

**Agenda:**

1. Bias and variance for linear regression
2. Summary: linear regression vs. nearest neighbors algorithm

## 1 Bias and variance for linear regression

Suppose now that the true relationship between  $y$  and  $\mathbf{x}$  is linear and that there is noise

$$y = \mathbf{x}^T \boldsymbol{\theta} + n$$

where  $n \sim \mathcal{N}(0, \sigma^2)$ . In this lecture we will assume that the input variables  $\mathbf{x}$  come from a random distribution, for example, they may be uniformly distributed in a box  $[-1, 1]^p$  in  $p$ -dimensional space.

Let  $\mathbf{x}_0$  be an arbitrary test point. Then the corresponding output is  $y_0 = \mathbf{x}_0^T \boldsymbol{\theta} + n_0$ . Note that

$$\mathbb{E}[y_0] = \mathbb{E}[\mathbf{x}_0^T \boldsymbol{\theta} + n_0] = \mathbf{x}_0^T \boldsymbol{\theta}.$$

Therefore,  $\mathbf{x}_0^T \boldsymbol{\theta}$  is our ideal answer at test time. As before, our estimate at test time is

$$\hat{y}_0 = \mathbf{x}_0^T \hat{\boldsymbol{\theta}}$$

where

$$\hat{\boldsymbol{\theta}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}.$$

Above,  $\mathbf{X}$  is the input data matrix composed of rows  $(\mathbf{x}^{(1)})^T, \dots, (\mathbf{x}^{(n)})^T$ , the data points in the training set;  $\mathbf{y} = [y^{(1)}, \dots, y^{(n)}]^T$  are the output variables in the training set;  $n$  is the size of the training set. Clearly  $\hat{\boldsymbol{\theta}}$  is random, even though we do not use our special font to denote that.

Now compute expected prediction error in this case:

$$\text{EPE}(\mathbf{x}_0) = \mathbb{E}_{y_0 | \mathbf{x}_0 = \mathbf{x}_0} \mathbb{E}_T (y_0 - \hat{y}_0)^2.$$

Above,  $\mathbb{E}_T$  denotes expectation over the randomness in the training data generation, i.e. over noise and over the randomness in the data points  $\mathbf{X}$ . The expectation  $\mathbb{E}_{y_0 | \mathbf{x}_0 = \mathbf{x}_0}$  is with respect to noise at test time.

Let's simplify the expression for the expected prediction error. First add and subtract our ideal answer,  $\mathbb{E}[y_0] = \mathbf{x}_0^\top \boldsymbol{\theta}$ :

$$\begin{aligned} \mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (y_0 - \hat{y}_0)^2 &= \mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (y_0 - \mathbf{x}_0^\top \boldsymbol{\theta} + \mathbf{x}_0^\top \boldsymbol{\theta} - \hat{y}_0)^2 \\ &= \text{Var}(y_0|\mathbf{x}_0 = \mathbf{x}_0) + \mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (\mathbf{x}_0^\top \boldsymbol{\theta} - \hat{y}_0)^2 + \underbrace{2 \mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (y_0 - \mathbf{x}_0^\top \boldsymbol{\theta})(\mathbf{x}_0^\top \boldsymbol{\theta} - \hat{y}_0)}_0. \end{aligned}$$

To see that the last term is zero note that  $y_0 - \mathbf{x}_0^\top \boldsymbol{\theta} = m_0$  is the test-time noise, which is zero mean and independent of  $\mathbf{x}_0^\top \boldsymbol{\theta} - \hat{y}_0$ , which contains only the randomness of the training set.

Next, let's consider the second term:

$$\begin{aligned} \mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (\mathbf{x}_0^\top \boldsymbol{\theta} - \hat{y}_0)^2 &= \mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (\hat{y}_0 - \mathbb{E}_T \hat{y}_0 + \mathbb{E}_T \hat{y}_0 - \mathbf{x}_0^\top \boldsymbol{\theta})^2 \\ &= \mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (\hat{y}_0 - \mathbb{E}_T \hat{y}_0)^2 + \mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (\mathbf{x}_0^\top \boldsymbol{\theta} - \mathbb{E}_T \hat{y}_0)^2 \\ &\quad + \underbrace{2 \mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (\hat{y}_0 - \mathbb{E}_T \hat{y}_0)(\mathbb{E}_T \hat{y}_0 - \mathbf{x}_0^\top \boldsymbol{\theta})}_0. \end{aligned} \quad (1)$$

Above,  $\mathbf{x}_0^\top \boldsymbol{\theta} - \mathbb{E}_T \hat{y}_0$  is nonrandom and  $\mathbb{E}_T (\hat{y}_0 - \mathbb{E}_T \hat{y}_0) = \mathbb{E}_T (\hat{y}_0) - \mathbb{E}_T (\hat{y}_0) = 0$ ; therefore the last term is zero.

Again, since  $\mathbf{x}_0^\top \boldsymbol{\theta} - \mathbb{E}_T \hat{y}_0$  is nonrandom, the second term in (1) can be written as:

$$\mathbb{E}_{y_0|\mathbf{x}_0=\mathbf{x}_0} \mathbb{E}_T (\mathbf{x}_0^\top \boldsymbol{\theta} - \mathbb{E}_T \hat{y}_0)^2 = (\mathbf{x}_0^\top \boldsymbol{\theta} - \mathbb{E}_T \hat{y}_0)^2 = (\text{Bias}(\hat{y}_0))^2.$$

This is the squared bias term. It measures how far  $\mathbb{E}_T(\hat{y}_0)$  is from the ideal estimate,  $\mathbf{x}_0^\top \boldsymbol{\theta}$ . In this case, the squared bias is zero. This can be seen as follows. First observe that:

$$\hat{y}_0 = \mathbf{x}_0^\top \hat{\boldsymbol{\theta}} = \mathbf{x}_0^\top \boldsymbol{\theta} + \mathbf{x}_0^\top (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}).$$

Now,

$$\begin{aligned} \hat{\boldsymbol{\theta}} - \boldsymbol{\theta} &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} - \boldsymbol{\theta} \\ &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{X} \boldsymbol{\theta} + (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{m} - \boldsymbol{\theta} \\ &= \boldsymbol{\theta} + (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{m} - \boldsymbol{\theta} \\ &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{m}, \end{aligned}$$

where  $\mathbf{m}$  denotes the  $n$ -dimensional vector of noise in the training set. Therefore,

$$\mathbf{x}_0^\top (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) = \mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{m}$$

and so

$$\hat{y}_0 = \mathbf{x}_0^\top \boldsymbol{\theta} + \sum_{i=1}^n m_i \left[ \mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \right]_i. \quad (2)$$

Taking the expectation of the last expression and using that  $m_i$  is independent of the choice of training data points in  $\left[ \mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \right]_i$ , we conclude:

$$\mathbb{E}_T \hat{y}_0 = \mathbf{x}_0^\top \boldsymbol{\theta} + \sum_{i=1}^n \mathbb{E}_T [m_i] \mathbb{E}_T \left[ \mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \right]_i = \mathbf{x}_0^\top \boldsymbol{\theta}$$

so that

$$(\text{Bias}(\hat{y}_0))^2 = (\mathbf{x}_0^\top \boldsymbol{\theta} - \mathbb{E}_T \hat{y}_0)^2 = 0 \quad (3)$$

as claimed above.

Putting pieces together:

$$\text{EPE}(\mathbf{x}_0) = \underbrace{\text{Var}(y_0 | \mathbf{x}_0 = \mathbf{x}_0)}_{\text{Var}(\mathbf{x}_0^\top \boldsymbol{\theta} + \eta_0) = \sigma^2} + \underbrace{\mathbb{E}_{y_0 | \mathbf{x}_0 = \mathbf{x}_0} \mathbb{E}_T (\hat{y}_0 - \mathbb{E}_T(\hat{y}_0))^2}_{\text{Var}_T(\hat{y}_0)} + \underbrace{(\text{Bias}(\hat{y}_0))^2}_0.$$

The first term is the additional variance  $\sigma^2$  in the expected prediction error since our target is not deterministic. Let us analyze the second term, which is the variance in our predictions due to the randomness in the training set. Starting from (2) and using  $\mathbb{E}_T(\hat{y}_0) = \mathbf{x}_0^\top \boldsymbol{\theta}$ , we have:

$$\begin{aligned} \text{Var}_T(\hat{y}_0) &= \mathbb{E}_T \left( \sum_{i=1}^n m_i \left[ \mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \right]_i \right)^2 \\ &= \sum_{i,j} \underbrace{\mathbb{E}_T [m_i m_j]}_{0 \text{ if } i \neq j} \mathbb{E}_T \left[ \mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \right]_i \left[ \mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \right]_j \\ &= \sigma^2 \mathbb{E}_T \|\mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top\|_2^2 \\ &= \sigma^2 \mathbb{E}_T \mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{X} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{x}_0 \\ &= \sigma^2 \mathbb{E}_T \mathbf{x}_0^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{x}_0. \end{aligned}$$

We observe that the variance depends on the point  $\mathbf{x}_0$ .

To analyze this expression consider the case where we have a lot of training data, i.e.  $n$  is large and for simplicity assume  $\mathbb{E}[\mathbf{x}] = \mathbf{0}$ .

Under these assumptions,

$$\mathbf{X}^\top \mathbf{X} \rightarrow n \text{Cov}(\mathbf{x}).$$

To see this consider components of this matrix.

For example,

$$\frac{1}{n} [\mathbf{X}^\top \mathbf{X}]_{1,1} = \frac{1}{n} \sum_i \mathbb{X}_{i1} \mathbb{X}_{i1} \rightarrow \mathbb{E}[\mathbf{x}_1^2], \quad n \rightarrow \infty. \quad (4)$$

Above,  $\mathbb{X}_i$  denotes the  $i$ th component of the vector  $\mathbf{x}$  and the last step follows from independence of  $\mathbb{X}_{1i}$  over  $i$ , and the law of large numbers.

Similarly,

$$\frac{1}{n} [\mathbf{X}^\top \mathbf{X}]_{1,2} = \frac{1}{n} \sum_i \mathbb{X}_{i1} \mathbb{X}_{i2} \rightarrow \mathbb{E}[\mathbf{x}_1 \mathbf{x}_2], \quad n \rightarrow \infty. \quad (5)$$

Next, assume that the test point  $\mathbf{x}_0$  is drawn randomly from the same distribution as the points in

the test set. Then we can average the expected prediction error over this choice:

$$\begin{aligned}
 \mathbb{E}_{\mathbf{x}_0} \text{EPE}(\mathbf{x}_0) &\rightarrow \sigma^2 + \mathbb{E}_{\mathbf{x}_0} \mathbf{x}_0^\top \text{Cov}(\mathbf{x})^{-1} \mathbf{x}_0 \frac{\sigma^2}{n} \\
 &= \sigma^2 + \frac{\sigma^2}{n} \mathbb{E}_{\mathbf{x}_0} \text{tr}(\mathbf{x}_0^\top \text{Cov}(\mathbf{x})^{-1} \mathbf{x}_0) \\
 &= \sigma^2 + \frac{\sigma^2}{n} \mathbb{E}_{\mathbf{x}_0} \text{tr}(\text{Cov}(\mathbf{x})^{-1} \mathbf{x}_0 \mathbf{x}_0^\top) \\
 &= \sigma^2 + \frac{\sigma^2}{n} \text{tr}(\text{Cov}(\mathbf{x})^{-1} \mathbb{E}_{\mathbf{x}_0}[\mathbf{x}_0 \mathbf{x}_0^\top]) \\
 &= \sigma^2 + \frac{\sigma^2}{n} \text{tr}(\underbrace{\text{Cov}(\mathbf{x})^{-1} \text{Cov}(\mathbf{x})}_{\mathbf{I}_p}) \\
 &= \sigma^2 \frac{p}{n} + \sigma^2
 \end{aligned}$$

where we used that  $\text{tr}(\mathbf{I}_p) = p$ .

We see that the expected prediction error increases linearly as a function of  $p$  with the slope  $\sigma^2/n$ . Hence the expected prediction error is small if  $n$  is large or  $\sigma^2$  is small. We have avoided the curse of dimensionality by putting heavy restrictions on the model class. We have no bias at all. However, if the model is wrong all bets are off.

## 2 Summary: linear regression vs. nearest neighbors algorithm

We conclude that linear regression has very desirable properties. If the true underlying model is linear, the linear regression has no bias and its variance increases only mildly with the number of variables in the model,  $p$ . However, if the underlying model is nonlinear, we might generate very large estimation errors because of the model mismatch. The nearest neighbors algorithm can flexibly adapt to any model. However, its bias is huge even in moderately small dimensions,  $p = 10$ . The art of machine learning is to design algorithms between these two extremes, well adapted to the problem at hand.