CSCI 145 Problem 10: When One Threshold Is Not Enough

August 27, 2026

Submission Instructions

Please upload your work to Gradescope by **11:59pm on DATE**.

- Turn in a single PDF of *handwritten* solutions. Typed (including LaTeX) solutions will not be accepted.
- Problems are grouped into three-week **cycles**. Once per cycle, you will meet with the grader and present the solution to *one* problem from that cycle, selected at random when you arrive — so bring your work and be ready to explain any of them.

Problem 10: When One Threshold Is Not Enough 2 points

A scalar measurement X comes from one of two equally likely classes. The class-conditional distributions are

$$X \mid Y = 0 \sim \mathcal{N}(0, 1), \qquad X \mid Y = 1 \sim \mathcal{N}(0, 4).$$

A researcher proposes fitting ordinary one-feature logistic regression, $\Pr(Y = 1 \mid X = x) = \sigma(wx+b)$, and argues that the sigmoid can represent the Bayes-optimal probability because it maps every real score into $(0, 1)$.

Audit the claim. Derive the exact posterior probability from Bayes' rule and express it as a sigmoid of a score. Characterize the region in which the Bayes classifier predicts class 1, prove that no choice of w and b in the proposed model has the same decision region, and give the smallest natural feature map that makes logistic regression represent the exact posterior. Explain what this example says about the roles of the sigmoid and the features in determining a classifier's decision boundary.

Results you may use. You may quote the Gaussian density and the posterior-odds identity

$$\frac{\Pr(Y = 1 \mid X = x)}{\Pr(Y = 0 \mid X = x)} = \frac{\Pr(Y = 1)}{\Pr(Y = 0)} \frac{p(x \mid Y = 1)}{p(x \mid Y = 0)}.$$

If $z = \log(p/(1 - p))$, then $p = \sigma(z) = 1/(1 + e^{-z})$.