

Probabilistic Modelling and Reasoning: Assignment Part 1/2*

This is Part 1/2 of the assignment with Questions 1 and 2 out of four, containing 50/100 marks.

Deadline: Wednesday March 20 2019 at 16:00

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

Group work: You may work alone or in pairs. If the work is done in pairs, both students have to work together on all questions in both Part 1 and Part 2.

Submission instructions: ~~After completing Part 1 and Part 2 you should submit the solutions to both parts as a single document manually to the ITO office by the deadline.~~ There are different deadlines for the two parts of the assignment. When completed, submit Part 1 manually to the ITO office by the deadline. Handwritten submissions are acceptable if the handwriting is *neat and legible*. If we cannot read something we cannot give you marks for it. Do not put your name but only your student ID on the copy.

If the work is done in pairs, only a single report needs to be submitted. Note both student IDs on the copy and *add a statement confirming that both students contributed equally to answering all questions.*

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

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

In particular, do not show your code, answers or write-up to anyone else.

Late submissions: We follow the policy of the School of Informatics: <http://web.inf.ed.ac.uk/infweb/student-services/ito/admin/coursework-projects/late-coursework-extension-requests>

*Based on material courtesy of Dr. Graham and Prof. Storkey.

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 and $\Phi(x)$ is the standard Gaussian cumulative density function (cdf), i.e.

$$\mathcal{N}(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad \Phi(x) = \int_{-\infty}^x \mathcal{N}(u; 0, 1) du.$$

$\{x_i\}_{i=1}^N$ indicates a collection of N variables i.e. $\{x_i\}_{i=1}^N = \{x_1, x_2, \dots, x_N\}$. \mathbb{R} is the set of real numbers and \mathbb{N} is the set of natural numbers $\{1, 2, 3, \dots\}$.

Question 1: Modelling player skills (20 marks)

We here consider the problem of modelling the skills of Go players.

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 the *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 directed 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 graph in Figure 1 models the outcome of one game only. The skills of the black and white players in the k^{th} game are represented by 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: Directed graph for a 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 in game k . The skill of the black player is s_{b_k} , the skill of the white player is s_{w_k} 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.

[5 MARKS]

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 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 graph.]

- (b) Using the directed graph from part (a) and the rules for assessing conditional independence in directed graphical models, state and explain whether s_1 and s_2 are conditionally independent if we

[5 MARKS]

know the value of $r^{(2)}$. Similarly state and explain whether s_1 and s_2 are conditionally independent if we know $r^{(2)}$ and $r^{(3)}$. Interpret the results in terms of the Go game.

[Expected response: Two conditional independency statements, justifications, and interpretation]

- (c) Draw an undirected minimal I-map for the directed graphical model specified by the DAG in part (a). [2 MARKS]

[Expected response: Labelled undirected graph.]

- (d) Fill in a table like that shown in Table 1. Use a tick or cross to indicate whether the independence relation shown in the table row generally holds under the given model (**UGM** is the Undirected Graphical Model, **DGM** is the Directed Graphical Model). Provide a short justification for your answer, e.g. UGM: trail a-b-c is not blocked, DGM: blocked by conditioning on d. [8 MARKS]

[Expected response: A filled out table like that provided]

Independence	UGM	DGM	Justification
$r^{(1)} \perp\!\!\!\perp \{r^{(2)}, r^{(3)}\} \mid \{s_1, s_2, s_3\}$	✓	✓	UGM: cond. on s_1 and s_2 blocks all trails from $r^{(1)}$. DGM: ditto because s_1 and s_2 in diverging configuration.
$r^{(1)} \perp\!\!\!\perp r^{(2)}$			
$r^{(1)} \perp\!\!\!\perp r^{(2)} \mid s_2$			
$r^{(1)} \perp\!\!\!\perp r^{(2)} \mid \{s_1, s_2\}$			
$r^{(1)} \perp\!\!\!\perp r^{(2)} \mid \{s_2, r^{(3)}\}$			

Table 1: The table to fill out for Question (d)

Question 2: Inference in a discrete-valued skill model (30 marks)

For this question, we assume that the skill of the players is represented as a discrete random variable taking values between $1, 2, \dots, 5$, where 1 is the lowest skill level and 5 the highest.

- (a) How does the joint probability mass function (pmf) $\mathbb{P} [s_1, s_2, s_3, r^{(1)}, r^{(2)}, r^{(3)}]$ over the skills s_1, s_2, s_3 and the three match outcomes in the table of Question 1 (a) factorise? [2 MARKS]

[Expected response: Factorisation of the joint pmf with explanation of the terms involved.]

- (b) Observe the undirected factor graph given in Figure 2 and understand what it represents. Define each factor in terms of a (conditional) probability. [3 MARKS]

[Expected response: Six equations, one for each factor.]




Figure 2: Undirected factor graph representing the joint pmf $\mathbb{P} [s_1, s_2, s_3, r^{(1)}, r^{(2)}, r^{(3)}]$, i.e. three games between three players.

- (c) Assume you observe the following outcomes for the three games [3 MARKS]

$$r^{(1)} = 1, \quad r^{(2)} = 0, \quad r^{(3)} = 1. \quad (1)$$

Draw the factor graph for $\mathbb{P} [s_1, s_2, s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1]$ and define or provide equations for all factors in the graph.

[Expected response: Labelled undirected factor graph with definitions or equations for the factors.]

We now consider the situation where we have reasonable knowledge about the skills of the first two players, but know nothing about the skill of the third player. Our knowledge about the skills is represented by the probability mass functions $\mathbb{P} [s_i]$ whose values are shown below as vectors such that the k -th element of the vector equals $\mathbb{P} [s_i = k]$, i.e. the probability of having skill level k , with $k \in \{1, \dots, 5\}$.

$$\mathbb{P} [s_1] = \begin{pmatrix} 0.01 \\ 0.01 \\ 0.08 \\ 0.2 \\ 0.7 \end{pmatrix}, \quad \mathbb{P} [s_2] = \begin{pmatrix} 0.02 \\ 0.02 \\ 0.06 \\ 0.3 \\ 0.6 \end{pmatrix}, \quad \mathbb{P} [s_3] = \begin{pmatrix} 0.2 \\ 0.2 \\ 0.2 \\ 0.2 \\ 0.2 \end{pmatrix}. \quad (2)$$

For example, $\mathbb{P} [s_1 = 3] = 0.08$. The goal of the following questions is to use probabilistic inference to update the belief about the third player's skill given the outcomes of the three games. In other words, we will compute the posterior pmf

$$\mathbb{P} [s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1].$$

We assume that the conditional probability for black winning in any game is given by

$$\mathbb{P} [r^{(k)} = 1 | s_{b_k}, s_{w_k}] = \frac{1}{10}(s_{b_k} - s_{w_k}) + 0.5. \quad (3)$$

This conditional probability only depends on the difference between the skill of the black and white player.

- (d) Can (standard) message passing be used to compute $\mathbb{P}[s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1]$? [1 MARKS]

[Expected response: Yes/no answer with justification.]

- (e) We can marginalise over s_1 to obtain [5 MARKS]

$$\mathbb{P}[s_2, s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1]$$

from

$$\mathbb{P}[s_1, s_2, s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1].$$

Show that summing out s_1 eliminates some nodes in the factor graph and creates a new factor $\tilde{\phi}_1(s_2, s_3)$. Derive the equation defining $\tilde{\phi}_1(s_2, s_3)$ and draw the factor graph representing $\mathbb{P}[s_2, s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1, r^{(3)} = 1]$.

[Expected response: Step-by-step derivation of the equation with explanation of the manipulations performed (e.g. “used sum rule” etc). Labelled undirected factor graph]

- (f) The desired posterior $\mathbb{P}[s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1]$ can be represented with the factor graph below where ψ_1 is purely determined by the prior pmf of s_3 . [6 MARKS]



Derive expressions for $\psi_1(s_3)$ and $\psi_2(s_3)$, and compute their numerical values for $s_3 \in \{1, 2, \dots, 5\}$. Rescale both vectors such that the maximal value of the elements is one (this facilitates marking since there may be differences in scale).

[Expected response: Two vectors with the values of $\psi_1(s_3)$ and $\psi_2(s_3)$. Derivation of the result and code snippet showing numerical computation.]

- (g) Compute the numerical values of the posterior $\mathbb{P}[s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1]$ for $s_3 \in \{1, 2, \dots, 5\}$. Make a plot that compares $\mathbb{P}[s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1]$ with $\mathbb{P}[s_3]$, explain the differences between the prior and posterior. [4 MARKS]

[Expected response: Vector with values of the posterior. Explanation and figure with the plot.]

- (h) We now consider an alternative model for how the probability of black winning depends on the difference in the skill-levels. We replace $\mathbb{P}[r^{(k)} = 1 | s_{b_k}, s_{w_k}]$ in Equation (3) with [6 MARKS]

$$\mathbb{P}[r^{(k)} = 1 | s_{b_k}, s_{w_k}] = \frac{1}{2} (1 + \tanh(s_{b_k} - s_{w_k})). \quad (4)$$

Re-compute the numerical values of the posterior $\mathbb{P}[s_3 | r^{(1)} = 1, r^{(2)} = 0, r^{(3)} = 1]$ for $s_3 \in \{1, 2, \dots, 5\}$ and make a plot that compares the obtained posterior with $\mathbb{P}[s_3]$. Explain how the observed differences between this posterior and the one obtained before relates to the change in the model of $\mathbb{P}[r^{(k)} = 1 | s_{b_k}, s_{w_k}]$.

[Expected response: Vector with values of the posterior. Figure with the plot. Differences in the obtained posterior. Explanation.]