Logistic Regression practice problems

(find and work with a partner)

1. Say I train a binary logistic regression model (i.e. outcomes $\in \{0, 1\}$) and end up with $\hat{\boldsymbol{w}} = [\hat{w}_0, \hat{w}_1]^T = [-4, -5]^T$. What is the decision boundary? Sketch a graph of this logistic model and label the decision boundary. How would you classify a new point $x_{\text{test}} = -2$?

2. The graph below shows the cost for logistic regression as a function of the hypothesis $h_{\boldsymbol{w}}(\boldsymbol{x})$, for one example $\boldsymbol{x}$. Which curve corresponds to the true label $y = 0$ and which corresponds to $y = 1$?

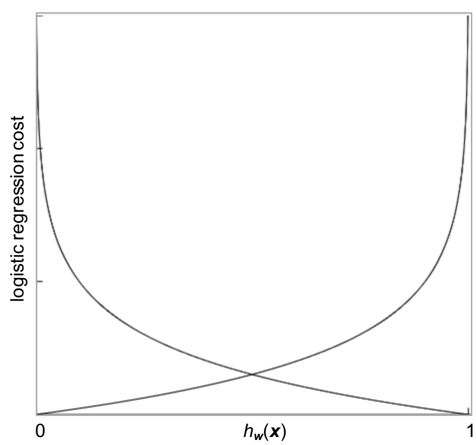

3. A key step in our derivation of the SGD updates for logistic regression was the fact that $g'(z) = g(z)(1 - g(z))$, where $g(z) = \frac{1}{1+e^{-z}}$. This allowed us to cancel out the terms in the denominators. Compute the derivative of $g(z)$ to demonstrate this fact. What does $g'(z)$ tend to as $z \to \infty$? As $z \to -\infty$?

4. Say we have $p = 1$ and two training examples: $(x_1, y_1) = (3, 0)$ and $(x_2, y_2) = (7, 1)$, and we would like to fit a logistic model to this dataset.

   (a) Draw these two examples on a coordinate system and sketch a logistic function that would fit them (roughly). What is the optimal decision boundary? How does this give us a family of solutions for $\hat{w}_0$ and $\hat{w}_1$?

   (b) Say in our SGD method, we choose to analyze (x_2, y_2) first. Before starting SGD, we set $w_0 = 0$ and $w_1 = 0$. After analyzing (x_2, y_2), what are w_0 and w_1? Choose $\alpha = 0.1$. Plot the decision boundary your graph above.

   (c) Next we consider (x_1, y_1). What are w_0 and w_1 be after this second data point? Plot this decision boundary on your graph above. At this point we have finished *one* iteration of SGD.