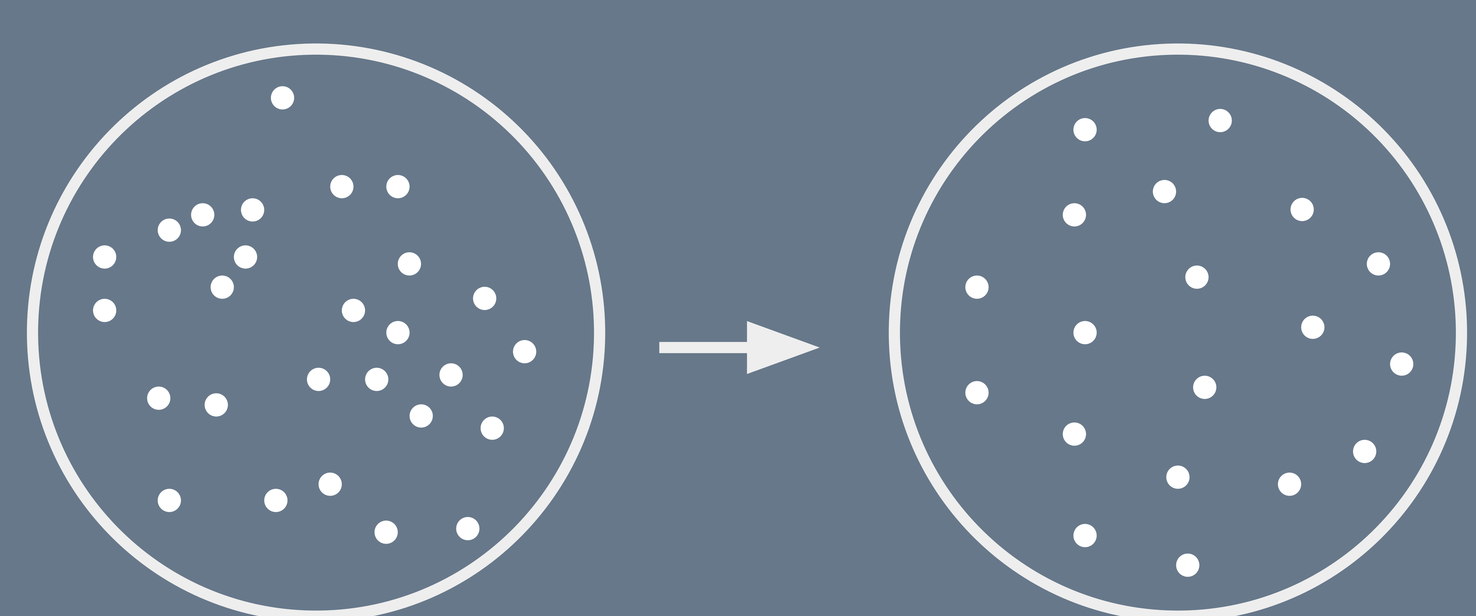

Random Sampling

DPP-based Sampling
(negatively correlated)

Sortition for Platform Governance

Soham De, Ariel Procaccia, Max Nickel, Smitha Milli
sohamde@uw.edu

Summary

- Community-driven approaches to platform governance often require users to collectively decide on key platform issues
  - Wikipedia (editors), Community Notes (contributors), Weibo, Nextdoor (digital juries) ...
- For online platforms, effective user sampling process (aka sortition) should satisfy the following criteria:
  - **Diverse viewpoints:** A selected subset should include users with a wide range of perspectives
  - **Fairness in participation:** Every user should have an equal chance of being chosen
  - **Relevance:** The subset of users should be relevant to the specific decision at hand
- Using Determinantal Point Processes (DPPs)[1,2] we propose a sortition method that allows a practitioner to sample diverse and relevant sets of users while maintaining fair marginal selection probabilities.
  - DPPs sample subsets in proportion to their diversity (volume spanned by their feature vectors)
  - We optimize the DPP such that it ensures fair marginal selection probabilities for all users

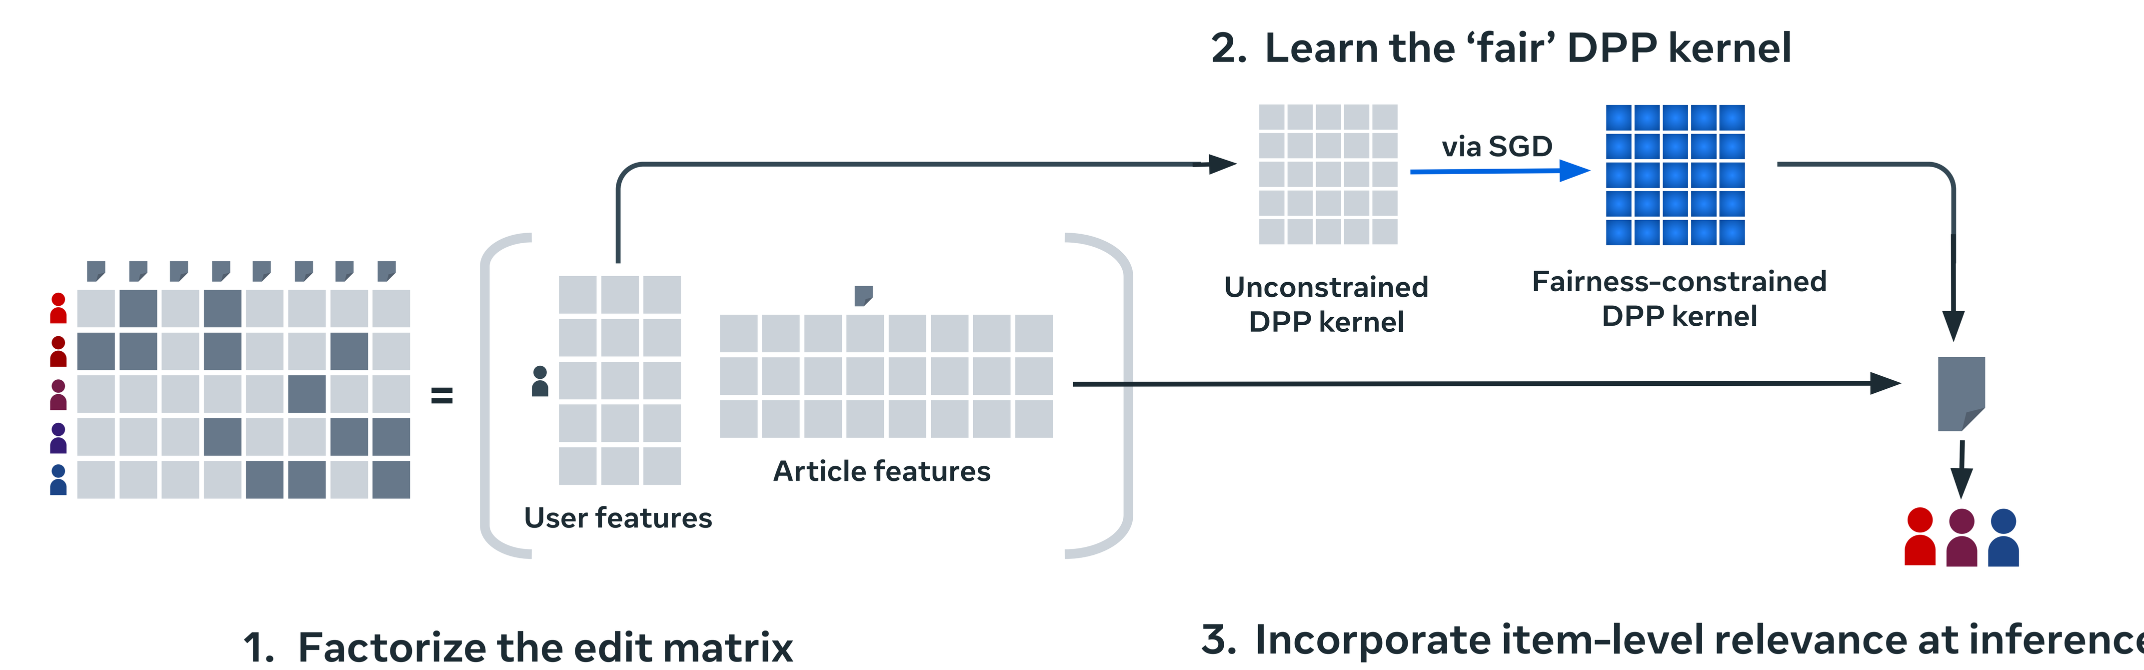

Results

- On a Wikipedia dataset[3], we demonstrate the tradeoffs involved in this approach (Figure 1)
- At inference time, we may incorporate relevance by applying a softmax transformation to the marginal probabilities (Figures 2,3)

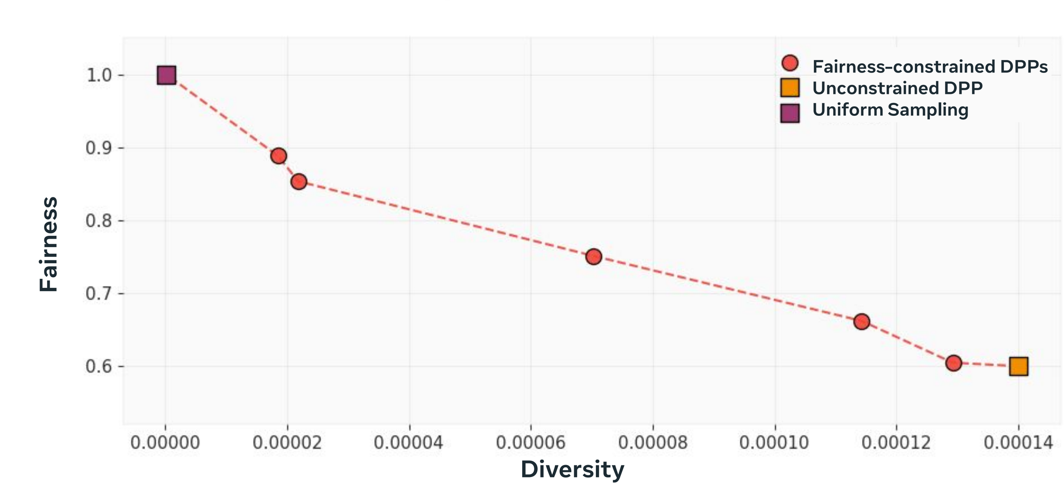

Figure 1: Fairness vs diversity tradeoff

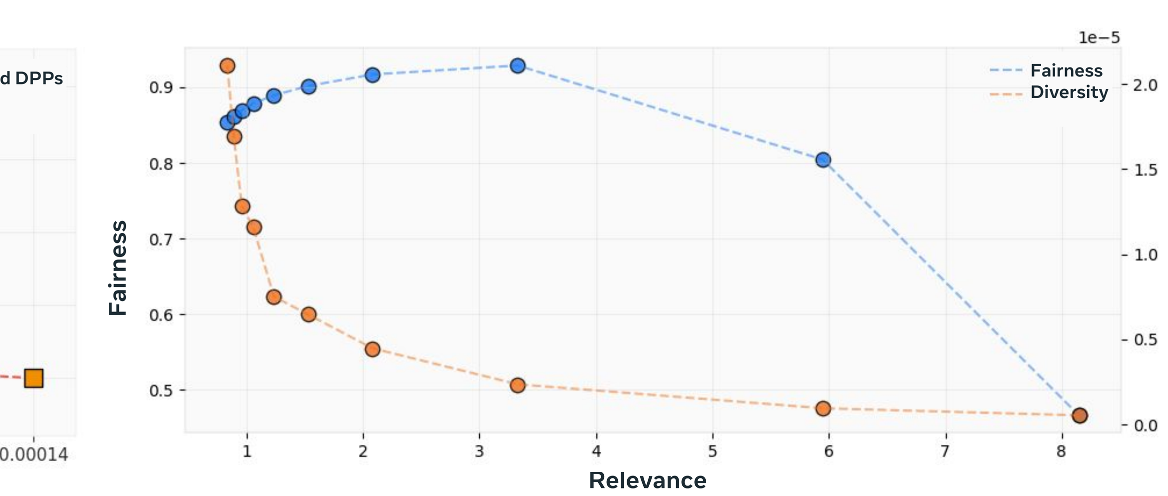

Figure 2: Tradeoffs when incorporating relevance

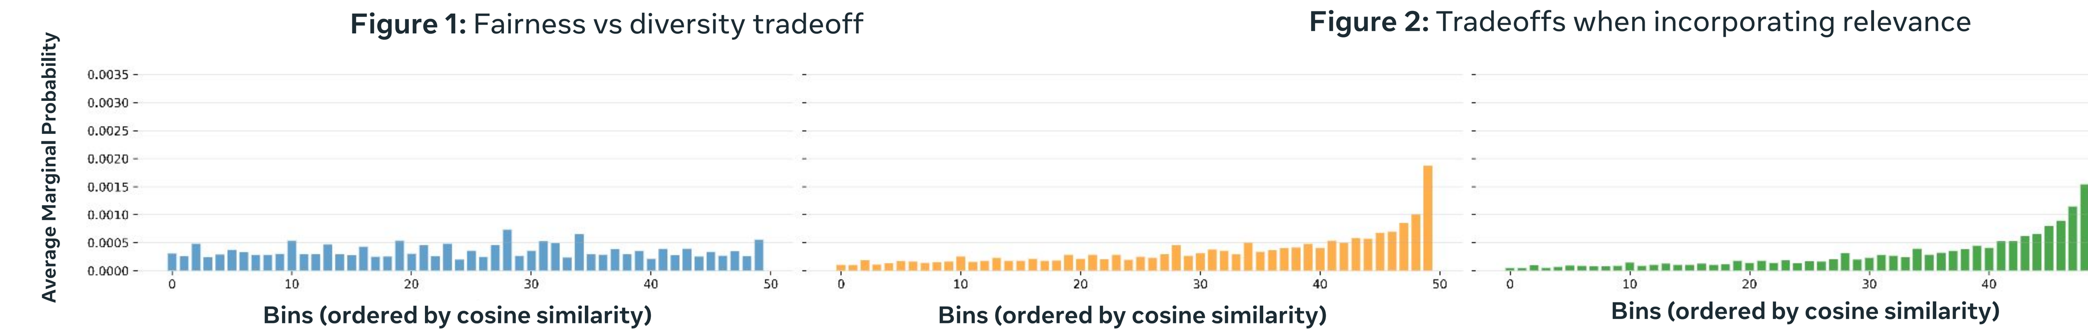

Figure 3: (Close to) Fair marginal selection probabilities → Higher marginal selection probabilities for more relevant users

Determinantal Point Processes

- DPPs[1] are designed to sample diverse sets
- Probability of a sampling a subset U is proportional to its diversity

$$\mathcal{P}_L(\mathbf{U} = U) \propto \det(L_U)$$

where L is a PSD kernel that defines a DPP and captures relationships between users.

- We factorize L as $L = X^T X$ where X represents user features

Fairness in DPPs

- Regular DPPs do not provide fairness guarantees – unique users may have disproportionately high chances of selection
- To achieve fairness in DPPs, we minimize the loss

$$\mathcal{L}(L, L^*) = \delta(L, L^*) - \lambda f(p_1(L), \ldots, p_n(L))$$

where δ is a distance function and $p_1(L), \ldots, p_n(L)$ are marginal selection probabilities of all users

- We make use of SGD and the Gram factor decomposition of DPPs[1] to implement this optimization scalably

Limitations & Future Work

- Learning the DPP kernel (L) involves a non-convex optimization problem. While we are able to scale this method up to millions of users, the global optimality of the learned solutions is not guaranteed
- Future work may present an evaluation of different choices of fairness and distance functions and also validate our measure of diversity (based on determinants) against more interpretable measures using explicit user characteristics

References

[1] Alex Kulesza and Ben Taskar. Determinantal Point Processes for Machine Learning. Now Publishers Inc, 2012.

[2] Mark Wilhelm, Ajith Ramanathan, Alexander Bonomo, Sagar Jain, Ed H. Chi, and Jennifer Gillenwater. Practical diversified recommendations on youtube with determinantal point processes. ACM CIKM, 2018

[3] Feng Shi, Misha Teplitskiy, Eamon Duede, and James A Evans. The wisdom of polarized crowds. Nature human behaviour, 2019

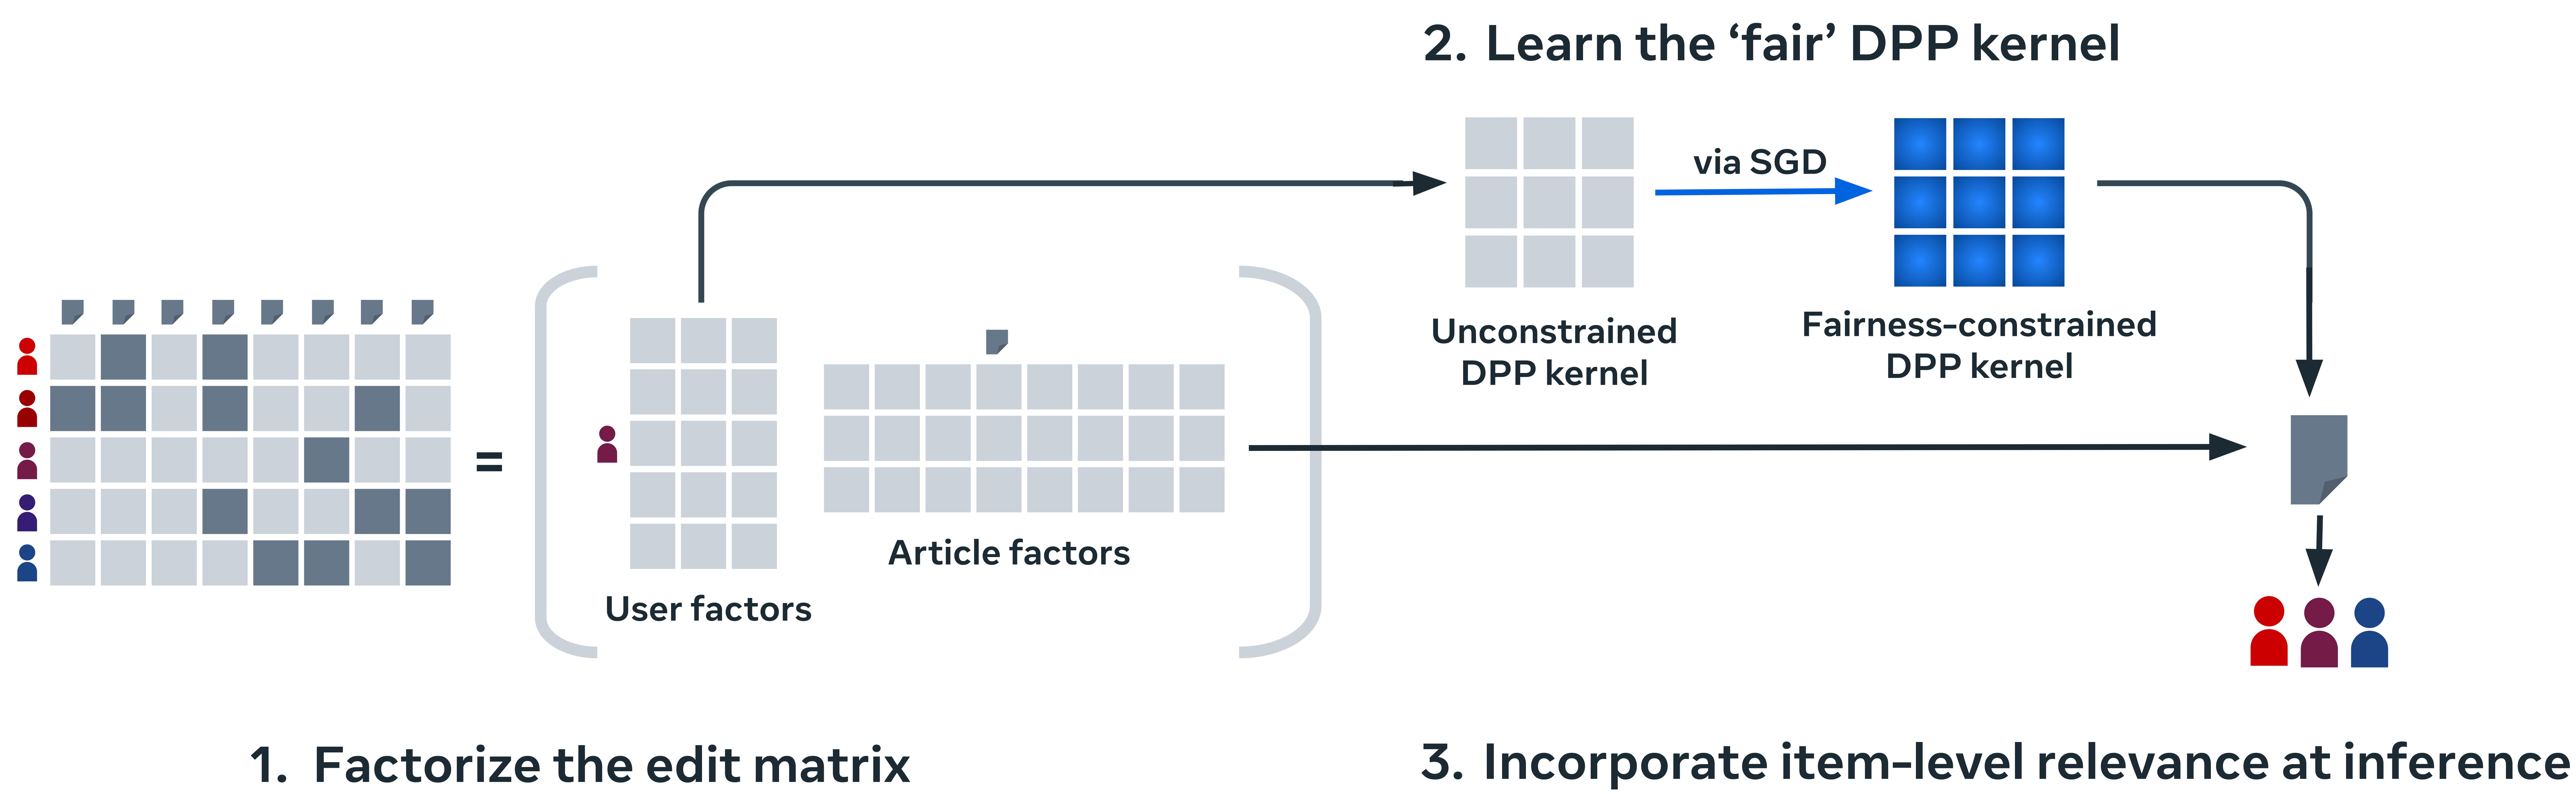
2. Learn the 'fair' DPP kernel
via SGD
Unconstrained DPP kernel
Fairness-constrained DPP kernel
=
User factors
Article factors
1. Factorize the edit matrix
3. Incorporate item-level relevance at inference