A Graphical Model for Skill Ratings in Competitive Games

April 8, 2013

1 Introduction

Skill ratings in games and sports have three main functions:

1. Players of similar skill can be matched, what leads to interesting, balanced matches.
2. Public ratings stimulate interest and competition.
3. Ratings can be used as criteria of qualification for tournaments.

In this exercise we will describe a probabilistic model that can be used to model the outcomes of different matches among multiple players in terms of the players' skill. We will assume a scenario in which several players can form a team and competition can be among different teams of players. The model that we will describe is based on the TrueSkill™ system [1].

2 Data Description

We assume there is a population of n players with ids $1, \ldots, n$ and k teams compete all of them in several matches. For each match, the teams are formed by k disjoint subsets $A_1, \ldots, A_k \subset \{1, \ldots, n\}$ of players, where $A_j \cap A_i = \emptyset$ if $i \neq j$. The outcome of a match is a k-dimensional vector $\mathbf{r} = (r_1, \ldots, r_k)^{\mathrm{T}}$ such that $r_i \in \{1, \ldots, k\}$ specifies a rank for the i-th team, with $r_i = 1$ if this is the winning team and $r_j = k$ if j-th is the losing team. The data available for the m-th match is $\mathcal{D}_m = \{P_m = \{A_1^m, \ldots, A_k^m\}, \mathbf{r}^m\}$, where P_m represents the partition of the players into k disjoint teams, that is, P_m is a set with the sets containing the players in each team and $\mathbf{r}^m$ is the k-dimensional vector encoding the outcome of the match. The complete data available is a dataset $\mathcal{D} = \{\mathcal{D}_1, \ldots, \mathcal{D}_M\}$ with the results of a total of M matches.

Our objective is to formulate a probabilistic graphical model that can (i) explain the observed data $\mathcal{D}$ and (ii) make predictions about an unknown outcome vector $\mathbf{r}_{m+1}$ for a match that has not taken place yet and in which the players will be split in teams according to the partition $P_{m+1} = \{A_1^{m+1}, \ldots, A_k^{m+1}\}$.

3 Probabilistic Model

In this section we propose a generative model for the data $\mathcal{D}$. This requires that we formalize in terms of random variables the process by which teams of players compete with each other in different games.

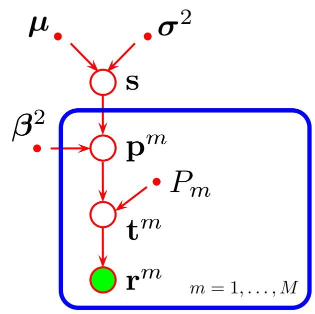

Figure 1: Probabilistic graphical model for skill rating in competitive games.

For this, we assume that for each player i there is a scalar number $s_i \in \mathbb{R}$ representing the true skill of that player. However, for different reasons, the player's performance may deviate slightly from s_i on a particular game. For example, in a sport game, player i could be injured. Therefore, we introduce the scalar p_i^m to represent the performance of the i-th player on the m-th game. The n-dimensional vector $\mathbf{p}^m = (p_1^m, \ldots, p_n^m)$ contains the performances of all the players on the m-th game. We assume that $\mathbf{p}^m$ is obtained from $\mathbf{s} = (s_1, \ldots, s_n)^\mathrm{T}$ as follows:

$$p(\mathbf{p}^m|\mathbf{s}) = \prod_{i=1}^{n} \mathcal{N}(p_i^m|s_i, \beta^2)\,,$$

where β is an estimate of the expected change in the performance of a player across different matches. Furthermore, we assume that the performance t_j^m of team j during the m-th match is given by the sum of the performances of its team members on that match. Let $\mathbf{t}^m = (t_1^m, \ldots, t_k^m)^\mathrm{T}$ be the k-th dimensional vector with the performance of each team on match m, then

$$p(\mathbf{t}^m|\mathbf{p}^m, P_m) = \prod_{j=1}^{k} \delta(t_j - \sum_{i \in A_j^m} p_i^m)\,,$$

where $\delta(\cdot)$ is a point probability mass at the origin. We assume that the probability of the match outcome $\mathbf{r}^m = (r_1^m, \ldots, r_k^m)^\mathrm{T}$ is determined by the k-th dimensional vector $\mathbf{t}^m = (t_1^m, \ldots, t_k^m)^\mathrm{T}$ as follows:

$$\begin{aligned} p(\mathbf{r}^m|\mathbf{t}^m) &= \begin{cases} 1 & \text{if} \quad t^m_{r_1^m} > \ldots > t^m_{r_k^m} \\ 0 & \text{otherwise} \end{cases} \\ &= \prod_{i=1}^{k-1} \Theta[t^m_{r_i^m} - t^m_{r_{i+1}^m}]\,, \end{aligned}$$

where $\Theta[x] = 0.5(\mathrm{sign}(x) + 1)$ is the Heaviside step function. Finally, we select a factorizing Gaussian prior on the vector of true skills $\mathbf{s} = (s_1, \ldots, s_n)^\mathrm{T}$, that is,

$$p(\mathbf{s}) = \prod_{i=1}^{n} \mathcal{N}(s_i|\mu_i, \sigma_i^2)\,.$$

Figure 1 shows the resulting graphical model.

4 Posterior and Predictive Distributions

The posterior distribution for $\mathbf{s}$ given $\mathcal{D}$ is

$$p(\mathbf{s}|\mathcal{D}) = \frac{[\prod_{m=1}^{M} p(\mathbf{r}^m|\mathbf{t}^m)p(\mathbf{t}^m|\mathbf{p}^m, A_j^m)p(\mathbf{p}^m|\mathbf{s})]p(\mathbf{s})}{p(\mathcal{D})}\,.$$

The predictive distribution for $\mathbf{r}_{m+1}$ given $P_{m+1} = \{A_1^{m+1}, \dots, A_k^{m+1}\}$ and $\mathcal{D}$ is

$$p(\mathbf{r}_{m+1}|P_{m+1}, \mathcal{D}) \quad = \quad \int p(\mathbf{r}^{m+1}|\mathbf{t}^{m+1})p(\mathbf{t}^{m+1}|\mathbf{p}^{m+1}, P_{m+1})p(\mathbf{p}^{m+1}|\mathbf{s})p(\mathbf{s}|\mathcal{D})\, d\mathbf{t}^{m+1} d\mathbf{s}\,.$$

References

[1] Ralf Herbrich, Tom Minka, Thore Graepel. TrueSkill:::TM:::: A Bayesian Skill Rating System. In Bernhard Schölkopf, John C. Platt, Thomas Hoffman, editors, Advances in Neural Information Processing Systems 19, Proceedings of the Twentieth Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 4-7, 2006. pages 569-576, MIT Press, 2006