

Probabilistic Modelling and Reasoning: Assignment

Modelling the skills of Go players

Instructor: Dr. Amos Storkey

Due in 16:00 Thursday 30th March 2017.

Mechanics

Marks: This assignment is out of 100 marks and forms 20% of your final grade for the course.

Academic conduct: Assessed work is subject to University regulations on academic conduct:

<http://web.inf.ed.ac.uk/infweb/admin/policies/academic-misconduct>

Do not show your code, answers or write-up to anyone else.

Never copy-and-paste material that is not your own into your assignment and edit it.

Submission instructions: You should submit this assignment manually to the ITO office in Forrest Hill by the deadline. Handwritten submissions are acceptable if the handwriting is neat and legible.

Late submissions: The School of Informatics policy is that late coursework normally gets a mark of zero. See <http://tinyurl.com/edinflate> for exceptions to this rule. Any requests for extensions should go to the ITO, either directly or via your Personal Tutor.

Data files: Associated data files for the assignment are available as a compressed TAR archive file `go_player_skill_model.tar.gz` from the course website.

Introduction

A central theme of this course is constructing probabilistic models to express the relationships we believe exist between variables in real world problems. Given a model and observed data we can then make principled predictions of the answers to questions we have and importantly also get some indication of our degree of certainty in those answers. In this coursework assignment we will consider the problem of inferring information about the skills of players from game outcome data, and using the resulting model to predict the results of future games.

In particular you will be constructing a model to infer the skill of professional Go players. As many of you will have seen in the news last year, a computer-based Go agent AlphaGo, developed by Google DeepMind, beat the reigning 3-time European Go champion Fan Hui in October 2015. This was the first time a computer program has beaten a professional level Go player in an even (non-handicapped) game. AlphaGo then went on to compete against Lee Sedol, a South Korean Go player who is regarded as one of the best players in the world in the last decade, in a five game match in March 2016. Of the five games in the match, AlphaGo won four and Lee Sedol won one.

In April 2017, AlphaGo is expected to play the current world number one, Ke Jie, in a further series of games. Part of your task in this assignment will be to use a probabilistic model fitted to a dataset of over 50 000 professional Go game results, kindly provided by <http://Go4Go.net>, as well as the recent AlphaGo match results to find out what the available data suggests the relative skill of AlphaGo is compared to the best human player and to make predictions about the outcomes of the Ke Jie games. Note, no specialist knowledge of the game Go beyond the information given in the assignment is required!

Notation

$\mathbb{P}[\cdot]$ indicates a probability mass function on discrete random variables and $\mathbb{p}[\cdot]$ a probability density function on real-valued random variables. $\mathbb{1}[\text{condition}]$ is an indicator function which is equal to 1 if `condition` is true and 0 otherwise. $\mathcal{N}(x; \mu, \sigma^2)$ is a Gaussian probability density function on x with mean μ and variance σ^2 while $\Phi(x)$ is the standard Gaussian cumulative density function i.e. $\Phi(x) = \int_{-\infty}^x \mathcal{N}(u; 0, 1) du$. $\{x_i\}_{i=1}^N$ indicates a collection of N variables i.e. $\{x_1, x_2, \dots, x_N\}$. \mathbb{R} is the set of real number and \mathbb{N} is the set of natural numbers $\{1, 2, 3 \dots\}$.

Question 1: Player skill graphical models (15 marks)

As a first step we will consider a basic probabilistic graphical model relating the skills of two Go players and the result of the game between them.

In Go the two players lay either black or white stones on the playing board, hence we will refer to the two players as the *black player* and *white player*. The black player always lays the first stone; to adjust for this there is usually a point offset called *komi* given to the white player in compensation. The *komi* is usually set to have a fractional component (e.g. 6.5 is common).

As the main scoring methods only allow integer scores this generally means games can only result in a win or loss with draws very infrequent. Therefore we will model the result as being a binary outcome of either the black player winning or the white player winning.

In figure 1 a simple factor graph is given to describe the relationship between the skills of two players and the result of a game between them. The game is indexed by $k \in \mathbb{N}$, and every player is also assigned a unique ID $\in \mathbb{N}$, with the ID of the black player in the k^{th} game being given by b_k and the ID of the white player by w_k . In general there may be more than one (or zero) games between each pair of players. The skills of the black and white players in the k^{th} game are represented by real values s_{b_k} and s_{w_k} respectively. The result of the game is denoted $r^{(k)}$ and is either equal to 1 if the black player wins or 0 if the white player wins.

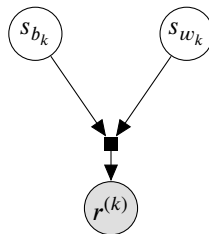


Figure 1: Factor graph of simple player skill model when observing one game between two players. Games are indexed by $k \in \mathbb{N}$, $b_k \in \mathbb{N}$ is the ID of the black player and $w_k \in \mathbb{N}$ the ID of the white player. The skill of the black player is $s_{b_k} \in \mathbb{R}$, the skill of the white player is $s_{w_k} \in \mathbb{R}$ and the result of the game is $r^{(k)} \in \{0, 1\}$, with 0 indicating a white win and 1 a black win.

- (a) Consider the case where we have three players and games between each pair of players.

Game index k	Black player ID b_k	White player ID w_k
1	1	2
2	2	3
3	3	1

Draw a directed factor graph to represent the joint distribution over the skills of the three players s_1, s_2 and s_3 and the three match outcomes $r^{(1)}, r^{(2)}$ and $r^{(3)}$.

[Expected response: Labelled directed factor graph diagram.]

- (b) Using the factor graph from part (a) and the rules for assessing conditional independence from factor graphs (covered in [lecture 3](#)), state and explain whether s_1 and s_2 are conditionally independent if we know the value of $r^{(2)}$. Similarly state and explain if s_1 and s_2 are conditionally independent if we know $r^{(2)}$ and $r^{(3)}$.

[Expected response: Two conditional independency statements and justifications.]

- (c) Convert the directed factor graph from part (a) into an undirected graphical model.

[Expected response: Labelled undirected graphical model diagram.]

- (d) Using the u -separation criteria and symmetry considerations on the undirected graphical model you constructed in part (c) what conditional independence statements can you make about the variables in the graphical model? Is the same set of conditional independence statements described by the directed factor graph in part (a)?

[Expected response: Set of conditional independency statements. Comparison to directed factor graph.]

Question 2: Gaussian player skill model (30 marks)

We will now consider a more concrete model relating the skills of M Go players to the results of N games between them. In particular we will consider the model described by the factor graph in figure 2. The real-valued skill variables are modelled as having zero-mean Gaussian prior distributions with fixed variance σ^2 . As well as the underlying per-player skill variables, we will also model the players in a particular game as having per-game performances ($p_b^{(k)}$ and $p_w^{(k)}$ for the black and white players in game k respectively) which are Gaussian distributed with a constant variance 1 and means equal to the player skills.

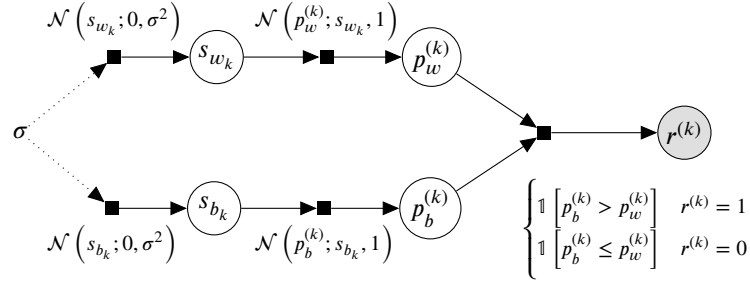


Figure 2: Factor graph of Gaussian player skill model for a single game. The variables $r^{(k)}$, s_{b_k} and s_{w_k} are as defined for figure 1 and $p_b^{(k)}$ and $p_w^{(k)}$ are the performances of the black and white player respectively in game k .

Our simple model is that the player whose performance is higher in the game is the winner. Therefore if $p_b^{(k)} > p_w^{(k)}$ then the black player is the winner and $r^{(k)} = 1$. If $p_b^{(k)} \leq p_w^{(k)}$ then the white player is the winner and $r^{(k)} = 0$. Note the slight asymmetry in defining the case when $p_b^{(k)} = p_w^{(k)}$ as a win for white will not matter in practice as the probability of this occurring exactly is zero.

We will assume that the result of each game is independently and identically distributed given the player skills and performances. To begin with in parts (a)–(d) we will consider a single game and so for the sake of clarity and brevity will drop the game index k from notation.

- (a) By taking the relevant factors from the factor graph in figure 2, write down an expression for $\mathbb{P}[p_b, p_w | s_b, s_w]$, the probability density across the player performance levels given known player skill levels.

[Expected response: Factorised expression for $\mathbb{P}[p_b, p_w | s_b, s_w]$.]

- (b) We will now define a change of variables using

$$\psi = \frac{(p_w - s_w) - (p_b - s_b)}{\sqrt{2}} \quad \theta = \frac{(p_w - s_w) + (p_b - s_b)}{\sqrt{2}}. \quad (1)$$

Calculate the conditional expectations $\mathbb{E}[\theta | s_b, s_w]$, $\mathbb{E}[\psi | s_b, s_w]$, $\mathbb{E}[\theta^2 | s_b, s_w]$, $\mathbb{E}[\psi^2 | s_b, s_w]$ and $\mathbb{E}[\theta\psi | s_b, s_w]$ with in all cases the expectations being with respect to $\mathbb{P}[p_b, p_w | s_b, s_w]$. Hence write down the conditional density $\mathbb{P}[\theta, \psi | s_b, s_w]$.

[Expected response: Values for each of the five conditional expectations including working to show how you got results. An expression for $\mathbb{P}[\theta, \psi | s_b, s_w]$ and explanation of how you arrived at it.]

- (c) Using the change of variables in part (b) write down an expression for $\mathbb{P}[r = 1 | \theta, \psi, s_b, s_w]$.

[Expected response: Expression for $\mathbb{P}[r = 1 | \theta, \psi, s_b, s_w]$ - this should be in terms of s_b, s_w, ψ and θ only.]

- (d) Using your answers from the two previous parts show that

$$\mathbb{P}[r = 1 | s_b, s_w] = \Phi\left[\frac{s_b - s_w}{\sqrt{2}}\right]. \quad (2)$$

[Expected response: Derivation of result shown with explanation of steps taken.]

- (e) Now considering all N games and M players, write down an expression for $\mathbb{P}\left[\left\{r^{(k)}\right\}_{k=1}^N \mid \left\{s_i\right\}_{i=1}^M\right]$, the probability of observing a set of game results $\left\{r^{(k)}\right\}_{k=1}^N$ between M players with known skills $\left\{s_i\right\}_{i=1}^M$.

[Expected response: Expression for $\mathbb{P}\left[\left\{r^{(k)}\right\}_{k=1}^N \mid \left\{s_i\right\}_{i=1}^M\right]$ with any assumptions used to derive this result explicitly stated and explained.]

- (f) Show that the posterior density on player skills given observed game outcomes can be written in the form of a probit regression problem with Gaussian prior on the ‘regression weights’ (here player skills which are grouped into a vector $\mathbf{s} = [s_1 \ s_2 \ \dots \ s_M]^T$), i.e.

$$\mathbb{P}\left[\left\{s_i\right\}_{i=1}^M \mid \left\{r^{(k)}\right\}_{k=1}^N\right] \propto \prod_{k=1}^N \left\{\Phi\left[\mathbf{s}^T \mathbf{x}^{(k)}\right]\right\} \mathcal{N}\left(\mathbf{s}; \mathbf{m}_0, \mathbf{V}_0\right) \quad (3)$$

giving expressions for $\mathbf{x}^{(k)}$, \mathbf{m}_0 and \mathbf{V}_0 . You may find the identity $\Phi(-x) = 1 - \Phi(x)$ useful.

[Expected response: Derivation showing how to get from your answer to part (e) to an expression for the posterior density in the form given and identification of how to express the new variables $\mathbf{x}^{(k)}$, \mathbf{m}_0 and \mathbf{V}_0 in terms of the variables defined in the factor graph in figure 2.]

Question 3: Approximate inference with Go game data (30 marks)

The posterior density on the player skills we derived in the previous question from our model cannot be evaluated exactly. In particular we can only express the posterior density up to an unknown normalisation constant as the integral

$$\int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \mathbb{P}\left[\left\{s_i\right\}_{i=1}^M, \left\{r^{(k)}\right\}_{k=1}^N\right] ds_1 \dots ds_M \quad (4)$$

does not have an analytic solution. If we want to perform Bayesian inference with our model, for example to predict the probability of the result r^* of a new game between two of the players b_* and w_* , we need to be able to calculate values for integrals like

$$\mathbb{P}\left[r^* = 1 \mid \left\{r^{(k)}\right\}_{k=1}^N\right] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbb{P}\left[r^* = 1 \mid s_{b_*}, s_{w_*}\right] \mathbb{P}\left[s_{b_*}, s_{w_*} \mid \left\{r^{(k)}\right\}_{k=1}^N\right] ds_{b_*} ds_{w_*} \quad (5)$$

which are in terms of a (marginal) normalised posterior density $\mathbb{P}\left[s_{b_*}, s_{w_*} \mid \left\{r^{(k)}\right\}_{k=1}^N\right]$.

As we cannot perform exact inference in the model, we will therefore use an *approximate inference* method which will allow us to estimate integrals like (5). In particular here we will make use of *expectation propagation* [Minka, 2001]. All of the details you need to know about expectation propagation for the purposes of this assignment are covered in appendix A, however if you wish to find out more about the method, possible textbook references are [Bishop, 2006, pg. 505-517], [Koller and Friedman, 2009, pg. 442-448], [Murphy, 2012, pg. 787-799] and [Barber, 2012, pg. 639-643].

In the provided archive file `go_player_skill_model.tar.gz` there are python `numpy .npz` and MATLAB `.mat` data files containing the parameters of expectation propagation based Gaussian approximations to the posterior density over the skills of 1049 Go players (1048 of the top professional players and AlphaGo) given the results of 51 413 games between them. These can be loaded using either the `numpy.load` function in python or the `load` function in MATLAB.

The parameters of two different approximate densities are provided. The parameters in the `full_covar.*` files are for a Gaussian approximate density parametrised with a full $M \times M$ covariance matrix \mathbf{V} . The parameters in the `diag_covar.*` files are for Gaussian approximation using a diagonal covariance matrix $\mathbf{V} = \text{diag } \mathbf{v}$ $\mathbf{v} \in \mathbb{R}^M$. In both a full mean vector $\mathbf{m} \in \mathbb{R}^M$ was used. The full list of variables defined in each of the files and mappings between notation in this document and variable names is given in table 1.

Description	Symbolic representation	Variable name	Dimensions
Number of players	M	n_players	scalar
Number of games	N	n_games	scalar
Overall mean vector	\mathbf{m}	approx_mean	1049
Overall covariance matrix	\mathbf{v} or \mathbf{V}	approx_covar	1049 or 1049 × 1049
Skill prior variance	σ^2	skill_prior_var	scalar

Table 1: Relationships between variable symbols used in this document and variable names in .npz / .mat files.

- (a) The true and approximate factors $f_k(s) = \Phi[s^T \mathbf{x}^{(k)}]$ and $\tilde{f}_k(s) = \mathcal{N}(s; \mathbf{m}_k, \mathbf{V}_k)$ are defined over the 1049 dimensional player skill space. However the true factor only varies along one direction in that high-dimensional space - in particular parallel to $\mathbf{x}^{(k)}$. To visualise the fit between the approximate and true factors we can therefore plot the variation in the true and approximate factors along this line.

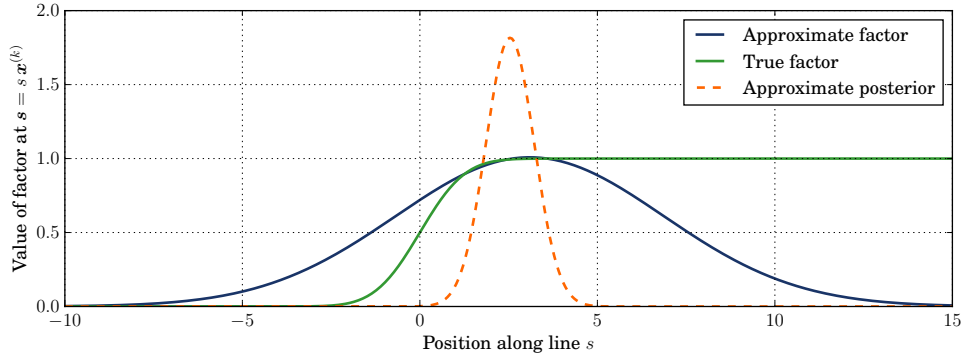


Figure 3: Values of corresponding approximate and true factors and scaled value of approximate posterior density along line $s = s \mathbf{x}^{(k)}$ for a single factor in the full covariance approximate Gaussian density.

Figure 3 shows one such plot, with three values plotted

- Blue: value of the approximate factor $c_k \mathcal{N}(s; \mathbf{m}_k, \mathbf{V}_k)$
- Green: value of the true factor $\Phi(s^T \mathbf{x}^{(k)})$
- Orange: value of overall approximation density $\mathcal{N}(s; \mathbf{m}, \mathbf{V})$ scaled to make it visible

along the line $s = s \mathbf{x}^{(k)} : s \in [-10, 15]$ for a single factor in the full-covariance Gaussian approximate posterior skill density corresponding to one of the AlphaGo–Fan Hui games. Interpret what is shown in the plot in the context of how expectation propagation iteratively fits each of the approximate factors.

[Expected response: An explanation of the relationships shown between the curves in the provided plot.]

- (b) In the data files containing the fitted model parameters, there are also three integer variables alpha_go_id, lee_sedol_id and ke_jie_id defined. These contain the player IDs / indices corresponding to AlphaGo, Lee Sedol and Ke Jie respectively. Plot the approximate posterior marginal skill densities for each of these three players on the same axis for both the full covariance and diagonal covariance models. Interpret what you see in the plots.

[Expected response: Two plot figures - one for each of the full covariance approximation and diagonal covariance approximation, with the three specified approximate posterior marginal distributions plotted on the same axis in each plot. A code snippet showing how you produced the plots. An interpretation of what the distributions shown mean.]

- (c) Using (9), calculate the predicted probability of AlphaGo winning when playing as black in a single game against Ke Jie, using both the full and diagonal covariance approximate Gaussian posterior densities.

[Expected response: Two predicted win probabilities i.e. values in [0, 1], one for each of the two approximate densities. A code snippet showing how you calculated these values.]

- (d) Suggest one advantage and one disadvantage of using expectation propagation with a Gaussian approximating density that is parametrised with a full covariance matrix rather than a diagonal covariance matrix.

[Expected response: A positive and negative aspect of using the full covariance Gaussian approximation to the posterior density compared to the diagonal covariance approximation, for example comparing in terms of accuracy or computational cost, including justifications for any assertions made.]

Question 4: Model evaluation (25 marks)

An important part of any probabilistic modelling task is evaluating the assumptions used in setting up the model and performing inference with it. In this question we will consider some of the limitations of the model and inference method used here.

Note that though you should provide sufficient detail to make yourself clear, long answers are not required here. For each of the following question parts, 4–5 sentences plus one or two accompanying figures and/or formulae would be an acceptable response.

- (a) In Go, the black player always makes the first move. Although *komi* points are designed to compensate for this advantage there is still potentially some asymmetry between playing as black or white. The model used in Q2 and Q3 is symmetric with respect to the player colour - it does not allow any (dis)advantage to playing as black or white to be accounted for. How might the model be extended to include this possibility?

*[Expected response: Proposition of a reasonable method (i.e. one which would do what was asked for) for modelling player colour asymmetry with sufficient detail in explanation to allow someone to implement this extension themselves. You do **not** need to implement the proposed new model.]*

- (b) We have assumed that all game outcomes are independently and identically distributed given player skills and conversely that a-priori every game result is equally important in determining a player's current skill. In reality the player skills will be dynamic quantities which vary over time and more recent game results will be more informative of the current player skill than older results. Discuss how we might extend our model to allow for time varying player skills.

*[Expected response: Example of a possible model for dynamic player skills. You do not need to provide all the details needed to implement inference with the proposed model but you should provide sufficient detail to make the structure of your proposed model clear - you could do this with a graphical model diagram for example. If you wish to use an existing model from the literature this is fine providing you appropriately reference it and explain how the model works in your own words. You do **not** need to implement the proposed new model.]*

- (c) As well as the limitations imposed by our modelling assumptions, our choice of approximate inference method also imposes some constraints on the results we get from our model. Discuss the limitations of the inference method used in Q3 and compare and contrast with an alternative approximate inference method that could be used.

*[Expected response: Description of what approximations the inference method used made. Brief descriptions of one other approximate inference method which could be applied to this problem and a discussion of the relative advantages and disadvantages of each of the methods, for example in terms of accuracy of results, computational cost, implementation complexity. You do **not** need to implement the proposed alternative inference method for this problem.]*

A Introduction to expectation propagation

For compactness in the description that follows we use the shorthand $p(\mathbf{s})$ for the true posterior density we wish to approximate.

Expectation propagation is a form of approximate message passing which generalises loopy belief propagation. The parameters of an approximating density $q(\mathbf{s})$, which is chosen from the exponential family, are iteratively updated so as to minimise a measure of the difference between the approximate and true posterior densities.

We can express our true posterior density as a product of factors $p(\mathbf{s}) \propto \prod_{k=0}^N \{f_k(\mathbf{s})\}$ with here $f_0(\mathbf{s}) = \mathcal{N}(\mathbf{s}; \mathbf{m}_0, \mathbf{V}_0)$ and $f_k(\mathbf{s}) = \Phi[\mathbf{s}^T \mathbf{x}^{(k)}] \quad \forall k \in \{1 \dots N\}$. In expectation propagation we assume our approximate density factorises equivalently

$$q(\mathbf{s}) = \prod_{k=0}^N \{\tilde{f}_k(\mathbf{s})\} \quad (6)$$

with each of the approximate density factors $\tilde{f}_k(\mathbf{s})$ an (unnormalised) exponential family distribution itself, with a property of exponential family meaning that the product of exponential family distributions is always another exponential family distribution.

Here we will choose for each of the approximate terms to be scaled Gaussian densities $\tilde{f}_k(\mathbf{s}) = c_k \mathcal{N}(\mathbf{s}; \mathbf{m}_k, \mathbf{V}_k)$, with in this case we being able to exactly incorporate the prior using $\tilde{f}_0(\mathbf{s}) = f_0(\mathbf{s}) = \mathcal{N}(\mathbf{s}; \mathbf{m}_0, \mathbf{V}_0)$ and the overall approximate density itself being Gaussian

$$q(\mathbf{s}) = \mathcal{N}(\mathbf{s}; \mathbf{m}, \mathbf{V}) \quad \text{with } \mathbf{V} = \left[\sum_{k=0}^N (\mathbf{V}_k^{-1}) \right]^{-1} \quad \text{and } \mathbf{m} = \mathbf{V} \sum_{k=0}^N (\mathbf{V}_k^{-1} \mathbf{m}_k) \quad (7)$$

One way we might fit our approximate distribution is to individually fit each of our approximate factors \tilde{f}_k to the corresponding true factor f_k . In practice this does not work very well as even if each of the individual factors are well approximated this does not mean the overall approximation will be good.

Expectation propagation instead optimises each approximate factor to fit to the true factor in the context of all the remaining factors. Iteratively for each factor $j \in \{1 \dots N\}$ it sets the approximate factor parameters c_j , \mathbf{m}_j and \mathbf{V}_j so as to minimise a ‘distance’¹ between the densities proportional to $\tilde{f}_j(\mathbf{s}) \prod_{k=0, k \neq j}^N \tilde{f}_k(\mathbf{s})$ and $f_j(\mathbf{s}) \prod_{k=0, k \neq j}^N \tilde{f}_k(\mathbf{s})$, thus making the approximate factor \tilde{f}_j be a good fit to the true factor f_j in regions of the state space where the remaining approximate density factors currently suggest most of the posterior probability mass is.

Although this is not guaranteed to lead to an optimal overall fit between the true and approximate posterior densities, in practice this method usually leads to a good fit with the algorithm converging to a fixed point after a few sequential passes across all the factors.

Given an approximate posterior density fitted using expectation propagation we can then use the identity

$$\int_{\mathbb{R}^M} \Phi[\mathbf{s}^T \mathbf{x}] \mathcal{N}(\mathbf{s}; \mathbf{m}, \mathbf{V}) d\mathbf{s} = \Phi \left[\frac{\mathbf{m}^T \mathbf{x}}{\sqrt{\mathbf{x}^T \mathbf{V} \mathbf{x} + 1}} \right] \quad (8)$$

to approximately compute predictive probabilities of the form given in (5)

$$\mathbb{P} \left[r^* = 1 \mid \{r^{(k)}\}_{k=1}^N \right] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbb{P} \left[r^* = 1 \mid s_{b^*}, s_{w^*} \right] \mathbb{P} \left[s_{b^*}, s_{w^*} \mid \{r^{(k)}\}_{k=1}^N \right] ds_{b^*} ds_{w^*} \quad (9)$$

$$= \int_{\mathbb{R}^M} \mathbb{P} \left[r^* = 1 \mid \mathbf{s} \right] \mathbb{P} \left[\mathbf{s} \mid \{r^{(k)}\}_{k=1}^N \right] d\mathbf{s} \quad (10)$$

$$\approx \int_{\mathbb{R}^M} \Phi[\mathbf{s}^T \mathbf{x}^*] \mathcal{N}(\mathbf{s}; \mathbf{m}, \mathbf{V}) d\mathbf{s} \quad (11)$$

$$= \Phi \left[\frac{\mathbf{m}^T \mathbf{x}^*}{\sqrt{\mathbf{x}^{*T} \mathbf{V} \mathbf{x}^* + 1}} \right] \quad (12)$$

¹The ‘distance’ measure is the Kullback–Leibler divergence or relative entropy which is not a true distance measure as it is non-symmetric.

References

- D. Barber. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.
- C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT press, 2009.
- T. P. Minka. Expectation propagation for approximate Bayesian inference. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 362–369. Morgan Kaufmann Publishers Inc., 2001.
- K. P. Murphy. *Machine Learning: a Probabilistic Perspective*. MIT press, 2012.