

Exercises

§6.2,6.3: Logistic Regression

Python: There is an optional starter script that supports Python-based problem(s) on this sheet. Use it as scaffolding: read the comments, complete the lines marked **TODO**, and compare the output with the mathematics requested in the problem(s).

Problem 1 Compute $\sigma(-2)$, $\sigma(0)$, $\sigma(2)$. Verify numerically that $\sigma(-2) = 1 - \sigma(2)$.

Problem 2 Suppose $f(x) = -4 + 1.5x$. (a) Find the value of x for which $P(Y = +1 \mid x) = 0.5$. (b) Compute the probability of class +1 when $x = 2$. (c) Compute the probability when $x = 4$.

Problem 3 A logistic model has $w_1 = \log 3$. By what factor do the odds of class +1 change when x_1 increases by one unit, with all other features held fixed? What if x_1 increases by two units?

Problem 4 For labels $y \in \{-1, +1\}$, compute logistic loss $\ell(y, f) = \log(1 + e^{-yf})$ for the following pairs: (a) $y = +1, f = 2$, (b) $y = +1, f = -2$, (c) $y = -1, f = -2$, (d) $y = -1, f = 0$. Which cases are correctly classified?

Problem 5 Show directly that $P(Y = y \mid \mathbf{x}) = \sigma(yf)$ works for both $y = +1$ and $y = -1$. Then derive $-\log P(Y = y \mid \mathbf{x}) = \log(1 + e^{-yf})$.

Problem 6 Let $\ell(m) = \log(1 + e^{-m})$. Differentiate with respect to m and show that $\ell'(m) = -\sigma(-m)$. Explain why this derivative has the correct sign for a classification loss.

Problem 7 Use Python to compute probabilities and predicted classes for

```
scores = np.array([-3., -1., 0., 0.5, 2.])
```

using the sigmoid and a 0.5 threshold. Then compute the same classes using a zero threshold on the scores. Verify that the predictions agree.

Problem 8 Consider the separable data from class $x = (1, 2, 3, 4, 5, 6)$, $y = (-1, -1, -1, +1, +1, +1)$. Compare the average logistic loss for $f_1(x) = x - 3.5$ and $f_2(x) = 2(x - 3.5)$. Which has smaller logistic loss? Do they have different classification error rates? What does this illustrate?