Calibrated Recommendations for Users with Decaying Attention

Jon Kleinberg[1] Emily Ryu[1] Éva Tardos[1]

[1]Cornell University

Diversity in recommendations

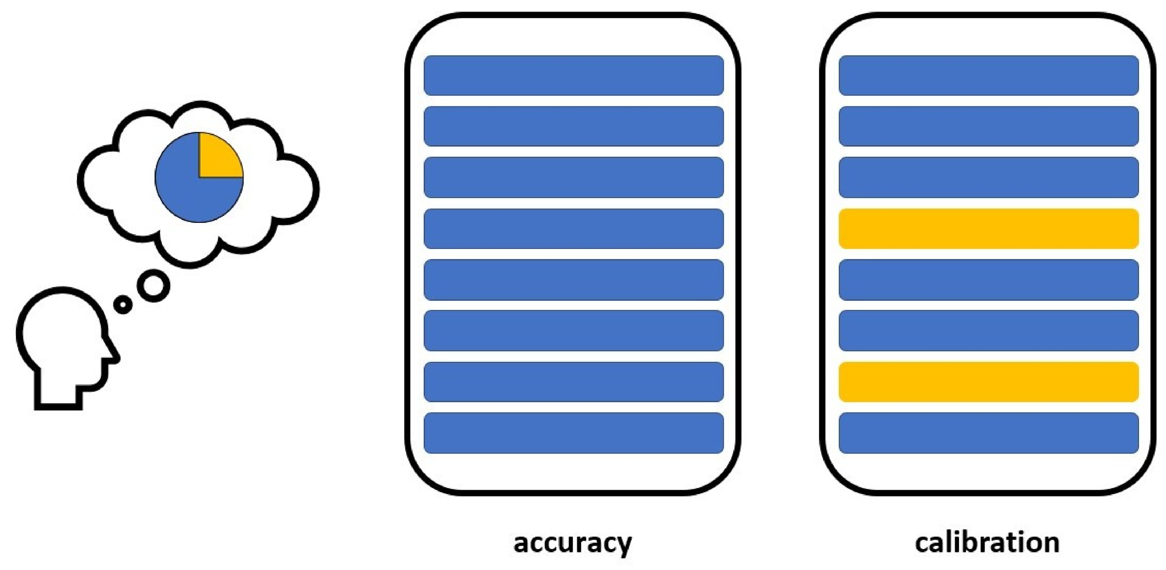

- Modeling recommendations: items j are **distributions** q_j **over genres**
- **Calibration** as a goal: user has a target genre distribution p; average of recommended items should (approximately) **match user's distribution**

Order matters

- Traditionally formulated as *subset selection* problem and approached using tools such as *submodularity*
- Key assumption: all items contribute equally to objective
- **Not true** in reality – *e.g.* decreasing attention, patience, trust moving down a ranked list means that **higher items contribute more**
- **Not modeled** by previous frameworks that attempt to generalize submodularity to sequences
- A more realistic model with **position-based weights**: a recommendation list $\pi = \pi_1\pi_2 \dots \pi_k$ induces genre distribution

$$q(\pi)(g) = \sum_{j=1}^{k} w_j \cdot q_{\pi_j}(g)$$

Main results

1. **Distributional genres:** reduction to **submodular optimization over a matroid** obtains a $(1-1/e)$-approximation to the optimal calibration
2. **Discrete genres:** bin-packing analysis of simple **greedy algorithm** to establish approximation guarantee of $2/3$
3. Introduction of algorithmic framework of **ordered submodularity** and divergence-based **overlap measures** of calibration

Ordered submodularity and overlap measures

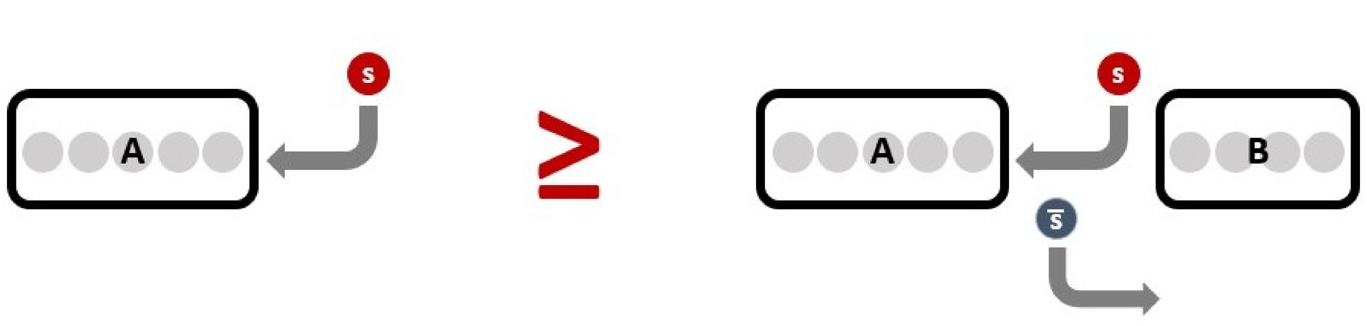

A sequence function f is ***ordered-submodular*** if for all sequences A and B, the following property holds for all elements s and $\bar{s}$:

$$f(A||s) - f(A) \geq f(A||s||B) - f(A||\bar{s}||B).$$

Remark: Generalizes monotonicity and submodularity for set functions

Given a target distribution p and a recommendation list π, an ***overlap measure*** $G(\pi) \coloneqq G(p, q(\pi))$ captures the similarity of p to the induced $q(\pi)$. It satisfies: $G(p,q) \geq 0$, $\arg\max_q G(p,q) = p$.

We study overlap measures that come from f-divergences, such as the squared Hellinger distance:

$$G_{H^2}(p,q) = \sum_x \sqrt{p(x) \cdot q(x)}$$

Distributional genres: continuous greedy algorithm

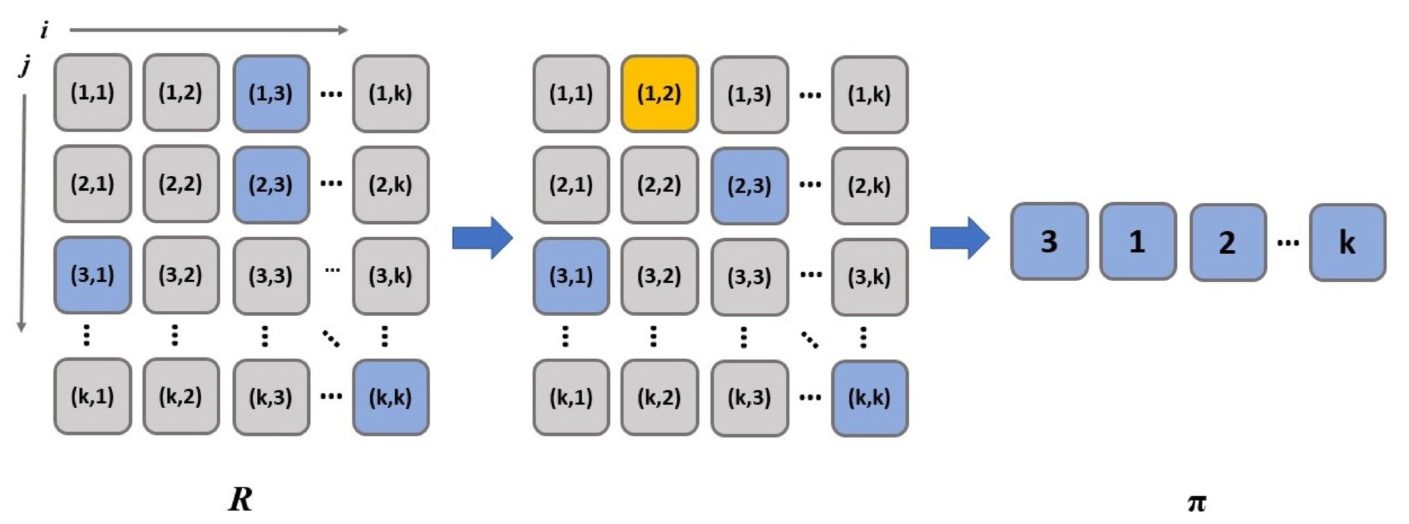

Key ideas (extension of Asadpour et al., EC '22)

- Laminar matroid on "exploded" ground set $V = \{i_j\}$ (pairs of item i, slot j)
- Independent sets in $\mathcal{I}$ have at most ℓ items in the first ℓ slots ($\forall \ell \in [k]$)
- From overlap measure G, construct monotone submodular set function F_G on V with a translation from every independent set R to a sequence π s.t. $G(\pi) \geq F_G(R)$
- Find $\bar{R}$ via continuous greedy such that $F_G(\bar{R}) \geq (1-1/e)\max_{R \in \mathcal{I}} F_G(R) \geq (1-1/e)\max G(\pi)$, then convert to $\bar{\pi} \implies G(\bar{\pi}) \geq (1-1/e)\max G(\pi)$

Discrete genres: bin-packing analysis of simple greedy algorithm

Interpret as assigning *slots* to *genres*; use squared Hellinger-based overlap:

$$f(S) = \sum_{\text{genres } g} \sqrt{p(g)} \sqrt{\sum_{i \in [k]: s_i = g} w_i}$$

Analyze each iteration of greedy algorithm: what if greedy assigns w_i to g' instead of optimal g^*?

(1) "Correct the mistake" if possible: move $z \geq \frac{w_i}{2}$ from g' to g^*
(2) If not possible, mistake must have been "not so bad"

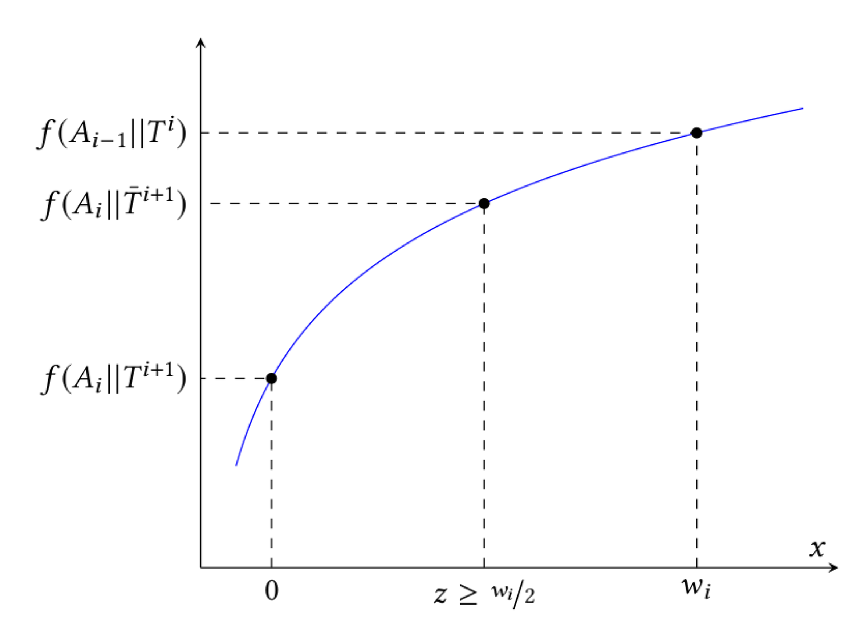

Inductive invariant $f(A_i||\bar{T}^{i+1}) \geq f(A_{i-1}||T^i) - \frac{1}{2}(f(A_i) - f(A_{i-1}))$

$$\implies f(A) \geq \frac{2}{3} OPT$$

Conclusions

- First performance guarantees for near-optimal calibration of recommendation *lists*, with ordering effects due to attention decay
- In the discrete model, $2/3$-approximation surpasses the $(1-1/e)$ barrier imposed by NP-hardness of general submodular optimization
- Future directions:
  1. Discrete model: Is $2/3$ tight for the greedy algorithm?
  2. Improved approximation algorithms for either model
  3. Parametrization of special "easy" cases
  4. Broader characterization of general overlap measures

References

[1] Jon Kleinberg, Emily Ryu, and Éva Tardos. Calibrated recommendations for users with decaying attention, 2023. URL `https://arxiv.org/abs/2302.03239`.