Bias-Variance Trade-off

Introduction

The bias-variance trade-off is a fundamental concept in machine learning and statistical modeling that describes the balance between two sources of error in predictive models. Understanding this trade-off is essential for building models that generalize well to new data.

1 Definitions

Consider a dataset with inputs X and outputs Y, where we assume Y is generated by an unknown function $f(X)$ with added noise ϵ, i.e.,

$$Y = f(X) + \epsilon,$$

where ϵ is a random error term with zero mean and variance σ^2. Given a model $\hat{f}(X)$ that approximates $f(X)$, the goal is to minimize the error when predicting new data.

1.1 Mean Squared Error (MSE)

The Mean Squared Error (MSE) between the true values Y and predicted values $\hat{f}(X)$ is given by

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

Expanding this equation yields

$$\text{MSE} = \underbrace{[f(X) - \mathbb{E}(\hat{f}(X))]^2}_{\text{Bias}^2} + \underbrace{\mathbb{E}[(\hat{f}(X) - \mathbb{E}(\hat{f}(X)))^2]}_{\text{Variance}} + \sigma^2,$$

where Bias^2 and Variance are two components of the error.

2 Bias and Variance

2.1 Bias

Bias measures how far the average prediction $\mathbb{E}[\hat{f}(X)]$ is from the true function $f(X)$. Formally,

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

A high-bias model makes strong assumptions about the data and is often too simple to capture the underlying structure, leading to underfitting.

2.2 Variance

Variance measures the variability of the model predictions $\hat{f}(X)$ for different training sets. Formally,

$$\text{Variance} = \mathbb{E}[(\hat{f}(X) - \mathbb{E}[\hat{f}(X)])^2].$$

A high-variance model is sensitive to fluctuations in the training data, which can lead to overfitting.

3 The Trade-off

The bias-variance trade-off illustrates the inverse relationship between bias and variance. Increasing model complexity typically decreases bias but increases variance, while reducing model complexity increases bias but decreases variance.

3.1 Implications

- **Underfitting:** High bias and low variance.
- **Overfitting:** Low bias and high variance.
- **Optimal Model:** Achieves a balance where the sum of bias and variance is minimized, leading to good generalization on new data.

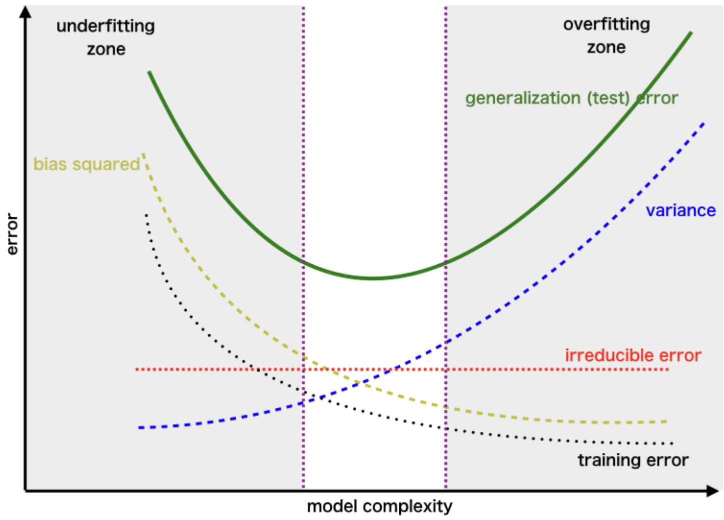

Figure 1: Illustration of the Bias-Variance Trade-off

4 Conclusion

The bias-variance trade-off is a critical concept for selecting and tuning models. In practice, cross-validation is often used to evaluate models and identify the point where the trade-off achieves the lowest MSE on unseen data. This balance is crucial to prevent underfitting and overfitting, thus enabling better generalization.