

1 Logistic Regression

Assume that we have n i.i.d. data points $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$, where each y_i is a binary label in $\{0, 1\}$. We model the posterior probability as a Bernoulli distribution and the probability for each class is the sigmoid function, i.e., $p(y|\mathbf{x}; \mathbf{w}) = q^y(1 - q)^{1-y}$, where $q = s(\mathbf{w}^\top \mathbf{x})$ and $s(\zeta) = \frac{1}{1+e^{-\zeta}}$ is the sigmoid function.

- (a) Show that for a given data point \mathbf{x} , the log ratio of the conditional probabilities, or *log odds*, is linear in \mathbf{x} . More specifically, show that

$$\log \frac{p(y = 1 | \mathbf{x}; \mathbf{w})}{p(y = 0 | \mathbf{x}; \mathbf{w})} = \mathbf{w}^\top \mathbf{x}.$$

- (b) Write out the likelihood and log likelihood functions for the n data points.

- (c) Show that finding maximum likelihood estimate of \mathbf{w} is equivalent to the following optimization problem:

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\operatorname{argmin}} \left[\sum_{i=1}^n (1 - y_i) \mathbf{w}^\top \mathbf{x}_i + \log(1 + \exp\{-\mathbf{w}^\top \mathbf{x}_i\}) \right]$$

- (d) Comment on whether it is possible to find a closed form maximum likelihood estimate of \mathbf{w} , and describe an alternate approach.

2 Gaussian Classification

Let $P(x | \omega_i) \sim \mathcal{N}(\mu_i, \sigma^2)$ for a two-category, one-dimensional classification problem with classes ω_1 and ω_2 , $P(\omega_1) = P(\omega_2) = 1/2$, and $\mu_2 > \mu_1$.

A common procedure to classify a data point x is to assign it to the class ω_i with the highest posterior probability $P(\omega_i | x)$. We will justify this choice later in class.

In this problem, the *decision boundary* is the line that separates the two classes, that is, the line where the posterior probabilities $P(\omega_1 | x)$ and $P(\omega_2 | x)$ are equal.

(a) Find the optimal decision boundary and the corresponding decision rule.

(b) The probability of misclassification (error rate) is:

$$P_e = P(\text{misclassified as } \omega_1 | \omega_2) P(\omega_2) + P(\text{misclassified as } \omega_2 | \omega_1) P(\omega_1).$$

Show that the probability of misclassification (error rate) associated with this decision rule is

$$P_e = \frac{1}{\sqrt{2\pi}} \int_a^\infty e^{-z^2/2} dz, \quad \text{where } a = \frac{\mu_2 - \mu_1}{2\sigma}.$$

(c) What is the limit of P_e as σ goes to 0?

3 Softmax Regression

Logistic regression directly models the probability of a point x belonging to class C as $P(Y = C|X = x) = s(w^\top x)$ where s is the sigmoid function $s(\gamma) = 1/(1 + e^{-\gamma})$. (If you want a bias term α , assume it is included as the last component of w , and each data point x has a 1 appended.) It is limited to binary classification problems. While logistic regression can be applied to multi-class classification with many-to-one or one-to-one approaches, there is a more elegant method.

Rather than modeling only $P(Y = C|X = x)$, softmax regression models the entire categorical distribution over k classes, $P(Y = 1|X = x), P(Y = 2|X = x), \dots, P(Y = k|X = x)$. It does so by learning a linear model with weight vector w_i for each of the k classes and turning them into probabilities with the softmax function, $s_i(z) = e^{-z_i} / (\sum_{j=1}^k e^{-z_j})$, giving the predictions

$$P(Y = i|X = x) = \frac{e^{-w_i^\top x}}{\sum_{j=1}^k e^{-w_j^\top x}}.$$

The idea is that class i 's probability is proportional to $e^{-w_i^\top x}$. To make all the probabilities sum to 1, the denominator is the sum of the numerators.

Note: For future reference, the most general form of the softmax function is $s_i(z) = e^{\beta z_i} / (\sum_{j=1}^k e^{\beta z_j})$, where $\beta \in [-\infty, \infty]$ is a user-defined hyperparameter. If β is negative, smaller inputs are mapped to larger outputs. If β is positive, larger inputs are mapped to larger outputs. Here, we only consider $\beta = -1$.

- Show that where $k = 2$, the softmax regression function has the same form as the logistic regression function.
- In the default form we gave above, softmax regression is overparameterized—there are more parameters than needed for the model. This should be evident in your answer to part (a). Reformulate softmax regression so it requires fewer parameters.
- Recall the logistic (aka binary cross-entropy) loss function (applied to a single training point),

$$L(\hat{y}, y) = -y \log \hat{y} - (1 - y) \log(1 - \hat{y}).$$

How would you design the analogous loss function $L(\hat{y}_1, \hat{y}_2, \dots, \hat{y}_k, y_1, y_2, \dots, y_k)$ for softmax regression? Assume that for each training point, the input includes k class labels in $[0, 1]$ satisfying $\sum_{i=1}^k y_i = 1$, and the k class predictions are $\hat{y}_i = P(Y = i|X = x)$ (which also sum to 1 because we forced them to sum to 1).