CSCI 145 Problem 10: A Classifier That Refuses to Learn

October 27, 2026

Submission Instructions

Please upload your work to Gradescope by **11:59pm on Wednesday, October 14**.

- Turn in a single PDF of *handwritten* solutions. Typed (including LaTeX) solutions will not be accepted.
- A complete, good-faith attempt earns submission credit. The grader may select any recent assigned problem for a brief conversation, so keep your work and be prepared to explain each solution and answer a related question. Scheduling details will be announced in class.

Problem 10: A Classifier That Refuses to Learn

Consider these four equally weighted training examples:

i	x_{i1}	x_{i2}	y_i
1	1	1	1
2	1	−1	0
3	−1	1	0
4	−1	−1	1

Consider a logistic regression model

$$z = b + w_1 x_1 + w_2 x_2, \qquad \sigma(z) = \frac{1}{1 + e^{-z}}.$$

With $n = 4$, the model is learned by minimizing the average cross-entropy

$$\mathcal{L}(b, w_1, w_2) = -\frac{1}{n} \sum_{i=1}^{n} \left[y_i \log \sigma(z_i) + (1 - y_i) \log(1 - \sigma(z_i)) \right].$$

The model predicts class 1 when $\sigma(z) > 1/2$ and class 0 otherwise.

1. Show that full-batch gradient descent initialized at $b = w_1 = w_2 = 0$ never moves. Is it stuck at a bad local minimum?
2. Find a linear classifier that correctly labels three of the four examples, and explain why this does not contradict your conclusions above.
3. Then add one feature that allows all four examples to be classified correctly. Give weights for the expanded model and verify its predictions.

Hint: Notice that

$$-y \log \sigma(z) - (1 - y) \log(1 - \sigma(z)) = \log(1 + e^{-z}) + (1 - y)z.$$

Takeaway: *Cross-entropy and classification accuracy can favor different weights. Here, gradient descent finds the best loss but gets half the labels wrong. Adding a feature lets the model classify all four points correctly.*