

1 Derivation of method

The derivation follows largely the program given in [1].

For our model, each player i is assigned a category $c(i) \in \{1, \dots, C\}$ and a skill vector

$$\Theta_i^{(k)} = (\zeta_i^{(k)}, \eta_i^{(k),1}, \dots, \eta_i^{(k),C})$$

at each period k . The superscript (k) will mostly be suppressed in the following.

Each element is assumed normally distributed with a given mean and variance, independent from each other:

$$\zeta_i \sim N(\zeta_i | \mu(\zeta_i), \sigma(\zeta_i)^2), \quad \eta_i^c \sim N(\eta_i^c | \mu(\eta_i^c), \sigma(\eta_i^c)^2).$$

The interpretation being that a player players with strength $\zeta + \eta^c$ against another player of category c . Thus, ζ is a general indicator of skill, and η^c is a category modifier. We will assume the normalization

$$\sum_c \eta^c = 0.$$

The rating model we use differs from the one in [1] in that we use a normal distribution rather than a logistic. Thus, when two players i and j play, we have an outcome s with probability distribution

$$P(s = 1) = \Phi(\Delta_{i,j}), \quad P(s = -1) = \Phi(-\Delta_{i,j}),$$

where Φ is the cumulative standard normal distribution function, and

$$\Delta_{i,j} = \zeta_i - \zeta_j + \eta_i^{c(j)} - \eta_j^{c(i)}$$

is the difference of the playing strengths.

Assume now that a player with prior playing strength Θ and category c plays a set of games against various opponents, with the result of game k against opponent j denoted as $s_{j,k} \in \{-1, 1\}$. As given in [1], we approximate the posterior likelihood of Θ as

$$f(\Theta | s) \propto \bar{\varphi}(\Theta | \bar{\mu}, \bar{\sigma}^2) \prod_{j,k} \int \Phi(s_{j,k} \Delta_j) \bar{\varphi}(\Theta_j | \bar{\mu}_j, \bar{\sigma}_j^2) d\Theta_j = \bar{\varphi}(\Theta | \bar{\mu}, \bar{\sigma}^2) \prod_{j,k} I_{j,k},$$

where $\bar{\varphi}$ denotes the distribution functions of Θ , Θ_j (which is a product of the normal distributions of each element), $\Delta_j = \zeta - \zeta_j + \eta^{c(j)} - \eta_j^c$ is the strength gap, and each integral runs over \mathbb{R}^{C+1} .

We will now evaluate the integrals $I_{j,k}$.

First, we note that the integrand in each of the variables η_j^d for $d \neq c$ is just a normal distribution, thus we can integrate them away to get

$$I_{j,k} = \int_{\mathbb{R}} \int_{\mathbb{R}} \int_{-\infty}^{s_{j,k} \Delta_j} \varphi(\alpha | 0, 1) \varphi(\zeta_j | \mu(\zeta_j), \sigma(\zeta_j)^2) \varphi(\eta_j^c | \mu(\eta_j^c), \sigma(\eta_j^c)^2) d\alpha d\zeta_j d\eta_j^c.$$

This integral can be evaluated in a CAS (Computer Algebra System) such as SAGE, and we find

$$I_{j,k} \propto \begin{cases} \Phi(\zeta + \eta^{c(j)} | \bar{\mu}_j, \bar{\sigma}_j^2), & s_{j,k} = 1, \\ 1 - \Phi(\zeta + \eta^{c(j)} | \bar{\mu}_j, \bar{\sigma}_j^2), & s_{j,k} = -1, \end{cases}$$

with the modified means and variances

$$\bar{\mu}_j = \mu(\zeta_j) + \mu(\eta_j^c), \quad \bar{\sigma}_j^2 = 1 + \sigma(\zeta_j)^2 + \sigma(\eta_j^c)^2.$$

Now, to simplify notation, assume that there are in total W_j wins and L_j losses against player j . We then find

$$f(\Theta|s) \propto \varphi(\Theta | \bar{\mu}, \bar{\sigma}^2) \prod_j \Phi_j(M_j)^{W_j} (1 - \Phi_j(M_j))^{L_j}, \quad (1)$$

where $M_j = \zeta + \eta^{c(j)}$.

As in [1], we now aim to approximate the product in (1) with a normal distribution in each variable. To this end, we take the logarithm, and rely on numerical optimization methods to find the maximum $\hat{\Theta}$. The gradient and Hessian is available, and it is our experience that this optimization is generally quite quick. We have used the NCG algorithm, but other, such as BFGS, Powell's iteration and a simple downhill simplex method work as well.

Thus, the product in (1) is approximated by a normal distribution with mean $\hat{\Theta}$ and variances $\hat{\sigma}^2$ given by the negative reciprocals of the diagonal elements in the Hessian evaluated at $\hat{\Theta}$.

The ratings are then updated exactly as in [1], namely

$$(\hat{\sigma}^2)' = \left(\frac{1}{\hat{\sigma}^2} + \frac{1}{\bar{\sigma}^2} \right)^{-1}$$

and

$$\bar{\mu}' = (\hat{\sigma}^2)' \left(\frac{\hat{\Theta}}{\hat{\sigma}^2} + \frac{\bar{\mu}}{\bar{\sigma}^2} \right)$$

where all algebraic operations are understood to be element-wise.

2 Notes

1. To keep the optimization well-defined, the normalization condition for all η^c must be enforced.
2. Even so, it is not guaranteed that in the updated mean, this condition remains true. Thus, after the update, we subtract the mean of the η^c from each η^c , and add it to the general rating ζ . This does not change the playing strengths.
3. If a player does not play during a period, his or her variance is updated according to the formula given in [1]. This also applies to category-specific rating variances $\sigma(\eta^d)^2$ if a player does not play any games against a player of category d .

4. To prevent ratings from becoming stale, a rating variance floor should be enforced. Thus, after updating the ratings

$$\vec{\sigma}' \leftarrow \max(\vec{\sigma}', \sigma_{\text{floor}}).$$

3 Parameters

The system should be fitted in each instance with the parameters:

- The rating variance floor σ_{floor} .
- The initial rating variances σ_{init} for new players.
- The period length.

References

- [1] M. Glickman. Parameter estimation in large dynamic paired comparison experiments. *Applied Statistics*, 48:377–394, 1999.