Trueskill^{TM}: A Bayesian skill rating system

Herbrich, Ralf, Tom Minka, and Thore Graepel. *NIPS*. 2006.

Kuhwan Jeong [1]

[1]Department of Statistics, Seoul National University, South Korea

January, 2017

Introduction

- TrueSkill ranking system is a skill based ranking system for Xbox Live developed at Microsoft Research.
- The purpose is to both identify and track the skills of gamers in order to be able to match them into competitive matches.
- TrueSkill ranking system only uses the final standings of all teams in a game in order to update the skill estimates of all gamers playing in this game.
- Ranking systems have been proposed for many sports but possibly the most prominent ranking system in use today is ELO.

ELO

- In 1959, Arpad Elo developed a statistical rating system for Chess.

- ELO models the probability of the game outcomes as a function of the two players' skill ratings s_1 and s_2 :

$$P(\text{Player 1 wins}|s_1, s_2) = \Phi\left(\frac{s_1 - s_2}{\sqrt{2}\beta}\right).$$

- Let $y = 1$ if player 1 wins, $y = -1$ if player 2 wins and $y = 0$ if a draw occurs.

- Then the ELO update is given by $s_1 \leftarrow s_1 + y\Delta, s_2 \leftarrow s_2 - y\Delta$ and

$$\Delta = \alpha\beta\sqrt{\pi}\left(\frac{y+1}{2} - \Phi\left(\frac{s_1 - s_2}{\sqrt{2}\beta}\right)\right).$$

TrueSkill

- Skill s is characterized by two numbers:
  - The average μ skill of the gamer.
  - The degree of uncertainty σ in the gamer's skill.

- From among a population of n players $\{1,\dots,n\}$ in a game let k teams compete a match.

- The team assignments are specified by k non-overlapping subsets

$$A_j \subset \{1,\dots,n\}, A_i \cap A_j = \text{ if } i \neq j.$$

TrueSkill

- Assume an independent normal prior $p(s) = \prod_{i=1}^{n} \mathcal{N}(s_i; \mu_i, \sigma_i^2)$.
- Each player i exhibits a performance $p_i \sim \mathcal{N}(p_i; s_i, \beta^2)$.
- The performance t_j of team j is modeled as the sum of the performances of its members $t_j = \sum_{i \in A_j} p_i$.
- The outcome $\mathbf{r} = (r_1, \ldots, r_k) \in \{1, \ldots, k\}^k$ is specified by a rank r_j for each team j.
- Disregarding draws, the probability of a game outcome r is modeled as

$$P(r|t_1, \ldots, t_k) = P(t_{r_{(1)}} > t_{r_{(2)}} > \cdots > t_{r_{(k)}}).$$

- If draws are permitted the wining outcome $r_{(j)} < r_{(j+1)}$ requires $t_{r_{(j)}} > t_{r_{(j+1)}} + \epsilon$ and the draw outcome $r_{(j)} = r_{(j+1)}$ requires $|t_{r_{(j)}} - t_{r_{(j+1)}}| \leq \epsilon$, where ϵ is a draw margin.

TrueSkill

- The posterior distribution $P(\mathbf{s}|\mathbf{r}, A)$ is approximated by a normal distribution $Q(\mathbf{s})$.

- μ_i and σ_i^2 are updated to the mean and variance of the marginal distribution $Q(s_i) = \int_{\mathbf{s}_{-i}} Q(\mathbf{s}) d\mathbf{s}_{-i}$, respectively.

Example

- Consider a game with $k = 3$ teams with team assignments $A_1 = \{1\}, A_2 = \{2, 3\}$ and $A_3 = \{4\}$.
- Assume that team 1 is the winner and that teams 2 and 3 draw, i.e., $\mathbf{r} = (1, 2, 2)$.
- Then the conditional probability of $\mathbf{s}, \mathbf{p}$ is

$$\begin{aligned} P(\mathbf{s}, \mathbf{p} | \mathbf{r}, A) \propto & \prod_{i=1}^{4} \left\{ \phi\big(s_i; \mu_i, \sigma_i^2\big) \phi\big(p_i; s_i, \beta^2\big) \right\} \\ & \times I(p_1 - p_2 - p_3 > \epsilon) I(|p_2 + p_3 - p_4| \leq \epsilon) \\ = & \ \Phi_1(\mathbf{s}, \mathbf{p}) I(p_1 - p_2 - p_3 > \epsilon) I(|p_2 + p_3 - p_4| \leq \epsilon). \end{aligned}$$

Example

- If we introduce auxiliary variables $d_1 = p_1 - p_2 - p_3$ and $d_2 = p_2 + p_3 - p_4$, then

$$
\begin{aligned}
P(\mathbf{s}, \mathbf{p}, \mathbf{d}|\mathbf{r}, A) \propto\ & \Phi_1(\mathbf{s}, \mathbf{p}) I(d_1 > \epsilon) I(d_1 = p_1 - p_2 - p_3) \\
& \times I(|d_2| \leq \epsilon) I(d_2 = p_2 + p_3 - p_4)
\end{aligned}
$$

and $\int P(\mathbf{s}, \mathbf{p}, \mathbf{d}|\mathbf{r}, A) d\mathbf{d} = P(\mathbf{s}, \mathbf{p}|\mathbf{r}, A)$.

- By integrating $\mathbf{s}$ out, we get

$$
\begin{aligned}
& P(\mathbf{p}, \mathbf{d}|\mathbf{r}, A) \\
= & \int P(\mathbf{s}, \mathbf{p}, \mathbf{d}|\mathbf{r}, A) d\mathbf{s} \\
\propto & \prod_{i=1}^{4} \left\{ \phi\left(p_i; \mu_i, \sigma_i^2 + \beta^2\right) \right\} I(d_1 > \epsilon) I(d_1 = p_1 - p_2 - p_3) \\
& \times I(|d_2| \leq \epsilon) I(d_2 = p_2 + p_3 - p_4) \\
= & \Phi_2(\mathbf{p}) I(d_1 > \epsilon) I(d_1 = p_1 - p_2 - p_3) I(|d_2| \leq \epsilon) I(d_2 = p_2 + p_3 - p_4).
\end{aligned}
$$

Example

- The marginal distribution of d_1 is proportional to

$$
\begin{aligned}
&P(d_1|\mathbf{r}, A) \\
\propto &\int\int \Phi_2(\mathbf{p}) I(d_1 > \epsilon) I(d_1 = p_1 - p_2 - p_3) I(|d_2| \leq \epsilon) I(d_2 = p_2 + p_3 - p_4) dd_2 d\mathbf{p} \\
= &I(d_1 > \epsilon) \int \Phi_2(\mathbf{p}) I(d_1 = p_1 - p_2 - p_3) I(|p_2 + p_3 - p_4| \leq \epsilon) d\mathbf{p}.
\end{aligned}
$$

- We assume that $I(d_1 > \epsilon)$ and $I(|d_2| \leq \epsilon)$ are proportional to $\phi\left(d_1; \hat{m}_1, \hat{s}_1^2\right)$ and $\phi\left(d_2; \hat{m}_2, \hat{s}_2^2\right)$, respectively.

Example

- Given $\hat{m}_2$ and $\hat{s}_2^2$,

$$P(d_1|\mathbf{r},A) \propto I(d_1 > \epsilon) \int I(d_1 = p_1 - p_2 - p_3)\Phi_2(\mathbf{p})\phi\big(p_2 + p_3 - p_4; \hat{m}_2, \hat{s}_2^2\big) d\mathbf{p}$$

- We can find $\tilde{m}_1, \tilde{s}_1^2$ such that

$$\int I(d_1 = p_1 - p_2 - p_3)\Phi_2(\mathbf{p})\phi\big(p_2 + p_3 - p_4; \hat{m}_2, \hat{s}_2^2\big) d\mathbf{p} \propto \phi\big(d_1; \tilde{m}_1, \tilde{s}_1^2\big).$$

- Therefore, $P(d_1|\mathbf{r},A)$ is a truncated normal distribution.
- We approximate $P(d_1|\mathbf{r},A)$ by $\mathcal{N}\big(d_1; m_1, s_1^2\big)$ via moment matching.
- Then we have

$$\phi\big(d_1; m_1, s_1^2\big) \propto I(d_1 > \epsilon)\phi\big(d_1; \tilde{m}_1, \tilde{s}_1^2\big).$$

- So we update $\hat{m}_1$ and $\hat{s}_1^2$ to

$$\hat{m}_1 = \frac{m_1\tilde{s}_1^2 - \tilde{m}_1 s_1^2}{s_1^2 + \tilde{s}_1^2}, \hat{s}_1^2 = \frac{s_1^2\tilde{s}_1^2}{s_1^2 + \tilde{s}_1^2}.$$

Example

- Likewise, we update $\hat{m}_2$ and $\hat{s}_2^2$ given $\hat{m}_1$ and $\hat{s}_1^2$.
- We iterate these updates until the convergence of m_1, m_2, s_1^2, s_2^2.
- Then $P(\mathbf{s}, \mathbf{p}, \mathbf{d}|\mathbf{r}, A)$ is approximated by $Q(\mathbf{s}, \mathbf{p}, \mathbf{d})$ where

$$
\begin{aligned}
Q(\mathbf{s}, \mathbf{p}, \mathbf{d}) &\propto \Phi_1(\mathbf{s}, \mathbf{p}) I(d_1 = p_1 - p_2 - p_3) I(d_2 = p_2 + p_3 - p_4) \\
&\quad \times \phi\big(d_1; \hat{m}_1, \hat{s}_1^2\big) \phi\big(d_2; \hat{m}_2, \hat{s}_2^2\big) \\
Q(\mathbf{s}, \mathbf{p}) &\propto \Phi_1(\mathbf{s}, \mathbf{p}) \phi\big(p_1 - p_2 - p_3; \hat{m}_1, \hat{s}_1^2\big) \phi\big(p_2 + p_3 - p_4; \hat{m}_2, \hat{s}_2^2\big) \\
Q(s_i) &\propto \phi\big(s_i; \mu_i, \sigma_i^2\big) \int \phi\big(p_i; s_i, \beta^2\big) \prod_{j \neq i} \phi\big(p_j; \mu_j, \sigma_j^2 + \beta^2\big) \\
&\quad \times \phi\big(p_1 - p_2 - p_3; \hat{m}_1, \hat{s}_1^2\big) \phi\big(p_2 + p_3 - p_4; \hat{m}_2, \hat{s}_2^2\big) d\mathbf{p}
\end{aligned}
$$

- We update μ_i and σ_i^2 by the mean and variance of $Q(s_i)$.