Reasoning Under Uncertainty: Marginal and Conditional Independence

CPSC 322 – Uncertainty 3

Textbook §6.2

Lecture Overview

Conditioning

- Probabilistic conditioning specifies how to revise beliefs based on new information.
- You build a probabilistic model taking all background information into account. This gives the prior probability.
- All other information must be conditioned on.
- If evidence e is all of the information obtained subsequently, the conditional probability $P(h|e)$ of h given e is the posterior probability of h.

Conditional Probability

The conditional probability of formula h given evidence e is

$$P(h|e) = \frac{P(h \wedge e)}{P(e)}$$

Chain rule:

$$P(f_1 \wedge f_2 \wedge \ldots \wedge f_n) = \prod_{i=1}^{n} P(f_i | f_1 \wedge \cdots \wedge f_{i-1})$$

Bayes' theorem:

$$P(h|e) = \frac{P(e|h) \times P(h)}{P(e)}.$$

Lecture Overview

Marginal independence

Definition (marginal independence)

Random variable X is marginally independent of random variable Y if, for all $x_i \in dom(X)$, $y_j \in dom(Y)$ and $y_k \in dom(Y)$,

$$\begin{aligned} & P(X = x_i | Y = y_j) \\ = \; & P(X = x_i | Y = y_k) \\ = \; & P(X = x_i). \end{aligned}$$

That is, knowledge of Y's value doesn't affect your belief in the value of X.

Examples of marginal independence

- The probability that the Canucks will win the Stanley Cup is independent of whether light $l1$ is lit.
    - remember the diagnostic assistant domain—the picture will recur in a minute!
- Whether there is someone in a room is independent of whether a light $l2$ is lit.
- Whether light $l1$ is lit is *not* independent of the position of switch $s2$.

Lecture Overview

Conditional Independence

- Sometimes, two random variables might not be marginally independent. However, they can *become* independent after we observe some third variable.

Definition

Random variable X is conditionally independent of random variable Y given random variable Z if, for all $x_i \in dom(X)$, $y_j \in dom(Y)$, $y_k \in dom(Y)$ and $z_m \in dom(Z)$,

$$\begin{aligned} & P(X = x_i | Y = y_j \wedge Z = z_m) \\ = \ & P(X = x_i | Y = y_k \wedge Z = z_m) \\ = \ & P(X = x_i | Z = z_m). \end{aligned}$$

- That is, knowledge of Y's value doesn't affect your belief in the value of X, given a value of Z.

Conditional Independence Example

- Kevin separately phones two students, Alice and Bob.
- To each, he tells the same number, $n_k \in \{1, \ldots, 10\}$.
- Due to the noise in the phone, Alice and Bob each imperfectly (and independently) draw a conclusion about what number Kevin said.
- Let the numbers Alice and Bob think they heard be n_a and n_b respectively.
- Are n_a and n_b marginally independent?

Conditional Independence Example

- Kevin separately phones two students, Alice and Bob.
- To each, he tells the same number, $n_k \in \{1, \ldots, 10\}$.
- Due to the noise in the phone, Alice and Bob each imperfectly (and independently) draw a conclusion about what number Kevin said.
- Let the numbers Alice and Bob think they heard be n_a and n_b respectively.
- Are n_a and n_b marginally independent?
  - No: we'd expect (e.g.) $P(n_a = 1 | n_b = 1) > P(n_a = 1)$.

Conditional Independence Example

- Kevin separately phones two students, Alice and Bob.
- To each, he tells the same number, $n_k \in \{1, \ldots, 10\}$.
- Due to the noise in the phone, Alice and Bob each imperfectly (and independently) draw a conclusion about what number Kevin said.
- Let the numbers Alice and Bob think they heard be n_a and n_b respectively.
- Are n_a and n_b marginally independent?
  - No: we'd expect (e.g.) $P(n_a = 1 | n_b = 1) > P(n_a = 1)$.
- Why are n_a and n_b conditionally independent given n_k?

Conditional Independence Example

- Kevin separately phones two students, Alice and Bob.
- To each, he tells the same number, $n_k \in \{1, \ldots, 10\}$.
- Due to the noise in the phone, Alice and Bob each imperfectly (and independently) draw a conclusion about what number Kevin said.
- Let the numbers Alice and Bob think they heard be n_a and n_b respectively.
- Are n_a and n_b marginally independent?
  - No: we'd expect (e.g.) $P(n_a = 1 | n_b = 1) > P(n_a = 1)$.
- Why are n_a and n_b conditionally independent given n_k?
  - Because if we know the number that Kevin actually said, the two variables are no longer correlated.
  - e.g., $P(n_a = 1 | n_b = 1, n_k = 2) = P(n_a = 1 | n_k = 2)$

Example domain (diagnostic assistant)

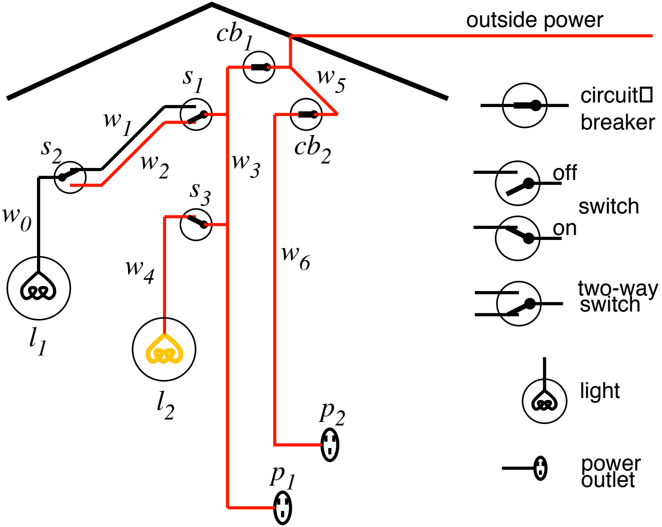

More examples of conditional independence

- Whether light $l1$ is lit is independent of the position of light switch $s2$ given whether there is power in wire w_0.
  - two random variables that are not marginally independent can still be conditionally independent
- Every other variable may be independent of whether light $l1$ is lit given whether there is power in wire w_0 and the status of light $l1$ (if it's ok, or if not, how it's broken).

More examples of conditional independence

- The probability that the Canucks will win the Stanley Cup is independent of whether light $l1$ is lit given whether there is outside power.
  - sometimes, when two random variables are marginally independent, they're also conditionally independent given a third variable.
- But not always...
  - Let C_1 be the proposition that coin 1 is heads; let C_2 be the proposition that coin 2 is heads; let B be the proposition that coin 1 and coin 2 are both either heads or tails.
  - $P(C_1|C_2) = P(C_1)$: C_1 and C_2 are marginally independent.
  - But $P(C_1|C_2, B) \neq P(C_1|B)$: if I know both C_2 and B, I know C_1 exactly, but if I only know B I know nothing.
  - Hence C_1 and C_2 are *not* conditionally independent given B.