Generalized Results for the Existence and Consistency of the MLE in the BTL Model

Heejong Bong and Alessandro Rinaldo

Department of Statistics and Data Sciences
Carnegie Mellon University
Pittsburgh, PA 15213

Partial ordering observation datasets by pairwise comparisons are getting more widely available

Global ranking

	Team
1	Heat
2	Celtics
3	Bucks
4	76ers
5	Raptors
6	Bulls
7	Nets
8	Hawks
9	Cavaliers
10	Hornets
⋮	

[1]

[1] National basketball Association (2022). NBA. *Google*. https://www.google.com/search?q=nba&source=lmns&bih=746&biw=1536&hl=en&sa=X&ved=2ahUKEwilpYjQjbr4AhXdA50JHflUDKUQ_AUoAHoECAEQAA

Partial ordering observation datasets by pairwise comparisons are getting more widely available

Pairwise comparison

⋮

Pacers	126	Wizards	108
Nets	134 ◂	Hornets	124 ◂
Pistons	106	Heat	111
76ers	118 ◂	Magic	125 ◂
Raptors	94	Celtics	139 ◂
Knicks	105 ◂	Grizzlies	110

⋮

Global ranking

	Team
1	Heat
2	Celtics
3	Bucks
4	76ers
5	Raptors
6	Bulls
7	Nets
8	Hawks
9	Cavaliers
10	Hornets

⋮

[1]

[1] National basketball Association (2022). NBA. *Google*. https://www.google.com/search?q=nba&source=lmns&bih=746&biw=1536&hl=en&sa=X&ved=2ahUKEwilpYjQjbr4AhXdA50JHflUDKUQ_AUoAHoECAEQAA

The Maximum Likelihood Estimator (MLE) of the Bradley-Terry-Luce (BTL) model is used for ranking

Data: n pairwise comparison results

$$i_k \text{ defeated } j_k \text{ at comparison } k$$

The Maximum Likelihood Estimator (MLE) of the Bradley-Terry-Luce (BTL) model is used for ranking

Data: n pairwise comparison results

$$i_k \text{ defeated } j_k \text{ at comparison } k$$

BTL model: $w^* = (w_1^*, \ldots, w_d^*)$: unknown quality scores of d competitors so that

$$\mathbb{P}[i \text{ defeats } j] = \frac{1}{1+e^{w_j^* - w_i^*}}$$

The Maximum Likelihood Estimator (MLE) of the Bradley-Terry-Luce (BTL) model is used for ranking

Data: n pairwise comparison results

i_k defeated j_k at comparison k

BTL model: $w^* = (w_1^*, \ldots, w_d^*)$: unknown quality scores of d competitors so that

$$\mathbb{P}[i \text{ defeats } j] = \frac{1}{1+e^{w_j^* - w_i^*}}$$

MLE: $\widehat{w} = (\widehat{w}_1, \ldots, \widehat{w}_d)$ satisfying

$$\operatorname{argmin}_{w \in \mathbb{R}^d} -l(w) \quad \text{w.r.t.} \quad \sum_{k=1}^{n} w_i = 0,$$

where $l(w)$ is the log-likelihood function.

The Maximum Likelihood Estimator (MLE) of the Bradley-Terry-Luce (BTL) model is used for ranking

Data: n pairwise comparison results

$$i_k \text{ defeated } j_k \text{ at comparison } k$$

BTL model: $w^* = (w_1^*, \dots, w_d^*)$: unknown quality scores of d competitors so that

$$\mathbb{P}[i \text{ defeats } j] = \frac{1}{1+e^{w_j^* - w_i^*}}$$

MLE: $\hat{w} = (\hat{w}_1, \dots, \hat{w}_d)$ satisfying

$$\operatorname{argmin}_{w \in \mathbb{R}^d} - l(w) \quad \text{w.r.t.} \quad \sum_{k=1}^{n} w_i = 0,$$

where $l(w)$ is the log-likelihood function.

Rank the competitors from the greatest to the lowest estimate $\hat{w}_i$.

Our paper develops a state-of-the-art understanding of the model in the spirit of two key concepts.

Dynamic range: the maximal performance gap among d competitors

$$B := \max_{i,j} \left| w_i^* - w_j^* \right|$$

Our paper develops a state-of-the-art understanding of the model in the spirit of two key concepts.

Dynamic range: the maximal performance gap among d competitors

$$B := \max_{i,j} \left| w_i^* - w_j^* \right|$$

Comparison graph: undirected weighted graph with weights to be the number of comparisons between the adjacent nodes

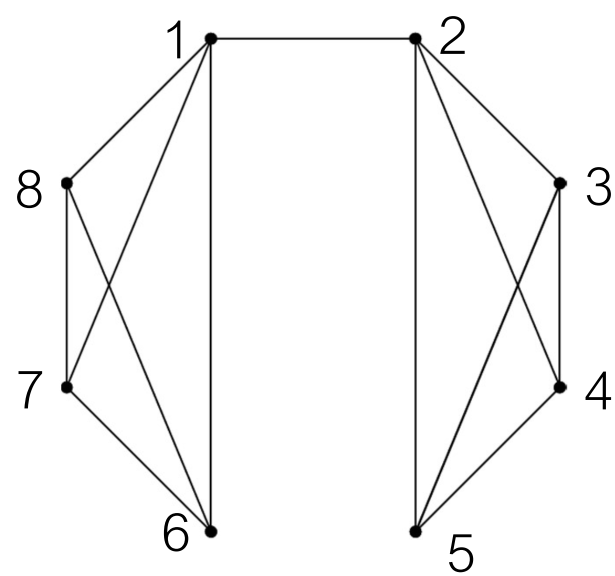

Our paper develops a state-of-the-art understanding of the model in the spirit of two key concepts.

Dynamic range: the maximal performance gap among d items

$$B := \max_{i,j} \left| w_i^* - w_j^* \right|$$

Comparison graph: undirected weighted graph with weights to be the number of comparisons between the adjacent nodes

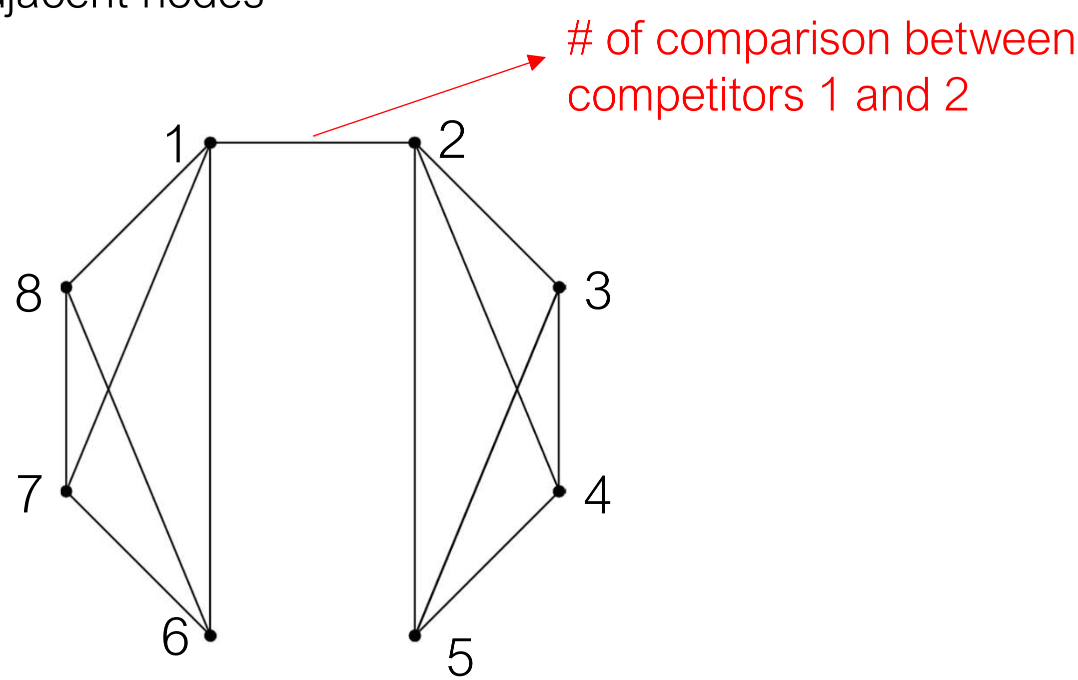

Related works on BTL model

Existence of the MLE:

- Simons & Yao (1999): complete graphs, bounded dynamic range
- Yan et al. (2012) and Han et al. (2020): Erdös-Rényi graphs, bounded dynamic range

Related works on BTL model

Existence of the MLE:

- Simons & Yao (1999): complete graphs, bounded dynamic range
- Yan et al. (2012) and Han et al. (2020): Erdös-Rényi graphs, bounded dynamic range

Consistency:

- Shah et al.(2016): ℓ_2-consistency, arbitrary graphs, bounded dynamic range
- Hendrickx et al. (2019;2020): sine error-consistency, arbitrary graphs, bounded dynamic range
- Chen et al. (2019;2020): ℓ_∞-consistency, Erdös-Rényi graphs, bounded dynamic range

Recent pairwise comparison dataset has (1) large dynamic range and (2) sparse comparison graph

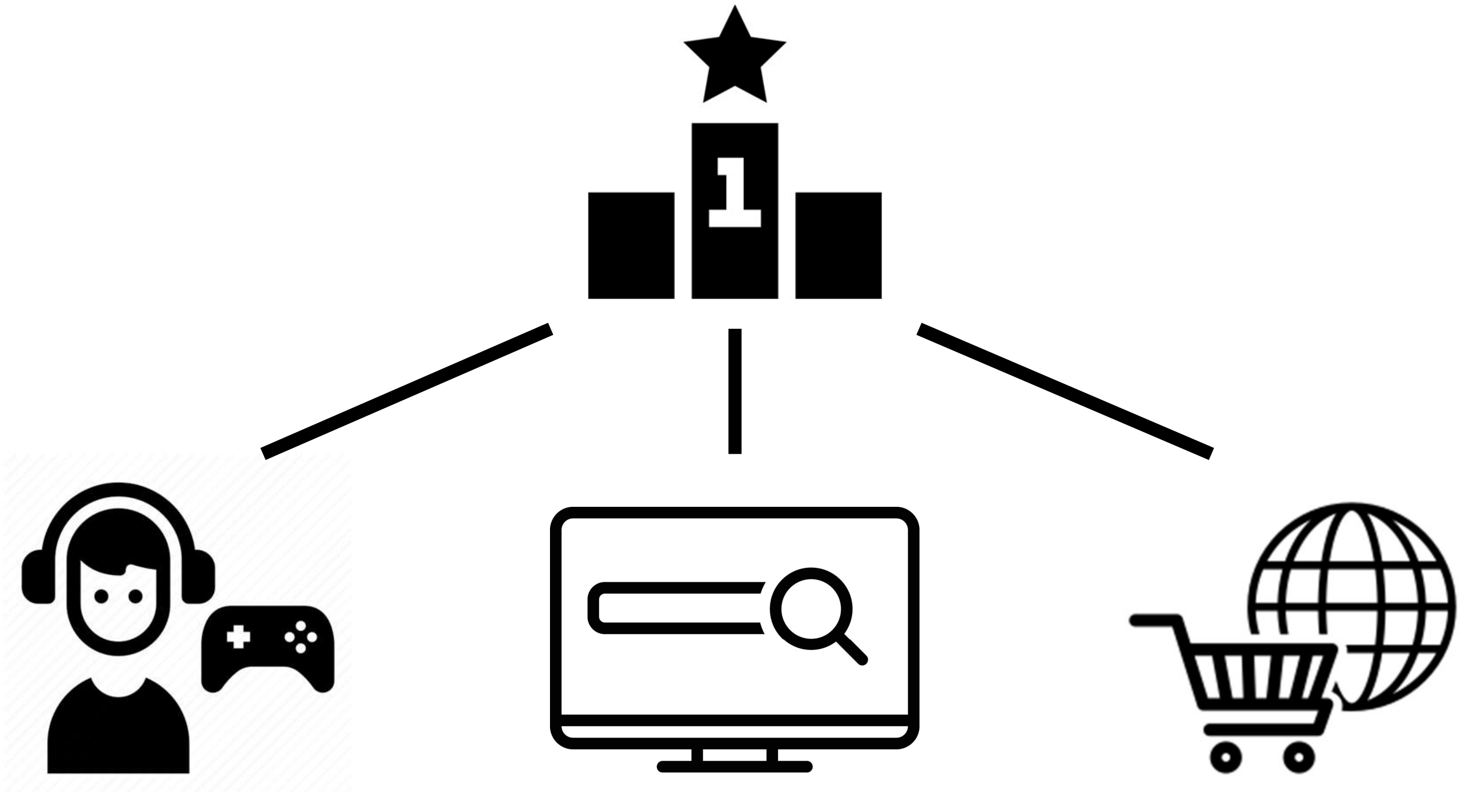

Our main contribution

Novel necessary condition for MLE existence

- Our condition first encompasses arbitrary graph topologies and infinitely increasing dynamic range.

Novel ℓ_2-consistency result for the unregularized MLE

- Our result first establish the consistency of unregularized MLE under arbitrary graph topologies and infinitely increasing dynamic range.

Fisher information matrix incorporates comparison graph and distribution of performance across items

log-likelihood:

$$l(w) = -\frac{1}{n}\sum_{k=1}^{n} \log(1 + e^{w_{j_k} - w_{i_k}})$$

Fisher information matrix incorporates comparison graph and distribution of performance across items

log-likelihood:

$$l(w) = -\frac{1}{n}\sum_{k=1}^{n} \log(1 + e^{w_{j_k} - w_{i_k}})$$

Fisher information matrix:

$$\mathcal{J}_{ij}^* = [\mathcal{J}(w^*)]_{ij} = \mathbb{E}[-\nabla^2 l(w^*)]_{ij}$$
$$= \begin{cases} \frac{L_{ij}}{\left(e^{\left(w_i^* - w_j^*\right)/2} + e^{\left(w_j^* - w_i^*\right)/2}\right)^2}, & \text{if } i \neq j, \\ -\sum_{k:k\neq i} \mathcal{J}_{ik}^*, & \text{if } i = j. \end{cases}$$

where L is the normalized graph Laplacian of the comparison graph.

Fisher information matrix incorporates comparison graph and distribution of performance across items

log-likelihood:

$$l(w) = -\frac{1}{n}\sum_{k=1}^{n} \log(1 + e^{w_{j_k} - w_{i_k}})$$

Fisher information matrix:

$$\mathcal{J}^*_{ij} = [\mathcal{J}(w^*)]_{ij} = \mathbb{E}[-\nabla^2 l(w^*)]_{ij}$$

$$= \begin{cases} \dfrac{L_{ij}}{\left(e^{(w_i^* - w_j^*)/2} + e^{(w_j^* - w_i^*)/2}\right)^2}, & \text{if } i \neq j, \\ -\sum_{k:k\neq i} \mathcal{J}^*_{ik}, & \text{if } i = j. \end{cases}$$

graph topology

performance distribution

where L is the graph Laplacian of the comparison graph.

Theorem 3.2.

If the smallest positive eigenvalue $\lambda_2(\mathcal{J}^*) \geq 2\frac{\log d}{n}$, then $\mathbb{P}[\text{MLE exist}] \geq 1 - \frac{2}{\sqrt{d}}$.

The established existence condition for the MLE extends previous works

For the Erdös-Rényi graphs, $ER(p)$, with bounded dynamic range, the condition $\lambda_2(\mathcal{J}^*) \geq 2\frac{\log d}{n}$ in the theorem holds with high probability if

$$e^B = o\left(\frac{dp}{\log d}\right).$$

It copies the results of Simons & Yao (1999); Yan et al. (2012); and Han et al. (2020).

For consistency, we construct a surrogate objective which is strongly convex against a proxy function

An example negative log-likelihood of BTL model:

2 items ($w_1^* = w_2^*$) tied in comparison

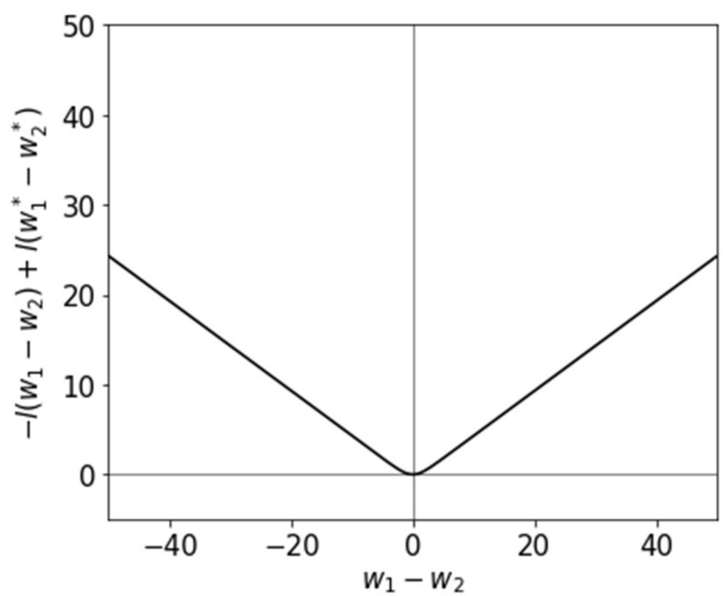

For consistency, we construct a surrogate objective which is strongly convex against a proxy function

An example negative log-likelihood of BTL model:

2 items ($w_1^* = w_2^*$) tied in comparison

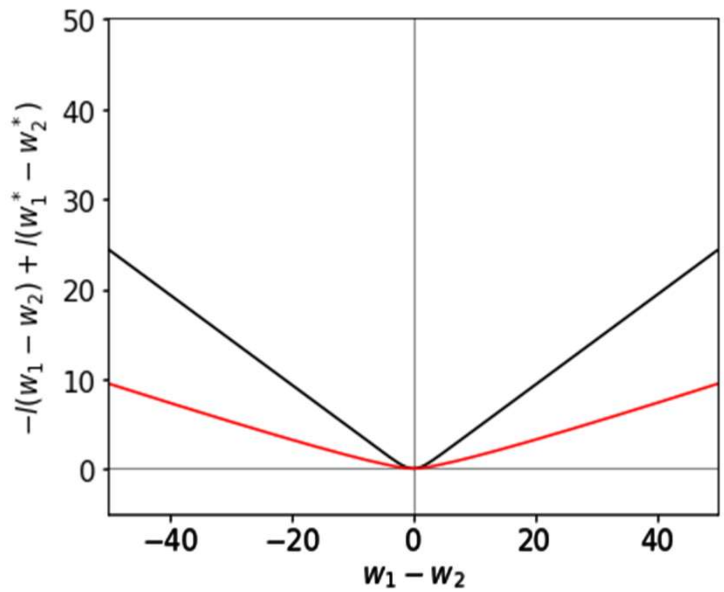

$$\mathcal{J}(w_1^* - w_2^*) \cdot h(w_1 - w_2)^2$$

where $h(x) = \operatorname{sgn}(x)\left(\sqrt{|x| + 1} - 1\right)$.

Theorem 4.1.

Under some regularity condition, the MLE $\widehat{w}$ exists, and the error $\Delta = \widehat{w} - w^*$ satisfies

$$\|h_d(\Delta)\|_{\mathcal{J}^*}^2 \leq c_1 \lambda_d\left(L^{1/2} \mathcal{J}^{*+} L^{1/2}\right) \frac{td}{n},$$

with probability at least $1 - e^{-t} - \frac{1}{d}$ for some universal constant $c_1 > 0$,

where h_d: $[h_d(x)]_i = h(x_i)$.

$h_d(\Delta)^\top \mathcal{J}^* h_d(\Delta)$

Comparison to an existing result under bounded dynamic range B

Shah et al. (2016) analyzed ℓ_∞-regularized MLE

$$\tilde{w} = \text{argmin}_{w \in \mathbb{R}^d} - l(w) \quad \text{w.r.t.} \ \|w\|_\infty \le B \ \text{and} \ \sum_{k=1}^{n} w_i = 0.$$

and provided ℓ_2-consistency for $\tilde{\Delta} = \tilde{w} - w^*$:

$$\left\|\tilde{\Delta}\right\|_L^2 \le c'(e^{-B} + e^{B})^4 \frac{td}{n}, \text{ w. p. at least } 1 - e^{-p}.$$

Comparison to an existing result under bounded dynamic range B

Shah et al. (2016) analyzed ℓ_∞-regularized MLE

$$\widetilde{w} = \operatorname{argmin}_{w \in \mathbb{R}^d} - l(w) \quad \text{w.r.t. } \|w\|_\infty \leq B \text{ and } \textstyle\sum_{k=1}^n w_i = 0.$$

and provided ℓ_2-consistency for $\widetilde{\Delta} = \widetilde{w} - w^*$:

$$\|\widetilde{\Delta}\|_L^2 \leq c'(e^{-B} + e^B)^4 \frac{td}{n}, \text{ w. p. at least } 1 - e^{-p}.$$

Our result: $\|h_d(\Delta)\|_L^2 \leq c_1(e^{-B} + e^B)^4 \frac{td}{n}$, w.p. at least $1 - e^{-t} - \frac{1}{d}$

asymptotically linear as $\Delta \to 0$

The resulting consistency theorem provides a better rate if the dynamic range grows with d

- w_i^*'s are evenly distributed along the dynamic range

$$w_i^* := \frac{2i-d}{d} B \text{ for } i \in [d].$$

- T comparisons only between items i and j having ranking difference smaller than W

The resulting consistency theorem provides a better rate if the dynamic range grows with d

- w_i^*'s are evenly distributed along the dynamic range

$$w_i^* := \frac{2i-d}{d} B \text{ for } i \in [d].$$

- T comparisons only between items i and j having ranking difference smaller than W

- The graph Laplacian has banded form: we call it a banded comparison graph case

$$L_{ij} = \begin{cases} -\frac{T}{n}, & \text{if } 0 < |i-j| \leq W, \\ -\sum_{k:k\neq i} L_{ik}, & \text{if } i = j, \\ 0, & \text{elsewhere.} \end{cases}$$

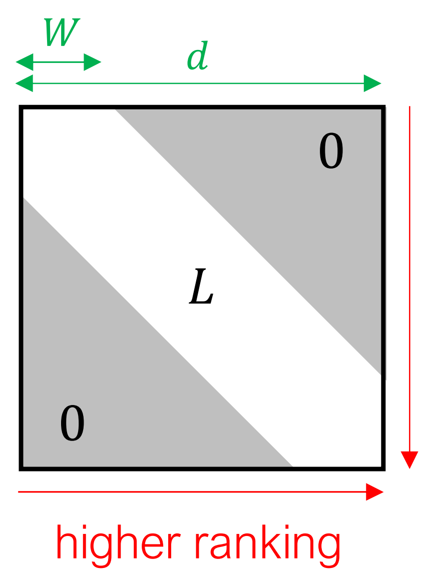

The resulting consistency theorem provides a better rate if the dynamic range grows with d

Shah et al. (2016): $\|\widetilde{\Delta}\|_L^2 \leq c'(e^{-B} + e^{B})^4 \frac{td}{n}$, w. p. at least $1 - e^{-p}$

Our result: $\|h_d(\Delta)\|_L^2 \leq \frac{c_1}{4}\left(e^{-BW/d} + e^{BW/d}\right)^4 \frac{td}{n}$, w.p. at least $1 - e^{-t} - \frac{1}{d}$

better dependence when $\frac{W}{d} \to 0$ as $d \to 0$

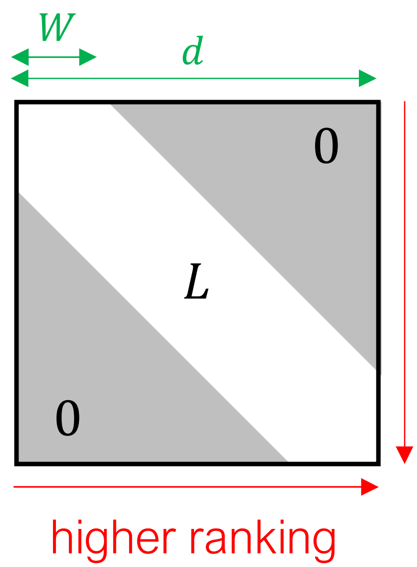

Shah et al. (2016) and our theorem were compared to MLEs on simulated data

- w_i^*'s are evenly distributed along the dynamic range

$$w_i^* := \frac{2i-d}{d} B \text{ for } i \in [d].$$

- T comparisons only between items i and j having ranking difference smaller than W

Shah et al. (2016) and our theorem were compared to MLEs on simulated data

- w_i^*'s are evenly distributed along the dynamic range

$$w_i^* := \frac{2i-d}{d} B \text{ for } i \in [d].$$

- T comparisons only between items i and j having ranking difference smaller than W

Simulation setting:

- $d = 100$ and $T = 5$
- B varies from $4e^{-1}\sqrt{\log d}$ to $4e\sqrt{\log d}$, and W from $\mathrm{e}^{-1}\frac{d}{4\sqrt{\log d}}$ to $\mathrm{e}\frac{d}{4\sqrt{\log d}}$.
- For each combination of the parameters, we performed 100 simulations.

Simulation results demonstrated the better rate of our result under banded comparison graph cases

(a) $W = \frac{d}{\sqrt{log\ d}}$, B changes

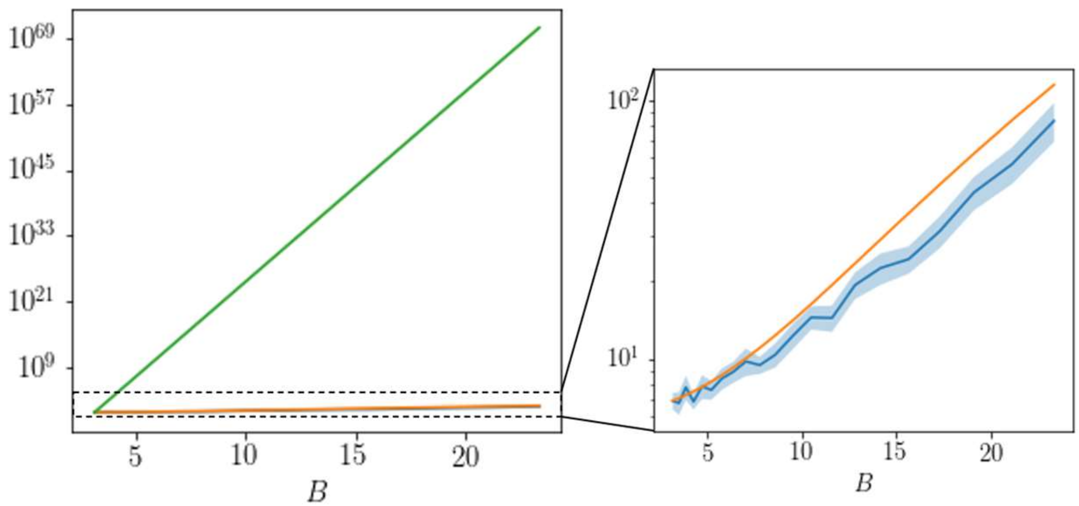

(b) $B = \sqrt{log\ d}$, W changes

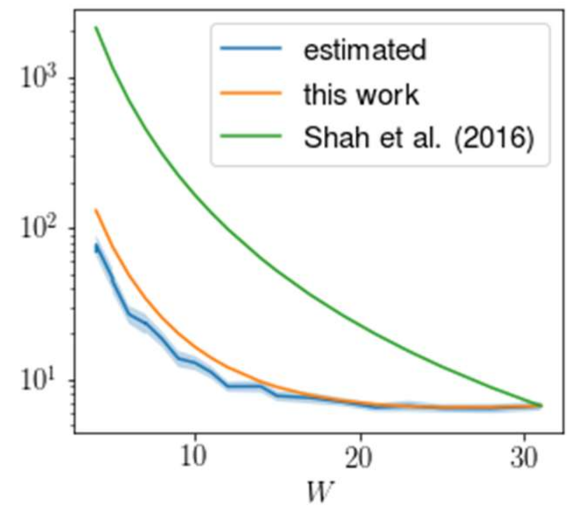

- empirical mean (blue line) and point-wise 95%-confidence region (light blue shade) of ℓ_2-loss were obtained from 100 MLEs on simulated data
- compared to the theoretical upperbounds up to constant scales (which are constant shifts in the plots)

Simulation codes are available in `github.com/HeejongBong/mmpc`.

References

Simons, G. and Yao, Y.-C. Asymptotics when the number of parameters tends to infinity in the Bradley-Terry model for paired comparisons. *The Annals of Statistics*, 27(3): 1041–1060, 1999.

Yan, T., Yang, Y., and Xu, J. Sparse paired comparisons in the Bradley-Terry model. *Statistica Sinica*, pp. 1305–1318, 2012.

Han, R., Ye, R., Tan, C., Chen, K., et al. Asymptotic theory of sparse Bradley–Terry model. *Annals of Applied Probability*, 30(5):2491–2515, 2020.

Shah, N. B., Balakrishnan, S., Bradley, J., Parekh, A., Ramchandran, K., and Wainwright, M. J. Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence. *The Journal of Machine Learning Research*, 17(1):2049–2095, 2016.

Hendrickx, J., Olshevsky, A., and Saligrama, V. Graph resistance and learning from pairwise comparisons. In *International Conference on Machine Learning*, pp. 2702–2711. PMLR, 2019.

Hendrickx, J., Olshevsky, A., and Saligrama, V. Minimax rate for learning from pairwise comparisons in the BTL model. In *International Conference on Machine Learning*, pp. 4193–4202. PMLR, 2020.

Chen, Y., Fan, J., Ma, C., and Wang, K. Spectral method and regularized MLE are both optimal for top-K ranking. *Ann. Statist.*, 47(4):2204–2235, 2019. ISSN 0090-5364. doi: 10.1214/18-AOS1745.

Chen, P., Gao, C., and Zhang, A. Y. Partial recovery for top-K ranking: Optimality of MLE and sub-optimality of spectral method. *arXiv preprint* arXiv:2006.16485, 2020.