ECE531 Screencast 1.2: Assessing Estimator Performance

D. Richard Brown III

Worcester Polytechnic Institute

The Cost of Making Bad Estimates

In order to understand whether an estimator is any good, we must assign a "cost" to making bad estimates.

- Intuitively, an optimal estimator will find the "best" guess of the true parameter θ.
- The solution to this problem depends on how we define "best" and how we penalize any deviation from "best".
- **Cost assignment**: $C_\theta(\hat{\theta}) : \Lambda \times \Lambda \mapsto \mathbb{R}$ is the cost of the parameter estimate $\hat{\theta} \in \Lambda$ given the true parameter $\theta \in \Lambda$.

Note that the cost $C_\theta(\hat{\theta})$ depends indirectly on the observation y.

Example

Suppose

$$Y[n] = \theta + W[n]$$

for $n = 0, \dots, N-1$ with

$$W[n] \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2).$$

Your estimator is specified as

$$\hat{\theta}(y) = \frac{1}{N} \sum_{n=0}^{N-1} y[n]$$

and your cost assignment is

$$C_\theta(\hat{\theta}) = \left(\hat{\theta}(y) - \theta\right)^2.$$

If $\theta = 1$ and we observe $y[n] = -1, 0, 1$ for $n = 0, 1, 2$, what is the estimate and the associated cost?

Average Cost

Note that the observations are random. This implies that the estimates and their associated costs are also random. Hence, it is of limited value to consider the cost for one particular observation.

We typically consider the **average cost** when the unknown parameter is θ:

$$\begin{aligned} R_\theta(\hat{\theta}) &:= \mathrm{E}_\theta\left[C_\theta(\hat{\theta}(Y))\right] \\ &= \int_{\mathcal{Y}} C_\theta(\hat{\theta}(y)) p_Y(y\,;\,\theta)\,dy \end{aligned}$$

Note that the expectation here is over the random observations. The parameter θ is not random.

For non-random parameter estimation (our current focus), $R_\theta(\hat{\theta})$ will usually be a function of θ.

Estimation Error and Common Cost Assignments

The **estimation error** is defined as

$$\epsilon := \hat{\theta}(y) - \theta.$$

Many cost assignments can be written as $C_\theta(\hat{\theta}) = C(\epsilon)$.

- Squared error: $C_\theta(\hat{\theta}) = \epsilon^2$.
- Absolute error: $C_\theta(\hat{\theta}) = |\epsilon|$.
- Uniform error ("hit or miss"):

$$C_\theta(\hat{\theta}) = \begin{cases} 0 & |\epsilon| \le \frac{\Delta}{2} \\ 1 & \text{otherwise} \end{cases}$$

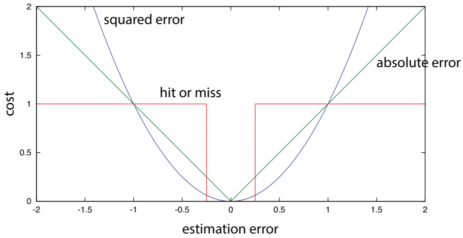

The squared error cost assignment is the most common and will be the focus of our study of "Minimum Variance Unbiased Estimators".