A *Propos* of Soft Condorcet Optimization and The Elo Rating System

Manfred Diaz

1 Introduction

DISCLAIMER: This are contemporary notes I used to answer the question on what are the main differences between SCO and Elo while we were working towards Lanctot et al. (2024). They are still very much work in progress but nonetheless, contain interesting material to understand SCO and Elo.

Recently, we introduced in Lanctot et al. (2024) *Soft Condorcet Optimization* (SCO). Frequently, we got asked the question on how SCO differs from *Elo* (Elo, 1967), the *de facto* rating system in the evaluation of LLM agents (Chiang et al., 2024), and how it could overcome the limitations of *Elo* as a rating method widely discussed in the literature (Herbrich et al., 2006; Shah et al., 2015; Balduzzi et al., 2018; Bertrand et al., 2023; Lanctot et al., 2023).

In this document, I offer an analysis of the structure of pairwise comparison data leveraged by SCO and Elo, the relationship between Elo and SCO objectives, including a characterization of Elo stationary points and the complete understanding of their properties at convergence explaining (predicting) the results on *Penthatlon* data from VasE and the AAMAS paper example in Section 4.1. Similarly, I present a characterization of SCO stationary points and their relationship with the margin matrix and the covariance of the predictions, and what are the main differences between a stochastic gradient descent update between Elo and SCO. Lastly, I briefly discuss the problem of inconsistent estimation of Elo through MLE.

2 Soft-Condorcet Optimization: The Origins

The main idea behind Soft Condorcet Optimization emanates from the Kemeny-Young voting rule (Kemeny, 1959; Young and Levenglick, 1978). This voting rule computes a function that minimizes the sum of Kendall-tau distances between the ranking and all the votes in the preference profile:

$$\min_{f_w([\succeq])\in\Pi(A)} \sum_{v\in[\succeq]} K_d(v, f_w([\succeq])).$$

Given a preference profile and parameters $\boldsymbol{\theta}$ (ratings), we define a loss function:

$$L([\succeq], A, V, \boldsymbol{\theta}) = \sum_{v\in[\succeq]} \sum_{i,j\in\{0,1,\cdots,|v|-1\}, i<j} D_v(\theta_{v[i]}, \theta_{v[j]}), \tag{1}$$

2.1 Distance-based Ranking Models

The problem of selecting an ordering that satisfies voters preferences has also been studied in statistics and psychology through probabilistic models defined over the space of permutations (Diaconis, 1988; Marden, 1995; Alvo and Yu, 2014). In particular, distance-based models (Fligner and Verducci, 1986) assume that voters have a true, collective preference over alternatives v^*, and the probability of observing any individual preference $v \in V$ is proportional to its distance to this ranking. Thus, the probability of any individual vote $v \in V$ is

$$p(v) \propto \exp(-K(v, v^*)) \tag{2}$$

where K is a distance between rankings.

The objective in Equation 1 is a particular instance of maximum likelihood estimation (MLE) of the true collective preference v^*, when the choice of distance is the Kendall-tau distance K_d and the ranking among alternatives is represented by scalar ratings $\boldsymbol{\theta}$.

2.2 Derivation

The properties of these models under several choices of K, including the Kendall-tau distance K_d, have been studied in the literature (Critchlow et al., 1991).Let $\pi \in \Pi$ be a ranking of order k. The generalized distance-based probabilistic model over the space of ranks Fligner and Verducci (1986) is given by: The natural interpretation of this model is that the probability of observing a rank π is related to its distance from the (Kemeny-Young) true/consensus/ rank $\pi^* \in \Pi$ (Fligner and Verducci, 1986; Marden, 1995).

Let $p^*(\pi)$ be the true distribution of rankings under the consensus π^*, and let $p_\theta(\pi)$ be the distance-based model under the parameterized consensus distribution π_θ, implicitly represented by each player parameter $\theta = [\theta_1, \dots, \theta_n]$, such that:

$$p_\theta(\pi) \propto \exp(-d(\pi, \pi_\theta)) \tag{3}$$

By the KL divergence formulation of the MLE:

$$\min_\theta D_{KL}(p^*(\pi)||p_\theta(\pi)) \tag{4}$$

we optimize the straightforward objective:

$$\min_\theta \mathbb{E}_{\pi \sim p^*(\pi)} \left[d(\pi, \pi_\theta)\right] \tag{5}$$

2.3 A Family of Methods

Equation 5 describes a family of ranking methods, for which SCO is obtained by setting $d = K_d$ the Kendal-tau distance or its continuous approximation such that:

$$\hat{K}_d(\pi, \pi_\theta) = \sum_{(i,j)\in\binom{k}{2}} \sigma(\theta_{\pi[i]} - \theta_{\pi[j]}) \tag{6}$$

Substituting in d_{Kem} and $p^*(\pi)$ by the sketch (dataset) D.

$$\min_\theta \mathbb{E}_{\pi \sim D} \left[\sum_{(i,j)\in\binom{k}{2}} \sigma(\theta_{\pi[i]} - \theta_{\pi[j]}) \right] \tag{7}$$

3 The Elo Rating System and Rankings

We adapt the Elo rating to operating over rankings to pairwise comparisons. First, we note that a complete rank (total order) π contains $\binom{k}{2}$ pairwise comparisons. Notationally, let $\{X_1, \dots, X_{\binom{k}{2}}\}$ be a set of $\binom{k}{2}$ independent Bernoulli distributed random variables $X_i \sim Bern(p)$. This is equivalent to X_i Bradley-Terry models (Bradley and Terry, 1952; Hunter, 2004) the probability $P(a \succ b)$ where $P(X_i = 1) = p$ represents $P(a \succ b)$ and $P(X_i = 0) = 1 - p$ represents $P(b \succ a)$ (no ties). Furthermore, assume a set of rating vectors $\theta = [\theta_1, \dots, \theta_k]$ such that:

$$P(a \succ b) = p = \sigma(\theta_a - \theta_b) = \frac{e^{\theta_a}}{e^{\theta_a} + e^{\theta_b}}$$

Let π^* be the true ranking among k players and π_θ our estimated ranking. By the D_{KL} formulation of MLE, we have again that:

$$\min_\theta D_{KL}\left(p^*(X_1,\ldots,X_{\binom{k}{2}})\middle\|p(X_1,\ldots,X_{\binom{k}{2}}|\pi_\theta)\right) \tag{8}$$

$$\min_\theta \sum_{i=1}^{\binom{k}{2}} D_{KL}\left(p^*(X_i)\middle\|p_\theta(X_i)\right) \tag{9}$$

$$\max_\theta \sum_{i=1}^{\binom{k}{2}} \mathbb{E}_{x_i\sim p^*(X_i)}\left[\log p_\theta(x_i)\right] \tag{10}$$

$$\max_\theta \sum_{i=1}^{\binom{k}{2}} p^*(x_i=0)\log(1-\sigma(\theta_a-\theta_b)) + p^*(x_i=1)\log\sigma(\theta_a-\theta_b) \tag{11}$$

$$\max_\theta \sum_{i=1}^{\binom{k}{2}} p^*(x_i=0)\log\sigma(\theta_b-\theta_a) + p^*(x_i=1)\log\sigma(\theta_a-\theta_b) \tag{12}$$

$$\min_\theta \sum_{i=1}^{\binom{k}{2}} -(1-y_i)\log\sigma(\theta_b-\theta_a) - y_i\log\sigma(\theta_a-\theta_b) \tag{13}$$

From (10) to (11) we used $1-\sigma(x)=\sigma(-x)$, $p^*(x_i=1)=y$ as the $\{0,1\}$ outcomes from $\binom{k}{2}$ results in the total order, and $\max_\theta f(\theta)=\min_\theta -f(\theta)$ to get the binary cross entropy loss.

4 Unify Representation of Pairwise Comparisons

In matrix form, the data is a collection of vectors $X_k=[0,\ldots,1_i,\ldots,-1_j,\ldots]$, where the *i-th* and *j-th* entries are $X_{ki}=1, X_{kj}=-1$ if alternative $i\succ j$. We denote by $X\in \mathrm{R}^{m\times n}$ of m pairwise comparisons among all the n alternatives. Also, let $\Theta=[\theta_1,\ldots,\theta_i,\ldots,\theta_j,\ldots,\theta_n]$ denotes each alternative ranking such that $X_k\Theta=\theta_i-\theta_j$.

4.1 Properties of The Design Matrix

We can write the design matrix X as the composition of a win and a loss matrix:

$$X = X_+ + X_- \tag{14}$$

where $X_+\in\mathrm{R}^{m\times n}$ contains only the wins ($X_{ki}=1$), and X_- holds only the negative entries ($K_{jj}=-1$). Therefore, the *Gramian* matrix $LL=X^{\mathrm{T}}X$ can be decomposed as:

$$L = X^{\mathrm{T}}X = (X_+ + X_-)^{\mathrm{T}}(X_+ + X_-) \tag{15}$$

$$= X_+^TX_+ + X_+^{\mathrm{T}}X_- + X_-^{\mathrm{T}}X_+ + X_-^{\mathrm{T}}X_- \tag{16}$$

$$= W - P^{\mathrm{T}} - P + D \tag{17}$$

where each of the four matrices W, M, M^{T} and D have the following interpretations.

First, the **wins matrix** $W=X_+^TX_+$ is a diagonal matrix $W\in\mathbb{Z}^{n\times n}$ where each entry W_{ii} represents the total number of pairwise comparisons won by alternative i. Conversely, the **losses matrix** $D=X_-^{\mathrm{T}}X_-$ is a diagonal matrix $D\in\mathbb{Z}^{n\times n}$ where each entry D_{ii} contains the number of comparisons alternative i lost. And finally, the matrix $P=X_-^{\mathrm{T}}X_+$, or more precisely $-M$, is the **pairwise matrix**, where each entry $-P_{ij}$ counts the number of times alternative i is preferred to alternative j. Moreover, note that $P^{\mathrm{T}}=(X_-^{\mathrm{T}}X_+)^T=X_-^{\mathrm{T}}X_+$ is the transpose of the matrix P. The **margin matrix** is $M=P-P^{\mathrm{T}}$.

4.2 Outcomes

The pairwise comparison problem is a *imbalanced* binary classification problem where only the positive class is present. As such, the target output is always the constant vector $Y = \mathbf{1}_m$. A vector will play an essential role next is $C = X^{\mathrm{T}}Y = X^{\mathrm{T}}\mathbf{1}_m \in \mathbb{R}^n$. For every alternative i, this vector contains in $C_i \in \mathbb{Z}$ the sum of pairwise comparison won and lost by alternative i. Note that, following the decomposition of the design matrix $X = X_+ + X_-$, we have that:

$$C = X^{\mathrm{T}}\mathbf{1}_m = (X_+ + X_-)^{\mathrm{T}}\mathbf{1}_m = X_+^{\mathrm{T}}\mathbf{1}_m + X_-^{\mathrm{T}}\mathbf{1}_m = C_+ + C_- \tag{18}$$

where $X^T\mathbf{1}_m$ computes the sum of the columns of the design matrix X, or its positive-negative decomposition.

4.3 Working Examples

Example 1. *In the Penthatlon example on Lanctot et al. (2023), we have the following:*

$$X^T X = \begin{bmatrix} 10 & -5 & -5 \\ -5 & 10 & -5 \\ -5 & -5 & 10 \end{bmatrix} = \begin{bmatrix} 6 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 6 \end{bmatrix} - \begin{bmatrix} 0 & 4 & 2 \\ 1 & 0 & 2 \\ 3 & 3 & 0 \end{bmatrix} + \begin{bmatrix} 4 & 0 & 0 \\ 0 & 7 & 0 \\ 0 & 0 & 4 \end{bmatrix} - \begin{bmatrix} 0 & 1 & 3 \\ 4 & 0 & 3 \\ 2 & 2 & 0 \end{bmatrix} \tag{19}$$

and

$$C = X^T\boldsymbol{1}_m = \begin{bmatrix} 2 & -4 & 2 \end{bmatrix} \tag{20}$$

Example 2. *In the AAMAS paper example*

$$X^T X = \begin{bmatrix} 10 & -5 & -5 \\ -5 & 10 & -5 \\ -5 & -5 & 10 \end{bmatrix} = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 8 & 0 \\ 0 & 0 & 4 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 3 \\ 5 & 0 & 3 \\ 2 & 2 & 0 \end{bmatrix} + \begin{bmatrix} 7 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 6 \end{bmatrix} - \begin{bmatrix} 0 & 5 & 2 \\ 0 & 0 & 2 \\ 3 & 3 & 0 \end{bmatrix} \tag{21}$$

and

$$C = X^T\boldsymbol{1}_m = \begin{bmatrix} 4 & -6 & 2 \end{bmatrix} \tag{22}$$

Example 3. *For the Rock-Paper-Scissors game, we have the following:*

$$X^T X = \begin{bmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} - \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} - \begin{bmatrix} 0 & -1 & 0 \\ 0 & 0 & -1 \\ -1 & 0 & 0 \end{bmatrix} \tag{23}$$

and

$$C = X^T\mathbf{1}_m = \begin{bmatrix} 0 & 0 & 0 \end{bmatrix} \tag{24}$$

5 Relationship Between The Objectives

We compare Elo and SCO by understanding the connection between their objectives. First, for a set of m pairwise comparisons, the Elo objective is the negative log-likelihood:

$$\sum_{k=1}^{m} -(1 - y_k)\log(1 - \sigma(X_k\Theta)) - y_k \log\sigma(X_k\Theta) \tag{25}$$

and note that, for the pairwise comparison problem, all $y_k = 1$, so the traditional *NLL* objective reduces to:

$$\mathrm{NLL}(\Theta) = \sum_{k=1}^{m} -\log\sigma(X_k\Theta) \tag{26}$$

which can be written in matrix form as:

$$\mathrm{NLL}(\Theta) = -\log\sigma(X\Theta) \tag{27}$$

where the *logarithm* and the *sigmoid* fucntions are applied entrywise to the vector $X\Theta$. And similarly, the SCO objective can be rewritten such as:

$$\text{SCO}(\Theta) = \mathbf{1}_m - \sigma(X\Theta) \tag{28}$$

by leveraging the identity $\sigma(-x) = 1 - \sigma(x)$ and where $\mathbf{1}_m$ is as before.

5.1 Stationary Points

5.1.1 Elo.

The gradient of the NLL objective, written in matrix form, is:

$$\nabla_\Theta \text{NLL}(\theta) = -X^{\mathrm{T}}(\mathbf{1}_m - \sigma(X\Theta)) \tag{29}$$

Therefore, the stationary points of the Elo objective are characterized as:

$$\nabla_\Theta \text{NLL}(\Theta) = -X^{\mathrm{T}}[\mathbf{1}_m - \sigma(X\Theta)] = 0 \tag{30}$$

Recall, from our analysis of the pairwise data representation, the outcome matrix $C = X^{\mathrm{T}}\mathbf{1}_m$ is the vector of cumulative wins and losses for every alternative. Therefore, Elo stationary points are those where the *sum of* predictions match each alternative cumulative wins and losses.

$$X^{\mathrm{T}}\sigma(X\Theta) = X^{\mathrm{T}}\mathbf{1}_m \tag{31}$$

This result explains why, in the Penthatlon data, A and C receive the exact same gradients and have the same ratings, as their sum of wins and losses is the same. It also explains why, in the new example on the AAMAS paper, A receives a higher rating than C.

Proposition 1. *Let a and b be two alternatives with ratings θ_i and θ_j, respectively. If i and i have the same win-loss aggregated statistics $X_a^T \boldsymbol{1}_n = X_b^T \boldsymbol{1}_n$, then alternatives a and b will have equal ratings $\theta_a = \theta_b$ when minimizing Elo by gradient descent.*

Proof. Straightforward from Equation 31. □

5.1.2 SCO

In contrast, the gradient of the SCO objective in matrix form is computed as:

$$\nabla_\Theta \text{SCO}(\Theta) = -X^{\mathrm{T}}[(\mathbf{1}_m - \sigma(X\Theta)) \circ \sigma(X\Theta)] \tag{32}$$

where $\circ$ represents Hadamard's product. The stationary points of SCO are then characterized by:

$$X^{\mathrm{T}}[(\mathbf{1}_m - \sigma(X\Theta)) \circ \sigma(X\Theta)] = 0 \tag{33}$$

Understanding the stationary points SCO requires a more involved algebraic manipulation and noticing that:

$$\sigma(\theta_i - \theta_j)[1 - \sigma(\theta_i - \theta_j)] = \sigma(\theta_j - \theta_i)(1 - \sigma(\theta_j - \theta_i)) \tag{34}$$

is the covariance of the predictions, and as such, every comparison between alternatives i, j will result in entries:

$$(w_{ij} - l_{ij})\sigma(\theta_i - \theta_j)[1 - \sigma(\theta_i - \theta_j)]$$

where w_{ij} and l_{ij} are the wins and losses for alternative i over j for every index $i > j$. Therefore, the stationary point condition of SCO reduces to:

$$[M \circ \Sigma]\mathbf{1}_p = 0 \tag{35}$$

where $M \in \mathbb{R}^{n\times n}$ is the *skew-symmetric* **margin matrix** and $\Sigma \in \mathbb{R}^{n\times n}$ is a *symmetric* covariance matrix $\Sigma_{ij} = \Sigma_{ji} = [\sigma(\theta_i - \theta_j)[1 - \sigma(\theta_i - \theta_j)]]$.

NOTE: more analysis is needed here. Why these stationary points are better than Elo's? I sense a connection here with Maximal Lotteries?

Example 4. *For the AAMAS paper example, one can verify that Equation 33 reduces to*

$$\begin{bmatrix} 0 & 5 & -1 \\ 0 & 0 & -1 \\ 1 & 1 & 0 \end{bmatrix} \circ \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \Sigma_{13} \\ \Sigma_{12} & \Sigma_{22} & \Sigma_{23} \\ \Sigma_{13} & \Sigma_{23} & \Sigma_{33} \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \tag{36}$$

where $\Sigma_{ij} = \Sigma_{ji} = \sigma(\theta_i - \sigma_j)[1 - \sigma(\theta_i - \theta_j)]$

The expression in (22) is equivalent to

$$\mathrm{diag}(MI_n\Sigma^T) = 0 \tag{37}$$
$$\mathrm{diag}(M\Sigma) = 0 \tag{38}$$

5.2 Stochastic Gradient Step

We briefly rewrite the SCO objective to contrast it with the NLL objective underpinning Elo.

$$= \sum_{v\in[\succeq]} \sum_{i,j\in I_2(v)} \sigma(\theta_j - \theta_i) \tag{39}$$
$$= \sum_{v\in[\succeq]} \sum_{i,j\in I_2(v)} 1 - \sigma(\theta_i - \theta_j) \tag{40}$$
$$\mathrm{SCO}(\theta) = \sum_{v\in[\succeq]} \sum_{i,j\in I_2(v)} y_{ij} - \sigma(\theta_i - \theta_j) \tag{41}$$

Therefore, for every entry X_k in the design matrix containing the pairwise comparison $i \succ j$, the ratings $\theta_i, \theta_j \in \Theta$ receive from the SCO objective the gradients:

$$\nabla_{\theta_i}\mathrm{SCO}(\theta) = -\nabla_\theta \sigma(\theta_i - \theta_j) \tag{42}$$
$$= -\sigma(\theta_i - \theta_j)(1 - \sigma(\theta_i - \theta_j)) \tag{43}$$
$$\nabla_{\theta_j}\mathrm{SCO}(\theta) = -\nabla_\theta \sigma(\theta_i - \theta_j) \tag{44}$$
$$= \sigma(\theta_i - \theta_j)(1 - \sigma(\theta_i - \theta_j)) \tag{45}$$

Then, from the Elo objective each involved rating receive the gradients:

$$\nabla_{\theta_i}\mathrm{NLL}(\theta) = -\nabla_\theta \log\sigma(\theta_i - \theta_j) \tag{46}$$
$$= -(1 - \sigma(\theta_i - \theta_j)) \tag{47}$$
$$\nabla_{\theta_j}\mathrm{NLL}(\theta) = -\nabla_{\theta_j} \log\sigma(\theta_i - \theta_j) \tag{48}$$
$$= 1 - \sigma(\theta_i - \theta_j) \tag{49}$$

Therefore, one stochastic gradient descent (SGD) update of SCO is performed as:

$$\begin{bmatrix} \theta_i \\ \theta_j \end{bmatrix}_k = \begin{bmatrix} \theta_i \\ \theta_j \end{bmatrix}_{k-1} - \eta\sigma(\theta_i - \theta_j) \begin{bmatrix} -\sigma(\theta_j - \theta_i) \\ \sigma(\theta_j - \theta_i) \end{bmatrix}_{k-1} \tag{50}$$

while the NLL update through SGD looks like:

$$\begin{bmatrix} \theta_i \\ \theta_j \end{bmatrix}_k = \begin{bmatrix} \theta_i \\ \theta_j \end{bmatrix}_{k-1} - \eta \begin{bmatrix} -\sigma(\theta_j - \theta_i) \\ \sigma(\theta_j - \theta_i) \end{bmatrix}_{k-1} \tag{51}$$

Conclusion 1. A damping coefficient $\epsilon = \sigma(\theta_i - \theta_j)$ is the difference between SCO and Elo objectives.

6 Problems with Elo

Inconsistent Estimator. Chen et al. (1999) established that an MLE estimator is a strongly consistent estimator of a Generalized Linear Model (GLM) if the eigenvalues of the *Gramian* $X^{\mathrm{T}}X$ are bounded away from 0 (i.e., $X^{\mathrm{T}}X$ is non-singular). But, in the data representation Elo leverages, the Gramian $X^{\mathrm{T}}X$ is **always** *singular*. Therefore, MLE may not be the appropriate procedure to estimate the parameters. We confirmed this empirically. Both GLM parameterizations of the Bernoulli distribution (*logit* and *probit*) present the same issue with the Pentathlon and AAMAS example data.

7 Future Work

There seems to be a connection between SCO & Elo and diffusion processes on graphs.

References

M. Alvo and P. L. H. Yu. *Statistical Methods for Ranking Data.* Springer New York, Sept. 2014.

D. Balduzzi, K. Tuyls, J. Perolat, and T. Graepel. Re-evaluating evaluation. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, pages 3272–3283, Red Hook, NY, USA, Dec. 2018. Curran Associates Inc.

Q. Bertrand, W. M. Czarnecki, and G. Gidel. On the Limitations of the Elo, Real-World Games are Transitive, not Additive. In F. Ruiz, J. Dy, and J.-W. van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 2905–2921. PMLR, 2023.

R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324, Dec. 1952.

K. Chen, I. Hu, and Z. Ying. Strong consistency of maximum quasi-likelihood estimators in generalized linear models with fixed and adaptive designs. *Annals of statistics*, 27(4):1155–1163, Aug. 1999.

W.-L. Chiang, L. Zheng, Y. Sheng, A. N. Angelopoulos, T. Li, D. Li, H. Zhang, B. Zhu, M. Jordan, J. E. Gonzalez, and I. Stoica. Chatbot Arena: An Open Platform for Evaluating LLMs by Human Preference. *arXiv [cs.AI]*, Mar. 2024.

D. E. Critchlow, M. A. Fligner, and J. S. Verducci. Probability models on rankings. *Journal of mathematical psychology*, 35(3):294–318, Sept. 1991.

P. Diaconis. *Group representations in probability and statistics*, volume 11 of *Lecture Notes-Monograph Series.* Institute of Mathematical Statistics, Jan. 1988.

A. E. Elo. The Proposed USCF Rating System: Its Development, Theory, and Application. *Chess Life*, XXII(8):242–249, Aug. 1967.

M. A. Fligner and J. S. Verducci. Distance based ranking models. *Journal of the Royal Statistical Society. Series B, Statistical methodology*, 48(3):359–369, 1986.

R. Herbrich, T. Minka, and T. Graepel. TrueSkill™: A Bayesian Skill Rating System. *Advances in Neural Information Processing Systems*, 19, 2006.

D. R. Hunter. MM algorithms for generalized Bradley-Terry models. *The Annals of Statistics*, 32 (1):384–406, Feb. 2004.

J. G. Kemeny. Mathematics without Numbers. *Daedalus*, 88(4):577–591, 1959.

M. Lanctot, K. Larson, Y. Bachrach, L. Marris, Z. Li, A. Bhoopchand, T. Anthony, B. Tanner, and A. Koop. Evaluating Agents using Social Choice Theory. *arXiv [cs.AI]*, Dec. 2023.

M. Lanctot, K. Larson, M. Kaisers, Q. Berthet, I. Gemp, M. Diaz, R.-R. Maura-Rivero, Y. Bachrach, A. Koop, and D. Precup. Soft Condorcet Optimization for Ranking of General Agents. *arXiv [cs.MA]*, Oct. 2024.

J. I. Marden. *Analyzing and Modeling Rank Data*. Chapman and Hall, 1995.

N. B. Shah, S. Balakrishnan, A. Guntuboyina, and M. J. Wainwright. Stochastically Transitive Models for Pairwise Comparisons: Statistical and Computational Issues. *arXiv [stat.ML]*, Oct. 2015.

H. P. Young and A. Levenglick. A consistent extension of Condorcet's election principle. *SIAM journal on applied mathematics*, 35(2):285–300, Sept. 1978.

A Scrapbook

Quasi-Newton secant condition:

$$B_k[\theta_k - \theta_{k-1}] = \nabla f(\theta_k) - \nabla f(\theta_{k-1}) \tag{52}$$

Let's compute $\Delta_{f_k} = \nabla f(\theta_k) - \nabla f(\theta_{k-1})$ for the NLL loss.

$$\begin{aligned}
\Delta_f &= \begin{bmatrix} -\sigma_{ji}^k \\ \sigma_{ji}^k \end{bmatrix} - \begin{bmatrix} -\sigma_{ji}^{k-1} \\ \sigma_{ji}^{k-1} \end{bmatrix} && (53)\\
&= \begin{bmatrix} -\sigma_{ji}^k + \sigma_{ji}^{k-1} \\ \sigma_{ji}^k - \sigma_{ji}^{k-1} \end{bmatrix} && (54)\\
&= \begin{bmatrix} \sigma_{ji}^{k-1} - \sigma_{ji}^k \\ \sigma_{ji}^k - \sigma_{ji}^{k-1} \end{bmatrix} && (55)
\end{aligned}$$

Then, we know that $\theta^{k+1} = \theta^k + s^k$. Thus we have the following:

$$\begin{aligned}
\sigma_{ji}^k &= \sigma(\theta_j^k - \theta_i^k) && (56)\\
&= \sigma(\theta_j^{k-1} + s_j^k - (\theta_i^{k-1} + s_i^k)) && (57)\\
&= \sigma(\theta_j^k - \theta_i^k + (s_j^k - s_i^k)) && (58)
\end{aligned}$$

And thus,

$$\sigma_{ji}^{k+1} - \sigma_{ji}^k = \sigma(\theta_j^k - \theta_i^k + (s_j^k - s_i^k)) - \sigma(\theta_j^k - \theta_i^k) \tag{59}$$

We need the following identity:

$$\sigma(x+c) - \sigma(x) = \frac{e^{x+c}}{1+e^{x+c}} - \frac{e^x}{1+e^x} \tag{60}$$
$$= \frac{e^{x+c}(1+e^x) - e^x(1+e^{x+c})}{(1+e^{x+c})(1+e^x)} \tag{61}$$
$$= \frac{e^{x+c} - e^x}{(1+e^{x+c})(1+e^x)} \tag{62}$$
$$= \frac{e^x(e^c - 1)}{(1+e^{x+c})(1+e^x)} \tag{63}$$
$$= \sigma(x)\frac{e^c - 1}{1+e^{x+c}} \tag{64}$$
$$= \sigma(x)\frac{1}{1+e^{x+c}}[e^c - 1] \tag{65}$$
$$= \sigma(x)\sigma(-x-c)[e^c - 1] \tag{66}$$
$$= \sigma(x)[1 - \sigma(x+c)][e^c - 1] \tag{67}$$

Thus,

$$\Delta_f = \begin{bmatrix} \sigma(\theta_j^k - \theta_i^k)[1 - \sigma(\theta_j^k - \theta_i^k + (s_j^k - s_i^k))][e^{s_j^k - s_i^k} - 1] \\ -\sigma(\theta_j^k - \theta_i^k)[1 - \sigma(\theta_j^k - \theta_i^k + (s_j^k - s_i^k))][e^{s_j^k - s_i^k} - 1] \end{bmatrix} \tag{68}$$
$$= [e^{s_j^k - s_i^k} - 1] \begin{bmatrix} \sigma(\theta_j^k - \theta_i^k)[1 - \sigma(\theta_j^k - \theta_i^k + (s_j^k - s_i^k))] \\ -\sigma(\theta_j^k - \theta_i^k)[1 - \sigma(\theta_j^k - \theta_i^k + (s_j^k - s_i^k))] \end{bmatrix} \tag{69}$$

A.1 Negative Log-Likelihood

Hessian. The Hessian of the NLL objective is given by:

$$H_{\text{NLL}} = X^{\mathrm{T}}[\hat{Y}(Y - \hat{Y})^{\mathrm{T}}]X \tag{70}$$

A.2 SCO

B How to Fix Your Elo

Intepretation Of A Complete Ranking.. We can leverage the intuition behind the Plackett-Luce model to understand a ranking, e.g., $R = A > B > C$, as a tournament when first A won and everyone else lost so $[1, -1, -1]$, next B won and everyone else (C) lost so $[0, 1, -1]$.

$$p(i \text{ wins}) = \frac{1}{1 + \sum_{k \neq i} e^{\theta_k - \theta_i}} \tag{71}$$

B.1 Second Order Analysis

B.1.1 Negative Log Likelihood

We start by the NLL(θ)

$$\mathrm{H}_{\text{NLL}} = \begin{bmatrix} \nabla_{\theta_{ii}} & \nabla_{\theta_{ij}} \\ \nabla_{\theta_{ji}} & \nabla_{\theta_{jj}} \end{bmatrix} \tag{72}$$

$$
\begin{aligned}
\nabla_{\theta_{ii}} &= \nabla_{\theta_i}\nabla_{\theta_i}\text{NLL}(\theta) && (73)\\
&= \nabla_{\theta_i}\left[-(1-\sigma(\theta_i-\theta_j))\right] && (74)\\
&= \nabla_{\theta_i}\sigma(\theta_i-\theta_j)-1 && (75)\\
&= \sigma(\theta_i-\theta_j)[1-\sigma(\theta_i-\theta_j)] && (76)\\
&= \sigma(\theta_i-\theta_j)\sigma(\theta_j-\theta_i) && (77)\\
\nabla_{\theta_{ij}} &= \nabla_{\theta_j}\nabla_{\theta_i}\text{NLL}(\theta) && (78)\\
&= \nabla_{\theta_j}\sigma(\theta_i-\theta_j)-1 && (79)\\
&= -\sigma(\theta_i-\theta_j)[1-\sigma(\theta_i-\theta_j)] && (80)\\
&= -\sigma(\theta_i-\theta_j)\sigma(\theta_j-\theta_i) && (81)\\
\nabla_{\theta_{ji}} &= \nabla_{\theta_i}\nabla_{\theta_j}\text{NLL}(\theta) && (82)\\
&= \nabla_{\theta_i}(1-\sigma(\theta_i-\theta_j)) && (83)\\
&= -\sigma(\theta_i-\theta_j)[1-\sigma(\theta_i-\theta_j)] && (84)\\
&= -\sigma(\theta_i-\theta_j)\sigma(\theta_j-\theta_i) && (85)\\
\nabla_{\theta_{jj}} &= \nabla_{\theta_j}\nabla_{\theta_j}\text{NLL}(\theta) && (86)\\
&= \nabla_{\theta_j}(1-\sigma(\theta_i-\theta_j)) && (87)\\
&= \sigma(\theta_i-\theta_j)[1-\sigma(\theta_i-\theta_j)] && (88)\\
&= \sigma(\theta_i-\theta_j)\sigma(\theta_j-\theta_i) && (89)
\end{aligned}
$$

$$
\begin{aligned}
\text{H}_{\text{NLL}} &= \begin{bmatrix} \sigma(\theta_i-\theta_j)\sigma(\theta_j-\theta_i) & -\sigma(\theta_i-\theta_j)\sigma(\theta_j-\theta_i) \\ -\sigma(\theta_i-\theta_j)\sigma(\theta_j-\theta_i) & \sigma(\theta_i-\theta_j)\sigma(\theta_j-\theta_i) \end{bmatrix} && (90)\\
&= \sigma(\theta_i-\theta_j)\begin{bmatrix} \sigma(\theta_j-\theta_i) & -\sigma(\theta_j-\theta_i) \\ -\sigma(\theta_j-\theta_i) & \sigma(\theta_j-\theta_i) \end{bmatrix} && (91)
\end{aligned}
$$

Singular Hessian. The Hessian is singular in this case.

$$\det(\text{H}_{\text{NLL}}) = 0 \tag{92}$$

Thus, we could compute the pseudo inverse H^+:

$$
\begin{aligned}
H^+_{NLL} &= (H^T H)^{-1} H^T && (93)\\
&= \frac{1}{2}\begin{bmatrix} \frac{1}{\sigma_{ij}\sigma_{ji}} & 0 \\ 0 & \frac{1}{\sigma_{ij}\sigma_{ji}} \end{bmatrix} && (94)
\end{aligned}
$$

Then,

$$H^+_{\text{NLL}}(\theta)\nabla_\theta \text{NLL}(\theta) = \frac{1}{2}\begin{bmatrix} -\frac{1}{\sigma_{ij}} \\ \frac{1}{\sigma_{ij}} \end{bmatrix} \tag{95}$$

B.2 Connection With Graph Diffussion

$$L = X^T X = (W - M) + (L - M^T) \tag{96}$$

B.2.1 Connection with The Comparison Laplacian

SCO, Elo & Comparison Multigraph. The Gramian matrix $L = X^{\mathrm{T}}X$ is the Laplacian matrix of the comparison multigraph. In this equivalence, X is the gradient operator (incidence matrix), and X^{T} is the divergence operator (the transpose or conjugate).