



Linear Least Square Regression of a line



A simple and trivial looking problem, but a good illustration



Assume X to be the independent variable with no errors.

Errors are only with Y .

Reduce deviations, $d = \Delta Y$.

METHOD - I

Cost Function:

$$E = \sum_{i=1}^N d_i^2 = \sum_{i=1}^N (Y_i - \tilde{Y}_i)^2$$

$$= \sum_{i=1}^N (Y_i - mX_i - C)^2$$

Minimize using derivatives:

$$\frac{\partial E}{\partial m} = -2 \sum_{i=1}^N (Y_i - mX_i - C)(-X_i) = 0;$$

$$\frac{\partial E}{\partial C} = - \sum_{i=1}^N 2(Y_i - mX_i - C) = 0;$$

**Re-arranging,
we get Normal
Equations:**

$$\begin{aligned} NC + \left(\sum_{i=1}^N X_i\right)m &= \sum_{i=1}^N Y_i; \\ \left(\sum_{i=1}^N X_i\right)C + \left(\sum_{i=1}^N X_i^2\right)m &= \sum_{i=1}^N (X_i Y_i); \end{aligned}$$

Solving, we get:

$$m = \frac{N \sum_{i=1}^N (X_i Y_i) - \sum_{i=1}^N X_i \sum_{i=1}^N Y_i}{DEN};$$

$$C = \frac{\sum_{i=1}^N Y_i \sum_{i=1}^N X_i^2 - \sum_{i=1}^N X_i \sum_{i=1}^N (X_i Y_i)}{DEN};$$

where,

$$DEN = N \sum_{i=1}^N X_i^2 - \left(\sum_{i=1}^N X_i \right)^2$$

In parametric form:

$$m = \frac{N^2 \sigma_{XY} - N^2 \mu_X \mu_Y}{N^2 \sigma_X^2 - N^2 \mu_X^2} = \frac{\sigma_{XY} - \mu_X \mu_Y}{\sigma_X^2 - \mu_X^2}$$

In parametric form:

$$m = \frac{N^2 \sigma_{XY} - N^2 \mu_X \mu_Y}{N^2 \sigma_X^2 - N^2 \mu_X^2} = \frac{\sigma_{XY} - \mu_X \mu_Y}{\sigma_X^2 - \mu_X^2}$$

$$C = \frac{\sum_{i=1}^N Y_i \sum_{i=1}^N X_i^2 - \sum_{i=1}^N X_i \sum_{i=1}^N (X_i Y_i)}{N \sum_{i=1}^N X_i^2 - (\sum_{i=1}^N X_i)^2};$$

$$= \frac{N^2 \mu_Y \sigma_X^2 - N^2 \mu_X \sigma_{XY}}{N^2 \sigma_X^2 - N^2 \mu_X^2} = \frac{\mu_Y \sigma_X^2 - \mu_X \sigma_{XY}}{\sigma_X^2 - \mu_X^2}$$

Check from above that the LSQ-line passes through the point: (μ_X, μ_Y) . Thus shift the origin to the point: (μ_X, μ_Y) .

The equation of the line in the transformed space:

$$m' = \sigma_{XY} / \sigma_X^2 ; C' = 0.$$

METHOD - II

Solving the same, $Y = mX + C \Rightarrow mX + C = Y;$
using matrix concepts:

$$\begin{bmatrix} X & 1 \end{bmatrix} \begin{bmatrix} m \\ C \end{bmatrix} = Y;$$

Any two points on the line, can give us the parameters:

$$\begin{bmatrix} X_1 & 1 \end{bmatrix} \begin{bmatrix} m \\ C \end{bmatrix} = Y_1; \quad \begin{bmatrix} X_2 & 1 \end{bmatrix} \begin{bmatrix} m \\ C \end{bmatrix} = Y_2;$$

$$\begin{bmatrix} X_1 & 1 \\ X_2 & 1 \end{bmatrix} \begin{bmatrix} m \\ C \end{bmatrix} = \begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix};$$

Thus:

$$\begin{bmatrix} m \\ C \end{bmatrix} = \begin{bmatrix} X_1 & 1 \\ X_2 & 1 \end{bmatrix}^{-1} \begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix} = \left(\frac{1}{X_1 - X_2} \right) \begin{bmatrix} 1 & -1 \\ -X_2 & X_1 \end{bmatrix} \begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix}$$

If you use this:

$$\begin{bmatrix} X_1 & 1 \end{bmatrix} \begin{bmatrix} m \\ C \end{bmatrix} = Y_1; \quad \begin{bmatrix} X_2 & 1 \end{bmatrix} \begin{bmatrix} m \\ C \end{bmatrix} = Y_2;$$

$$\begin{bmatrix} X_1 & 1 \\ X_2 & 1 \end{bmatrix} \begin{bmatrix} m \\ C \end{bmatrix} = \begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix};$$

N C₂ lines may be obtained for each pair.

In case of best fit:

**We are basically
trying to solve
an ill-posed problem,
where:**

$$\begin{bmatrix} X_1 & 1 \\ X_2 & 1 \\ \vdots & \vdots \\ X_N & 1 \end{bmatrix} \begin{bmatrix} m \\ C \end{bmatrix} = \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_N \end{bmatrix};$$

$$\begin{bmatrix} X_1 & 1 \\ X_2 & 1 \\ \vdots & \vdots \\ X_N & 1 \end{bmatrix} \begin{bmatrix} m \\ C \end{bmatrix} = \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_N \end{bmatrix};$$

Take this as: $A\beta = B$,
where A is a non-square
(or even singular square) matrix.

Use Pseudo-inverse in this case:

$$A\beta = B \Rightarrow A^T A\beta = A^T B;$$

$$(A^T A)\beta = A^T B \quad \Rightarrow \quad \beta = (A^T A)^{-1} A^T B;$$

$$\beta = A^+ B;$$

where, $A^+ = (A^T A)^{-1} A^T$; is the Pseudo - inverse.

$(A^T A)$ is square and assumed to be non-singular (generally).
If not, look for alternative formula (hang on, for now)

$A^+ A$ or AA^+ is not equal to I (except non-singular square A), but I_p .

$$A = \begin{bmatrix} X_1 & \mathbf{1} \\ X_2 & \mathbf{1} \\ \vdots & \vdots \\ X_N & \mathbf{1} \end{bmatrix};$$

$$B = \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_N \end{bmatrix};$$

$$\beta = [\mathbf{m} \ \mathbf{C}]^T$$

$$A\beta = B \Rightarrow A^T A\beta = A^T B;$$

$$\beta = (A^T A)^{-1} A^T B;$$

$$A^T A = \begin{bmatrix} X_1 & X_2 & \dots & X_N \\ \mathbf{1} & \mathbf{1} & \dots & \mathbf{1} \end{bmatrix} \begin{bmatrix} X_1 & \mathbf{1} \\ X_2 & \mathbf{1} \\ \vdots & \vdots \\ X_N & \mathbf{1} \end{bmatrix};$$

$$A^T A = \begin{bmatrix} \sum_{i=1}^N X_i^2 & \sum_{i=1}^N X_i \\ \sum_{i=1}^N X_i & N \end{bmatrix} = N \begin{bmatrix} \sigma_x^2 & \mu_x \\ \mu_x & \mathbf{1} \end{bmatrix}$$

$$A^T A = N \begin{bmatrix} \sigma_x^2 & \mu_x \\ \mu_x & 1 \end{bmatrix}; (A^T A)^{-1} = \left(\frac{1}{N \cdot D} \right) \begin{bmatrix} 1 & -\mu_x \\ -\mu_x & \sigma_x^2 \end{bmatrix}$$

$$A^T B = \begin{bmatrix} X_1 & X_2 & \dots & X_N \\ 1 & 1 & \dots & 1 \end{bmatrix} \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_N \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^N X_i Y_i \\ \sum_{i=1}^N Y_i \end{bmatrix} = N \begin{bmatrix} \sigma_{XY} \\ \mu_Y \end{bmatrix};$$

where, $D = \sigma_x^2 - \mu_x^2$

Thus:

$$\beta = A^+ B = (A^T A)^{-1} A^T B = \left(\frac{1}{D} \right) \begin{bmatrix} 1 & -\mu_X \\ -\mu_X & \sigma_x^2 \end{bmatrix} \begin{bmatrix} \sigma_{XY} \\ \mu_Y \end{bmatrix}$$

$$= \left(\frac{1}{D} \right) \begin{bmatrix} \sigma_{XY} - \mu_X \mu_Y \\ \sigma_x^2 \mu_Y - \sigma_{XY} \mu_X \end{bmatrix}$$

The solution is same as in LSQ-FIT.

Pseudo-inverse satisfies the Least-square criteria.

**So we have seen the relation between
LSQ and Pseudo-inverse.**

**Where and how does the
eigen-space (PCA)
help us?**

- Solve that analytical mystery yourselves.

Next:

***Ref: Chap. 2 of Hastie – ESL book,
for “HAT matrix” in regression;***

.

Notations in ESLT – Hastie:

$$\mathbf{X} = \begin{bmatrix} x_{11} & \cdots & x_{1p} \\ \vdots & \ddots & \vdots \\ x_{N1} & \cdots & x_{Np} \end{bmatrix} = [x_1 \cdots x_N]^T = [\mathbf{X}_1 \cdots \mathbf{X}_p]$$

i-th Observation (row-vector):

$$x_i^T = [x_{i1} \cdots x_{ip}]; i = 1, \dots, N$$

All observations of j-th variable (e.g. sensor):

$$\mathbf{X}_j = [x_{1j} \cdots x_{Nj}]^T; j = 1, \dots, p$$

$$\begin{bmatrix} 1 & X_1 \\ 1 & X_2 \\ \vdots & \vdots \\ 1 & X_N \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \end{bmatrix} = \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_N \end{bmatrix}; p = 1.$$

$$\begin{bmatrix} 1 & X_1^1 & X_1^2 & X_1^3 \\ 1 & X_2^1 & X_2^2 & X_2^3 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & X_N^1 & X_N^2 & X_N^3 \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix} = \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_N \end{bmatrix};$$

$p = 3$

$$\begin{bmatrix} 1 & X_1^1 & X_1^2 \\ 1 & X_2^1 & X_2^2 \\ \vdots & \vdots & \vdots \\ 1 & X_N^1 & X_N^2 \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \end{bmatrix} = \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_N \end{bmatrix}; p = 2.$$

$$\begin{bmatrix} 1 & \mathbf{X}_1 & \dots & \mathbf{X}_p \end{bmatrix} \begin{bmatrix} \beta_0 \\ \vdots \\ \beta_p \end{bmatrix} = [\mathbf{Y}]$$

$$(A^T A)X = A^T B$$

$$\Rightarrow \beta = (A^T A)^{-1} A^T B;$$


$$\beta = A^+ B;$$


$$\begin{bmatrix} X_1 & 1 \\ X_2 & 1 \\ & \\ & \\ X_N & 1 \end{bmatrix} \begin{bmatrix} m \\ c \end{bmatrix} = \begin{bmatrix} Y_1 \\ Y_2 \\ \\ \\ Y_N \end{bmatrix};$$

where, $A^+ = (A^T A)^{-1} A^T$; is the Pseudo-inverse.

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

where $\hat{y}_i = \hat{f}(x_i)$. The matrix $\mathbf{H} = \mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$ appearing in equation (3.7) is sometimes called the “hat” matrix because it puts the hat on \mathbf{y} .







***Now see chap. 2 of Hastie – ESL book,
for “HAT matrix” in regression;***

And come back for related concepts.

The vector of observed y does not lie in the plane X . OLS swaps out y for a vector that IS IN THE PLANE, but as close as possible to y . This vector is selected through "projection", not unlike shining a light from above y onto the plane. This calculation is achieved through the matrix H , or fitted-value maker.



e is the "error" between observed vector y and projected vector h that we plan to swap for y . Goal: choose h to make e as small as possible. We want to minimize the error between h and y . Choose h perpendicular to error vector e .



- when fit is good, angle is small, h is long
- when fit is poor, angle large, and h is short
- leverage points will have HIGH FIT:
- leverage points mean that, given an extreme X , fitted y will predict observed y

Column space of X (flat plane) (our mathematical model):
 Columns of X are vectors x_1, x_2 (in 3-space) .. x_1, \dots, x_k (in n -space)
 (The plane represents all possible linear combinations of the vectors x_1 and x_2 .)

$$\beta = (X^T X)^{-1} X^T y;$$

$$\begin{aligned}\text{var}(\beta) &= E(\beta\beta^T) = E[(X^T X)^{-1} X^T y)(X^T X)^{-1} X^T y)^T] \\ &= (X^T X)^{-1} X^T E(yy^T) X (X^T X)^{-1} \\ &= (X^T X)^{-1} \sigma^2 \quad \text{or} \quad \sigma^2 (X^T X)^{-1}\end{aligned}$$

In statistics and machine learning, the **bias–variance tradeoff** (or dilemma) is the problem of simultaneously minimizing two sources of error that prevent supervised learning algorithms from generalizing beyond their training set:

The **bias** is error from erroneous assumptions in the learning algorithm. High bias can cause an algorithm to miss the relevant relations between features and target outputs (underfitting).

The **variance** is error from sensitivity to small fluctuations in the training set. High variance can cause overfitting: modeling the random noise in the training data, rather than the intended outputs.

The **bias–variance decomposition** is a way of analyzing a learning algorithm's expected generalization error with respect to a particular problem as a sum of three terms, the bias, variance, and a quantity called the irreducible error, resulting from noise in the problem itself.

This tradeoff applies to all forms of supervised learning: classification, regression (function fitting), and structured output learning. It has also been invoked to explain the effectiveness of heuristics in human learning.

<Src: WIKI>

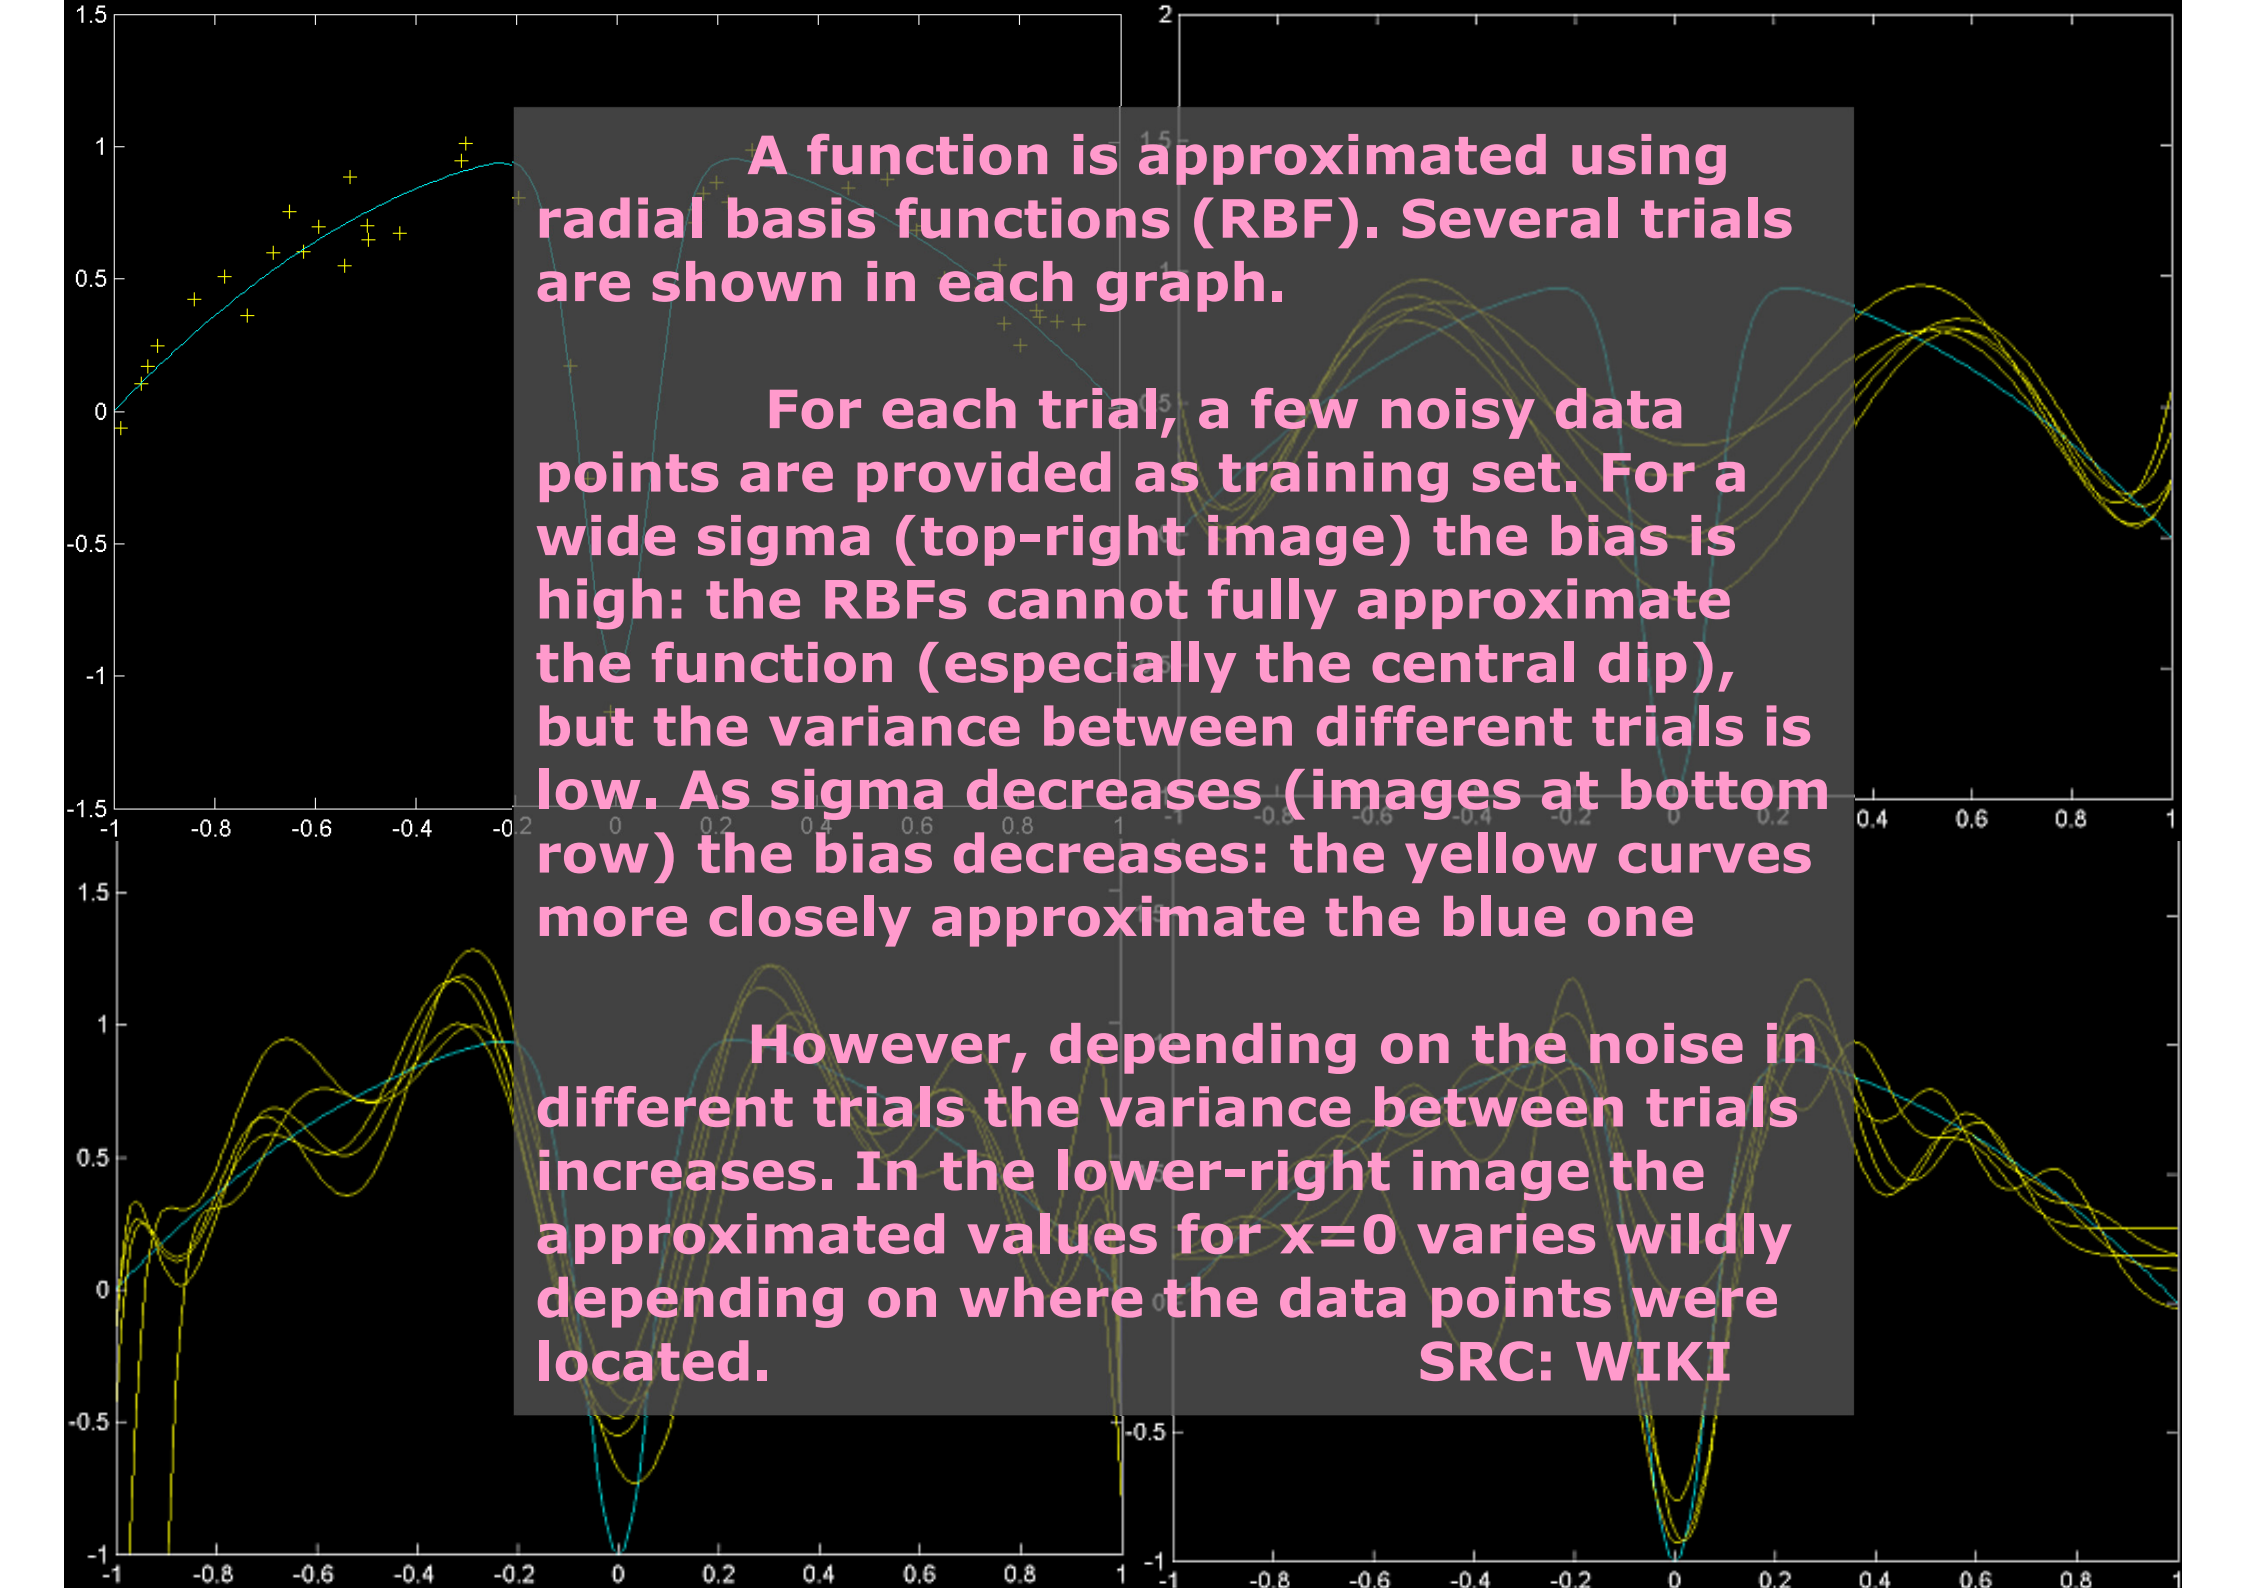
Motivation (WIKI):

The **bias-variance tradeoff** is a central problem in supervised learning. Ideally, one wants to choose a model that both accurately captures the regularities in its training data, but also generalizes well to unseen data. Unfortunately, it is typically impossible to do both simultaneously.

High-variance learning methods may be able to represent their training set well, but are at risk of overfitting to noisy or unrepresentative training data. In contrast, algorithms with **high bias** typically produce simpler models that don't tend to overfit, but may *underfit* their training data, failing to capture important regularities.

Models with **low bias** are usually more complex (e.g. higher-order regression polynomials), enabling them to represent the training set more accurately. In the process, however, they may also represent a large noise component in the training set, making their predictions less accurate - despite their added complexity.

In contrast, models with **higher bias (low variance)** tend to be relatively simple (low-order or even linear regression polynomials), but may produce lower variance predictions when applied beyond the training set.



The figure consists of four subplots arranged in a 2x2 grid. The top-left plot shows a smooth blue curve representing a target function, with several yellow '+' markers representing noisy training data points. The top-right plot shows a single yellow curve that is a smoothed version of the blue curve, representing a high-sigma approximation. The bottom-left plot shows multiple yellow curves that are more oscillatory and closely follow the blue curve, representing a low-sigma approximation with low noise. The bottom-right plot shows multiple yellow curves that are highly oscillatory and vary significantly in their values at x=0, representing a low-sigma approximation with high noise. A central text box contains three paragraphs explaining these observations.

A function is approximated using radial basis functions (RBF). Several trials are shown in each graph.

For each trial, a few noisy data points are provided as training set. For a wide sigma (top-right image) the bias is high: the RBFs cannot fully approximate the function (especially the central dip), but the variance between different trials is low. As sigma decreases (images at bottom row) the bias decreases: the yellow curves more closely approximate the blue one

However, depending on the noise in different trials the variance between trials increases. In the lower-right image the approximated values for $x=0$ varies wildly depending on where the data points were located.

SRC: WIKI




FIGURE 2.11. *Test and training error as a function of model complexity.*

2.3.1 Linear Models and Least Squares

The linear model has been a mainstay of statistics for the past 30 years and remains one of our most important tools. Given a vector of inputs $X^T = (X_1, X_2, \dots, X_p)$, we predict the output Y via the model


$$\hat{Y} = \hat{\beta}_0 + \sum_{j=1}^p X_j \hat{\beta}_j. \quad (2.1)$$

The term $\hat{\beta}_0$ is the intercept, also known as the *bias* in machine learning. Often it is convenient to include the constant variable 1 in X , include $\hat{\beta}_0$ in the vector of coefficients $\hat{\beta}$, and then write the linear model in vector form as an inner product


$$\hat{Y} = X^T \hat{\beta}, \quad (2.2)$$

Lets examine two learning techniques for prediction: the stable but biased linear model and the less stable but apparently less biased class of k-nearest-neighbor estimates.

Linear Regression of 0/1 Response



1-Nearest Neighbor Classifier



15-Nearest Neighbor Classifier




FIGURE 2.3. The same classification example in two dimensions as in Figure 2.1. The classes are coded as a binary variable (BLUE = 0, ORANGE = 1), and then predicted by 1-nearest-neighbor classification.

2. The same classification example in two dimensions as in Figure 2.1. The classes are coded as a binary variable (BLUE = 0, ORANGE = 1) and -nearest-neighbor averaging as in (2.8). The predicted class is hence majority vote amongst the 15-nearest neighbors.

The linear decision boundary from least squares is very smooth, and apparently stable to fit. It does appear to rely heavily on the assumption that a linear decision boundary is appropriate. It **has low variance and potentially high bias**.

On the other hand, the k-nearest-neighbor procedures do not appear to rely on any stringent assumptions about the underlying data, and can adapt to any situation. However, any particular sub-region of the decision boundary depends on a handful of input points and their particular positions, and is **thus wiggly and unstable—high variance and low bias**.

The MSE is generally broken into two components that will become familiar as we proceed: variance and squared bias. Such a decomposition is always possible and often useful, and is known as the **bias–variance decomposition**.

$$\begin{aligned}\text{MSE}(x_0) &= E_{\mathcal{T}}[f(x_0) - \hat{y}_0]^2 \\ &= E_{\mathcal{T}}[\hat{y}_0 - E_{\mathcal{T}}(\hat{y}_0)]^2 + [E_{\mathcal{T}}(\hat{y}_0) - f(x_0)]^2 \\ &= \text{Var}_{\mathcal{T}}(\hat{y}_0) + \text{Bias}^2(\hat{y}_0).\end{aligned}$$

REFER : ELT - HASTIE BOOK (2.25) PP (24, 239, 242)

Bias–variance decomposition of squared error [\[edit \]](#)

Suppose that we have a training set consisting of a set of points x_1, \dots, x_n and real values y_i associated with each point x_i . We assume that there is a functional, but noisy relation $y = f(x) + \epsilon$, where the noise, ϵ , has zero mean and variance σ^2 .

We want to find a function $\hat{f}(x)$, that approximates the true function $f(x)$ as well as possible, by means of some learning algorithm. We make "as well as possible" precise by measuring the mean squared error between y and $\hat{f}(x)$: we want $(y - \hat{f}(x))^2$ to be minimal, both for x_1, \dots, x_n and for points outside of our sample. Of course, we cannot hope to do so perfectly, since the y_i contain noise ϵ ; this means we must be prepared to accept an *irreducible error* in any function we come up with.

Finding an \hat{f} that generalizes to points outside of the training set can be done with any of the countless algorithms used for supervised learning. It turns out that whichever function \hat{f} we select, we can decompose its expected error on an unseen sample x as follows:^{[3]:34}^{[4]:223}

$$\mathbb{E}[(y - \hat{f}(x))^2] = \text{Bias}[\hat{f}(x)]^2 + \text{Var}[\hat{f}(x)] + \sigma^2$$

Where:

$$\text{Bias}[\hat{f}(x)] = \mathbb{E}[\hat{f}(x) - f(x)]$$

and

$$\text{Var}[\hat{f}(x)] = \mathbb{E}[\hat{f}(x)^2] - \mathbb{E}[\hat{f}(x)]^2$$

$$\mathbb{E}[(y - \hat{f}(x))^2] = \text{Bias}[\hat{f}(x)]^2 + \text{Var}[\hat{f}(x)] + \sigma^2$$

Where:

$$\text{Bias}[\hat{f}(x)] = \mathbb{E}[\hat{f}(x) - f(x)]$$

and

$$\text{Var}[\hat{f}(x)] = \mathbb{E}[\hat{f}(x)^2] - \mathbb{E}[\hat{f}(x)]^2$$

The expectation ranges over different choices of the training set $x_1, \dots, x_n, y_1, \dots, y_n$, all sampled from the same (conditional) distribution. The three terms represent:

- the square of the *bias* of the learning method, which can be thought of the error caused by the simplifying assumptions built into the method. E.g., when approximating a non-linear function $f(x)$ using a learning method for *linear models*, there will be error in the estimates $\hat{f}(x)$ due to this assumption;
- the *variance* of the learning method, or, intuitively, how much the learning method $\hat{f}(x)$ will move around its mean;
- the irreducible error σ^2 . Since all three terms are non-negative, this forms a lower bound on the expected error on unseen samples. [3]:34

The more complex the model $\hat{f}(x)$ is, the more data points it will capture, and the lower the bias will be. However, complexity will make the model "move" more to capture the data points, and hence its variance will be larger.

Derivation [\[edit\]](#)

The derivation of the bias–variance decomposition for squared error proceeds as follows.^{[5][6]} For notational convenience, abbreviate $f = f(x)$ and $\hat{f} = \hat{f}(x)$. First, note that for any random variable X , we have

$$\text{Var}[X] = \text{E}[X^2] - \text{E}[X]^2$$

Rearranging, we get:

$$\text{E}[X^2] = \text{Var}[X] + \text{E}[X]^2$$

Since f is [deterministic](#)

$$\text{E}[f] = f.$$

This, given $y = f + \epsilon$ and $\text{E}[\epsilon] = 0$, implies $\text{E}[y] = \text{E}[f + \epsilon] = \text{E}[f] = f$.

Also, since $\text{Var}[\epsilon] = \sigma^2$

$$\text{Var}[y] = \text{E}[(y - \text{E}[y])^2] = \text{E}[(y - f)^2] = \text{E}[(f + \epsilon - f)^2] = \text{E}[\epsilon^2] = \text{Var}[\epsilon] + \text{E}[\epsilon]^2 = \sigma^2$$

Thus, since ϵ and \hat{f} are independent, we can write

$$\begin{aligned}\text{E}[(y - \hat{f})^2] &= \text{E}[y^2 + \hat{f}^2 - 2y\hat{f}] \\ &= \text{E}[y^2] + \text{E}[\hat{f}^2] - \text{E}[2y\hat{f}] \\ &= \text{Var}[y] + \text{E}[y]^2 + \text{Var}[\hat{f}] + \text{E}[\hat{f}]^2 - 2f\text{E}[\hat{f}] \\ &= \text{Var}[y] + \text{Var}[\hat{f}] + (f - \text{E}[\hat{f}])^2 \\ &= \text{Var}[y] + \text{Var}[\hat{f}] + \text{E}[f - \hat{f}]^2 \\ &= \sigma^2 + \text{Var}[\hat{f}] + \text{Bias}[\hat{f}]^2\end{aligned}$$

Approaches [\[edit \]](#)

Dimensionality reduction and feature selection can decrease variance by simplifying models. Similarly, a larger training set tends to decrease variance. Adding features (predictors) tends to decrease bias, at the expense of introducing additional variance. Learning algorithms typically have some tunable parameters that control bias and variance, e.g.:

- (Generalized) linear models can be regularized to decrease their variance at the cost of increasing their bias.^[10]
- In artificial neural networks, the variance increases and the bias decreases with the number of hidden units.^[1] Like in GLMs, regularization is typically applied.
- In k-nearest neighbor models, a high value of k leads to high bias and low variance (see below).
- In Instance-based learning, regularization can be achieved varying the mixture of prototypes and exemplars.^[11]
- In decision trees, the depth of the tree determines the variance. Decision trees are commonly pruned to control variance.^{[3]:307}

One way of resolving the trade-off is to use mixture models and ensemble learning.^{[12][13]} For example, boosting combines many "weak" (high bias) models in an ensemble that has lower bias than the individual models, while bagging combines "strong" learners in a way that reduces their variance.

K-nearest neighbors [\[edit \]](#)

In the case of k -nearest neighbors regression, a closed-form expression exists that relates the bias–variance decomposition to the parameter k .^{[4]:37, 223}

$$\mathbb{E}[(y - \hat{f}(x))^2] = \left(f(x) - \frac{1}{k} \sum_{i=1}^k f(N_i(x)) \right)^2 + \frac{\sigma^2}{k} + \sigma^2$$

where $N_1(x), \dots, N_k(x)$ are the k nearest neighbors of x in the training set. The bias (first term) is a monotone rising function of k , while the variance (second term) drops off as k is increased. In fact, under "reasonable assumptions" the bias of the first-nearest neighbor (1-NN) estimator vanishes entirely as the size of the training set approaches infinity.^[1]

Bias, variance and mean squared error [edit]

Main article: Bias–variance tradeoff

See also: Accuracy (trueness and precision)

**For Parameter set,
say using LSQ**

While bias quantifies the *average* difference to be expected between an estimator and an underlying parameter, an estimator based on a finite sample can additionally be expected to differ from the parameter due to the randomness in the sample.

One measure which is used to try to reflect both types of difference is the mean square error,

$$\text{MSE}(\hat{\theta}) = \mathbf{E} [(\hat{\theta} - \theta)^2].$$

This can be shown to be equal to the square of the bias, plus the variance:

$$\begin{aligned}\text{MSE}(\hat{\theta}) &= (\mathbf{E}[\hat{\theta}] - \theta)^2 + \mathbf{E}[(\hat{\theta} - \mathbf{E}[\hat{\theta}])^2] \\ &= (\text{Bias}(\hat{\theta}, \theta))^2 + \text{Var}(\hat{\theta})\end{aligned}$$

When the parameter is a vector, an analogous decomposition applies:^[13]

$$\text{MSE}(\hat{\theta}) = \text{trace}(\text{Var}(\hat{\theta})) + \|\text{Bias}(\hat{\theta}, \theta)\|^2$$

where

$$\text{trace}(\text{Var}(\hat{\theta}))$$

is the trace of the covariance matrix of the estimator.

An estimator that minimises the bias will not necessarily minimise the mean square error.

Predictor [\[edit \]](#)

If \hat{Y} is a vector of n predictions, and Y is the vector of observed values corresponding to the inputs to the function which generated the predictions, then the MSE of the predictor can be estimated by

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (\hat{Y}_i - Y_i)^2$$

I.e., the MSE is the *mean* $\left(\frac{1}{n} \sum_{i=1}^n \right)$ of the *square of the errors* $((\hat{Y}_i - Y_i)^2)$. This is an easily computable quantity for a particular sample (and hence is sample-dependent).

Estimator [\[edit \]](#)

The MSE of an estimator $\hat{\theta}$ with respect to an unknown parameter θ is defined as

$$\text{MSE}(\hat{\theta}) = \mathbb{E} \left[(\hat{\theta} - \theta)^2 \right].$$

The MSE can be written as the sum of the variance of the estimator and the squared bias of the estimator, providing a useful way to calculate the MSE and implying that in the case of unbiased estimators, the MSE and variance are equivalent.^[2]

$$\text{MSE}(\hat{\theta}) = \text{Var}(\hat{\theta}) + \text{Bias}(\hat{\theta}, \theta)^2.$$

Proof of variance and bias relationship [\[edit \]](#)

$$\begin{aligned} \text{MSE}(\hat{\theta}) &= \mathbb{E} \left[(\hat{\theta} - \theta)^2 \right] \\ &= \mathbb{E} \left[\left(\hat{\theta} - \mathbb{E}[\hat{\theta}] + \mathbb{E}[\hat{\theta}] - \theta \right)^2 \right] \\ &= \mathbb{E} \left[\left(\hat{\theta} - \mathbb{E}[\hat{\theta}] \right)^2 + 2 \left(\hat{\theta} - \mathbb{E}[\hat{\theta}] \right) \left(\mathbb{E}[\hat{\theta}] - \theta \right) + \left(\mathbb{E}[\hat{\theta}] - \theta \right)^2 \right] \\ &= \mathbb{E} \left[\left(\hat{\theta} - \mathbb{E}[\hat{\theta}] \right)^2 \right] + \mathbb{E} \left[2 \left(\hat{\theta} - \mathbb{E}[\hat{\theta}] \right) \left(\mathbb{E}[\hat{\theta}] - \theta \right) \right] + \mathbb{E} \left[\left(\mathbb{E}[\hat{\theta}] - \theta \right)^2 \right] \\ &= \mathbb{E} \left[\left(\hat{\theta} - \mathbb{E}[\hat{\theta}] \right)^2 \right] + 2 \left(\mathbb{E}[\hat{\theta}] - \theta \right) \mathbb{E} \left[\hat{\theta} - \mathbb{E}[\hat{\theta}] \right] + \left(\mathbb{E}[\hat{\theta}] - \theta \right)^2 && \mathbb{E}[\hat{\theta}] - \theta = \text{const.} \\ &= \mathbb{E} \left[\left(\hat{\theta} - \mathbb{E}[\hat{\theta}] \right)^2 \right] + 2 \left(\mathbb{E}[\hat{\theta}] - \theta \right) \left(\mathbb{E}[\hat{\theta}] - \mathbb{E}[\hat{\theta}] \right) + \left(\mathbb{E}[\hat{\theta}] - \theta \right)^2 && \mathbb{E}[\hat{\theta}] = \text{const.} \\ &= \mathbb{E} \left[\left(\hat{\theta} - \mathbb{E}[\hat{\theta}] \right)^2 \right] + \left(\mathbb{E}[\hat{\theta}] - \theta \right)^2 \\ &= \text{Var}(\hat{\theta}) + \text{Bias}(\hat{\theta}, \theta)^2 \end{aligned}$$

$$\begin{aligned}
\text{MSE}(x_0) &= \mathbb{E}_{\mathcal{T}}[f(x_0) - \hat{y}_0]^2 \\
&= \mathbb{E}_{\mathcal{T}}[\hat{y}_0 - \mathbb{E}_{\mathcal{T}}(\hat{y}_0)]^2 + [\mathbb{E}_{\mathcal{T}}(\hat{y}_0) - f(x_0)]^2 \\
&= \text{Var}_{\mathcal{T}}(\hat{y}_0) + \text{Bias}^2(\hat{y}_0).
\end{aligned} \tag{2.25}$$

$$\begin{aligned}
\text{EPE}(x_0) &= \mathbb{E}_{y_0|x_0} \mathbb{E}_{\mathcal{T}}(y_0 - \hat{y}_0)^2 \quad Y = X^T \beta + \varepsilon \quad \varepsilon \sim N(0, \sigma^2) \\
&= \text{Var}(y_0|x_0) + \mathbb{E}_{\mathcal{T}}[\hat{y}_0 - \mathbb{E}_{\mathcal{T}}\hat{y}_0]^2 + [\mathbb{E}_{\mathcal{T}}\hat{y}_0 - x_0^T \beta]^2 \\
&= \text{Var}(y_0|x_0) + \text{Var}_{\mathcal{T}}(\hat{y}_0) + \text{Bias}^2(\hat{y}_0) \\
&= \sigma^2 + \mathbb{E}_{\mathcal{T}} x_0^T (\mathbf{X}^T \mathbf{X})^{-1} x_0 \sigma^2 + 0^2.
\end{aligned} \tag{2.27}$$

$$\begin{aligned}
\mathbb{E}_{x_0} \text{EPE}(x_0) &\sim \mathbb{E}_{x_0} x_0^T \text{Cov}(X)^{-1} x_0 \sigma^2 / N + \sigma^2 \\
&= \text{trace}[\text{Cov}(X)^{-1} \text{Cov}(x_0)] \sigma^2 / N + \sigma^2 \\
&= \sigma^2(p/N) + \sigma^2.
\end{aligned} \tag{2.28}$$

$\mathbb{E}(X) = 0$, then $\mathbf{X}^T \mathbf{X} \rightarrow N \text{Cov}(X)$

Equation 2.25 Derivation; E_τ is expectation over training set

$$\begin{aligned}MSE(x_0) &= E_\tau([f(x_0) - \hat{y}_0]^2) \\&= E_\tau [f(x_0)^2 - 2f(x_0)\hat{y}_0 + \hat{y}_0^2] \\&= E_\tau [\hat{y}_0^2] - 2E_\tau[f(x_0)]E_\tau[\hat{y}_0] + E_\tau [f(x_0)^2]\end{aligned}$$

Add and subtract $(E_\tau [\hat{y}_0])^2$

RHS =

$$E_\tau [\hat{y}_0^2] - (E_\tau [\hat{y}_0])^2 + (E_\tau [\hat{y}_0])^2 - 2E_\tau[f(x_0)]E_\tau[\hat{y}_0] + E_\tau [f(x_0)^2]$$

Since, f is deterministic (given a particular input, will always provide the same output), $E_\tau[f(x_0)] = f(x_0)$ and

$$Var(A) = E[(A - E[A])^2] = E[A^2] - (E[A])^2$$

$$\begin{aligned}\text{LHS} &= E_\tau [\hat{y}_0 - E_\tau[\hat{y}_0]]^2 + (E_\tau[\hat{y}_0] - f(x_0))^2 \\&= Var_\tau (\hat{y}_0) + Bias^2(\hat{y}_0)\end{aligned}$$

$$MSE(x_0) = Var_\tau (\hat{y}_0) + Bias^2(\hat{y}_0) \quad (2.25)$$

Equation 2.27 and 2.28 Derivation

Suppose the relation between Y and X is linear,

$$Y = X^T \beta + \epsilon$$

where $\epsilon \sim N(0, \sigma^2)$ and we fit a model by least squares to the training data.

For an arbitrary point x_0 , we have $\hat{y}_0 = x_0^T \hat{\beta}$, which can be written as $\hat{y}_0 = x_0^T \beta + \sum_{i=1}^N l_i(x_0) \epsilon_i$, where $l_i(x_0)$ is the i -th element of $\mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}_0$. Since under this model the least squares estimates are unbiased, we find that:

$$\begin{aligned} EPE(x_0) &= E_{y_0|x_0} E_{\tau} (y_0 - \hat{y}_0)^2 \\ &= E_{y_0|x_0} E_{\tau} [y_0^2 - 2y_0 \hat{y}_0 + \hat{y}_0^2] \\ &= E_{y_0|x_0} [y_0^2 - 2y_0 E_{\tau}(\hat{y}_0) + E_{\tau}(\hat{y}_0^2)] \\ &= E_{y_0|x_0} [y_0^2] - 2E_{\tau}(\hat{y}_0) E_{y_0|x_0}(y_0) + E_{\tau}(\hat{y}_0^2) \end{aligned}$$

$$\begin{aligned}
EPE(x_0) &= E_{y_0|x_0} E_{\tau} (y_0 - \hat{y}_0)^2 \\
&= E_{y_0|x_0} E_{\tau} [y_0^2 - 2y_0\hat{y}_0 + \hat{y}_0^2] \\
&= E_{y_0|x_0} [y_0^2 - 2y_0 E_{\tau}(\hat{y}_0) + E_{\tau}(\hat{y}_0^2)] \\
&= E_{y_0|x_0} [y_0^2] - 2E_{\tau}(\hat{y}_0) E_{y_0|x_0}(y_0) + E_{\tau}(\hat{y}_0^2)
\end{aligned}$$

Using $E(A^2) = Var(A) + [E(A)]^2$ on above, we have, $EPE(x_0)$

$$= Var(y_0|x_0) + (E_{y_0|x_0}[y_0])^2 - 2E_{\tau}(\hat{y}_0)E_{y_0|x_0}(y_0) + E_{\tau}(\hat{y}_0^2)$$

Add and subtract $(E_{\tau}[\hat{y}_0])^2$; and $E_{y_0|x_0}[y_0] = y_0 = x_0^T \beta$

$$\begin{aligned}
EPE(x_0) &= Var(y_0|x_0) + E_{\tau}(\hat{y}_0^2) - (E_{\tau}[\hat{y}_0])^2 + (E_{\tau}[\hat{y}_0])^2 - 2y_0 E_{\tau}(\hat{y}_0) + y_0^2 \\
&= Var(y_0|x_0) + Var_{\tau}(\hat{y}_0) + (E_{\tau}[\hat{y}_0] - y_0)^2
\end{aligned}$$

$$= Var(y_0|x_0) + Var_{\tau}(\hat{y}_0) + (E_{\tau}[\hat{y}_0] - x_0^T \beta)^2$$

$$= Var(y_0|x_0) + Var_{\tau}(\hat{y}_0) + Bias^2(\hat{y}_0)$$

$$EPE(x_0)$$

$$= Var(y_0|x_0) + E_\tau(\hat{y}_0^2) - (E_\tau[\hat{y}_0])^2 + (E_\tau[\hat{y}_0])^2 - 2y_0E_\tau(\hat{y}_0) + y_0^2$$

$$= Var(y_0|x_0) + Var_\tau(\hat{y}_0) + (E_\tau[\hat{y}_0] - y_0)^2$$

$$= Var(y_0|x_0) + Var_\tau(\hat{y}_0) + (E_\tau[\hat{y}_0] - x_0^T\beta)^2$$

$$= Var(y_0|x_0) + Var_\tau(\hat{y}_0) + Bias^2(\hat{y}_0)$$

Using $Var(y_0|x_0) = \sigma^2$ and substituting the expansion for \hat{y}_0

$$EPE(x_0) = \sigma^2 + Var_\tau\left(x_0^T\beta + \sum_{i=1}^N l_i(x_0)\epsilon_i\right) + \left(E_\tau\left[x_0^T\beta + \sum_{i=1}^N l_i(x_0)\epsilon_i\right] - x_0^T\beta\right)^2$$

$$= \sigma^2 + Var_\tau\left(x_0^T\beta + \sum_{i=1}^N l_i(x_0)\epsilon_i\right) + \left(x_0^T\beta + E_\tau\left[\sum_{i=1}^N l_i(x_0)\epsilon_i\right] - x_0^T\beta\right)$$

$$\text{Since } E[\epsilon] = 0, \quad E_\tau\left[\sum_{i=1}^N l_i(x_0)\epsilon_i\right] = 0$$

$$EPE(x_0) = \sigma^2 + Var_\tau\left(x_0^T\beta + \sum_{i=1}^N l_i(x_0)\epsilon_i\right) + 0$$

Since, $E[\epsilon] = 0$, $E_{\tau} [\sum_{i=1}^N l_i(x_0)\epsilon_i] = 0$

$$EPE(x_0) = \sigma^2 + Var_{\tau} \left(x_0^T \beta + \sum_{i=1}^N l_i(x_0)\epsilon_i \right) + 0$$

$$EPE(x_0) = \sigma^2 + Var_{\tau}(x_0^T \beta) + Var_{\tau} \left(\sum_{i=1}^N l_i(x_0)\epsilon_i \right)$$

Since $Var_{\tau}(x_0^T \beta) = 0$,

$$EPE(x_0) = \sigma^2 + \sum_{i=1}^N l_i^2(x_0) \sigma^2$$

$l_i(x_0)$ is the i -th element of $\mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}_0$

$$= \sigma^2 + \mathbf{x}_0^T (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}_0 \sigma^2$$

Since $\mathbf{X}^T \mathbf{X}$ is symmetric:

$$EPE(x_0) = \sigma^2 + \mathbf{x}_0^T (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}_0 \sigma^2 \quad (2.27)$$

Since $X^T X$ is symmetric

$$EPE(x_0) = \sigma^2 + x_0^T (X^T X)^{-1} x_0 \sigma^2 \quad (2.27)$$

Here, we have incurred an additional variance σ^2 in the prediction error, since our target is not deterministic. There is no bias, and the variance depends on x_0 .

If N is large and T were selected at random, and assuming $E(X) = 0$, then $X^T X \rightarrow N \text{Cov}(X)$ and

$$E_{x_0} EPE(x_0) \sim \frac{E_{x_0} \left(x_0^T (\text{Cov}(X))^{-1} x_0 \right) \sigma^2}{N} + \sigma^2$$

Using property, $E(U^T V U) = E(\text{tr}(V U U^T)) = \text{tr} \left(V \left(E(U U^T) \right) \right)$

$$= \frac{\text{trace}[\text{Cov}(X)^{-1} \text{Cov}(x_0)] \sigma^2}{N} + \sigma^2$$

$$E_{x_0} EPE(x_0) \sim \sigma^2 \left(\frac{p}{N} \right) + \sigma^2 \quad (2.28)$$

$$E_{x_0}EPE(x_0) = \frac{\text{trace}[Cov(X)^{-1}Cov(x_0)]\sigma^2}{N} + \sigma^2$$

$$E_{x_0}EPE(x_0) \sim \sigma^2 \left(\frac{p}{N}\right) + \sigma^2 \quad (2.28)$$

Here we see that the expected EPE increases linearly as a function of p , with slope σ^2/N . If N is large and/or σ^2 is small, this growth in variance is negligible (0 in the deterministic case).

By imposing some heavy restrictions on the class of models being fitted, we have avoided the curse of dimensionality.

model $Y = f(X) + \varepsilon$, with $E(\varepsilon) = 0$ and $\text{Var}(\varepsilon) = \sigma^2$.

$$\begin{aligned}\text{EPE}_k(x_0) &= E[(Y - \hat{f}_k(x_0))^2 | X = x_0] \\ &= \sigma^2 + [\text{Bias}^2(\hat{f}_k(x_0)) + \text{Var}_{\mathcal{T}}(\hat{f}_k(x_0))] \quad (2.46)\end{aligned}$$

$$= \sigma^2 + \left[f(x_0) - \frac{1}{k} \sum_{\ell=1}^k f(x_{(\ell)}) \right]^2 + \frac{\sigma^2}{k}. \quad (2.47)$$

$$\begin{aligned}\text{Err}(x_0) &= E[(Y - \hat{f}(x_0))^2 | X = x_0] \\ &= \sigma_{\varepsilon}^2 + [E\hat{f}(x_0) - f(x_0)]^2 + E[\hat{f}(x_0) - E\hat{f}(x_0)]^2 \\ &= \sigma_{\varepsilon}^2 + \text{Bias}^2(\hat{f}(x_0)) + \text{Var}(\hat{f}(x_0)) \\ &= \text{Irreducible Error} + \text{Bias}^2 + \text{Variance}. \quad (7.9)\end{aligned}$$

Equation 2.47 Derivation

The k-nearest-neighbor regression fit $\hat{f}_k(x_0)$ usually illustrates the competing forces that effect the predictive ability of such approximation. Suppose the data arise from a model $Y = f(X) + \epsilon$, with $E(\epsilon) = 0$ and $Var(\epsilon) = \sigma^2$.

For simplicity here we assume that the values of x_i in the samples are fixed in advance (nonrandom). For an input x_0 ,

$$\hat{y}_0 = \hat{f}_k(x_0) = \frac{1}{k} \sum_{l=1}^k y_{(l)} = \frac{1}{k} \sum_{l=1}^k [f(x_{(l)}) + \epsilon_{(l)}]$$

The subscript in the parentheses- (l) indicate the sequence of nearest neighbors to x_0 . The expected prediction error at x_0 , also known as *test* or *generation* error, can be decomposed:

$$EPE_k(x_0) = E \left[\left(Y - \hat{f}_k(x_0) \right)^2 \mid X = x_0 \right]$$

$$EPE_k(x_0) = E \left[\left(Y - \hat{f}_k(x_0) \right)^2 \middle| X = x_0 \right]$$

Using $EPE(x_0)$

$$\begin{aligned} &= Var(y_0|x_0) + E_\tau(\hat{y}_0^2) - (E_\tau[\hat{y}_0])^2 + (E_\tau[\hat{y}_0])^2 - 2y_0E_\tau(\hat{y}_0) + y_0^2 \\ &= Var(y_0|x_0) + Var_\tau(\hat{y}_0) + (E_\tau[\hat{y}_0] - y_0)^2 \\ &= Var(y_0|x_0) + Var_\tau(\hat{y}_0) + (E_\tau[\hat{y}_0] - x_0^T\beta)^2 \\ &= Var(y_0|x_0) + Var_\tau(\hat{y}_0) + Bias^2(\hat{y}_0) \end{aligned}$$

$$EPE(x_0) = \sigma^2 + \left[Bias^2 \left(\hat{f}_k(x_0) \right) + Var_\tau \left(\hat{f}_k(x_0) \right) \right] \quad (2.46)$$

$$= \sigma^2 + \left(y_0 - E_\tau[\hat{f}_k(x_0)] \right)^2 + Var_\tau[\hat{f}_k(x_0)]$$

$$= \sigma^2 + \left(y_0 - E_\tau \left[\frac{1}{k} \sum_{l=1}^k [f(x_{(l)}) + \epsilon_{(l)}] \right] \right)^2 + Var_\tau \left[\frac{1}{k} \sum_{l=1}^k [f(x_{(l)}) + \epsilon_{(l)}] \right]$$

$$\begin{aligned}
EPE(x_0) &= \sigma^2 + \left[Bias^2 \left(\hat{f}_k(x_0) \right) + Var_{\tau} \left(\hat{f}_k(x_0) \right) \right] \quad (2.46) \\
&= \sigma^2 + \left(y_0 - E_{\tau} [\hat{f}_k(x_0)] \right)^2 + Var_{\tau} [\hat{f}_k(x_0)] \\
&= \sigma^2 + \left(y_0 - E_{\tau} \left[\frac{1}{k} \sum_{l=1}^k [f(x_{(l)}) + \epsilon_{(l)}] \right] \right)^2 + Var_{\tau} \left[\frac{1}{k} \sum_{l=1}^k [f(x_{(l)}) + \epsilon_{(l)}] \right] \\
&= \sigma^2 + \left(f(x_0) - \frac{1}{k} \left(E_{\tau} \left[\sum_{l=1}^k f(x_{(l)}) \right] + E_{\tau} \left[\sum_{l=1}^k \epsilon_{(l)} \right] \right) \right)^2 \\
&\quad + \frac{1}{k^2} \left(Var_{\tau} \left[\sum_{l=1}^k f(x_{(l)}) \right] + Var_{\tau} \left(\sum_{l=1}^k \epsilon_{(l)} \right) \right)
\end{aligned}$$

Since,

$\sum_{l=1}^k f(x_{(l)})$ is a constant:

$$\begin{aligned}
EPE_k(x_0) &= \sigma^2 + \left[f(x_0) - \frac{1}{k} \sum_{l=1}^k f(x_{(l)}) \right]^2 + \frac{k\sigma^2}{k^2} \\
EPE_k(x_0) &= \sigma^2 + \left[f(x_0) - \frac{1}{k} \sum_{l=1}^k f(x_{(l)}) \right]^2 + \frac{\sigma^2}{k} \quad (2.47)
\end{aligned}$$

7.3 The Bias–Variance Decomposition

As in Chapter 2, if we assume that $Y = f(X) + \varepsilon$ where $E(\varepsilon) = 0$ and $\text{Var}(\varepsilon) = \sigma_\varepsilon^2$, we can derive an expression for the expected prediction error of a regression fit $\hat{f}(X)$ at an input point $X = x_0$, using squared-error loss:

$$\begin{aligned}\text{Err}(x_0) &= E[(Y - \hat{f}(x_0))^2 | X = x_0] \\ &= \sigma_\varepsilon^2 + [E\hat{f}(x_0) - f(x_0)]^2 + E[\hat{f}(x_0) - E\hat{f}(x_0)]^2 \\ &= \sigma_\varepsilon^2 + \text{Bias}^2(\hat{f}(x_0)) + \text{Var}(\hat{f}(x_0)) \\ &= \text{Irreducible Error} + \text{Bias}^2 + \text{Variance}.\end{aligned}\tag{7.9}$$

The first term is the variance of the target around its true mean $f(x_0)$, and cannot be avoided no matter how well we estimate $f(x_0)$, unless $\sigma_\varepsilon^2 = 0$. The second term is the squared bias, the amount by which the average of our estimate differs from the true mean; the last term is the variance; the expected squared deviation of $\hat{f}(x_0)$ around its mean. Typically the more complex we make the model \hat{f} , the lower the (squared) bias but the higher the variance.

For the k -nearest-neighbor regression fit, these expressions have the simple form

$$\begin{aligned}\text{Err}(x_0) &= E[(Y - \hat{f}_k(x_0))^2 | X = x_0] \\ &= \sigma_\varepsilon^2 + \left[f(x_0) - \frac{1}{k} \sum_{\ell=1}^k f(x_{(\ell)}) \right]^2 + \frac{\sigma_\varepsilon^2}{k}.\end{aligned}\tag{7.10}$$

Compare:

$$\begin{aligned}\text{EPE}(x_0) &= E_{y_0|x_0} E_{\mathcal{T}} (y_0 - \hat{y}_0)^2 \\ &= \text{Var}(y_0|x_0) + E_{\mathcal{T}} [\hat{y}_0 - E_{\mathcal{T}} \hat{y}_0]^2 + [E_{\mathcal{T}} \hat{y}_0 - x_0^T \beta]^2 \\ &= \text{Var}(y_0|x_0) + \text{Var}_{\mathcal{T}}(\hat{y}_0) + \text{Bias}^2(\hat{y}_0) \\ &= \sigma^2 + E_{\mathcal{T}} x_0^T (X^T X)^{-1} x_0 \sigma^2 + 0^2.\end{aligned}\tag{2.27}$$

$$\begin{aligned}\text{Err}(x_0) &= E[(Y - \hat{f}(x_0))^2 | X = x_0] \\ &= \sigma_{\varepsilon}^2 + [E \hat{f}(x_0) - f(x_0)]^2 + E[\hat{f}(x_0) - E \hat{f}(x_0)]^2 \\ &= \sigma_{\varepsilon}^2 + \text{Bias}^2(\hat{f}(x_0)) + \text{Var}(\hat{f}(x_0)) \\ &= \text{Irreducible Error} + \text{Bias}^2 + \text{Variance}.\end{aligned}\tag{7.9}$$

In statistics, the **bias** of an estimator (or bias function) is the difference between this estimator's expected value and the true value of the parameter being estimated.

An **estimator or decision rule with zero bias is called unbiased**.

$$\text{Bias}(\hat{\theta}, \theta) = \text{Bias}_{\theta}[\hat{\theta}] = E_{x|\theta}[\hat{\theta}] - \theta = E_{x|\theta}[\hat{\theta} - \theta],$$

An estimator is said to be unbiased if its bias is equal to zero for all values of parameter θ , or equivalently, if the expected value of the estimator matches that of the parameter.

In statistical theory, **unbiased estimation** of a standard deviation is the calculation from a statistical sample of an estimated value of the standard deviation (a measure of statistical dispersion) of a population of values, in such a way that the expected value of the calculation equals the true value.

$$\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i \quad S^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2 \quad \text{vs} \quad S^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X})^2$$

The expected value of the **uncorrected sample variance** does not equal the population **variance σ^2** , unless multiplied by a normalization factor. The sample mean, on the other hand, is an unbiased estimator of the population mean μ .

$$\bar{X} - \mu = \frac{1}{n} \sum_{i=1}^n X_i - \mu = \frac{1}{n} \sum_{i=1}^n X_i - \frac{1}{n} \sum_{i=1}^n \mu = \frac{1}{n} \sum_{i=1}^n (X_i - \mu).$$

$$\begin{aligned} \mathbb{E}[S^2] &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2 \right] = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \left((X_i - \mu) - (\bar{X} - \mu) \right)^2 \right] \\ &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \left((X_i - \mu)^2 - 2(\bar{X} - \mu)(X_i - \mu) + (\bar{X} - \mu)^2 \right) \right] \\ &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 - \frac{2}{n} (\bar{X} - \mu) \sum_{i=1}^n (X_i - \mu) + \frac{1}{n} (\bar{X} - \mu)^2 \sum_{i=1}^n 1 \right] \\ &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 - \frac{2}{n} (\bar{X} - \mu) \sum_{i=1}^n (X_i - \mu) + \frac{1}{n} (\bar{X} - \mu)^2 \cdot n \right] \end{aligned}$$

$$\mathbb{E}[S^2] = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 - \frac{2}{n} (\bar{X} - \mu) \sum_{i=1}^n (X_i - \mu) + (\bar{X} - \mu)^2 \right]$$

$$= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 - \frac{2}{n} (\bar{X} - \mu) \cdot n \cdot (\bar{X} - \mu) + (\bar{X} - \mu)^2 \right]$$

$$= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 - 2(\bar{X} - \mu)^2 + (\bar{X} - \mu)^2 \right]$$

$$= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 - (\bar{X} - \mu)^2 \right]$$

$$= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (X_i - \mu)^2 \right] - \mathbb{E} \left[(\bar{X} - \mu)^2 \right]$$

$$= \sigma^2 - \mathbb{E} \left[(\bar{X} - \mu)^2 \right] = \left(1 - \frac{1}{n} \right) \sigma^2 < \sigma^2.$$

Algebraically speaking, $\mathbf{E}[S^2]$ is unbiased because:

$$\begin{aligned}\mathbf{E}[S^2] &= \mathbf{E}\left[\frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X})^2\right] = \frac{n}{n-1} \mathbf{E}\left[\frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2\right] \\ &= \frac{n}{n-1} \left(1 - \frac{1}{n}\right) \sigma^2 = \sigma^2,\end{aligned}$$

where the transition to the second line uses the result derived above for the biased estimator. Thus $\mathbf{E}[S^2] = \sigma^2$, and

therefore $S^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X})^2$ is an unbiased estimator of the population variance, σ^2 .

Bias and Unbias Estimator

If the following holds:

$$E[u(X_1, X_2, \dots, X_n)] = \theta$$

then the statistic $u(X_1, X_2, \dots, X_n)$ is an **unbiased estimator** of the parameter θ .
Otherwise, $u(X_1, X_2, \dots, X_n)$ is a **biased estimator** of θ .

If X_i are normally distributed random variables with mean μ and variance σ^2 , then:

$$\hat{\mu} = \frac{\sum X_i}{n} = \bar{X} \text{ and } \hat{\sigma}^2 = \frac{\sum (X_i - \bar{X})^2}{n}$$

are the maximum likelihood estimators of μ and σ^2 , respectively. Are the MLEs unbiased for their respective parameters?

$E(X_i) = \mu$ and $\text{Var}(X_i) = \sigma^2$. Therefore:

$$E(\bar{X}) = E\left(\frac{1}{n} \sum_{i=1}^n X_i\right) = \frac{1}{n} \sum_{i=1}^n E(X_i) = \frac{1}{n} \sum_{i=1}^n \mu = \frac{1}{n} (n\mu) = \mu$$

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 = \frac{1}{n} \sum_{i=1}^n (x_i^2 - 2x_i\bar{x} + \bar{x}^2)$$

$$Var(X) = \sigma^2 = E(X^2) - \mu^2 \text{ and } Var(\bar{X}) = \frac{\sigma^2}{n} = E(\bar{X}^2) - \mu^2$$

$$= \frac{1}{n} \sum_{i=1}^n x_i^2 - \bar{x}^2$$

$$\begin{aligned} E(\hat{\sigma}^2) &= E \left[\frac{1}{n} \sum_{i=1}^n X_i^2 - \bar{X}^2 \right] = \left[\frac{1}{n} \sum_{i=1}^n E(X_i^2) \right] - E(\bar{X}^2) \\ &= \frac{1}{n} \sum_{i=1}^n (\sigma^2 + \mu^2) - \left(\frac{\sigma^2}{n} + \mu^2 \right) \\ &= \frac{1}{n} (n\sigma^2 + n\mu^2) - \frac{\sigma^2}{n} - \mu^2 \\ &= \sigma^2 - \frac{\sigma^2}{n} = \frac{n\sigma^2 - \sigma^2}{n} = \frac{(n-1)\sigma^2}{n} \end{aligned}$$

$$\begin{aligned}
 E(\hat{\sigma}^2) &= E \left[\frac{1}{n} \sum_{i=1}^n X_i^2 - \bar{X}^2 \right] = \left[\frac{1}{n} \sum_{i=1}^n E(X_i^2) \right] - E(\bar{X}^2) \\
 &= \frac{1}{n} \sum_{i=1}^n (\sigma^2 + \mu^2) - \left(\frac{\sigma^2}{n} + \mu^2 \right) \\
 &= \frac{1}{n} (n\sigma^2 + n\mu^2) - \frac{\sigma^2}{n} - \mu^2 \\
 &= \sigma^2 - \frac{\sigma^2}{n} = \frac{n\sigma^2 - \sigma^2}{n} = \frac{(n-1)\sigma^2}{n}
 \end{aligned}$$

$$E(\hat{\sigma}^2) \neq \sigma^2$$

the maximum likelihood estimator of σ^2 is a biased estimator.

OLS is linear and unbiased

First of all, note that $\hat{\beta}$ is linear in y . In fact, $\hat{\beta}$ is the product between the $K \times N$ matrix $(X^T X)^{-1} X^T$ and y , and matrix multiplication is a linear operation.

It can easily be proved that $\hat{\beta}$ is unbiased, both conditional on X , and unconditionally, that is,

$$E[\hat{\beta}|X] = \beta$$

We can use the definition of y to re-write the OLS estimator as follows:

$$\begin{aligned}\hat{\beta} &= (X^T X)^{-1} X^T y \\ &= (X^T X)^{-1} X^T (X\beta + \varepsilon) \\ &= (X^T X)^{-1} X^T X\beta + (X^T X)^{-1} X^T \varepsilon \\ &= \beta + (X^T X)^{-1} X^T \varepsilon\end{aligned}$$

When we condition on X , we can treat X as a constant matrix. Therefore, the conditional expectation of $\hat{\beta}$ is

$$\begin{aligned}E[\hat{\beta}|X] &= \beta + (X^T X)^{-1} X^T E[\varepsilon|X] \\ &= \beta + (X^T X)^{-1} X^T \cdot 0 \\ &= \beta\end{aligned}$$

The Law of Iterated Expectations implies that

$$E[\hat{\beta}] = E[E[\hat{\beta}|X]] = E[\beta] = \beta$$

BYE