CS109: Probability for Computer Scientists

Lecture 12 — General Inference

Feb 2

Problem 1: Groundhog Day

Did you know today is Groundhog's Day?

Sees shadow = 6 more weeks of winter.

Doesn't see shadow = early spring.

Based on historical data:

- When the groundhog sees shadow, 70% chance of 6 more weeks of winter.
- When the groundhog doesn't see shadow, 40% chance of 6 more weeks of winter.
- Before we observe the groundhog, we think it is equally likely to be an early spring or 6 more weeks of winter.

a) Compute the probability of 6 more weeks of winter *given that the groundhog sees his shadow*.

Problem 2: Lidar in 1D

Let T be the true distance. Your prior belief is: $T \sim N(\mu = 1,\ \sigma^2 = 3)$. Your sensor has uncertainty:

$$X \mid (T = t) \sim \mathcal{N}(\mu = t,\ \sigma^2 = 1.5).$$

You observe: $X = 4$.

a) Write Bayes' rule for the posterior density in the form

$$f(T = t \mid X = 4) \ \propto\ f(X = 4 \mid T = t) \cdot f(T = t).$$

(No need to simplify.)

b) (Optional) Compute the posterior distribution for T given $X = 4$. Specifically, give the mean and variance of the posterior distribution. Will require some algebra and completing the square.

c) **(Optional — Super challenge: 2D Tracking)**

Now the object is at an unknown location (X, Y).

Prior: you believe (X, Y) is centered around $(3, 3)$ with joint density

$$f(X = x, Y = y) = \frac{1}{8\pi} \exp\left(-\frac{(x-3)^2 + (y-3)^2}{8} \right).$$

Likelihood: you observe a noisy distance reading D from a sensor at $(0, 0)$. The sensor model is

$$D \mid (X = x, Y = y) \sim \mathcal{N}\left(\mu = \sqrt{x^2 + y^2},\ \sigma^2 = 1 \right),$$

and you observe $D = 4$. Write the unnormalized posterior density using $\propto$:

$$f(X = x, Y = y \mid D = 4) \ \propto\ f(D = 4 \mid X = x, Y = y) \cdot f(X = x, Y = y).$$

Problem 3: Size of a Joint Distribution

Suppose you have N binary random variables.

a) If $N = 9$, how many entries are in the full joint probability table?

b) For general N, how many entries are in the full joint probability table?

Problem 4: Bayes Net with Probabilities

Consider the Bayes net with binary variables:

$$Flu \to Fever, \qquad (Flu, U) \to Tired,$$

where U stands for "Undergrad".

Given probabilities:

$$P(Flu = 1) = 0.1, \qquad P(U = 1) = 0.8,$$

$$P(Fever = 1 \mid Flu = 1) = 0.9, \qquad P(Fever = 1 \mid Flu = 0) = 0.05,$$

$$P(Tired = 1 \mid Flu = 0, U = 0) = 0.1$$
$$P(Tired = 1 \mid Flu = 0, U = 1) = 0.8$$
$$P(Tired = 1 \mid Flu = 1, U = 0) = 0.9$$
$$P(Tired = 1 \mid Flu = 1, U = 1) = 1.0$$

Diagram:

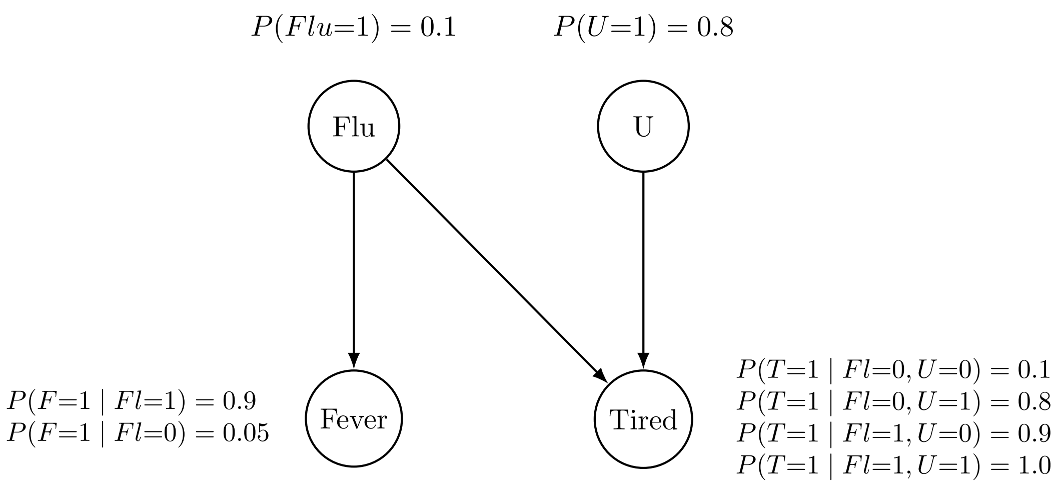

a) Compute $P(Fever = 0 \mid Flu = 1)$.

b) We want:

$$P(Flu = 1 \mid U = 1, Tired = 1).$$

A simulation-based estimate is:

$$P(Flu = 1 \mid U = 1, Tired = 1) \approx \frac{\#\text{ samples with } (Flu = 1, U = 1, Tired = 1)}{\#\text{ samples with } (U = 1, Tired = 1)}.$$

Explain *why* this ratio is a reasonable approximation.

Problem 6: The Cousin Problem

A simplified genetic model: each person has a binary variable indicating whether they have a recessive gene (1) or not (0). We observe that **Cousin 1 has the gene**.

$P(GP1{=}1) = \frac{1}{20}$ $P(GP2{=}1) = \frac{1}{20}$

$P(S1{=}1) = \frac{1}{20}$

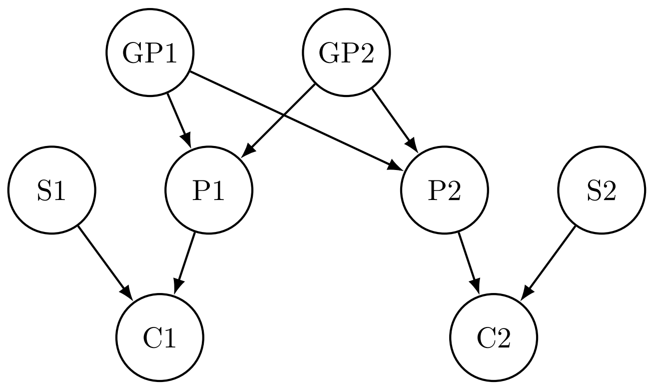

$P(S2{=}1) = \frac{1}{20}$

observed $C1 = 1$

We want (conceptually): $P(C2 = 1 \mid C1 = 1)$.

a) Describe in words: if you do **rejection sampling**, what samples do you **throw away**?

b) Describe in words: among the samples you keep, what do you **count**?

c) What steps would you include in a function `make_sample()` that generates one full assignment for all nodes in the Bayes net?