Ranking from sparse comparison data with non-uniformity

Yuepeng Yang

Department of Statistics, University of Chicago

Jan 6-7, 2025 @ CUHKSZ

- Data appears as comparisons:
  evaluating "Is A better than B" is easier than "how good is A"
- Goal: comparison data $\rightarrow$ item ranking

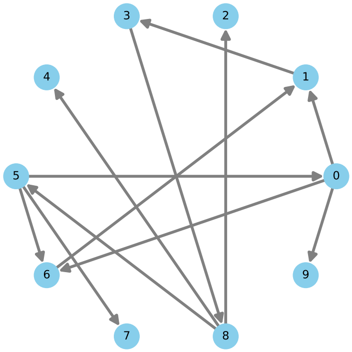

Chess player rating

Educational testing

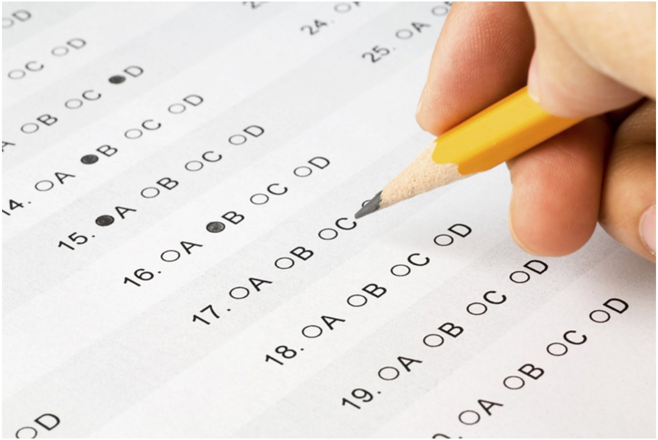

Reinforcement learning with human feedback (RLHF)

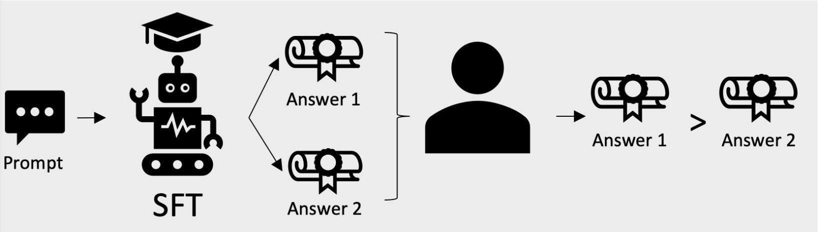

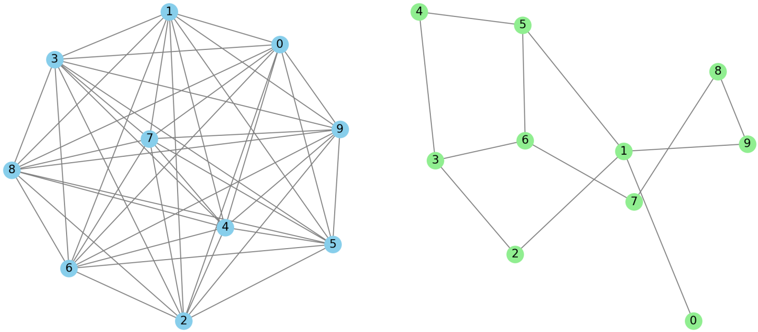

Dense vs sparse comparison graph

Can we do estimation for sparsely connected comparison graph?

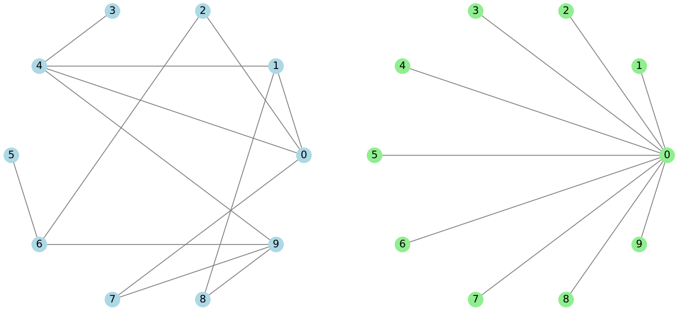

Uniform vs non-uniform comparison graph

Can we do estimation with non-uniformity in sampling?

- Top-K ranking in Bradley-Terry-Luce (BTL) model with uniform sampling
- Towards non-uniformity: result for general comparison graph
- Case 1: Heterogeneous sampling probability
- Case 2: Imbalanced bipartite structure

Top-K ranking in Bradley-Terry-Luce (BTL) model

- Latent scores $\boldsymbol{\theta}^\star = [\theta_1^\star, \ldots, \theta_n^\star]$, with $\Delta_K = \theta_{(K)} - \theta_{(K+1)}$
- Noisy binary pairwise comparisons from BTL (logistic) model:

$$\mathbb{P}\,[\, i \succ j \,] = \mathrm{sigmoid}(\theta_i^\star - \theta_j^\star) = \frac{e^{\theta_i^\star}}{e^{\theta_i^\star} + e^{\theta_j^\star}}$$

- Goal: identify the top-K items by latent score

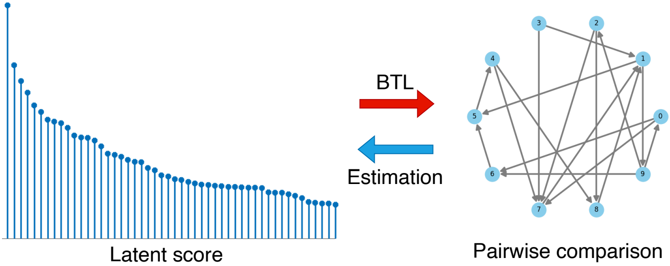

- Uniform sampling: each pair (i, j) observed with i.i.d. prob p
- Typical approach: rank by estimated the latent scores $\boldsymbol{\theta}^{\star}$
- Need to control ℓ_{∞} error to be at most $\Delta_K/2$ so that

$$\widehat{\theta}_{(K)} > \theta^{\star}_{(K)} - \Delta/2 \geq \theta^{\star}_{(K+1)} + \Delta/2 > \widehat{\theta}_{(K+1)}$$

MLE is optimal for uniform sampling

Theorem (Chen, Fan, Ma, Wang, AOS'19)

Assuming $np \gtrsim \log n$, *regularized MLE achieves error rate and*

$$\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^{\star}\|_{\infty} \lesssim \sqrt{\frac{\log n}{np}}$$

Optimal up to constant factors.

Towards a more general regmie

The ℓ_∞ analysis relies on

- independence between edges
- uniform sampling probability

Need a more general result to relax these assumptions

- $\boldsymbol{D}$: diagonal degree matrix $\boldsymbol{A}$: the adjacency matrix
- $\boldsymbol{L} := \boldsymbol{D} - \boldsymbol{A}$: (weighted) graph Laplacian
- $\lambda_{n-1}(\cdot)$: $(n-1)$-th largest eigenvalues (algebraic connectivity)
- $d_{\max}$: maximum degree

Theorem (Yang, Chen, Orecchia, Ma, COLT '24)

*Suppose $\lambda_{n-1}^5(\boldsymbol{L})/d_{\max}^4 \gtrsim \log^2(n)$, then weighted MLE on **general comparison graph** satisfies*

$$\left\| \widehat{\boldsymbol{\theta}}_{\mathrm{WMLE}} - \boldsymbol{\theta}^\star \right\|_\infty \lesssim \sqrt{\frac{\log n}{\lambda_{n-1}(\boldsymbol{L})}}$$

Good spectrum is enough for good estimation

- Non-uniform sampling: each pair (i,j) is sampled with unknown probability $p_{ij} \geq p$
- Uniform sampling has good spectrum:

$$d_{\max} \lesssim np, \qquad \lambda_{n-1}(\boldsymbol{L}) \gtrsim np$$

- Non-uniform sampling can have bad spectrum

$$d_{\max} \gtrsim n, \qquad \lambda_{n-1}(\boldsymbol{L}) \lesssim np$$

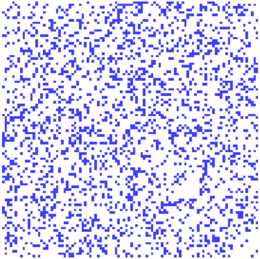

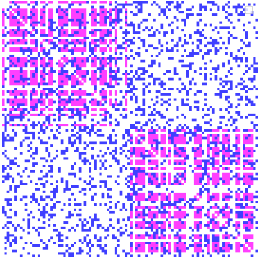

Reweighting to recover good spectrum

- Observation: comparison graph for non-uniform sampling always has a hidden Erdős–Rényi subgraph
- Select weights $\boldsymbol{W}$ by solving the semidefinite program

$$\max_{\boldsymbol{W}} \lambda_{n-1}(\boldsymbol{L}) \quad \text{s.t.} \quad d_{\max} \leq 2np$$

- Weight 1 on the Erdős–Rényi subgraph is a feasible solution
- The SDP always returns a weighted graph with a spectrum that is at least as good as the Erdős–Rényi subgraph

Result for heterogeneous sampling probability

The weighted comparison graph...

- has good spectrum
- has edge-dependent weights

... so we can invoke our result for general comparison graph to get

Theorem (Yang, Chen, Orecchia, Ma, COLT '24)

Suppose $np \gtrsim \log^3 n$. Weighted MLE with the SDP-based reweighting achieves error rate

$$\left\| \widehat{\boldsymbol{\theta}}_{\mathrm{WMLE}} - \boldsymbol{\theta}^{\star} \right\|_{\infty} \lesssim \sqrt{\frac{\log n}{np}}$$

Non-uniformity: imbalanced bipartite structure

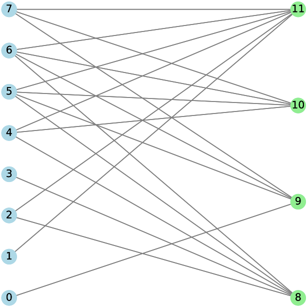

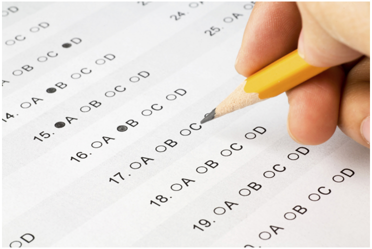

Rasch model

- m items with latent scores $\boldsymbol{\theta}^\star = [\theta_1^\star, \ldots, \theta_m^\star]$
- n users with latent scores $\boldsymbol{\zeta}^\star = [\zeta_1^\star, \ldots, \zeta_n^\star]$
- Each user-item pair (t, i) is sampled with i.i.d. probability p
- Observe outcome via logistic model

$$\mathsf{P}[\text{ user } t \prec \text{ item } i\ ] = \mathrm{sigmoid}(\theta_i^\star - \zeta_t^\star) = \frac{e^{\theta_i^\star}}{e^{\zeta_t^\star} + e^{\theta_i^\star}}$$

- Goal: estimate and identify top-K item parameters

Challenge for Rasch model with imbalanced groups

- When m, p are fixed and $n \to \infty$, Joint MLE on $(\boldsymbol{\theta}^\star, \boldsymbol{\zeta}^\star)$ is inconsistent for $\boldsymbol{\theta}^\star$ estimation

$$\limsup_{n \to \infty} \|\widehat{\boldsymbol{\theta}}_{\mathrm{JMLE}} - \boldsymbol{\theta}^\star\|_2 > 0$$

- Information gap: np samples / item vs mp samples / user
- Bad spectrum:

$$d_{\max} \gtrsim np \quad \text{and} \quad \lambda_{n+m-1}(\boldsymbol{L}) \lesssim mp$$

- Re-tabulate user-item comparisons to item-item comparisons

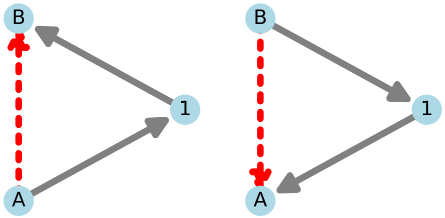

- Reduction to BTL (logistic) model

$$\frac{\mathbb{P}[i \succ t \succ j]}{\mathbb{P}[i \succ t \succ j \ \text{ or } \ i \prec t \prec j]} = \frac{e^{\theta_i^*}}{e^{\theta_i^*} + e^{\theta_j^*}}$$

- Reduce problem size by estimating item parameters $\boldsymbol{\theta}^\star$ directly

The resulting item-item comparison graph

- has dependent edges and heterogeneous sampling probability
- has good spectrum: $\lambda_{m-1}(\boldsymbol{L}) \gtrsim np$ and $d_{\max} \lesssim np$

So our results for general comparison graph comes in handy

Theorem (Yang and Ma, '24)

When $np \gtrsim \log^3 n$ and $mp \geq 2$. RP-MLE achieves error rate

$$\|\boldsymbol{\theta}_{\text{RP-MLE}} - \boldsymbol{\theta}^\star\|_\infty \lesssim \sqrt{\frac{\log n}{np}}$$

We study ranking with pairwise comparisons in sparse and non-uniform sampling regime

- Good spectrum = good estimation
- Weighted MLE for non-uniform sampling
- Random pairing MLE for Rasch model with imbalanced groups

Papers:

- Top-K ranking with a monotone adversary. COLT 2024.
- Random pairing MLE for estimation of item parameters in Rasch model. arXiv:2406.13989

Bonus: Analysis through preconditioned gradient descent

- Consider the *preconditioned* gradient descent: let $\boldsymbol{\theta}^0 = \boldsymbol{\theta}^\star$, run

$$\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta \nabla^2 \mathcal{L}_w(\boldsymbol{\theta}^\star)^\dagger \nabla \mathcal{L}_w(\boldsymbol{\theta}^t)$$

  and $\boldsymbol{\theta}^t \to \widehat{\boldsymbol{\theta}}_{\mathsf{WMLE}}$.

- Let $\boldsymbol{\delta}^t := \boldsymbol{\theta}^t - \boldsymbol{\theta}^\star$,

$$\boldsymbol{\delta}^{t+1} = (1-\eta)\,\boldsymbol{\delta}^t - \eta\left(\nabla^2 \mathcal{L}_w(\boldsymbol{\theta}^\star)^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}} - \nabla^2 \mathcal{L}_w(\boldsymbol{\theta}^\star)^\dagger \boldsymbol{r}^t\right)$$

Bonus: Analysis through preconditioned gradient descent

- Consider the *preconditioned* gradient descent: let $\boldsymbol{\theta}^0 = \boldsymbol{\theta}^\star$, run

$$\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta \nabla^2 \mathcal{L}_w(\boldsymbol{\theta}^\star)^\dagger \nabla \mathcal{L}_w(\boldsymbol{\theta}^t)$$

and $\boldsymbol{\theta}^t \to \widehat{\boldsymbol{\theta}}_{\mathsf{WMLE}}$.

- Let $\boldsymbol{\delta}^t := \boldsymbol{\theta}^t - \boldsymbol{\theta}^\star$,

$$\boldsymbol{\delta}^{t+1} = \underbrace{(1-\eta)\,\boldsymbol{\delta}^t}_{\substack{\text{coordinate-wise} \\ \text{error contraction}}} - \eta \left(\nabla^2 \mathcal{L}_w(\boldsymbol{\theta}^\star)^\dagger \boldsymbol{B} \widehat{\boldsymbol{\epsilon}} - \nabla^2 \mathcal{L}_w(\boldsymbol{\theta}^\star)^\dagger \boldsymbol{r}^t \right)$$

Bonus: Analysis through preconditioned gradient descent

- Consider the *preconditioned* gradient descent: let $\boldsymbol{\theta}^0 = \boldsymbol{\theta}^\star$, run

$$\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \eta \nabla^2 \mathcal{L}_w(\boldsymbol{\theta}^\star)^\dagger \nabla \mathcal{L}_w(\boldsymbol{\theta}^t)$$

  and $\boldsymbol{\theta}^t \to \widehat{\boldsymbol{\theta}}_{\mathsf{WMLE}}$.

- Let $\boldsymbol{\delta}^t := \boldsymbol{\theta}^t - \boldsymbol{\theta}^\star$,

$$\boldsymbol{\delta}^{t+1} = (1-\eta)\,\boldsymbol{\delta}^t - \eta\Big(\underbrace{\nabla^2 \mathcal{L}_w(\boldsymbol{\theta}^\star)^\dagger \boldsymbol{B}\widehat{\boldsymbol{\epsilon}}}_{\substack{\text{first-order}\\ \text{noise}}} - \underbrace{\nabla^2 \mathcal{L}_w(\boldsymbol{\theta}^\star)^\dagger \boldsymbol{r}^t}_{\substack{\text{second-order}\\ \text{residual}}}\Big)$$