A Bayesian rating system using W-Stein's identity

Ruby Chiu-Hsing Weng

Department of Statistics
National Chengchi University

2011.12.16

Joint work with C.-J. Lin

Outline

- The rating problem
- Review online rating systems
- Bayesian approximation methods
- Our approach
- Experiments on game data

Rating/Ranking

- paired comparison: chess, tennis
- multiple comparison: car racing, horse racing
- multiple teams/players comparison: double tennis, bridge, sports, online games
  - 3 teams: (A1,A2,A3), (B1,B2), (C1,C2,C3,C4)
- Questions: Who will win the next game? Top 10? A global ranking?

Online rating

Online rating: online method for rating
An online method learns cases sequentially.

- procedure:
  - predict the next outcome
  - soon the outcome is available
  - refine the prediction model
- discard cases after learning; large-scale data
- Online rating:
  - rank individuals/estimate skills (after each game)

Outline

- The rating problem
- Review online rating systems
- Bayesian approximation methods
- Our approach
- Experiments on game data

Online rating systems

- Elo (Elo (1960, 1986)): the first chess rating system with probabilistic underpinning
  - US Chess Federation (USCF), World Chess Federation (FIDE), world football league, etc
- Glicko (Glickman (1992)): Bayesian rating system
  - the free internet chess server, etc
- TrueSkill$^{\text{TM}}$ (Herbrich, Graepel and Minka (Microsoft Research 2006)): Bayesian rating system
  - Microsoft Xbox Live

  (multiple teams/players)

Elo vs Glicko

Elo (a physics professor):

player's skill characterized by strength θ

$$\theta_i \leftarrow \theta_i + K(s_i - P(\text{i wins}))$$

Glicko (a statistics professor):

$$\theta \sim Normal(\mu, \sigma^2)$$

player's skill characterized by (μ, σ^2)

A Bayesian framework

- player's skill $\theta \sim (\mu, \sigma^2)$
- before the game, i's skill is $N(\mu_i, \sigma_i^2)$
- after the game, i's skill is $N(\mu_i^{\mathrm{new}}, (\sigma_i^2)^{\mathrm{new}})$, the posterior mean and variance of θ_i

Outline

- The rating problem
- Review online rating systems
- Bayesian approximation methods
- Our approach
- Experiments on game data

The game model

- prior of stregnth θ_i: $N(\mu_i, \sigma_i^2)$, $i = 1, ..., k$
- D: game outcome
- likelihood: $L(\theta) = P(D|\theta)$
- the posterior density of $\theta = (\theta_1, ..., \theta_k)'$:

$$p(\theta|D) \propto p(\theta) \times P(D|\theta)$$

- goal: derive $E(\theta|D)$ and $\mathrm{Var}(\theta|D)$ fast and accurate

Models for ranked data

- R: observed rank of k objects of a game
- Thurstone-Mosteller (TM) model (1927)
  - X_i: unobserved actual performance of object i
  - $X_i \sim N(\theta_i, \sigma^2)$, θ_i is strength of i.
  - Example, $k = 2$ $\quad R = (1, 2)$ if and only if $X_1 > X_2$.

$$L(\theta) = P(X_1 > X_2) = \Phi(\frac{\theta_1 - \theta_2}{\sqrt{2\sigma^2}})$$

- Bradley-Terry (BT) model: $X_1 - X_2 \sim$ logistic

$$L(\theta) = P(X_1 > X_2) = \frac{\theta_1}{\theta_1 + \theta_2},\ \theta_i > 0$$

Likelihood for a multiple-team game

Suppose R=(1,2,3,4).

- Plackett-Luce (PL) model: (a generalization of BT)

$$L(\theta) = \left(\frac{\theta_1}{\theta_1+\theta_2+\theta_3+\theta_4}\right)\left(\frac{\theta_2}{\theta_2+\theta_3+\theta_4}\right)\left(\frac{\theta_3}{\theta_3+\theta_4}\right)$$

- Alternative: first decompose

$$\begin{aligned} L(\theta) &= P(X_1 > X_2)P(X_1 > X_3)P(X_1 > X_4)P(X_2 > X_3) \\ &\quad P(X_2 > X_4)P(X_3 > X_4) \end{aligned}$$

then use TM or BT for paired comparison.

It involves several terms!

Bayesian approximations

Bayesian inference often involves intractable integrations.
Popular approximation techniques:

- Markov Chain Monte Carlo (MCMC)
  - draw samples approximately from $p(\theta|D)$
  - slow for large-scale data
  - we need only $E(\theta|D)$ and $\mathrm{Var}(\theta|D)$
- Laplace method
- variational Bayes – bound integral by Jensen's ineq.

Glicko and TrueSkill$^{\mathrm{TM}}$

- Glicko: rating of chessplayers
  - two players
  - update skills after a rating period (e.g. a tournament)
  - better to have 5-10 games per player
  - based on CLT and some ad doc methods
    $p(\theta|D) \propto p(\theta) \times P(D|\theta)$

- TrueSkill: rating of players for Microsoft Xbox
  - multiple teams/players
  - update skills after a single game
  - based on EP (expectation propagation)

More on TrueSkill$^{\mathrm{TM}}$

For multiple-team games $L(\theta) = P(\theta|D) = P_1 P_2 \cdots P_m$

- expectation propagation:

1. employ Assumed Density Filtering (ADF)
   (it learns $P_1, P_2, \ldots$ sequentially; sensitive to order)
2. recursively refine it

- update rule:

- analytic form available for two-team game ($k = 2$)
- numerical integration for $k > 2$ (slower)

TrueSkill Ranking System - Microsoft Research

Microsoft Research

Videos Projects Publications People Downloads

Home | Our Research | Connections | Careers

Worldwide Labs | Research Areas | Research Groups

The making of research @ 20 years
Watch and subscribe

> Projects > TrueSkill Ranking System

TrueSkill™ Ranking System

The TrueSkill™ ranking system is a skill based ranking system for Xbox Live developed at Microsoft Research.

- Detailed Description
- Calculators
- FAQ

The TrueSkill ranking system is a skill based ranking system for Xbox Live developed at Microsoft Research. The purpose of a ranking system is to both identify and track the skills of gamers in a game (mode) in order to be able to match them into competitive matches. The TrueSkill ranking system only uses the final standings of all teams in a game in order to update the skill estimates (ranks) 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.

Ranking Players

So, what is so special about the TrueSkill ranking system? In short, the biggest difference to other ranking systems is that in the TrueSkill ranking system skill is characterised by **two** numbers:

- The average skill of the gamer (μ in the picture).
- The degree of uncertainty in the gamer's skill (σ in the picture).

The ranking system maintains a **belief** in every gamer's skill using these two numbers. If the

Outline

- The rating problem
- Review online rating systems
- Bayesian approximation methods
- Our approach
- Experiments on game data

Stein's lemma vs Stein's identity

- Stein's lemma (Stein 1961, 1981):
  - the expectations of normal distributions
  - W is standard normal iff $E(f'(W)) = E(Wf(W))$
  - applications: James-Stein estimator, empirical Bayes
- Stein's Identity (Woodroofe, 1989):
  - the expectations of distributions which are "nearly" normal: $d\Gamma(z) = f(z)\phi(z)dz$
  - coined as W-Stein's identity (Weng and Lin, 2011)
- essentially exchange the orders of integration

The basic equations

Z: a k-dim random vector with density

$$p(z) \propto \phi_p(z) f(z)$$

Let E be expectation w.r.t. Z. By W-Stein's Identity,

$$E(Z) = E\left[\frac{\nabla f(Z)}{f(Z)}\right], \quad E(Z_i Z_j) = \delta_{ij} + E\left[\frac{\nabla^2 f(Z)}{f(Z)}\right]_{ij}$$

∇: gradient w.r.t. z; $[\cdot]_{ij}$: ij component of a matrix
$\delta_{ij} = 1$ if $i = j$ and 0 otherwise.

The game model

- prior of θ_i: $N(\mu_i, \sigma_i^2)$, $i = 1, \ldots, k$
- let $Z_i = \frac{\theta_i - \mu_i}{\sigma_i}$; the posterior of $Z = (Z_1, \ldots, Z_k)$:

$$p(z|D) \propto \phi_k(z) \times f(z)$$

- Here $f(z) = P(D|z)$. If $f(z) = f_1(z) \cdots f_m(z)$, then

$$\begin{aligned} E(Z|D) &= E\left[\frac{\nabla f(Z)}{f(Z)}\right] \\ &= E\left[\frac{\nabla f_1(Z)}{f_1(Z)}\right] + E\left[\frac{\nabla f_2(Z)}{f_2(Z)}\right] + \cdots + E\left[\frac{\nabla f_m(Z)}{f_m(Z)}\right]. \end{aligned}$$

Team skill

$$\mu_i^{\text{new}} = \mu_i + \sigma_i E[Z_i|D] = \mu_i + \sigma_i E\left[\frac{\partial f(Z)/\partial Z_i}{f(Z)}\right]$$

$$\begin{aligned}(\sigma_i^{\text{new}})^2 &= \sigma_i^2 \mathrm{Var}[Z_i|D] = \sigma_i^2 \left(E[Z_i^2] - E[Z_i]^2\right) \\ &= \sigma_i^2 \left(1 + E\left[\frac{\nabla^2 f(Z)}{f(Z)}\right]_{ii} - E\left[\frac{\partial f(Z)/\partial Z_i}{f(Z)}\right]^2\right).\end{aligned}$$

Approximate these expectations at $z_i = 0$, i.e. $\theta_i = \mu_i$
i.e. as if distribution of θ is concentrated on μ.

How accurate is the approximation?

- Assess accuracy for $k = 2$ and Thurstone-Mosteller model $X_i \sim N(\theta_i, \beta_i^2)$
- Joint posterior pdf of (θ_1, θ_2)

$$\propto \phi\left(\frac{\theta_1 - \mu_1}{\sigma_1}\right)\phi\left(\frac{\theta_2 - \mu_2}{\sigma_2}\right)\Phi\left(\frac{\theta_1 - \theta_2}{\sqrt{\beta_1^2 + \beta_2^2}}\right),$$

- Exact $E(\theta_i|D)$ can be derived.
- Suggest correcting our approximation by

$$\beta_i^2 + \sigma_i^2 \leftarrow \beta_i^2.$$

Individual skill

After updating k team skills, we can apply our method to update individual player in each team. Here

- the ith team has n_i players
- the jth player in ith team has strength θ_{ij}
- prior of θ_{ij} is $N(\mu_{ij}, \sigma_{ij}^2)$
- assume $\theta_i = \sum_j \theta_{ij}$
- so, prior of θ_i is $N(\sum_j \mu_{ij}, \sum_j \sigma_{ij}^2)$

Reparametrize $f(z)$

$$\text{Let } Z_{ij} \equiv \frac{\theta_{ij} - \mu_{ij}}{\sigma_{ij}}. \text{ So, } Z_i \equiv \frac{\theta_i - \mu_i}{\sigma_i} = \frac{\sum_j \sigma_{ij} Z_{ij}}{\sigma_i}$$

$$z = [z_1, \ldots, z_k]^T, \quad \bar{z} = [z_{11}, \ldots, z_{1n_1}, \ldots, z_{k1}, \ldots, z_{kn_k}]^T.$$

Probability of game outcome $f(z)$ is:

$$f(z) = f\left(\sum_{j=1}^{n_1} \frac{\sigma_{1j} z_{1j}}{\sigma_1}, \ldots, \sum_{j=1}^{n_k} \frac{\sigma_{kj} z_{kj}}{\sigma_k} \right) = \bar{f}(\bar{z})$$

$$\text{Update rule: } \mu_{ij} \leftarrow \mu_{ij} + \sigma_{ij} \cdot E\left(\frac{\partial \bar{f}(z)/\partial z_{ij}}{\bar{f}} \right)$$

Our update rules

Team skill update:

$$\begin{aligned} \mu_i &\leftarrow \mu_i + \Omega_i \\ \sigma_i^2 &\leftarrow \sigma_i^2 \max(1 - \Delta_i, \kappa) \end{aligned}$$

where κ is a small positive value (0.0001).

Individual skill update:

$$\mu_{ij} \leftarrow \mu_{ij} + \frac{\sigma_{ij}^2}{\sigma_i^2} \Omega_i$$

$$\sigma_{ij}^2 \leftarrow \sigma_{ij}^2 \max\left(1 - \frac{\sigma_{ij}^2}{\sigma_i^2} \Delta_i, \kappa\right)$$

(adjustment to μ_{ij}) $\propto \sigma_{ij}^2$

Outline

Game Data

Data: generated by Bungie Studios during the beta testing of the Xbox title Halo 2.

Game type	# games	# players	description
Free for All	5,943	60,022	up to 8 players in a game
Small Teams	27,539	4,992	up to 12 players in 2 teams
Head to Head	6,227	1,672	2 players in a game
Large Teams	1,199	2,576	up to 16 players in 2 teams

Table: Data summary

Comparison with TrueSkill$^{\mathrm{TM}}$

	BT	PL	TM	TrueSkill
Free for All	*30.59%	31.74%	44.65%	30.82%
Small Teams	33.97%	*33.89%	36.46%	35.23%
Head to Head	32.53%	32.53%	*32.44%	*32.44%
Large Teams	*37.30%	37.67%	39.37%	38.15%

Table: Prediction error.

We apply our method for BT, PL, TM models.
The prediction uses $\mu - 3\sigma$.
red: accuracy beats TrueSkill
*: best among four methods

TM vs BT

- BT seems better than TM
- Thurstone-Mosteller: $X_1 - X_2$ follows normal.
  Bradley-Terry: $X_1 - X_2$ follows logistic.
- Most currently used Elo variants use logistic rather than normal, because it is argued that:

TM vs BT

- BT seems better than TM
- Thurstone-Mosteller: $X_1 - X_2$ follows normal.
  Bradley-Terry: $X_1 - X_2$ follows logistic.
- Most currently used Elo variants use logistic rather than normal, because it is argued that:

  weaker players have significantly greater winning chances than normal model predicts.

Normal vs logistic

Solid line: normal; Dashed line: logistic

x-axis: $\theta_1 - \theta_2$; y-axis: Prob(player 1 wins)

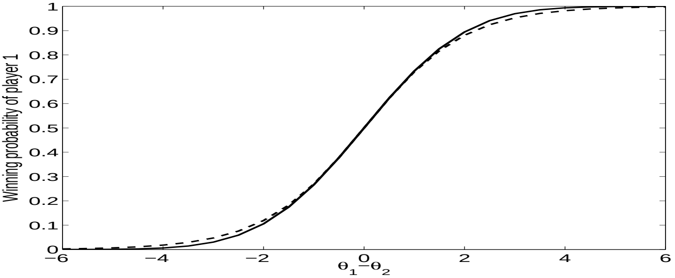

Normal vs logistic

Solid line: normal; Dashed line: logistic

With logistic, weaker player has bigger winning chance

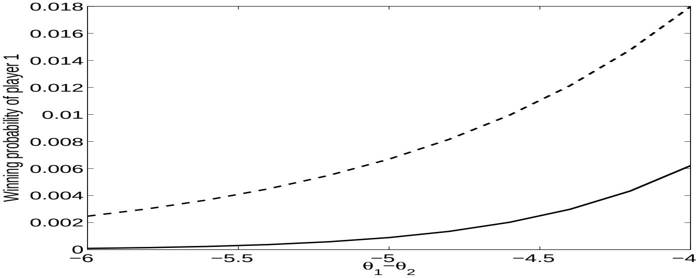

Computation time

- implementation: C and F#
- TrueSkill is written in F#
- "Free for All": using F#

  TrueSkill 13 seconds, ours 1.2 seconds
- competitive accuracy and shorter running time
- http://www.csie.ntu.edu.tw/~cjlin/papers/online_ranking

Comparison with Glicko

	Glicko	TrueSkill	ours
Model	BT	TM	BT, PL, TM
Game type	2 teams	multiple	multiple
Technique	CLT	EP	W-Stein
Update after	a rating period	one game	one game

* The single-game version of Glicko has prediction error 33.88% on Head to Head
* Results by ours and TrueSkill: 32.53% 32.53% 32.44% 32.44%
* Possible reason: Glicko is for "a rating period"

* Glicko considers the increase in variance with the elapse of time

A. E. Elo. *The Rating of Chessplayers, Past and Present*.
Arco Pub., New York, 2nd edition, 1986.

M. Glickman. Parameter estimation in large dynamic paired comparison experiments.
Applied Statistics, 48(3):377–394, 1999.

R. Herbrich, T. Minka, and T. Graepel. TrueSkill$^{\mathrm{TM}}$: A Bayesian skill rating system.
Advances in Neural Information Processing Systems 19, pages 569–576. MIT Press, Cambridge, MA, 2007.

R. C. Weng and C.-J. Lin. A Bayesian approximation method for online ranking.
Journal of Machine Learning Research, 12:267-300, 2011.