Gibbs sampling in TrueSkill

Carl Edward Rasmussen

October 28th, 2016

Key concepts

- In TrueSkill the joint distribution is intractable
- however, both the
    - performance differences given skills, and
    - skills given performances
  are tractable.
- we derive the Gibbs sampling updates in detail.

Gibbs sampling for the TrueSkill model

We have $g = 1, \ldots, G$ games where I_g: id of Player 1 and J_g: id of Player 2. The outcome of game g is $y_g = +1$ if I_g wins, $y_g = -1$ if J_g wins.

Gibbs sampling alternates between sampling skills $\mathbf{w} = [w_1, \ldots, w_M]^\top$ conditional on fixed performance differences $\mathbf{t} = [t_1, \ldots, t_N]^\top$, and sampling $\mathbf{t}$ conditional on fixed $\mathbf{w}$.

1. Initialise $\mathbf{w}$, e.g. from the prior $p(\mathbf{w})$.

2. Sample the *performance differences* from their conditional posteriors

$$p(t_g | w_{I_g}, w_{J_g}, y_g) \; \propto \; \delta(y_g - \text{sign}(t_g)) \mathcal{N}(t_g; w_{I_g} - w_{J_g}, 1)$$

3. Jointly sample the *skills* from the conditional posterior

$$p(\mathbf{w}|\mathbf{t}, \mathbf{y}) \; = \; \underbrace{p(\mathbf{w}|\mathbf{t})}_{\mathcal{N}(\mathbf{w}; \mu, \Sigma)} \; \propto \; \underbrace{p(\mathbf{w})}_{\mathcal{N}(\mathbf{w}; \mu_0, \Sigma_0)} \prod_{g=1}^{G} \underbrace{p(t_g | w_{I_g}, w_{J_g})}_{\propto \, \mathcal{N}(\mathbf{w}; \mu_g, \Sigma_g)}$$

4. Go back to step 2.

Gaussian identities

The distribution for the performance is both Gaussian in t_g and proportional to a Gaussian in $\mathbf{w}$

$$p(t_g | w_{I_g}, w_{J_g}) \propto \exp\left(-\tfrac{1}{2}(w_{I_g} - w_{J_g} - t_g)^2\right)$$
$$\propto \mathcal{N}\left(-\tfrac{1}{2}\begin{pmatrix} w_{I_g} - \mu_1 \\ w_{J_g} - \mu_2 \end{pmatrix}^\top \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \begin{pmatrix} w_{I_g} - \mu_1 \\ w_{J_g} - \mu_2 \end{pmatrix}\right)$$

with $\mu_1 - \mu_2 = t_g$. Notice that

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix} = \begin{pmatrix} t_g \\ -t_g \end{pmatrix}$$

Remember that for products of Gaussians precisions add up, and means weighted by precisions (natural parameters) also add up:

$$\mathcal{N}(\mathbf{w}; \mu_a, \Sigma_a)\mathcal{N}(\mathbf{w}; \mu_b, \Sigma_b) = z_c \, \mathcal{N}(\mathbf{w}; \mu_c, \Sigma_c)$$

where $\Sigma_c^{-1} = \Sigma_a^{-1} + \Sigma_b^{-1}$ and $\mu_c = \Sigma_c(\Sigma_a^{-1}\mu_a + \Sigma_b^{-1}\mu_b)$.

Conditional posterior over skills given performances

We can now compute the covariance and the mean of the conditional posterior.

$$\Sigma^{-1} = \Sigma_0^{-1} + \underbrace{\sum_{g=1}^{G} \Sigma_g^{-1}}_{\tilde{\Sigma}^{-1}} \qquad \mu = \Sigma\big(\Sigma_0^{-1}\mu_0 + \underbrace{\sum_{g=1}^{G} \Sigma_g^{-1}\mu_g}_{\tilde{\mu}}\big),$$

where each game precision Σ_g^{-1} contain only 4 non-zero entries. The combined precision is:

$$[\tilde{\Sigma}^{-1}]_{ii} = \sum_{g=1}^{G} \delta(i - I_g) + \delta(i - J_g)$$

$$[\tilde{\Sigma}^{-1}]_{i \neq j} = -\sum_{g=1}^{G} \delta(i - I_g)\delta(j - J_g) + \delta(i - J_g)\delta(j - I_g),$$

and for the mean we have

$$\tilde{\mu}_i = \sum_{g=1}^{G} t_g\big(\delta(i - I_g) - \delta(i - J_g)\big).$$

Implementing Gibbs sampling for the TrueSkill model

we have derived the conditional distribution for the performance differences in game g and for the skills. These are:

- the posterior conditional performance difference for t_g is a univariate truncated Gaussian. How can we sample from it?
  - by rejection sampling from a Gaussian, or
  - by the inverse transformation method (passing a uniform on an interval through the inverse cumulative distribution function).
- the conditional skills can be sampled jointly form the corresponding Gaussian (using the cholesky factorization of the covariance matrix).

Once samples have been drawn from the posterior, these can be used to make predictions for game outcomes, using the generative model.

How would you do this?