Spring 2023
Michael Gutmann &
Chris Williams

Exercises for the tutorials: 1 and 2.

The other exercises are for self-study and exam preparation. All material is examinable unless otherwise mentioned.

Exercise 1. *Conversion to factor graphs*

(a) Draw an undirected factor graph for the directed graphical model defined by the graph below.

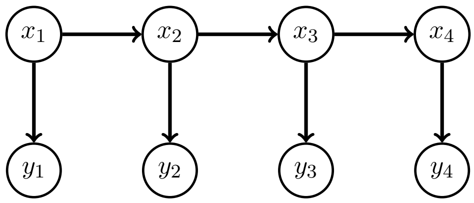

(b) Draw an undirected factor graph for directed graphical models defined by the graph below (this kind of graph is called a polytree: there are no loops but a node may have more than one parent).

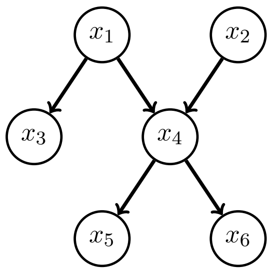

Exercise 2. *Variable elimination for the "car start" network*

Consider the belief network given in Figure 1, which gives the probability of a car starting (or not) (s), depending on the state of the battery (b), if it has fuel (f), the state of the fuel gauge (g), and if the car "turns over" (t).

Below we code the variables as 0/1, with:
$b = 0 \equiv b = bad$, and $b = 1 \equiv b = good$,
$f = 0 \equiv f = empty$, and $f = 1 \equiv f = not\ empty$,
$g = 0 \equiv g = empty$, and $g = 1 \equiv g = not\ empty$,
$t = 0 \equiv t = no$, and $t = 1 \equiv t = yes$,
$s = 0 \equiv s = no$, and $s = 1 \equiv s = yes$.

We have that $p(b = 0) = 0.02$, and $p(f = 0) = 0.05$. The other CPTs are given by

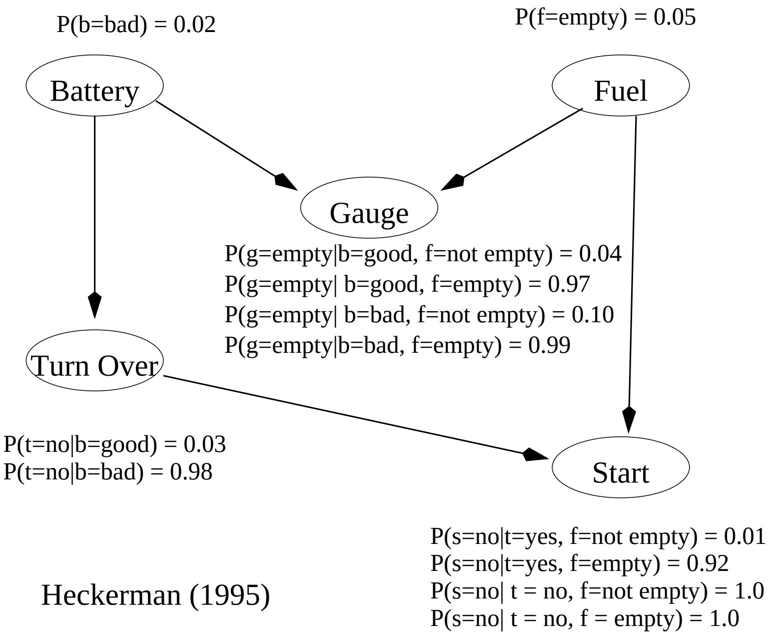

Figure 1: Simple network describing if a car's engine starts, due to Heckerman (1995).

b	f	$p(g = 0 \mid b, f)$
0	0	0.99
0	1	0.10
1	0	0.97
1	1	0.04

b	$p(t = 0 \mid b)$
0	0.98
1	0.03

t	f	$p(s = 0 \mid t, f)$
0	0	1.0
0	1	1.0
1	0	0.92
1	1	0.01

(a) From the graphical model we can read off

$$p(b, f, g, t, s) = p(b)p(f)p(g|b, f)p(t|b)p(s|t, f). \tag{1}$$

Now consider the potential representation

$$p(b, f, g, t, s) \overset{def}{=} \phi_{bt}(b, t)\phi_{stf}(s, t, f)\phi_{gbf}(g, b, f).$$

Assign the factors in eq. 1 to the potentials in a valid way.

(b) We wish to compute $p(b, f|s = 0)$ (i.e., the car does not start). This conditioning has the effect of turning

$$p(b, f, g, t, s) = \phi_{bt}(b, t)\phi_{stf}(s, t, f)\phi_{gbf}(g, b, f)$$

into

$$p(b, f, g, t|s = 0) \propto \phi_{bt}(b, t)\phi_{tf}^{s=0}(t, f)\phi_{gbf}(g, b, f).$$

Identify the modification needed to the factor(s) in the potential ϕ_{stf} to turn it into $\phi_{tf}^{s=0}$.

(c) First eliminate g to obtain a potential representation for $p(b, f, t|s = 0)$.

(d) Now eliminate t to obtain a potential representation for $p(b, f|s = 0)$. Evaluate this for all four possible combinations of b and f. Normalize this to obtain the posterior probability of all four configurations of b and f.

(e) Wellman and Henrion (1993)[1] state:

[1] IEEE Transactions on Pattern Analysis and Machine Intelligence, 15(3) pp 287–292 (1993).

"Explaining away" is a common pattern of reasoning in which the confirmation of one cause of an observed or believed event reduces the need to invoke alternative causes.

Think about the four posterior probabilites obtained above. Given that the car does not start, do these broadly meet with your expectations? Do you observe "explaining away"?

(f) In general we write computer programs to do inference in such models. See `https://github.com/vsimkus/pmr2023-pgm-demo` for a demo for this question prepared by the PMR TA Vaidotas Simkus using the `pgmpy` Python package. You can run the notebook on Google Colab directly via the link `http://colab.research.google.com/github/vsimkus/pmr2023-pgm-demo`.

Exercise 3. ***Limits of directed and undirected graphical models***

We here consider the probabilistic model $p(y_1, y_2, x_1, x_2) = p(y_1, y_2|x_1, x_2)p(x_1)p(x_2)$ where $p(y_1, y_2|x_1, x_2)$ factorises as

$$p(y_1, y_2|x_1, x_2) = p(y_1|x_1)p(y_2|x_2)\phi(y_1, y_2)n(x_1, x_2) \tag{2}$$

with $n(x_1, x_2)$ equal to

$$n(x_1, x_2) = \left(\int p(y_1|x_1)p(y_2|x_2)\phi(y_1, y_2)\mathrm{d}y_1\mathrm{d}y_2 \right)^{-1}. \tag{3}$$

In the lecture "Factor Graphs", we used the model to illustrate the setup where x_1 and x_2 are two independent inputs that each control the interacting variables y_1 and y_2 (see graph below).

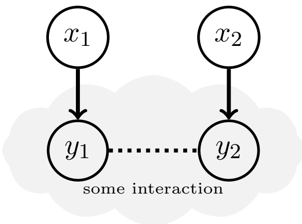

(a) Use the basic characterisations of statistical independence

$$u \perp\!\!\!\perp v|z \Longleftrightarrow p(u, v|z) = p(u|z)p(v|z) \tag{4}$$
$$u \perp\!\!\!\perp v|z \Longleftrightarrow p(u, v|z) = a(u, z)b(v, z) \qquad (a(u, z) \geq 0, b(v, z) \geq 0) \tag{5}$$

to show that $p(y_1, y_2, x_1, x_2)$ satisfies the following independencies

$$x_1 \perp\!\!\!\perp x_2 \qquad x_1 \perp\!\!\!\perp y_2 \mid y_1, x_2 \qquad x_2 \perp\!\!\!\perp y_1 \mid y_2, x_1$$

(b) *(optional, not examinable)* The following factor graph represents $p(y_1, y_2, x_1, x_2)$:

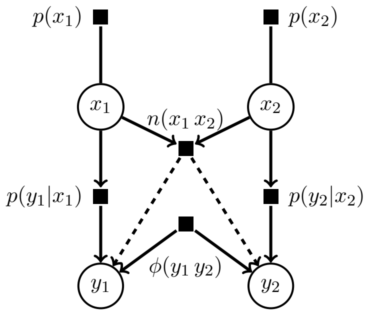

Use the separation rules for factor graphs to verify that we can find all independence relations. The separation rules are (see Barber, section 4.4.1, or the original paper by Brendan Frey: https://arxiv.org/abs/1212.2486):

"If all paths are blocked, the variables are conditionally independent. A path is blocked if one or more of the following conditions is satisfied:

1. One of the variables in the path is in the conditioning set.
2. One of the variables or factors in the path has two incoming edges that are part of the path (variable or factor collider), and neither the variable or factor nor any of its descendants are in the conditioning set."

Remarks:

- "one or more of the following" should best be read as "one of the following".
- "incoming edges" means directed incoming edges
- the descendants of a variable or factor node are all the variables that you can reach by following a path (containing directed or directed edges, but for directed edges, all directions have to be consistent)
- In the graph we have dashed directed edges: they do count when you determine the descendants but they do not contribute to paths. For example, y_1 is a descendant of the $n(x_1, x_2)$ factor node but $x_1 - n - y_2$ is not a path.