COMPUTER SCIENCE TRIPOS Part II – 2017 – Paper 7

3 Machine Learning and Bayesian Inference (SBH)

(*a*) For random variables (RVs) A_1, A_2 and B, define what it means for A_1 to be *conditionally independent of* A_2 *given* B, written $A_1 \perp A_2|B$. [1 mark]

(*b*) Given mutually disjoint sets X_1, X_2 and Y of random variables from some Bayesian network, define what it means to say that a path from $x_1 \in X_1$ to $x_2 \in X_2$ is *blocked* by Y. [5 marks]

(*c*) Given mutually disjoint sets X_1, X_2 and Y of random variables from some Bayesian network, define what it means for X_1 and X_2 to be *d-separated* by Y. What does this tell you about the probability distribution represented by the Bayesian network? [3 marks]

(*d*) Consider the following Bayesian network.

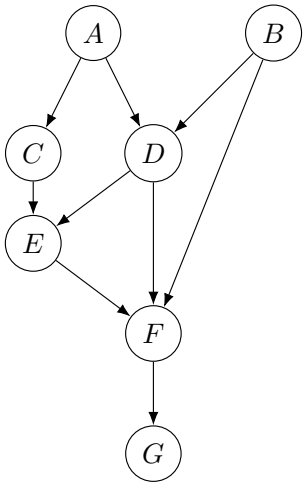

In each of the following cases, establish whether or not $X_1 \perp X_2|Y$:

(*i*) $X_1 = \{A, E\}$, $X_2 = \{G\}$ and $Y = \{F\}$. [2 marks]

(*ii*) $X_1 = \{A\}$, $X_2 = \{D\}$ and $Y = \{G\}$. [3 marks]

In each case explain your answer.

(*e*) Define how Bayesian networks and Markov random fields are used to represent probability distributions, and briefly describe the trade-offs involved in choosing one versus the other. [6 marks]