Ranking in rounds, with application to matchmaking

Hafedh El Ferchichi, Matthieu Lerasle, Vianney Perchet

December 14, 2023

CREST
CENTER FOR RESEARCH
IN ECONOMICS AND STATISTICS

INSTITUT POLYTECHNIQUE DE PARIS

ENSAE

Outline

Motivation

- Design a matchmaking system for 1vs1 games or sport (like Elo for chess or TrueSkill for videogames).
- At each iteration: **Every player must be matched.**
- $\{i, j\}$ is a good match if

$$|\mathbb{P}\left(i \textit{ defeats } j\right) - \frac{1}{2}| \leq \varepsilon^* \quad (1)$$

where $\varepsilon^* > 0$ is chosen by an operator.

Problem formulation

- $\varepsilon_{i,j} = \mathbb{P}(i \text{ defeats } j) - \frac{1}{2}$ the **"Skill Gap"**
- Choose M_t a maximal matching such that: if $\{i, j\} \in M_t$ then players i and j play against each other.
- The **cost** of a matching M:

$$C_M = \sum_{\{i,j\} \in M} \mathbb{1}_{|\varepsilon_{i,j}| > \varepsilon^*} \tag{2}$$

- The regret is:

$$C_M - C_{M^*} \tag{3}$$

where $M^* \in \underset{M \text{ matching}}{\arg\min} \ C_M$

Main assumptions

Assumption 1 (Strong Stochastic Transitivity SST)

If $\varepsilon_{i,j} \geq 0$ and $\varepsilon_{j,k} \geq 0$ then

$$\varepsilon_{i,k} \geq \max\{\varepsilon_{i,j}, \varepsilon_{j,k}\} \tag{4}$$

Assumption 2 (Stochastic Triangle inequality STI)

If $\varepsilon_{i,j} \geq 0$ and $\varepsilon_{j,k} \geq 0$ then

$$\varepsilon_{i,k} \leq \varepsilon_{i,j} + \varepsilon_{j,k} \tag{5}$$

- **Unique ranking R** :

$$R(i) > R(j) \iff \varepsilon_{i,j} \geq 0 \quad (6)$$

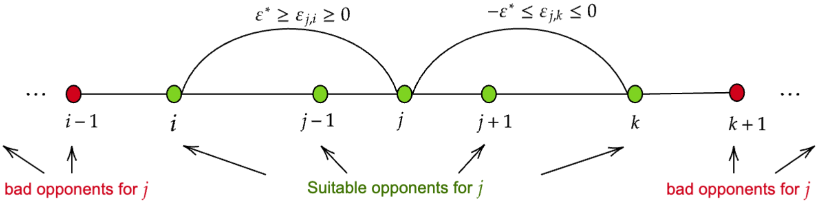

The optimal matching costs:

$$C_{M^*} = \#\{\textbf{odd-sized connected components of affinity graph}\} \quad (7)$$

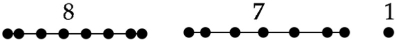

Definition 1 (ε-correct ranking)

a ranking $\hat{R}$ is said to be ε-correct if

$$\hat{R}(i) > \hat{R}(j) \implies \varepsilon_{i,j} \geq -\varepsilon \tag{8}$$

Definition 2 ((ε, δ)-PAC algorithm)

An algorithm that retrieve an ε-correct ranking with probability at least $1 - \delta$ is said to be (ε, δ) **Probably Approximately Correct (PAC)**

Without STI, there is an instance where any (ε, δ)-PAC algorithm requires $\Omega(n^2)$ comparisons (Falahatgar et al. [2017]).

Worst instance without (STI)

Consider an instance where $\varepsilon_{n,1} = 1/2$, and $|\varepsilon_{i,j}| = \eta$ when $\{i,j\} \neq \{1,n\}$.

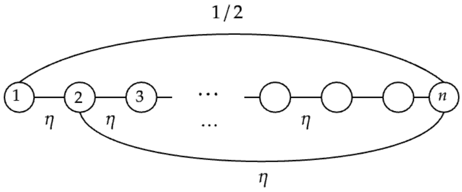

- (ε, δ)-PAC: Sample complexity $\Omega(\frac{n}{\varepsilon^2} \log \frac{n}{\delta})$ Falahatgar et al. [2017].
- Exact ranking requires Ren et al. [2019]

$$\Omega\left(\sum \frac{1}{\varepsilon_{i,i+1}^2}\left(\log \frac{n}{\delta} + \log \log \frac{1}{|\varepsilon_{i,i+1}|}\right)\right) \tag{9}$$

- Our method for (ε, δ)-PAC algorithm:
  - $O\left((\frac{n}{\varepsilon^2} \log \frac{n}{\delta}) \log^3 n\right)$ Sample complexity.
  - $O\left((\frac{1}{\varepsilon^2} \log \frac{n}{\delta}) \log^3 n\right)$ Time complexity
- Our method for exact ranking:
  - $O\left(\log^3 n \sum_i \frac{1}{\varepsilon_{i,i+1}^2}(\log(\frac{n}{\delta}) + \log \log \frac{1}{\varepsilon_{i,i+1}})\right)$ Sample complexity.
  - $O\left(\frac{\log^3 n}{\min_i \varepsilon_{i,i+1}^2}(\log \frac{n}{\delta} + \log \log \frac{1}{\min_i |\varepsilon_{i,i+1}|})\right)$ Time complexity.

(ε, δ)-PAC algorithm

- AKS sorting Network Ajtai et al. [1983]:
  - $O(\log n)$ matchings.
  - $O(n \log n)$ comparisons.
- COMPARE$(i, j, \varepsilon, \delta)$ substitutes for the comparators of AKS network: Repeat a comparison to gain confidence

Theoretical Performance

Lemma 1 (Theoretical Performance of COMPARE)

COMPARE terminates after $O(\varepsilon^{-2} \log(1/\delta))$ *comparisons and returns the stronger player with probability at least* $1/2$. *Further, if* $\varepsilon \leq |\varepsilon_{i,j}|$, *then COMPARE returns the stronger player with probability at least* $1 - \delta$.
If "Confident" is returned, that implies that $\varepsilon_{i,j} \geq \varepsilon$ *with probability at least* $1 - \delta$

Theoretical Performance of AKS(ε, δ)

Using COMPARE with parameter $\varepsilon' = \varepsilon/(2C^2 \log n)$ and $\delta' = \delta/(nC^2 \log^3 n)$ in AKS yield an (ε, δ)-PAC ranking algorithm.

Theorem 1 (Performance of AKS(ε, δ))

AKS(ε, δ) uses at most $O\left(\varepsilon^{-2} n \log^3(n) \log(n/\delta)\right)$ *comparisons, and* $O\left(\varepsilon^{-2} \log^3(n) \log(n/\delta)\right)$ *matchings.*

Is (ε, δ)-PAC sufficient to retrieve an optimal matching?

Unfortunately Not.

- Consider the following instance:

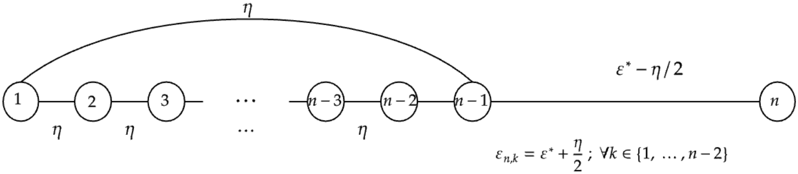

- Every ranking $\hat{R}$ where $\hat{R}(\{1, \ldots, n-1\}) < \hat{R}(n)$ is ε-correct.
- Need for $\frac{\eta}{2}$-correct ranking (i.e. an exact ranking).

Exact Ranking Retrieval (ERR)

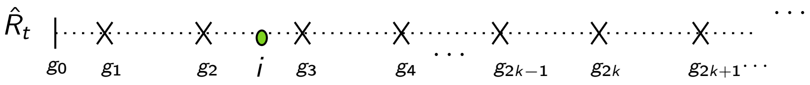

$\hat{R}_t$ is ε-correct

$\varepsilon_{g_{i+1},g_i} > \varepsilon$

$\times$: Frontiers

$\dots$: Cross at most one frontier

Exact Ranking Retrieval (ERR)

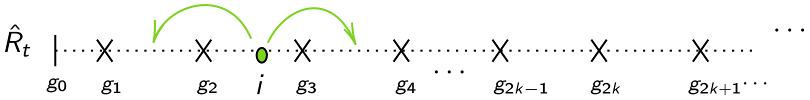

$\hat{R}_t$ is ε-correct

$\varepsilon_{g_{i+1},g_i} > \varepsilon$

$\times$: Frontiers

$\dots$: Cross at most one frontier

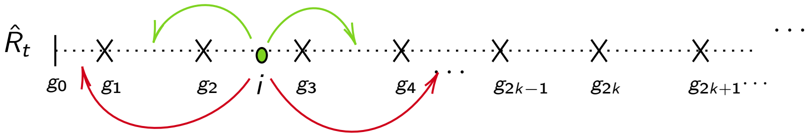

$\hat{R}_t$ is ε-correct

$\varepsilon_{g_{i+1},g_i} > \varepsilon$

$\times$: Frontiers

$\ldots$: Cross at most one frontier

Exact Ranking Retrieval (ERR)

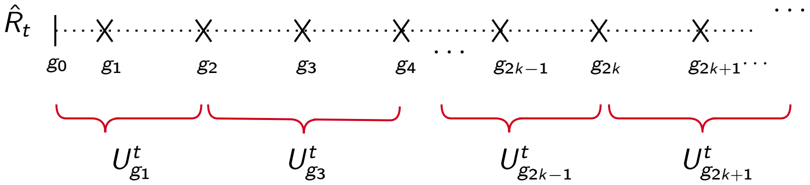

$\hat{R}_t$ is ε-correct

$\varepsilon_{g_{i+1},g_i} > \varepsilon$

$\times$: Frontiers

$\ldots$: Cross at most one frontier

Exact Ranking Retrieval (ERR)

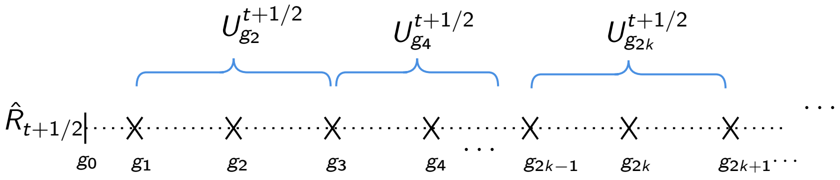

$\hat{R}_t$ is ε-correct

$\varepsilon_{g_{i+1},g_i} > \varepsilon$

$\times$: Frontiers

$\dots$: Cross at most one frontier

Properties of ERR

- Players that reached their true ranking are recognizable.
- These players don't have to participate in ranking anymore (moves to gap estimation right away).
- Performance: Let $\Delta_i = \min\{|\varepsilon_{i,i+1}|, |\varepsilon_{i,i-1}|\}$
  - Sample complexity:

$$O\left(log^3 n \sum_i \frac{1}{\Delta_i^2} (\log(n/\delta) + \log\log \frac{1}{\Delta_i}) \right) \tag{10}$$

  - Time complexity(for player i):

$$O\left(\frac{\log^3 n}{\Delta_i^2} (\log \frac{n}{\delta} + \log\log \frac{1}{\Delta_i}) \right) \tag{11}$$

Neighboring gap estimation

For this Phase: the exact ranking is provided.

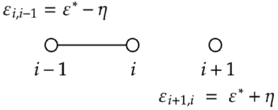

- Costs $O\left(\frac{1}{\eta^2}\log(\frac{1}{\delta})\right)$
- if $\eta_{i,i+1} = |\varepsilon^* - |\varepsilon_{i,i+1}||$, we have:
  - Sample complexity of retrieving an optimal matching is:

$$O\left(\sum_i \frac{1}{\eta_{i,i+1}^2}(\log(n/\delta) + \log\log\frac{1}{\eta_{i,i+1}})\right) \tag{12}$$

  - A cost of:

$$O\left(\sum_{i;|\varepsilon_{i,i+1}|>\varepsilon^*} \frac{1}{\eta_{i,i+1}^2}(\log(n/\delta) + \log\log\frac{1}{\eta_{i,i+1}})\right) \tag{13}$$

Summary

- (ε, δ)-PAC ranking:
  - Time complexity: $O\left(\frac{log^3 n}{\varepsilon^2} \log(n/\delta)\right)$
- Exact Ranking Retrieval:
  - Time complexity (for player i): $O\left(\frac{\log^3 n}{\Delta_i^2}(\log(\frac{n}{\delta}) + \log\log\frac{1}{\Delta_i})\right)$
  - Cost for player i: $O\left(\frac{\log^3 n}{\Delta_i^2}(\log(\frac{n}{\delta}) + \log\log\frac{1}{\Delta_i})\right)$
- Gap estimation:
  - Time complexity: $O\left(\frac{1}{\min_i \eta_i^2}(\log(\frac{n}{\delta}) + \log\log\frac{1}{\min_i \eta_i^2}\right)$
  - Cost: $O\left(\sum_{|\varepsilon_{i,i+1}|>\varepsilon^*} \frac{1}{\eta_{i,i+1}^2}(\log(n/\delta) + \log\log\frac{1}{\eta_{i,i+1}})\right)$

References

Miklós Ajtai, János Komlós, and Endre Szemerédi. An 0 (n log n) sorting network. In *Proceedings of the fifteenth annual ACM symposium on Theory of computing*, pages 1–9, 1983.

Moein Falahatgar, Yi Hao, Alon Orlitsky, Venkatadheeraj Pichapati, and Vaishakh Ravindrakumar. Maxing and ranking with few assumptions. *Advances in Neural Information Processing Systems*, 30, 2017.

Wenbo Ren, Jia (Kevin) Liu, and Ness Shroff. On sample complexity upper and lower bounds for exact ranking from noisy comparisons. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Questions ?