

Problem Sheet 2

1 Maximum Likelihood Estimation of σ

Fix the input datapoints $\mathbf{x}_1, \dots, \mathbf{x}_N$, we will not consider these as being generated by a random process. Let \mathbf{w} and σ be the parameters defining the linear model with Gaussian noise, *i.e.*,

$$y_i \sim \mathcal{N}(\mathbf{x}_i^\top \mathbf{w}, \sigma^2). \quad (1)$$

In class we showed that the maximum likelihood estimate for \mathbf{w} is the same as the least square estimator, $\mathbf{w}_{\text{ML}} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}$. Show that the MLE for σ^2 is given by

$$\sigma_{\text{ML}}^2 = \frac{1}{N} (\mathbf{y} - \mathbf{X} \mathbf{w}_{\text{ML}})^\top (\mathbf{y} - \mathbf{X} \mathbf{w}_{\text{ML}}). \quad (2)$$

2 Centering and Ridge Regression

Assume that $\frac{1}{N} \sum_{i=1}^N \mathbf{x}_i = \mathbf{0}$, *i.e.*, the data is centered. (In this question we will treat the constant term separately, as centering this would give us a column of 0s.) Let us denote the parameter for the leading constant term as b (for “bias”). So the linear model is $\hat{y} = b + \mathbf{x}^\top \mathbf{w}$. Consider minimizing the ridge objective:

$$\mathcal{L}_{\text{ridge}}(\mathbf{w}, b) = (\mathbf{X} \mathbf{w} + b \mathbf{1} - \mathbf{y})^\top (\mathbf{X} \mathbf{w} + b \mathbf{1} - \mathbf{y}) + \lambda \mathbf{w}^\top \mathbf{w} \quad (3)$$

Here $\mathbf{1}$ is the vector of all ones and note that b^2 is not regularized. Show that if \hat{b} and $\hat{\mathbf{w}}$ are the resulting solutions obtained by minimising the above objective, then

$$\begin{aligned} \hat{b} &= \frac{1}{N} \sum_{i=1}^N y_i \\ \hat{\mathbf{w}} &= (\mathbf{X}^\top \mathbf{X} + \lambda \mathbf{I}_D)^{-1} \mathbf{X}^\top \mathbf{y} \end{aligned}$$

What happens if you also center \mathbf{y} ?

3 Bias of the Least Squared Estimator

Suppose that the data $\mathcal{D} = \langle (\mathbf{x}_i, y_i) \rangle_{i=1}^N$ is truly generated from a linear model, *i.e.*,

$$\mathbb{E}[y \mid \mathbf{x}, \mathbf{w}^*] = \mathbf{x}^\top \mathbf{w}^* \quad (4)$$

for some fixed (but unknown) parameter vector \mathbf{w}^* . Recall that the least squares estimator is

$$\hat{\mathbf{w}}_{\text{LS}} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}. \quad (5)$$

1. Assume that $\mathbf{x}_1, \dots, \mathbf{x}_N$ are fixed and that $\mathbf{X}^\top \mathbf{X}$ is invertible. You can think of the data \mathcal{D} as a random variable (because of the possible noise in the y_i s). Thus, $\hat{\mathbf{w}}_{\text{LS}}(\mathcal{D})$ is itself a random variable. Show that the expectation of the estimator $\hat{\mathbf{w}}_{\text{LS}}(\mathcal{D})$ (with respect to \mathcal{D}) is \mathbf{w}^* . Such an estimator is called an *unbiased* estimator as its expectation equals the true parameter value.
2. Now suppose we have some other estimator $\hat{\mathbf{w}}$ which may not be unbiased. The bias is defined as

$$\text{Bias}(\hat{\mathbf{w}}) = \|\mathbb{E}_{\mathcal{D}} [\hat{\mathbf{w}}(\mathcal{D})] - \mathbf{w}^*\|. \quad (6)$$

Thus, the bias is the Euclidean distance between the expectation of the estimator and the true parameter. Suppose you are interested in minimizing the squared distance between the estimated parameters and true parameters, *i.e.*, to minimize $\|\hat{\mathbf{w}}(\mathcal{D}) - \mathbf{w}^*\|^2$. Show that the expected (with respect to \mathcal{D}) squared distance can be decomposed as follows:

$$\mathbb{E}_{\mathcal{D}} [\|\hat{\mathbf{w}}(\mathcal{D}) - \mathbf{w}^*\|^2] = \|\mathbb{E}_{\mathcal{D}} [\hat{\mathbf{w}}(\mathcal{D})] - \mathbf{w}^*\|^2 + \mathbb{E}_{\mathcal{D}} [\|\hat{\mathbf{w}}(\mathcal{D}) - \mathbb{E}_{\mathcal{D}} [\hat{\mathbf{w}}(\mathcal{D})]\|^2] \quad (7)$$

The first term above is just the squared bias and the second term above is the variance of the estimator. Thus, while being unbiased looks like a natural property to demand of estimators, it might sometimes be preferable to have a *biased* estimator if it has a much lower variance. This is what ridge regression or LASSO does.

4 Maximum Likelihood and Model Selection

Let the random variable $x \in \{0, 1\}$ model the outcome of an experiment, such that the event $x = 1$ occurs with probability θ_1 . Suppose that someone else observes the experiment and reports to you the outcome, y . But this person is unreliable and only reports the result correctly with probability θ_2 . That is, $p(y \mid x, \theta_2)$ is given by

	$y = 0$	$y = 1$
$x = 0$	θ_2	$1 - \theta_2$
$x = 1$	$1 - \theta_2$	θ_2

Assume that θ_2 is independent of x and θ_1 .

1. Write down the joint probability distribution $p(x, y \mid \boldsymbol{\theta})$ as a 2×2 table, in terms of $\boldsymbol{\theta} = (\theta_1, \theta_2)$.
2. Given the following dataset: $\mathbf{x} = (1, 1, 0, 1, 1, 0, 0)$, $\mathbf{y} = (1, 0, 0, 0, 1, 0, 1)$. What are the numerical values of the MLEs for θ_1 and θ_2 ? What is the numerical value $p(\mathcal{D} \mid \hat{\boldsymbol{\theta}}, M_2)$ where M_2 denotes this 2-parameter model? Justify your answer by including the derivations.
3. Now consider a model with 4 parameters, $\boldsymbol{\theta} = (\theta_{0,0}, \theta_{0,1}, \theta_{1,0}, \theta_{1,1})$, representing $p(x, y \mid \boldsymbol{\theta}) = \theta_{x,y}$. (Only 3 of these parameters are free to vary, since they must sum to one.) What is the MLE of $\boldsymbol{\theta}$? What is $p(\mathcal{D} \mid \hat{\boldsymbol{\theta}}, M_4)$ where M_4 denotes this 4-parameter model?



4. Suppose we are not sure which model is correct. We compute the leave-one-out cross-validated log-likelihood of the 2-parameter model and the 4-parameter model as follows:

$$L(M) = \sum_{i=1}^N \log p(x_i, y_i \mid M, \hat{\theta}(\mathcal{D}_{-i}))$$

and $\hat{\theta}(\mathcal{D}_{-i})$ denotes the MLE computed on \mathcal{D} excluding the i^{th} observation. Which model will CV pick and why?