On A Mallows-type Model For (Ranked) Choices

Yuxuan Tang
Institute of Operations Research and Analytics,
National University of Singapore

Joint Work with Yifan Feng, NUS Business School

A Motivating Example

A company designs

What's the preference over these n items?

Ask for top-k feedback

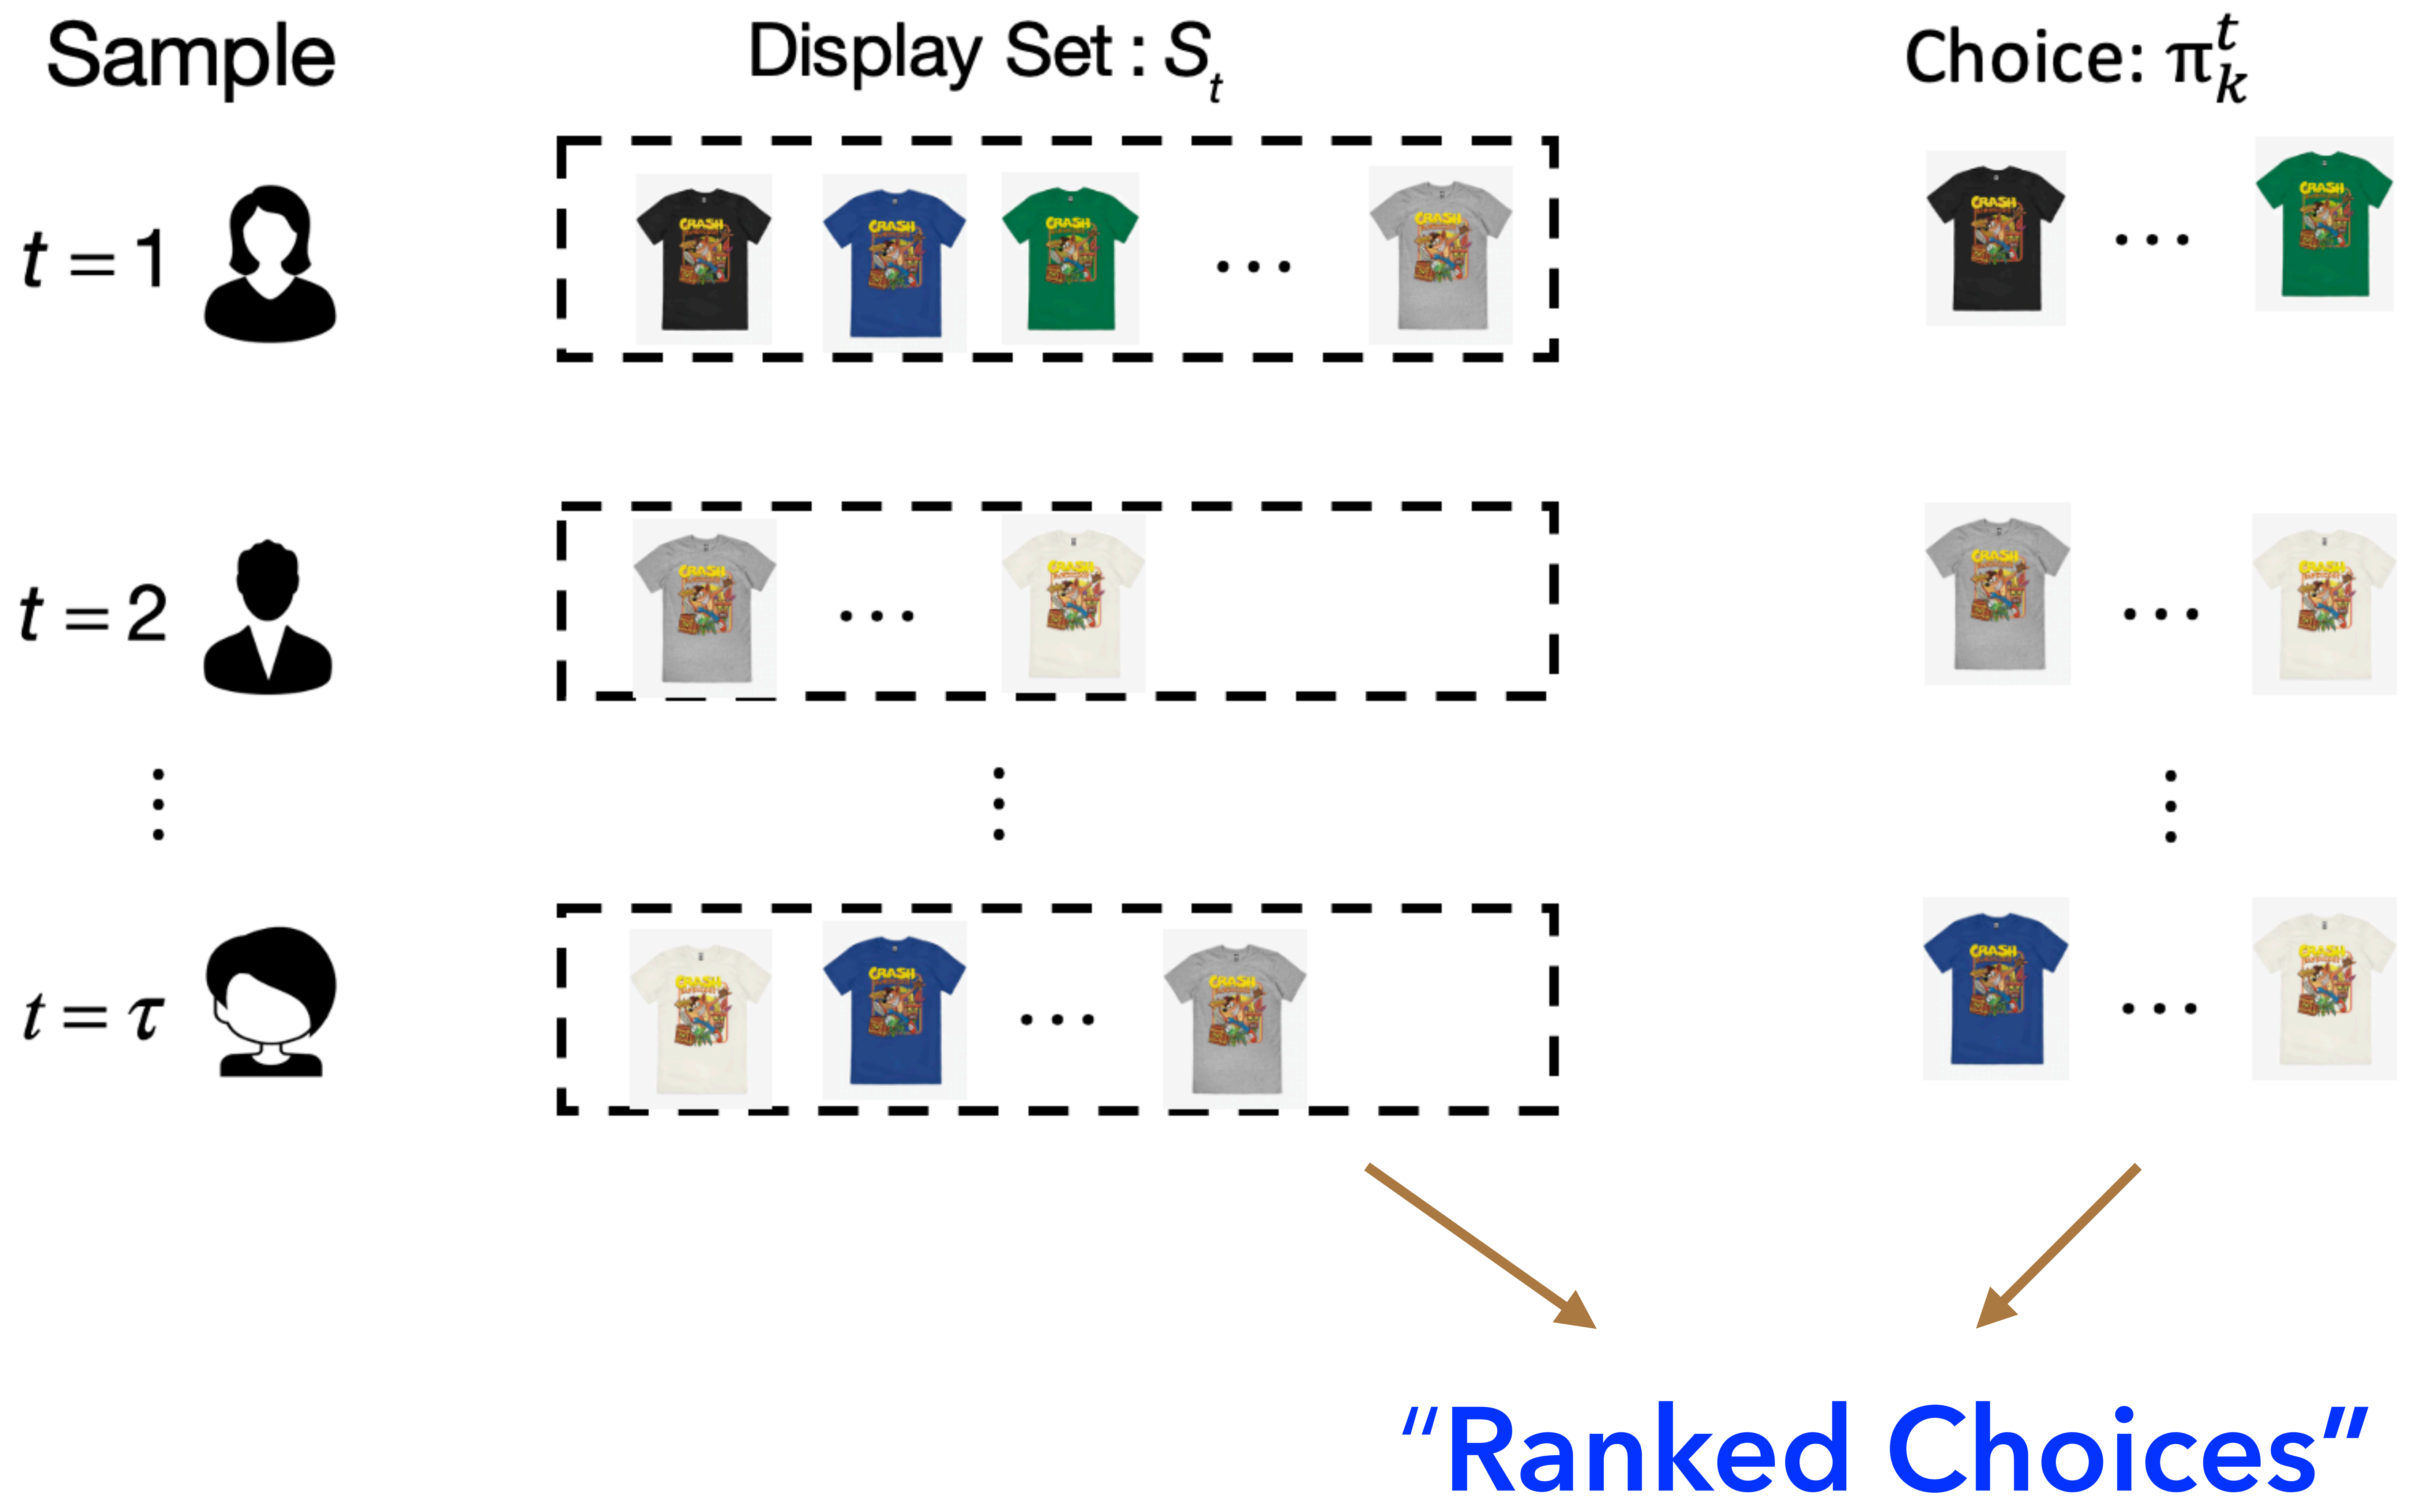

What is the Mallows Model?

Probability distribution of rankings/permutations

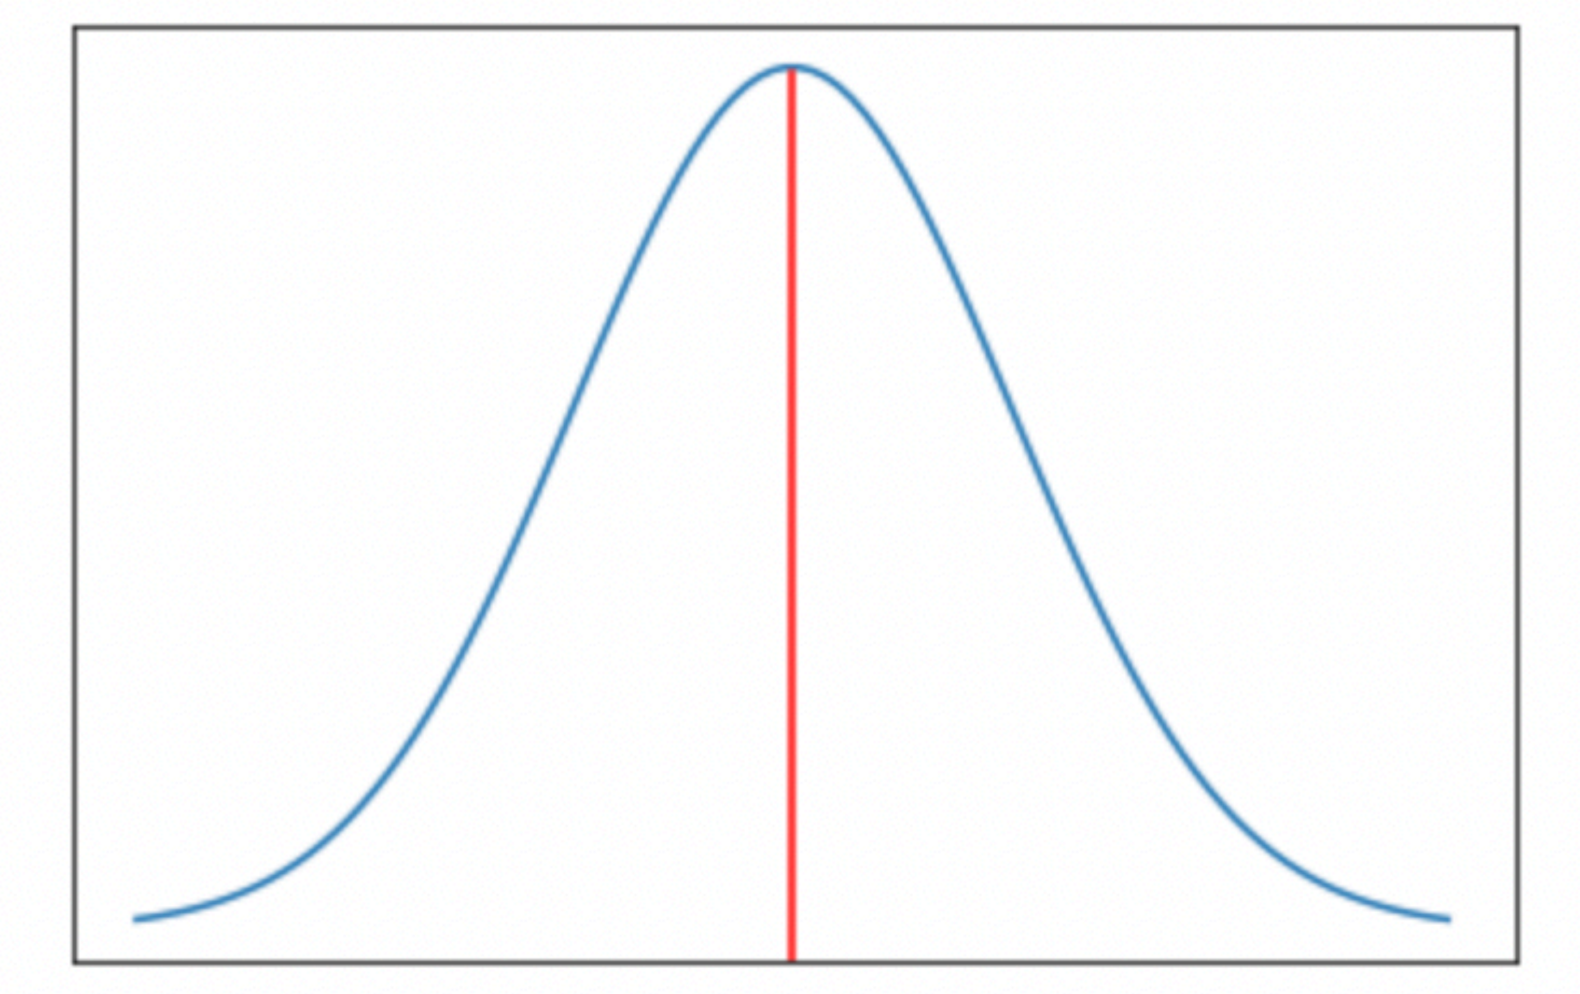

π^*: Central Ranking[1]

$$\lambda(\pi) = \frac{q^{d_K(\pi^*,\pi)}}{\sum_{\pi'} q^{d_K(\pi^*,\pi')}}$$

q: Dispersion parameter

Kendall's Tau Distance[1]:

$$d_K(\pi) = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} I\{\pi(i) > \pi(j)\}$$

[1] Assume $\pi^* = e$.

Apply Mallows model to ranked choices?

- For top-1 feedback (choice data), Désir, Goyal, Jagabathula and Segevd (Operations Research, 2021)[1] use mixture of Mallows model
  - Choice probabilities can be evaluated in $O(n^2 \log n)$.
  - Estimation is still hard, and they propose Mallows Smoothing to conduct estimation.
    - Numerically show better prediction power than sparse ranking-based choice model.

Theorem Even if all display sets S with |S|>2 are displayed infinite times, the estimator from the Mallows Smoothing heuristic is not consistent.

[1] Désir, A., Goyal, V., Jagabathula, S., & Segev, D. (2021). Mallows-smoothed distribution over rankings approach for modeling choice. *Operations Research*, *69*(4), 1206-1227.

Main Contributions

- **We develop a novel Mallows-type model and**
  - Characterize *simple closed-form* (ranked) choice probabilities,
  - *Learn the parameters* with guaranteed *consistency* relatively *easy*,
  - Learn in a *mixture* setting by Expectation Maximization (EM) algorithm,
  - Efficiently *sample* out a *top-k* list.

Our RMJ-based Ranking Model

- **A new distance function: Reverse Major Index (RMJ)[1]**

$$d_R(\pi) = \sum_{i=1}^{n-1} I\{\pi(i) > \pi(i+1)\} \cdot (n-i)$$

  - Adjacent disagreement
  - Linear decreasing weight

(4,2,1,3)		
Distance Function	RMJ	Kendall's Tau
Number of Disagreement	2 adjacent	4 pairwise
Disagreements	(4,2), (2,1)	(4,2), (4,1), (4,3), (2,1)
Weights	3,2	1,1,1,1
Distance	3+2=5	1+1+1+1=4

[1] Assume $\pi^* = e$.

Simple Closed-form Choice Probability

Theorem Given a display set $S = \{x_1, x_2, \ldots, x_M\}$ be such that $x_1 < x_2 < \cdots < x_M$. Then

$$\mathbb{P}(x_i \mid S) = \frac{q^{i-1}}{1 + q + \ldots + q^{M-1}}$$

- Relative ranking within display set matters.
- Choice probabilities decay exponentially fast.

Consistency

Parameter Learning

- Given historical data $H_T = (S_1, x_1, \ldots, S_T, x_T)$
- MLE formulation:

$$\sum_{t=1}^{T} \log \left(\frac{1-q}{1-q^{|S_t|}} \right) + \log q \sum_{(i,j):i \neq j} I\{j \succ_\pi i\} \cdot w_{ij}$$

where

- $w_{ij} := \sum_{t=1}^{T} I\{\{i,j\} \subseteq S_t \text{ and } x_t = i\}$
- Intuitively, a positive $w_{ij} - w_{ji}$ is an indication that item i should be preferred to item j.

$$\min \sum_{(i,j):i \neq j} w_{ij} x_{ji}$$

s.t.

$$x_{ij} + x_{ji} = 1 \quad \forall 1 \leq i,j \leq n$$
$$x_{ij} + x_{jr} + x_{ri} \leq 2 \quad \forall 1 \leq i,j,r \leq n$$
$$x_{ij} \in \{0,1\} \quad \forall 1 \leq i,j \leq n$$

$x_{ij} = 1$ means $i \succ_\pi j$.

- **Integer Programming Reformulation**: Well-studied weighted feedback arc set problem on tournaments.
- Given π and set $\alpha = -\log q$, MLE is a concave function of α.

Ranked Choice Probability

Theorem Given a display set S and a π_k such that $R(\pi_k) \subseteq S$, we have

$$\mathbb{P}(\pi_k \mid S) = q^{d_S(\pi_k)+L_S(\pi_k)} \cdot \frac{\psi(|S|-k, q)}{\psi(|S|, q)}$$

where $d_S(\pi_k) := \sum_{i=1}^{k-1} I\{\pi_k(i) > \pi_k(i+1)\} \cdot (|S|-i)$, $L_S(\pi_k) := |\{x \in R^c(\pi_k) \cap S : x < \pi_k(k)\}|$.

Learning on Ranked Choices

- Given historical data $H_T = (S_1, \pi_k^1, \ldots, S_T, \pi_k^T)$, where $\pi_k^t = (x_1^t, \ldots, x_k^t)$.

Theorem The MLE for the central ranking can be obtained from the same integer program with a generalized definition of w_{ij} below

$$w_{ij} = \sum_{t=1}^{T} \left[\sum_{h=1}^{k-1} (|S_t| - h) \cdot I\{x_h^t = i, x_{h+1}^t = j\} + I\{x_k^t = i\} \cdot I\Big\{\{i,j\} \subseteq S_t \backslash \{x_1^t, \ldots, x_{k-1}^t\}\Big\} \right]$$

Nice Properties

Lemma 1 *(Probability distribution of top-k rankings, $k \geq 1$)*

$$\lambda(\pi_k) = q^{d(\pi_k)+L(\pi_k)} \cdot \frac{\psi(n-k, q)}{\psi(n, q)}$$

Lemma 2 *(Sampling of Next Position)*

Given π_k such that $\pi_k(k) = z$, the conditional probability for the $(k+1)$-positioned item is

$$\mathbb{P}\left(\pi_{k+1} = \pi_k \oplus y \mid \pi_k\right) = \frac{q^{h(y|z)-1}}{1 + q + \cdots + q^{n-k-1}}$$

where $h(y \mid z) = \begin{cases} \sum_{x \in R^c(\pi_k)} \mathbb{I}\{z < x \leq y\} & \text{if } y > z \\ n - k - \sum_{x \in R^c(\pi_k)} \mathbb{I}\{y < x < z\} & \text{if } y < z \end{cases}.$

Efficient Sampling: $O(nk)$

- An example of 3 items, {1,2,3}.

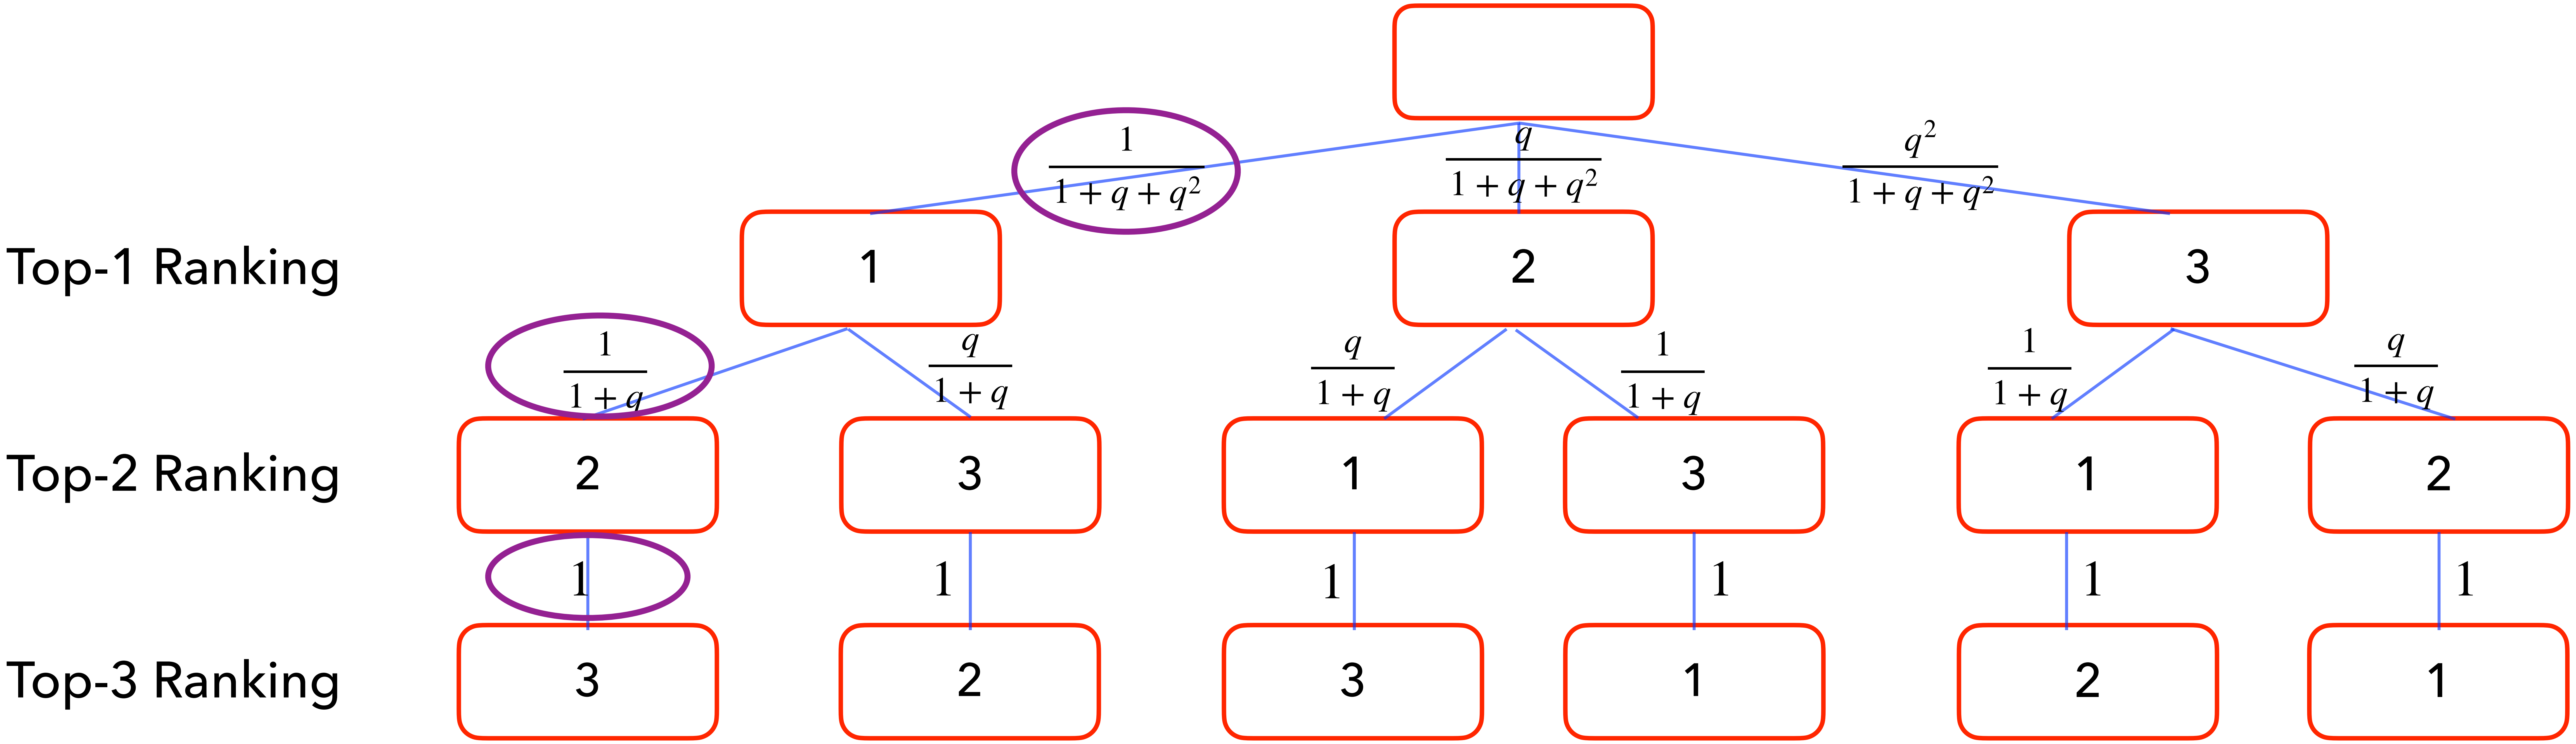

$$\lambda((1,2,3)) = \frac{1}{1+q+q^2} \times \frac{1}{1+q} \times 1$$

Experiments

- Experiment 1: Prediction Accuracy
- Experiment 2: Robustness Check
- Experiment 3: Efficient Estimation

Public Data Sets

- Two public ranking data sets about sushi preference[1]
  - 5000 complete rankings over 10 kinds of sushi.
  - 5000 top-10 rankings over 100 kinds of sushi.

[1] Source: Toshihiro Kamishima, Sushi Data, downloaded from PrefLib.org

Experiment 1: Prediction on Top-1 Choice

- Compare with Mallows model and Plackett-Luce model.
- Fit into mixture models.

Figure 1: Comparison of Explanation and Prediction Power for Top-1 choice

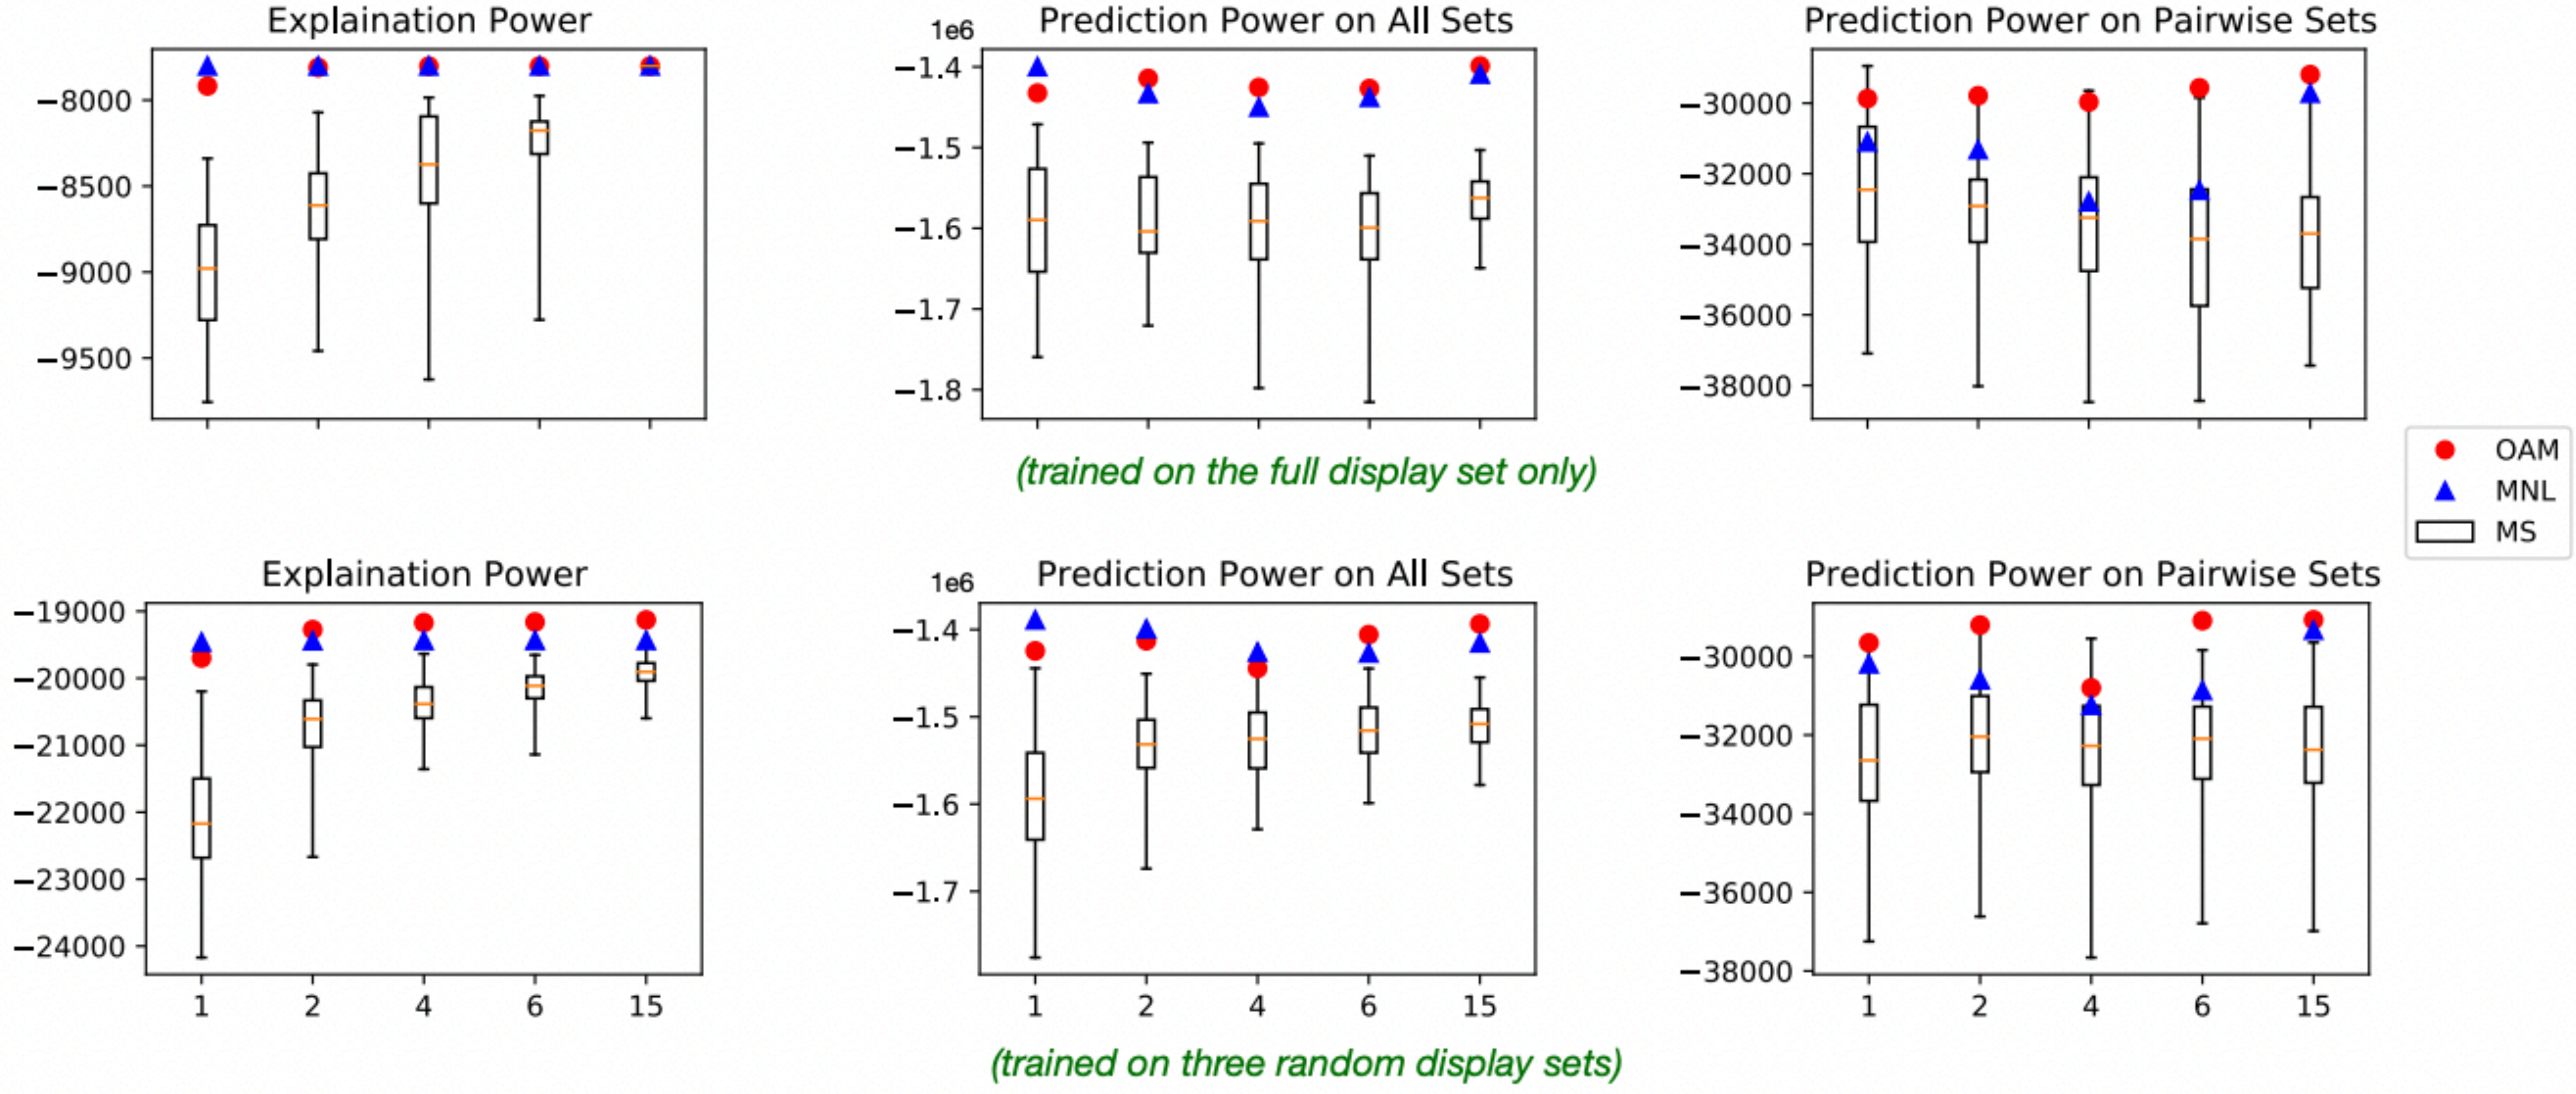

In each panel, the x-axis represents the number of clusters, and the y-axis represents the log-likelihood metric.

Experiment 2: Robustness Check on Top-k Choice

- The 10 Sushi data set. We conduct estimation on top-1, top-2 and top-3 ranked choices.
  - Compare with Borda Count[1] and Simple Count[2]
  - Discrepancy: Average pairwise Kendall' Tau distance
- **Setting 1**: Display Sets: {set with $size \geq k$} (Balanced Display)

	OAM	Borda Count	Simple Count
Top-3	(8,5,6,3,2,1,4,9,7,10)	(8,3,6,1,2,5,4,9,7,10)	(8,3,1,6,2,5,4,9,7,10)
Top-2	(8,5,6,3,2,1,4,9,7,10)	(8,3,5,6,1,2,4,9,7,10)	(8,3,6,1,5,2,4,9,7,10)
Top-1	(8,5,6,3,2,1,4,9,7,10)	(8,5,6,3,2,1,4,9,7,10)	(8,5,3,6,2,1,4,9,7,10)
Discrepancy	0	4	4

- **Setting 2**: Display Sets: {[10], {7,9,10}} (Unbalanced Display)

	OAM	Borda Count	Simple Count
Top-3	(8,5,3,2,6,1,4,9,7,10)	(8,3,5,6,9,2,1,7,4,10)	(9,7,10,8,3,5,6,2,1,4)
Top-2	(8,5,6,3,2,1,4,9,7,10)	(8,5,9,6,3,2,1,7,4,10)	(9,7,8,10,5,6,3,2,1,4)
Top-1	(8,5,2,6,1,3,4,9,7,10)	(8,5,9,7,2,6,1,3,4,10)	(9,7,8,5,2,6,10,1,3,4)
Discrepancy	2.67	7.33	6

1. Borda Count: the score of an item is linear decreasing with its rank.
2. Simple Count: simply count the occurrences of each sushi in the top-k choice.

Experiment 3: Efficiency When n is large

- Training data is 5000 top-10 choices out of 100 types of sushi.

- We use LP relaxation to speed up the IP.

- We bootstrap 10 times (each time drawing 10000 samples) and record the running times and optimality gaps[1].

	Model Building Time (min)	Model Solving Time (min)	Optimality Gap
Average	21.10	4.20	1.47%
Max	21.19	4.50	1.79%

$< 5min$ $< 2\%$

[1] Feng, Y., Caldentey, R., & Ryan, C. T. (2022). Robust learning of consumer preferences. *Operations Research*, *70*(2), 918-962.

On A Mallows-type Model For (Ranked) Choices

Yifan Feng
Department of Analytics and Operations
NUS Business School
yifan.feng@nus.edu.sg

Yuxuan Tang
Institute of Operations Research and Analytics
National University of Singapore
yuxuan.tang@u.nus.edu

Abstract

We consider a preference learning setting where every participant chooses an ordered list of k most preferred items among a displayed set of candidates. (The set can be different for every participant.) We identify a distance-based ranking model for the population's preferences and their (ranked) choice behavior. The ranking model resembles the Mallows model but uses a new distance function called Reverse Major Index (RMJ). We find that despite the need to sum over all permutations, the RMJ-based ranking distribution aggregates into (ranked) choice probabilities with simple closed-form expression. We develop effective methods to estimate the model parameters and showcase their generalization power using real data, especially when there is a limited variety of display sets.