

Homework 1

Belief and the failed search

Due: October 9, 2026, 11:59 PM PT

Student Name:

Instructor: John Lipor

Instructions

Submit your work as a PDF through the course submission system. If you need an accommodation or an alternate format, contact the instructor.

Note: For all questions asked in this course, unless the question is simply asking about definitions, please justify all answers. Stating the correct answer without justification will earn no credit. For example, if you write down a string of equations, each successive equation should be accompanied by some justification such as “by the additivity axiom” or “because the outcomes are equally likely” or “by the distributive property of set operations.” You do not need to justify algebra.

Problem 1: The wreck in the grid [5 pts each]

A vessel is lost in one of four seafloor regions with prior probabilities $p = (0.7, 0.15, 0.1, 0.05)$. A side-scan sonar pass over a region *containing* the wreck detects it with probability $1 - \alpha$, where $\alpha = 0.4$ is the overlook probability (rough terrain). A pass over an empty region never produces a detection.

- (a) The team searches region 1 and finds nothing. Compute the posterior over all four regions. Verify it sums to one.
- (b) Show that the searched region’s probability strictly decreased while every unsearched region’s probability increased *by the same multiplicative factor*. Interpret in one sentence: where does the probability “flow”?
- (c) The team searches region 1 again and again finds nothing. Compute the new posterior. After how many consecutive failed passes of region 1 does region 2 become the most probable region?
- (d) Session 2 proves that after n consecutive failed searches of a cell with prior p and overlook probability α , its posterior is $p\alpha^n/(p\alpha^n + 1 - p)$. For a general tolerance $\varepsilon \in (0, 1)$, derive in closed form (via logarithms) the smallest n for which this posterior falls below ε . State the inequality carefully: which way does the sign flip when you

divide by $\ln \alpha$, and why? Evaluate your bound for the numbers of part (a) with $\varepsilon = 0.01$. Finally, what happens to your n as $\alpha \rightarrow 1$, and why does that make sense?

Problem 2: The flagged channel [5 pts each]

A spectrum sensor monitors a channel that is occupied 4% of the time. When the channel is occupied, the detector flags it with probability 0.9; when idle, it false-alarms with probability 0.075.

- (a) *Before computing anything*: identify the latent variable, the observation, the prior, and the likelihood. Then compute the probability the channel is actually occupied given a flag.
- (b) Redo the computation in odds form: posterior odds = prior odds $\times$ likelihood ratio. Which form makes it more obvious *why* the answer is far from 0.9?

Problem 3: Three coins [5 pts]

You flip three fair coins. At least two are alike, and there is an even chance that the third is a head or a tail. Therefore the probability that all coins show the same side is equal to $1/2$. Do you agree? Provide both your intuited and computed answers.

Problem 4: One look, best look [5 pts]

With the prior and overlook probability of Problem 1, suppose each region i costs $c = (1, 1, 2, 4)$ hours to search and you have time for exactly one pass. If your objective is probability of detection per hour, which region do you search? Prove that for a single look, searching the region maximizing $p_i(1 - \alpha_i)/c_i$ is optimal among all single-look strategies. (Session 3 states this allocation rule without proof; supplying the proof is your task here.)

Problem 5: Posterior heat map [10 pts]

Simulate a 10×10 grid with a prior of your choosing (*not* uniform) and overlook probability $\alpha = 0.3$. Script a sequence of 25 greedy searches (always search the current-posterior-maximizing cell; the target's true cell is drawn once from the prior).

Before running: predict in one sentence what the posterior map will look like if the target happens to be in a low-prior cell.

Turn in: posterior heat maps after 5, 15, and 25 failed/successful looks, your prediction versus what happened, and your code.

Problem 6: Audit the TA-bot [5 pts]

An AI assistant was asked: “A monitoring channel is busy 10% of the time. A detector flags a busy channel with probability 0.85 and false-alarms on an idle channel with probability 0.15. Given a flag, how likely is the channel to be busy?” It answered:

TA-bot writes: We use the odds form of Bayes’ rule. The likelihood ratio is $L = \mathbb{P}(\text{flag} \mid \text{busy}) / \mathbb{P}(\text{flag} \mid \text{idle}) = 0.85 / 0.15 = 17/3$. The posterior odds of busy versus idle are therefore $17 : 3$, giving $\mathbb{P}(\text{busy} \mid \text{flag}) = 17/20 = 0.85$. Sanity check: the detector is highly reliable, so a high posterior makes sense.

This solution contains **exactly one error**.

- (a) Identify the first line that goes wrong.
- (b) Name the violated principle in one sentence, citing the session.
- (c) Give the corrected posterior.
- (d) One sentence: why is this a mistake a language model plausibly makes?

Problem 7: Reflection / Meta-Learning [5 pts]

Read Chapter 1 of [Solving Mathematical Problems](#) by Terence Tao, and answer the following questions **in the interesting-reading channel on Discord**.

- Among the problem solving strategies discussed, have you used any (whether implicitly or explicitly) in this course so far or in other courses?
- Among the problem solving strategies discussed, which do you think would be helpful to use going forward?