Comparison between OLS and Ridge in terms of ℓ_2 estimation error

Xin Bing

Department of Statistical Sciences
University of Toronto

A theoretical understanding of the role of regularization

Consider the linear regression

$$Y = X^\top \beta + \epsilon.$$

Suppose we have i.i.d. observations $(x_1, y_1), \ldots, (x_n, y_n)$. Further assume the design matrix $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_n)^\top \in \mathbb{R}^{n \times p}$ is deterministic and orthonormal, i.e.

$$\frac{1}{n}\mathbf{X}^\top \mathbf{X} = \mathbf{I}_p.$$

Consider the ridge estimator $\hat{\beta}_\lambda^R$ of β for any given regularization parameter $\lambda \geq 0$. Let $\hat{\beta}$ be the OLS estimator of β.

We now contrast the behaviour of the ridge estimator with that of the OLS estimator side by side.

- Criteria:

$$\hat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2$$

$$\hat{\boldsymbol{\beta}}_\lambda^R = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda\|\beta\|_2^2.$$

- Closed-form solutions:

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top\mathbf{y} = \frac{1}{n}\mathbf{X}^\top\mathbf{y}$$

$$\hat{\boldsymbol{\beta}}_\lambda^R = (\mathbf{X}^\top\mathbf{X} + \lambda\mathbf{I}_p)^{-1}\mathbf{X}^\top\mathbf{y} = \frac{1}{n+\lambda}\mathbf{X}^\top\mathbf{y}.$$

- We examine their statistical properties of estimating $\boldsymbol{\beta}$ in terms of
  - bias
  - variance
  - mean squared error

Bias of the OLS estimator

- OLS: unbiased

$$
\begin{aligned}
\mathbb{E}[\hat{\boldsymbol{\beta}}] &= \mathbb{E}\left[\frac{1}{n}\mathbf{X}^\top \mathbf{y}\right] && \text{by } \frac{1}{n}\mathbf{X}^\top\mathbf{X} = \mathbf{I}_p \\
&= \mathbb{E}\left[\frac{1}{n}\mathbf{X}^\top(\mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon})\right] && \text{by } \mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon} \\
&= \frac{1}{n}\mathbf{X}^\top\mathbf{X}\boldsymbol{\beta} + \mathbb{E}\left[\frac{1}{n}\mathbf{X}^\top\boldsymbol{\epsilon}\right] && \mathbf{X} \text{ and } \boldsymbol{\beta} \text{ are deterministic} \\
&= \frac{1}{n}\mathbf{X}^\top\mathbf{X}\boldsymbol{\beta} + \frac{1}{n}\sum_{i=1}^{n}\mathbf{x}_i\,\mathbb{E}\left[\epsilon_i\right] && \text{by } \mathbb{E}[\epsilon_i] = 0 \\
&= \boldsymbol{\beta}.
\end{aligned}
$$

Bias of the ridge estimator

By repeating similar arguments as before:

- Ridge: biased

$$\begin{aligned}
\mathbb{E}[\hat{\boldsymbol{\beta}}_\lambda^R] &= \mathbb{E}\left[\frac{1}{n+\lambda}\mathbf{X}^\top \mathbf{y}\right] \\
&= \mathbb{E}\left[\frac{1}{n+\lambda}\mathbf{X}^\top(\mathbf{X}\boldsymbol{\beta}+\boldsymbol{\epsilon})\right] \\
&= \frac{1}{n+\lambda}\mathbf{X}^\top\mathbf{X}\boldsymbol{\beta} + \mathbb{E}\left[\frac{1}{n+\lambda}\mathbf{X}^\top\boldsymbol{\epsilon}\right] \\
&= \frac{n}{n+\lambda}\boldsymbol{\beta} \\
&= \boldsymbol{\beta} - \frac{\lambda}{n+\lambda}\boldsymbol{\beta}.
\end{aligned}$$

Variance of the OLS and ridge estimators

- OLS:

$$\begin{aligned}\mathrm{Cov}(\hat{\boldsymbol{\beta}}) &= \frac{1}{n^2}\mathbf{X}^\top \mathrm{Cov}(\mathbf{y})\mathbf{X} \\ &= \frac{\sigma^2}{n^2}\mathbf{X}^\top\mathbf{X} && \text{by } \mathrm{Cov}(\mathbf{y}) = \sigma^2\mathbf{I}_n \\ &= \frac{\sigma^2}{n}\mathbf{I}_p.\end{aligned}$$

- Ridge:

$$\mathrm{Cov}(\hat{\boldsymbol{\beta}}_\lambda^R) = \frac{1}{(n+\lambda)^2}\mathbf{X}^\top \mathrm{Cov}(\mathbf{y})\mathbf{X} = \frac{\sigma^2 n}{(n+\lambda)^2}\mathbf{I}_p.$$

ℓ_2 estimation error of the OLS estimator

- OLS:

$$
\begin{aligned}
&\mathbb{E}[\|\hat{\beta}-\beta\|_2^2] \\
&= \mathbb{E}\left[(\hat{\beta}-\beta)^\top(\hat{\beta}-\beta)\right] \\
&= \mathbb{E}\left[\left(\hat{\beta}-\mathbb{E}[\hat{\beta}]+\mathbb{E}[\hat{\beta}]-\beta\right)^\top\left(\hat{\beta}-\mathbb{E}[\hat{\beta}]+\mathbb{E}[\hat{\beta}]-\beta\right)\right] \\
&= \mathbb{E}\left[\left(\hat{\beta}-\mathbb{E}[\hat{\beta}]\right)^\top\left(\hat{\beta}-\mathbb{E}[\hat{\beta}]\right)\right]+\mathbb{E}\left[\left(\mathbb{E}[\hat{\beta}]-\beta\right)^\top\left(\mathbb{E}[\hat{\beta}]-\beta\right)\right] \\
&= \text{trace}\,\mathbb{E}\left[\left(\hat{\beta}-\mathbb{E}[\hat{\beta}]\right)\left(\hat{\beta}-\mathbb{E}[\hat{\beta}]\right)^\top\right]+\left(\mathbb{E}[\hat{\beta}]-\beta\right)^\top\left(\mathbb{E}[\hat{\beta}]-\beta\right) \\
&= \text{trace}\left[\text{Cov}(\hat{\beta})\right]+\left\|\mathbb{E}[\hat{\beta}]-\beta\right\|_2^2 \\
&= \underbrace{\frac{\sigma^2 p}{n}}_{\textbf{Variance}}+\underbrace{0}_{\textbf{Bias}}.
\end{aligned}
$$

ℓ_2 estimation error of the ridge predictor

- Ridge:

$$\mathbb{E}[\|\hat{\beta}_\lambda^R - \beta\|_2^2] = \text{trace}\left[\text{Cov}(\hat{\beta}_\lambda^R)\right] + \left\|\mathbb{E}[\hat{\beta}_\lambda^R] - \beta\right\|_2^2$$
$$= \underbrace{\frac{\sigma^2 pn}{(n+\lambda)^2}}_{\textbf{Variance}} + \underbrace{\left(\frac{\lambda}{n+\lambda}\right)^2 \|\beta\|_2^2}_{\textbf{Bias}}.$$

- Recall for OLS:

$$\mathbb{E}[\|\hat{\beta} - \beta\|_2^2] = \underbrace{\frac{\sigma^2 p}{n}}_{\textbf{Variance}} + \underbrace{0}_{\textbf{Bias}}.$$

Remark: Ridge estimator has smaller variance by paying extra bias as the price. **This is the essential idea of regularization!** The balance between variance and bias of ridge is controlled by the magnitude of λ.

Same phenomenon for prediction

Since we predict $X = x$ by

- OLS:
$$\hat{y} = x^\top \hat{\beta}$$
- Ridge:
$$\hat{y}_\lambda^R = x^\top \hat{\beta}_\lambda^R$$

Regularization controlled by λ has the same effects on prediction MSE.

Same phenomenon for the Lasso

The same idea holds for the Lasso. But the analysis of the MSE estimation error of the Lasso is less straightforward than that of Ridge.