The Bias-Variance Tradeoff

PUBH 8442: Bayes Decision Theory and Data Analysis

Eric F. Lock
UMN Division of Biostatistics, SPH
elock@umn.edu

02/5/2025

Unbiased rules

- A rule $d(\mathbf{y})$ is unbiased if risk under a given $\theta \in \Theta$ is minimized at θ

$$E_{\mathbf{y} \mid \theta}\, l(\theta', d(\mathbf{y})) \geq E_{\mathbf{y} \mid \theta}\, l(\theta, d(\mathbf{y}))$$

  for all $\theta', \theta \in \Theta$.

- For $l(\theta, d(\mathbf{y})) = (\theta - d(\mathbf{y}))^2$, unbiased implies

$$E_{\mathbf{y} \mid \theta}\, d(\mathbf{y}) = \theta.$$

- The *bias* of d is

$$E_{\mathbf{y} \mid \theta} d(\mathbf{y}) - \theta.$$

Example: silly unbiased estimator

- Let X be the number of times an event occurs in an hour

$$X \sim \text{Poisson}(\lambda).$$

- Using X, estimate $P(A)$, where $A :=$ no events in two hours:

$$P(A \mid \lambda) = e^{-2\lambda}.$$

- Would like an unbiased estimate for $P(A)$ under squared loss:

$$E\left[d(X)\right] = e^{-2\lambda}.$$

Example: silly unbiased estimator

- d is unbiased only if $d(X) = (-1)^X$
    - A ridiculous estimate!

The bias-variance trade-off

- Let $y_1, \ldots, y_n$ be iid with mean μ and variance σ^2
- Consider estimates for μ of the form

$$d_B(\mathbf{y}) = B\mu_0 + (1 - B)\bar{y}$$

where $0 \leq B \leq 1$ is the "shrinkage factor".

- This estimate has variance

$$(1 - B)^2 \frac{\sigma^2}{n}$$

- And it has bias

$$B(\mu_0 - \mu)$$

The bias-variance trade-off

- The frequentist risk under squared error loss for d_B is

$$\begin{aligned} R(\mu, d_B) &= (1-B)^2\frac{\sigma^2}{n} + B^2(\mu_0 - \mu)^2 \\ &= Variance + Bias^2 \end{aligned}$$

- The "bias-variance trade-off"
- For $B = 0$, $d(\mathbf{y}) = \bar{y}$ is unbiased but maximizes variance
- For $B = 1$, $d(\mathbf{y}) = \mu_0$ has no variance but maximizes bias
- Ideal B depends on
  - σ^2 (larger $\sigma^2 \rightarrow B$ closer to 1)
  - n (larger $n \rightarrow B$ closer to 0)
  - Distance of μ from μ_0 (larger distance $\rightarrow B$ closer to 0)

The bias-variance trade-off

- Given prior $p(\mu)$ for μ with mean μ_0 and variance τ^2.
- The rule d_B that minimizes Bayes risk

$$r(p(\mu), d_B) = \int R(\mu, d_B)p(\mu)d\mu$$

has

$$B = \frac{\sigma^2}{\sigma^2 + n\tau^2}.$$

- Normal-normal model is a special case.
- Homework