Approximate Inference in Practice
Microsoft's Xbox TrueSkill^{TM}

Daniel Hernández-Lobato

Universidad Autónoma de Madrid

2014

Outline

Outline

Introduction

- **Competition is a key aspect of humans.**
  - Innate in most persons since a young age.
  - Used as a principle in most sports: soccer, basketball, etc.
- **Ratings used in games for fair competition.**
  - ELO system: Estimates the skill level of a chess player.
  - ATP system: Estimates the skill level of a tennis player.
  - Used for matchmaking in tournaments.
  - Generate a players ranking.
- **Online gaming poses additional challenges.**
  - Infer from a few match outcomes player skills.
  - Consider the possibility of teams with different number of players.

Questions that Arise in Online Gaming Skill Rating

- **Observed data:** Match outcomes of k teams with $n_1, \ldots, n_k$ players each, in the form of a ranking with potential ties between teams.

- **Information we would like to obtain:**
  - Skills of each player $s_1, \ldots, s_k$.

    If $s_i > s_j$ player i is expected to beat player j.
  - Global ranking among players.
  - Fair matches between players and teams of players.

Successfully achieved by Microsoft's Xbox TrueSkill^{TM}

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

Outline

TrueSkill^{TM} Observed and Latent Variables

Latent Variables:

- **Skill** s_i **for player** i: $p(s_i) = \mathcal{N}(s_i|\mu_i, \sigma_i^2)$.
- **Performance** p_i **of player** i: $p(p_i|s_i) = \mathcal{N}(p_i|s_i, \beta^2)$.
- **Performance** t_j **of team** j: $p(t_j|\{p_i : i \in \mathcal{A}_j\}) = \delta(t_j - \sum_{i \in \mathcal{A}_j} p_i)$.

Observed Variables:

- **Match outcome involving** k **teams**: Rank r_j for each team j.

$$p(r_1, \ldots, r_k | t_1, \ldots, t_k) = \mathbb{I}(t_{r_1} > t_{r_2} > \cdots > t_{r_k})$$

Thus, $r_j < r_{j+1}$ implies $t_j > t_{j+1}$.

The parameters of the model are β^2, μ_i and σ_i^2.

TrueSkill$^{\text{TM}}$ Example

We consider a game with 3 teams $\mathcal{A}_1 = \{1\}$, $\mathcal{A}_2 = \{2, 3\}$, $\mathcal{A}_3 = \{4\}$.

Results: $\mathbf{r} = (1, 2, 3)$. Team $\mathcal{A}_1$ wins followed by $\mathcal{A}_2$ and $\mathcal{A}_3$.

Define $d_1 = t_1 - t_2$ and $d_2 = t_2 - t_3$:

$$p(d_1|t_1, t_2) = \delta(d_1 - t_1 + t_2)\,, \qquad p(d_2|t_2, t_3) = \delta(d_2 - t_2 + t_3)\,.$$

The result implies $d_1 > 0$ and $d_2 > 0$. **Use transitivity!**

Thus, $\mathbf{r} = (1, 2, 3)$ is equivalent to observing $b_1 = 1$ and $b_2 = 1$ where:

$$p(b_1|d_1) = \begin{cases} 1 & \text{if } d_1 > 0\,, \\ 0 & \text{if } d_1 \leq 0\,. \end{cases} \qquad p(b_2|d_2) = \begin{cases} 1 & \text{if } d_2 > 0\,, \\ 0 & \text{if } d_2 \leq 0\,. \end{cases}$$

Bayesian Network for the Example

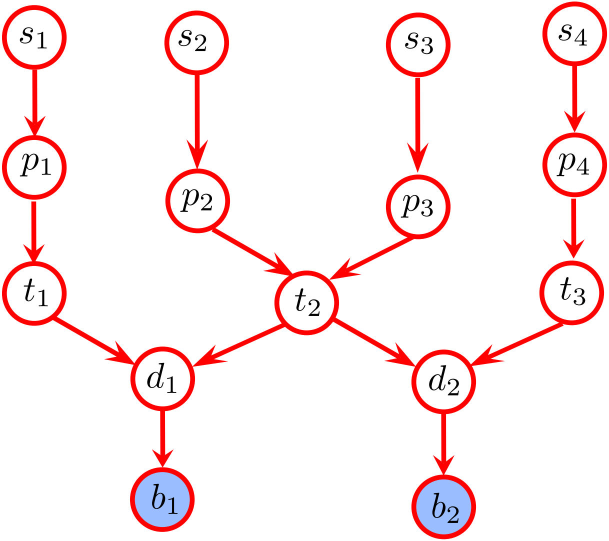

Note that we observe $b_1 = 1$ and $b_2 = 1$.

TrueSkill^{TM} Bayesian Online Learning

Inference involves computing a posterior given a match outcome:

$$p(\mathbf{s}, \mathbf{p}, \mathbf{t}, \mathbf{d}, |\mathbf{b}) = \frac{p(\mathbf{s}, \mathbf{b}, \mathbf{t}, \mathbf{p}, \mathbf{d})}{p(\mathbf{b})} = \frac{p(\mathbf{b}|\mathbf{d})p(\mathbf{d}|\mathbf{t})p(\mathbf{t}|\mathbf{p})p(\mathbf{p}|\mathbf{s})p(\mathbf{s})}{p(\mathbf{b})}$$

We marginalize out $\mathbf{t}$ and $\mathbf{p}$ and $\mathbf{d}$ to get the marginal posterior of $\mathbf{s}$.

Bayesian Online Learning

The posterior after one match is used as the prior for the next match.

$$\{\mu_i(0), \sigma_i^2(0)\} \xrightarrow{\text{Match } \#1} \{\mu_i(1), \sigma_i^2(1)\} \xrightarrow{\text{Match } \#2} \{\mu_i(2), \sigma_i^2(2)\}$$

Marginal Posterior is not Gaussian!

Outline

Bethe Cluster Graph for the Example

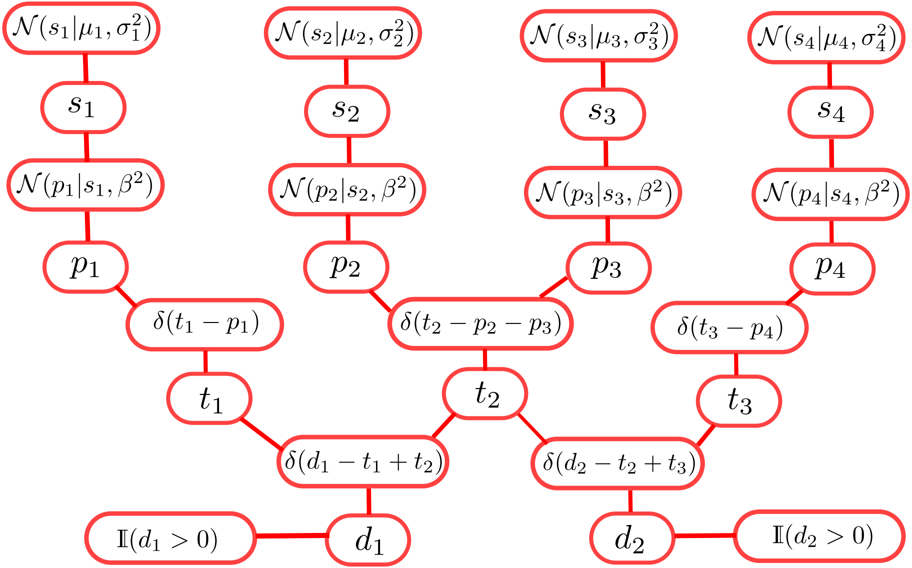

Message Passing Algorithm

We pass messages in the Bethe cluster graph until convergence to compute the posterior marginals of $s_1, \ldots, s_4$.

$$\delta_{i\to j}(s_{i,j}) = \left\{ \int \psi_i(\mathbf{c}_i) \left[\prod_{k\in \mathrm{Nb}_i} \delta_{k\to i}(s_{k,i}) \right] d\left(\mathbf{c}_i \setminus s_{i,j}\right) \right\} / \delta_{j\to i}(s_{i,j})\,.$$

where $s_{i,j}$ are the variables in the sepset of edge $i \to j$. This is the same rule as the one used in Belief Propagation.

The posterior marginals for $s_1, \ldots, s_4$ are given:

$$p(s_i|\mathbf{b}) \propto \psi_i(s_i) \prod_{k\in \mathrm{Nb}_i} \delta_{k\to i}(s_i) = \mathcal{N}(s_i|\mu_i, \sigma_i^2)\delta_{k\to i}(s_i)\,.$$

All messages are Gaussian except for the two bottom messages!

Bethe Cluster Graph: Message Passing

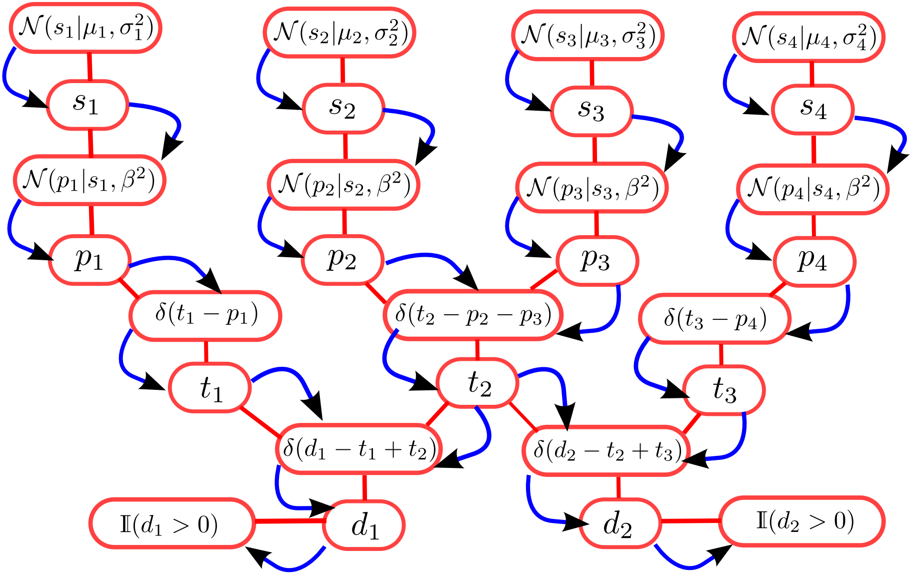

Bethe Cluster Graph: Message Passing

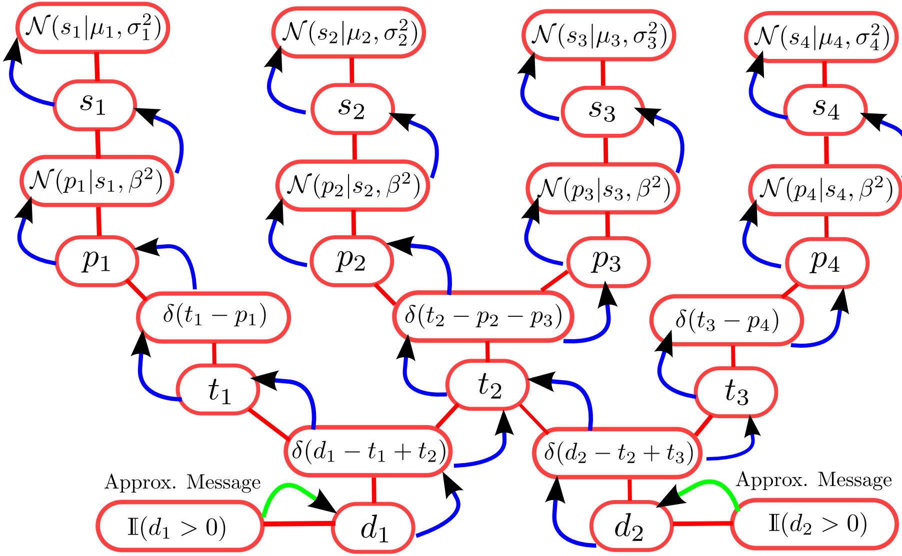

Approximate Messages: Projection Step

We consider the case of the first message. The first step is to approximate the marginal by projecting into the Gaussian family:

$$\psi_i(d_1)\delta_{j\to i}(d_1) = \mathbb{I}(d_1 > 0)\mathcal{N}(d_1|\hat{m}, \hat{v})\,.$$

For this, we compute the log of the normalization constant:

$$\log Z = \log \int \mathbb{I}(d_1 > 0)\mathcal{N}(d_1|\hat{m}, \hat{v}) dd_1 = \log \Phi\left(\frac{\hat{m}}{\sqrt{\hat{v}}}\right)\,.$$

We can obtain the mean and the variance of $\mathbb{I}(d_1 > 0)\mathcal{N}(d_1|\hat{m}, \hat{v})$ by computing the derivatives with respect to $\hat{m}$ and $\hat{v}$!

Assume m and v are the mean and the variance. The approximate message is then:

$$\delta_{i\to j}(d_1) \propto \frac{\mathcal{N}(d_1|m, v)}{\mathcal{N}(d_1|\hat{m}, \hat{v})}\,.$$

The computation of the other approximate message is equivalent.

Outline

Convergence Speed

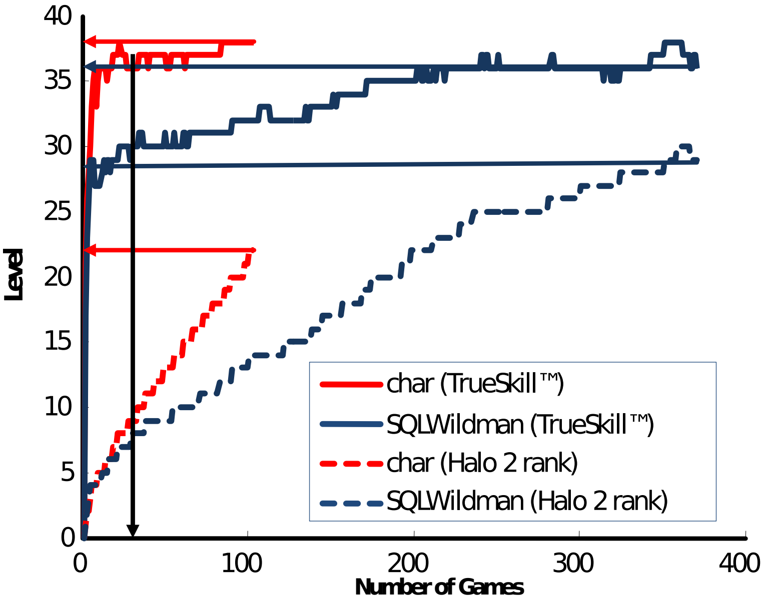

Applications to Online Gaming

- **Leader-board of players**:
  - Provides a global ranking of all players.
  - The rank is conservative: $\mu_i - 3\sigma_i$.

- **Matchmaking**:
  - For players: Most uncertain outcome is better.
  - For inference: Most uncertain outcome is most informative.
  - Good for both players and the model.

Rank	Division	Gamertag	Progress	Games	Rank	Division	Gamertag	Progress	Games	Rank	Division	Gamertag	Progress	Games
1	50	EMK PerFexionZx		6944	101	50	Dreadcide		7524	201	50	Spitfire Bolts		2604
2	50	wafflezonsunday		7388	102	50	ggravityz pull		3135	202	50	Scorpion Soulja		4013
3	50	Crow3112		11006	103	50	TuMuchShuMon3Y		2488	203	50	Mts Kl4r3		2848
4	50	YA ljublju Gera		4648	104	50	mio isurugi		3982	204	50	SykoRe		1480

Matchmaking: Probability of winning or loosing

We assume player i wants to play against player j. What is the probability of winning?

$$
\begin{aligned}
p(p_i > p_j) &= \int \mathbb{I}(p_i - p_j > 0) p(p_j|s_j) p(s_j) p(p_i|s_i) p(s_i) dp_i dp_j ds_i ds_j \\
&= \int \mathbb{I}(p_i - p_j > 0) \mathcal{N}(p_j|s_j, \beta^2) \mathcal{N}(p_i|s_i, \beta^2) \cdot \\
&\quad \cdot \mathcal{N}(s_j|\mu_j, \sigma_j^2) \mathcal{N}(s_i|\mu_i, \sigma_i^2) dp_i dp_j ds_i ds_j \\
&= \int \mathbb{I}(p_i - p_j > 0) \mathcal{N}(p_j|\mu_j, \sigma_i^2 + \beta^2) \mathcal{N}(p_i|\mu_i, \sigma_i^2 + \beta^2) dp_i dp_j \\
&= \Phi\left(\frac{\mu_i - \mu_j}{\sqrt{\sigma_i^2 + \sigma_j^2 + 2\beta^2}} \right).
\end{aligned}
$$

Game Over!