

Assignment 2 [version 2](#):

Stochastic Variational Inference in the TrueSkill Model

STA414/STA2014 and CSC412/CSC2506 Winter 2020

David Duvenaud and Jesse Bettencourt

Student Number

March 15, 2020

The goal of this assignment is to get you familiar with the basics of Bayesian inference in large models with continuous latent variables, and the basics of stochastic variational inference.

Background We'll implement a variant of the TrueSkill model, a player ranking system for competitive games originally developed for Halo 2. It is a generalization of the Elo rating system in Chess. For the curious, the original 2007 NIPS paper introducing the trueskill paper can be found here: <http://papers.nips.cc/paper/3079-trueskilltm-a-bayesian-skill-rating-system.pdf>

This assignment is based on one developed by Carl Rasmussen at Cambridge for his course on probabilistic machine learning: <http://mlg.eng.cam.ac.uk/teaching/4f13/1920/>

0.1 Model definition

We'll consider a slightly simplified version of the original trueskill model. We assume that each player has a true, but unknown skill $z_i \in \mathbb{R}$. We use N to denote the number of players.

The prior. The prior over each player's skill is a standard normal distribution, and all player's skills are *a priori* independent.

The likelihood. For each observed game, the probability that player i beats player j , given the player's skills z_A and z_B , is:

$$p(\text{A beat B} | z_A, z_B) = \sigma(z_i - z_j)$$

where

$$\sigma(y) = \frac{1}{1 + \exp(-y)}$$

There can be more than one game played between a pair of players, and in this case the outcome of each game is independent given the players' skills. We use M to denote the number of games.

The data. The data will be an array of game outcomes. Each row contains a pair of player indices. The first index in each pair is the winner of the game, the second index is the loser. If there were M games played, then the array has shape $M \times 2$.

1 Implementing the model [10 points]

- (a) [2 points] Implement a function `log_prior` that computes the log of the prior over all player's skills. Specifically, given a $K \times N$ array where each row is a setting of the skills for all N players, it returns a $K \times 1$ array, where each row contains a scalar giving the log-prior for that set of skills.
- (b) [3 points] Implement a function `logp_a_beats_b` that, given a pair of skills z_a and z_b evaluates the log-likelihood that player with skill z_a beat player with skill z_b under the model detailed above. To ensure numerical stability, use the function `log1pexp` that computes $\log(1 + \exp(x))$ in a numerically stable way. This function is provided by `StatsFuns.jl` and imported already, and also by Python's `numpy`.
- (c) [3 points] Assuming all game outcomes are i.i.d. conditioned on all players' skills, implement a function `all_games_log_likelihood` that takes a batch of player skills `zs` and a collection of observed games `games` and gives a batch of log-likelihoods for those observations. Specifically, given a $K \times N$ array where each row is a setting of the skills for all N players, and an $M \times 2$ array of game outcomes, it returns a $K \times 1$ array, where each row contains a scalar giving the log-likelihood of all games for that set of skills. Hint: You should be able to write this function without using for loops, although you might want to start that way to make sure what you've written is correct. If A is an array of integers, you can index the corresponding entries of another matrix B for every entry in A by writing `B[A]`.
- (d) [2 points] Implement a function `joint_log_density` which combines the log-prior and log-likelihood of the observations to give $p(z_1, z_2, \dots, z_N, \text{all game outcomes})$

2 Examining the posterior for only two players and toy data [10 points]

To get a feel for this model, we'll first consider the case where we only have 2 players, A and B . We'll examine how the prior and likelihood interact when conditioning on different sets of games.

Provided in the starter code is a function `skillcontour!` which evaluates a provided function on a grid of z_A and z_B 's and plots the isocontours of that function. As well there is a function `plot_line_equal_skill!`. We have included an example for how you can use these functions.

We also provided a function `two_player_toy_games` which produces toy data for two players. I.e. `two_player_toy_games(5,3)` produces a dataset where player A wins 5 games and player B wins 3 games.

- (a) [2 points] For two players A and B , plot the isocontours of the joint prior over their skills. Also plot the line of equal skill, $z_A = z_B$. Hint: you've already implemented the `log` of the likelihood function.
- (b) [2 points] Plot the isocontours of the likelihood function. Also plot the line of equal skill, $z_A = z_B$.
- (c) [2 points] Plot isocontours of the joint posterior over z_A and z_B given that player A beat player B in one match. Since the contours don't depend on the normalization constant, you can simply plot the isocontours of the log of joint distribution of $p(z_A, z_B, A \text{ beat } B)$ Also plot the line of equal skill, $z_A = z_B$.
- (d) [2 points] Plot isocontours of the joint posterior over z_A and z_B given that 10 matches were played, and player A beat player B all 10 times. Also plot the line of equal skill, $z_A = z_B$.
- (e) [2 points] Plot isocontours of the joint posterior over z_A and z_B given that 20 matches were played, and each player beat the other 10 times. Also plot the line of equal skill, $z_A = z_B$.

For all plots, label both axes.

3 Stochastic Variational Inference on Two Players and Toy Data [18 points]

One nice thing about a Bayesian approach is that it separates the model specification from the approximate inference strategy. The original Trueskill paper from 2007 used message passing. Carl Rasmussen's assignment uses Gibbs sampling, a form of Markov Chain Monte Carlo. We'll use gradient-based stochastic variational inference, which wasn't invented until around 2014.

In this question we will optimize an approximate posterior distribution with stochastic variational inference to approximate the true posterior.

- (a) [5 points] Implement a function `elbo` which computes an unbiased estimate of the evidence lower bound. As discussed in class, the ELBO is equal to the KL divergence between the true posterior $p(z|\text{data})$, and an approximate posterior, $q_\phi(z|\text{data})$, plus an unknown constant. Use a fully-factorized Gaussian distribution for $q_\phi(z|\text{data})$. This estimator takes the following arguments:

- `params`, the parameters ϕ of the approximate posterior $q_\phi(z|\text{data})$.
- A function `logp`, which is equal to the true posterior plus a constant. This function must take a batch of samples of z . If we have N players, we can consider B -many samples from the joint over all players' skills. This batch of samples `zs` will be an array with dimensions (N, B) .
- `num_samples`, the number of samples to take.

This function should return a single scalar. Hint: You will need to use the reparameterization trick when sampling `zs`.

- (b) [2 points] Write a loss function called `neg_toy_elbo` that takes variational distribution parameters and an array of game outcomes, and returns the negative elbo estimate with 100 samples.
- (c) [5 points] Write an optimization function called `fit_toy_variational_dist` which takes initial variational parameters, and the evidence. Inside it will perform a number of iterations of gradient descent where for each iteration :
- (a) Compute the gradient of the loss with respect to the parameters using automatic differentiation.
 - (b) Update the parameters by taking an `lr`-scaled step in the direction of the descending gradient.
 - (c) Report the loss with the new parameters (using `@info` or print statements)
 - (d) On the same set of axes plot the target distribution in red and the variational approximation in blue.

Return the parameters resulting from training.

- (d) [2 points] Initialize a variational distribution parameters and optimize them to approximate the joint where we observe player A winning 1 game. Report the final loss. Also plot the optimized variational approximation contours (in blue) and the target distribution (in red) on the same axes.
- (e) [2 points] Initialize a variational distribution parameters and optimize them to approximate the joint where we observe player A winning 10 games. Report the final loss. Also plot the optimized variational approximation contours (in blue) and the target distribution (in red) on the same axes.
- (f) [2 points] Initialize a variational distribution parameters and optimize them to approximate the joint where we observe player A winning 10 games and player B winning 10 games. Report the final loss. Also plot the optimized variational approximation contours (in blue) and the target distribution (in red) on the same axes.

For all plots, label both axes.

4 Approximate inference conditioned on real data [24 points]

Load the dataset from `tennis_data.mat` containing two matrices:

- W is a 107 by 1 matrix, whose i 'th entry is the name of player i .
- G is a 1801 by 2 matrix of game outcomes (actually tennis matches), one row per game. The first column contains the indices of the players who won. The second column contains the indices of the player who lost.

Compute the following using your code from the earlier questions in the assignment, but conditioning on the tennis match outcomes:

- (a) [1 point] For any two players i and j , $p(z_i, z_j | \text{all games})$ is always proportional to $p(z_i, z_j, \text{all games})$. In general, are the isocontours of $p(z_i, z_j | \text{all games})$ the same as those of $p(z_i, z_j | \text{games between } i \text{ and } j)$? That is, do the games between other players besides i and j provide information about the skill of players i and j ? A simple yes or no suffices.
- Hint: One way to answer this is to draw the graphical model for three players, i , j , and k , and the results of games between all three pairs, and then examine conditional independencies. If you do this, there's no need to include the graphical models in your assignment.
- (b) [5 points] Write a new optimization function `fit_variational_dist` like the one from the previous question except it does not plot anything. Initialize a variational distribution and fit it to the joint distribution with all the observed tennis games from the dataset. Report the final negative ELBO estimate after optimization.
- (c) [2 points] Plot the approximate mean and variance of all players, sorted by skill. For example, in Julia, you can use: `perm = sortperm(means); plot(means[perm], yerror=exp.(logstd[perm]))`. There's no need to include the names of the players.
- (d) [2 points] List the names of the 10 players with the highest mean skill under the variational model.
- (e) [3 points] Plot the joint approximate posterior over the skills of Roger Federer and Rafael Nadal. Use the approximate posterior that you fit in question 4 part b.
- (f) [5 points] Derive the exact probability under a factorized Gaussian over two players' skills that one has higher skill than the other, as a function of the two means and variances over their skills. Express your answer in terms of the cumulative distribution function of a one-dimensional Gaussian random variable.
- Hint 1: Use a linear change of variables $y_A, y_B = z_A - z_B, z_B$. What does the line of equal skill look like after this transformation?
 - Hint 2: If $X \sim \mathcal{N}(\mu, \Sigma)$, then $AX \sim \mathcal{N}(A\mu, A\Sigma A^T)$ where A is a linear transformation.
 - Hint 3: Marginalization in Gaussians is easy: if $X \sim \mathcal{N}(\mu, \Sigma)$, then the i th element of X has a marginal distribution $X_i \sim \mathcal{N}(\mu_i, \Sigma_{ii})$
- (g) [2 points] Using the formula from part c, compute the exact probability under your approximate posterior that Roger Federer has higher skill than Rafael Nadal. Then, estimate it using simple Monte Carlo with 10000 examples, again using your approximate posterior.
- (h) [2 points] Using the formula from part c, compute the probability that Roger Federer is better than the player with the lowest mean skill. Compute this quantity exactly, and then estimate it using simple Monte Carlo with 10000 examples, again using your approximate posterior.
- (i) [2 points] Imagine that we knew ahead of time that we were examining the skills of top tennis players, and so changed our prior on all players to $\text{Normal}(10, 1)$. Which answers in this section would this change? No need to show your work, just list the letters of the questions whose answers would be different in expectation.