COMPUTER SCIENCE TRIPOS Part IB – mock – Paper 6

4 Foundations of Data Science (DJW)

I am playing a game of solitaire, which involves repeatedly tossing a fair coin. If I get three heads in a row I win, if I get two tails in a row I lose.

(*a*) Devise a Markov chain to represent the state of the game, and draw the state space. The state space diagram should have eight states, including

- a state $\varnothing$ to represent "not yet tossed any coins",
- a state TT to represent "lost", with a single outgoing transition back to state TT,
- a state HHH to represent "won", with a single outgoing transition back to state HHH.

[5 marks]

(*b*) I wish to compute the probability of winning. Let ρ_x be the probability that I will win, when starting from state x. Clearly $\rho_{TT} = 0$ and $\rho_{HHH} = 1$. Show that for any other state x

$$\rho_x = \sum_y \mathbb{P}\big(\text{will win} \mid \text{start at } y\big) P_{xy}$$

for a suitable matrix P, which you should define. Explain your reasoning carefully. [6 marks]

(*c*) Write out a set of equations that could be solved to find $\rho_\varnothing$. You do not need to solve them. [3 marks]

(*d*) Explain what is meant by *stationary distribution*. [2 marks]

(*e*) Let $\lambda \in [0, 1]$, and define a distribution π by

$$\pi_x = \lambda 1_{x=TT} + (1 - \lambda) 1_{x=HHH}\,.$$

Show that π is a stationary distribution for your Markov chain. [4 marks]

Answer: Part (a). The state space diagram is like this. All edges have probability 1/2, apart from the self-loops on TT and HHH which have probability 1.

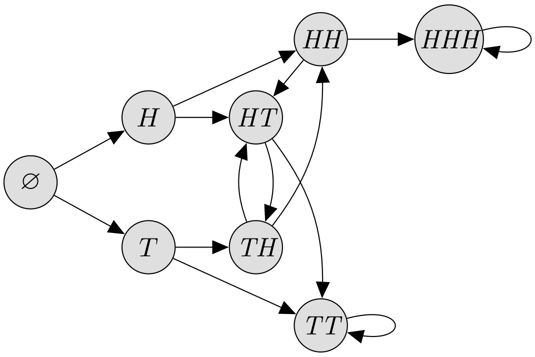

Part (b). Let X_n be the state at time n. Then

$$\begin{aligned}\rho_x &= \mathbb{P}\big(\text{will win} \mid X_0 = x\big)\\ &= \sum_y \mathbb{P}\big(\text{will win} \mid X_1 = y, X_0 = x\big)\,\mathbb{P}(X_1 = y \mid X_0 = x) \quad \text{by the law tot. prob.}\\ &= \sum_y \mathbb{P}\big(\text{will win} \mid X_1 = y\big) P_{xy} \quad \text{from the causal diagram}\\ &= \sum_y \mathbb{P}\big(\text{will win} \mid \text{start from } y\big) P_{xy}.\end{aligned}$$

The last equation is because the Markov chain has the same procedure for generating the next state from every timestep, and so it doesn't make any difference whether we start at timestep 0 or timestep 1. Here, P is the transition matrix, $P_{xy} = \mathbb{P}(X_1 = y \mid X_0 = x)$.

Part (c). All we need to do is write out the equation for ρ_x for every state ρ.

$$\begin{aligned}\rho_\varnothing &= \rho_H/2 + \rho_T/2\\ \rho_H &= \rho_{HH}/2 + \rho_{HT}/2\\ \rho_T &= \rho_{TH}/2\\ \rho_{HH} &= 1/2 + \rho_{HT}/2\\ \rho_{HT} &= \rho_{TH}/2\\ \rho_{TH} &= \rho_{HH}/2 + \rho_{HT}/2\\ \rho_{TT} &= 0 \quad \text{from the question}\\ \rho_{HHH} &= 1 \quad \text{from the question}\end{aligned}$$

It's actually fairly easy to solve these, starting from $\{TH, HH, HT\}$ and working backwards. The answer is $\pi_\varnothing = 3/10$.

Part (d). A distribution π is said to be *stationary* if "if the chain starts in distribution π it remains in distribution π", i.e.

$$\mathbb{P}(X_0 = x) = \pi_x \text{ for all } x \qquad \Rightarrow \qquad \mathbb{P}(X_n = x) = \pi_x \text{ for all } x \text{ and } n.$$

Part (e). From general theory, we know that if a distribution π satisfies

$$\pi_x = \sum_y \pi_y P_{yx} \quad \text{for all } x$$

then it is a stationary distribution. So, we need to verify that the formula for π given in the question satisfies this equation, for every x. Start with $x = HHH$:

$$\sum_y \pi_y P_{yHHH} = \frac{\pi_{HH}}{2} + \pi_{HHH} = \pi_{HHH}$$

where the first equality comes from enumerating all the edges that point into *HHH*, and the second equality comes because $\pi_{HH} = 0$ for the distribution π as defined in the question. So, the equation is satisfied at $x = HHH$. Likewise, $x = TT$. And likewise also all the other cases: for every other state x, all the states y that point into x are assigned $\pi_y = 0$.

Hence the equation is satisfied for all x, hence the given distribution is a stationary distribution.