Ranking from Pairwise Comparisons

Alex Mangiapane, UNC | Sean Stuhlsatz, UVA
Mentor: Cheng Mao, GA Tech

♛ Pairwise comparisons ♛

Goal: Rank simulated and real players based on pairwise comparisons

Models

Assume each player has a skill θ_i, and when players i and j compete,

$$\mathbb{P}(i \text{ beats } j) = f(\theta_i - \theta_j)$$

Different functions f yield different models. The most common models are Bradley-Terry and Thurstone-Mosteller:

Bradley-Terry

Bradley-Terry uses a logistic curve:

$$f(x) = \frac{1}{1 + e^{-cx}}$$

Thurstone-Mosteller

Thurstone-Mosteller uses a normal CDF:

$$f(x) = \Phi\left(\frac{x}{\beta}\right)$$

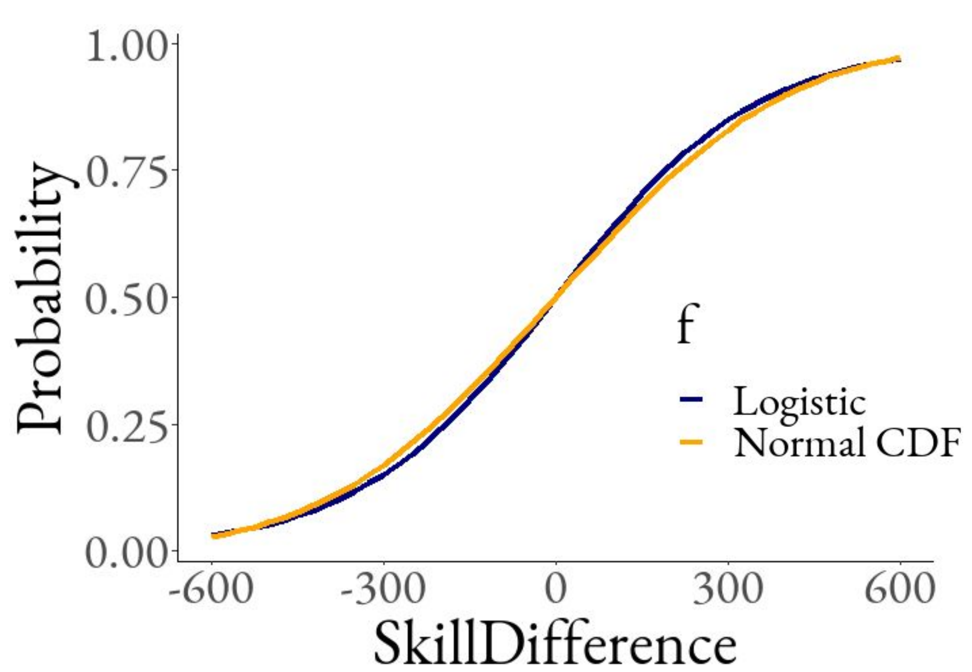

$$c = \frac{\ln(10)}{400}$$

$$\beta = \frac{400\pi}{\ln(10)\sqrt{3}}$$

Applications

This kind of modeling has a variety of applications, in areas like:

- ***Online games***
- ***Econometrics***
- ***Educational testing and psychometrics***

♛ Ranking methods ♛

Elo ratings

Elo is the oldest online method. If player i plays player j, it updates scores by:

$$\hat{\theta}_i \leftarrow \hat{\theta}_i + k(s - p_{ij})$$

where s is 1 if player i won, 0 if i lost, and p_{ij} is the estimated probability of i winning. Under the Bradley-Terry model Elo updates are gradient descent on the likelihood, and can be improved by having k decay over time, yielding:

$$\hat{\theta}_i \leftarrow \hat{\theta}_i + \frac{k}{t^d}(s - p_{ij})$$

Bayesian methods

Bayesian methods treat ratings as random variables, and store probability distributions rather than single ratings. In practice, we usually require that these distributions be normal. Given the Thurstone-Mosteller model, one can compute the exact update formulas, which are:

$$\mu_i \leftarrow \mu_i + (-1)^s \frac{\sigma_i^2}{c} \frac{\varphi}{Z}$$

$$\sigma_i^2 \leftarrow \sigma_i^2 - (-1)^s \frac{\sigma_i^2 \varphi}{c^3 Z}(\mu_i(\sigma_j^2 + \beta^2) + \mu_j \sigma_i^2) - \left(\frac{\varphi \sigma_i^2}{Zc}\right)^2$$

One can also approximate the posterior, as Trueskill does for Thurstone-Mosteller:

$$\mu_i \leftarrow \mu_i - (-1)^s \frac{\sigma_i^2}{c} V\left(\frac{\mu_{\text{winner}} - \mu_{\text{loser}}}{c}, \frac{\varepsilon}{c}\right)$$

$$\sigma_i^2 \leftarrow \sigma_i^2 \max\left(1 - \left(\frac{\sigma_i^2}{c}\right) W\left(\frac{\mu_{\text{winner}} - \mu_{\text{loser}}}{c}, \frac{\varepsilon}{c}\right), \kappa\right)$$

For Bradley-Terry, Glicko approximates the updates, yielding:

$$\mu_i \leftarrow \mu_i + \frac{c}{\sigma_i^{-2} + \delta^{-2}} \sum_k g(\sigma_{j_k}^2)(s_k - E_{j_k})$$

$$\sigma_i^2 \leftarrow \frac{1}{\sigma_i^{-2} + \delta^{-2}}$$

♛ Results ♛

Simulation

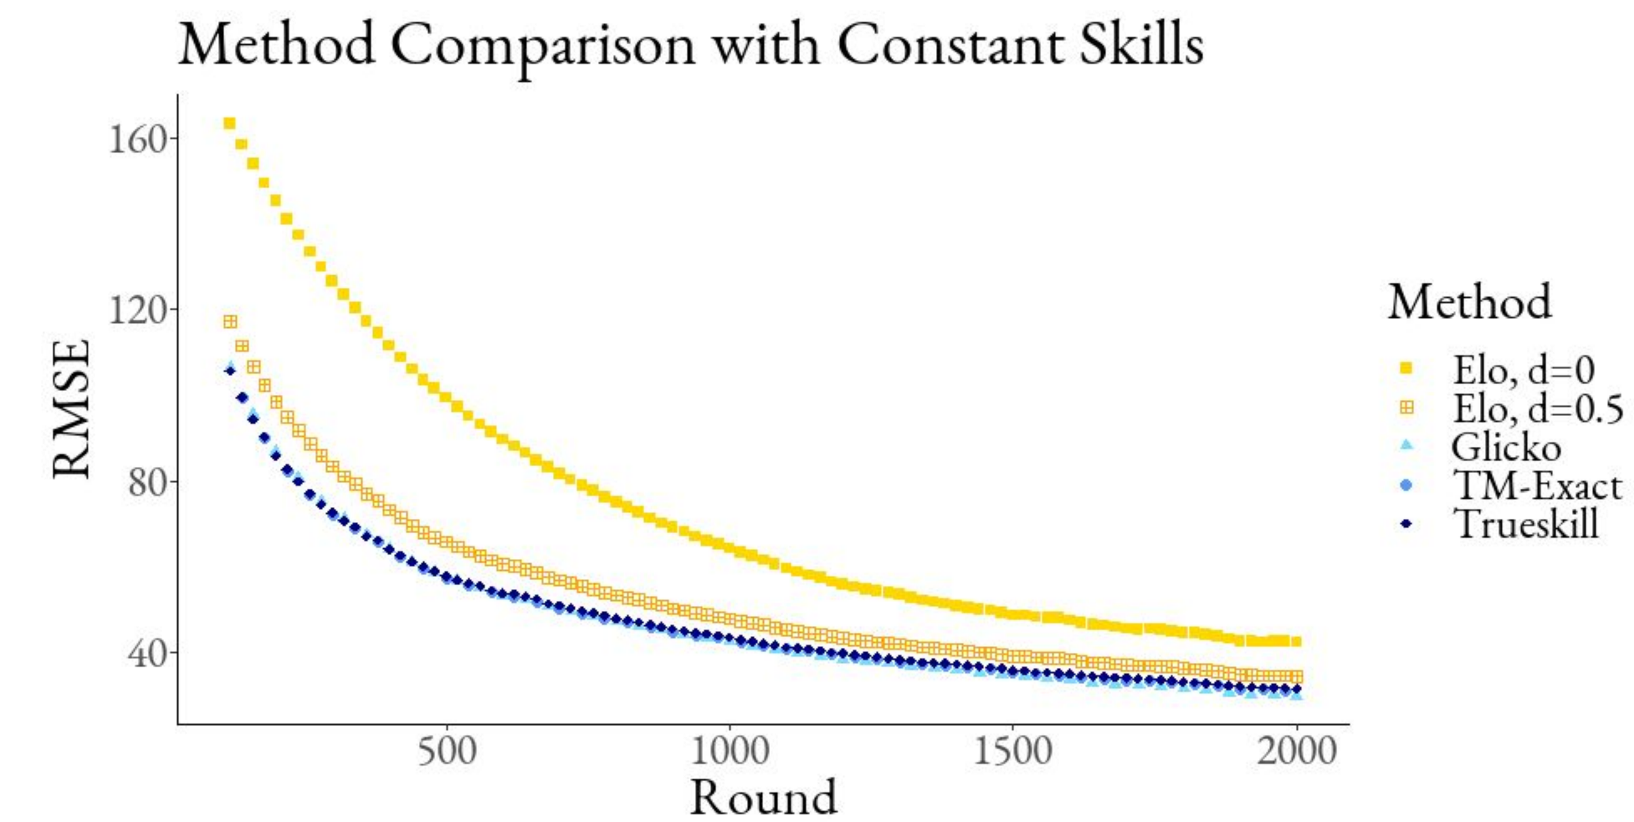

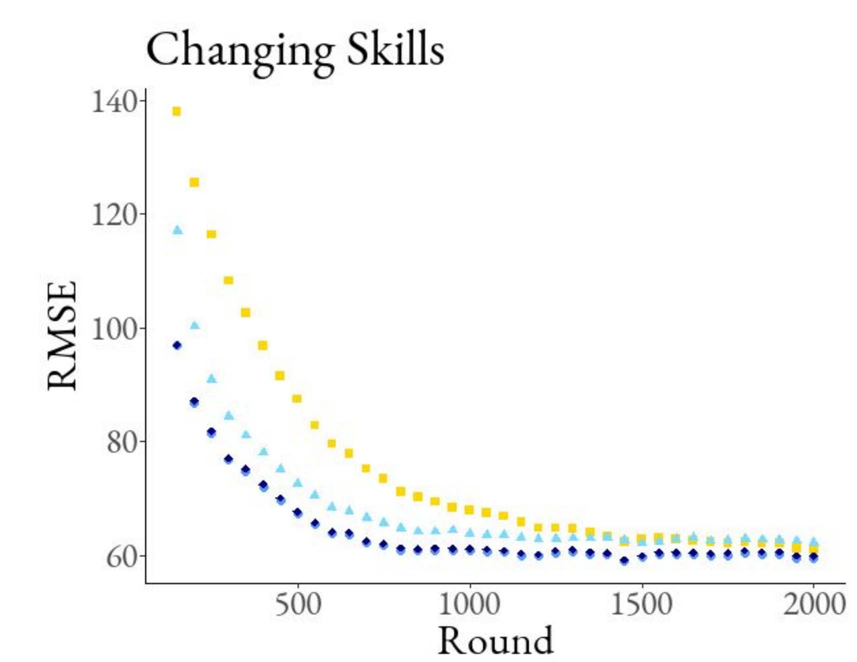

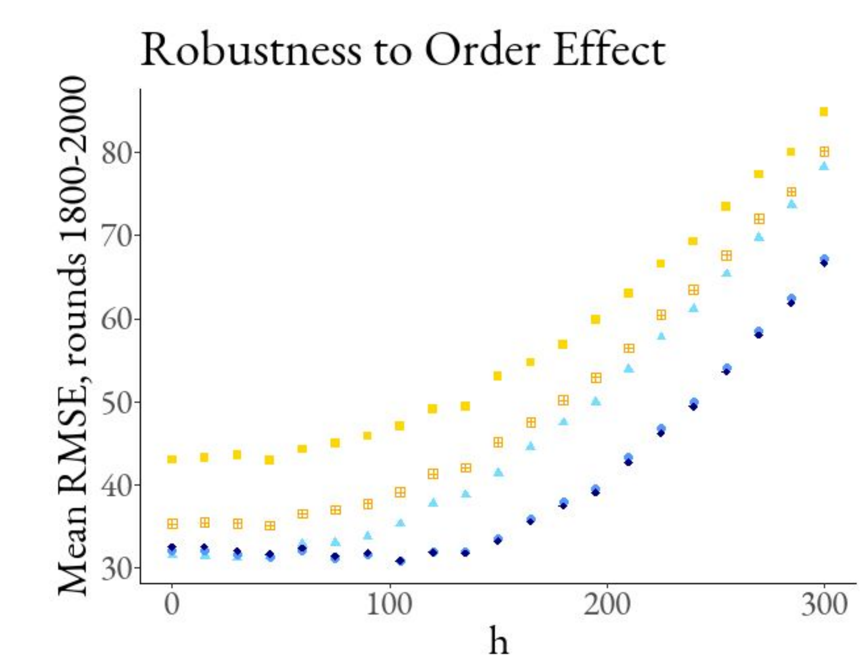

Real Data

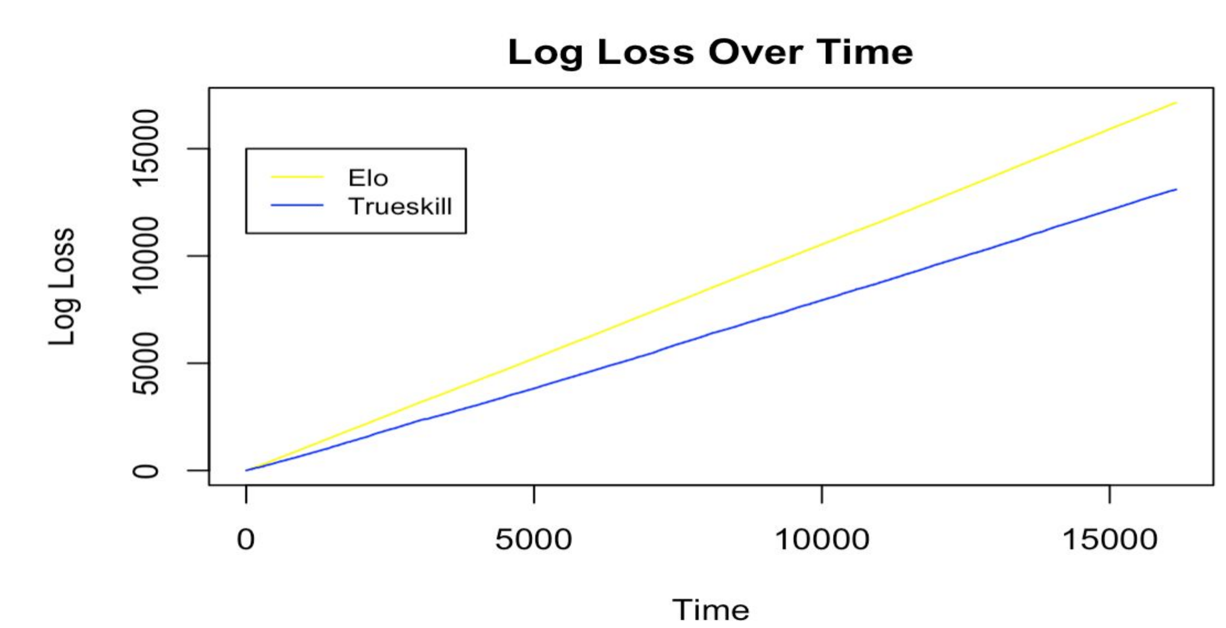

Acknowledgements

We would like to thank our advisor, Dr. Cheng Mao, as well as Georgia Tech and the National Science Foundation.