

9/19/13

1

BIAS-VARIANCE

① How do we pick d in ^(1D) polynomial regression?

One Idea:

$$\hat{y}^{(0)} = w_0 \quad x \in \mathbb{R}^1$$


$$\hat{y}^{(1)} = w_0 + w_1 x'$$

$$\hat{y}^{(d)} = w_0 + w_1 x' + \dots + w_d x^d$$

$$\hat{w}^{(d)} = \underset{w \in \mathbb{R}^d}{\operatorname{argmin}} \sum_{\text{training data}} (y_i - \hat{y}_i^{(d)})^2$$

$$\hat{d} = \underset{d \in \{0, 1, \dots, 100\}}{\operatorname{argmin}} \sum_{\text{training data}} (y_i - \hat{y}_i^{(d)})^2$$

Won't work too well. Model classes are nested,
 $\hat{d} = 100$ will always give lowest training error.



This is called the problem of "Model Selection"
 → How do I pick a model class to search in?

② Types of Error

→ larger Model classes will always do better on training error, but that's not what we care about.

what we really care about → Expected Loss / Error
 $X, Y \sim P(X, Y)$ [unknown]

$$E_{P(X, Y)} [L(y, g(x; w))]$$

↑
all parameters to hyperparameters

what we ideally want to do

$$\min_w \iint_{x, y} L(y, g(x; w)) p(x, y) dx dy$$

Two problems


→ Integral hard to compute

→ $p(x, y)$ unknown

So let's approximate integral with samples (x_i, y_i)
 $\sim P(X, Y)$

$$E_{\text{approx}}(w) = \frac{1}{N} \sum_{i=1}^N L(y_i, g(x_i; w))$$

So here's how we optimize various quantities:



E_{train}
Used to fit model parameters
 $\hat{w} = \underset{w}{\operatorname{argmin}} E_{train}$

E_{val}

E_{test}


Used to choose model classes
 $\hat{d} = \underset{d \in \{0, \dots, 100\}}{\operatorname{argmin}} E_{val}$

Used to estimate Expected Error/Loss

No tweaking, No learning on this, otherwise becomes a biased estimate.

If not enough data to do: (train, val) split we do cross-validation.

③ Overfitting vs Under Fitting



$E_{val} \approx E_{train} = \text{high} \Rightarrow \text{Underfitting}$
 $E_{val} \gg E_{train} \Rightarrow \text{Overfitting}$


Overfitting \Rightarrow Model Class too large, too expressive
 Need more assumptions, Need smaller model class

Underfitting \Rightarrow Model Class too small, too simple.
 Need fewer assumptions



④ Bias - Variance

Eval is high for both overfitting & underfitting
 Can we differentiate?



Back to basics: Coin Toss.

$D = \{x_1, \dots, x_n\} \sim \text{Ber}(\theta^*)$ (Say $\theta^* = 0.5$)
 IID

D	$\hat{\theta}_{MLE} = \frac{1}{N} \sum x_i$	$\hat{\theta}_{lazy} = x_1$	$\hat{\theta}_{silly} = 1$
{H, H, T}	2/3	1	1
{T, T, H}	1/3	0	1
{H, T, H}	2/3	1	1
⋮	⋮	⋮	⋮

$$\begin{aligned}
 E[\hat{\theta}_{MLE}] &= \frac{1}{N} \sum_{i=1}^N E[x_i] = \frac{1}{N} \sum_{i=1}^N [1 \cdot \theta^* + 0 \cdot (1 - \theta^*)] \\
 &= \frac{1}{N} \sum_{i=1}^N \theta^* = \theta^*
 \end{aligned}$$

bias = $E[\hat{\theta}] - \theta^*$

Difference between what your estimator reports on average & the Truth.

$\hat{\theta}_{MLE}$ is UNBIASED $\because E[\hat{\theta}_{MLE}] = \theta^*$

What about others?


$E[\theta_{lazy}] = E[X_1] = \theta^*$ (Also Unbiased)

$E[\theta_{silly}] = E[1] = 1$ (Biased if $\theta^* \neq 1$)

high variance

No variance

Variance = $E[(\hat{\theta} - E[\hat{\theta}])^2]$



⑤ Bias-Variance Decomposition for Squared-Loss

Say we want to estimate θ^*

← your estimator's arg value

$$\text{Let } \bar{\theta} \equiv E[\hat{\theta}]$$

$$\begin{aligned} \text{Now } E[\text{Loss}] &= E[(\hat{\theta} - \theta^*)^2] \\ &= E[(\hat{\theta} - \bar{\theta} + \bar{\theta} - \theta^*)^2] \end{aligned}$$

$$= E[(\hat{\theta} - \bar{\theta})^2 + (\bar{\theta} - \theta^*)^2 - 2(\hat{\theta} - \bar{\theta})(\bar{\theta} - \theta^*)]$$

$$= E[(\hat{\theta} - \bar{\theta})^2] + E[(\bar{\theta} - \theta^*)^2] - 2(\bar{\theta} - \theta^*) \underbrace{E[(\hat{\theta} - \bar{\theta})]}_{\text{constants}}$$

$$= \text{Var}(\hat{\theta}) + E[\text{bias}^2] - 2(\bar{\theta} - \theta^*) \underbrace{(E[\hat{\theta}] - \bar{\theta})}_{=0}$$

$$= \text{Var}(\hat{\theta}) + \text{bias}^2$$

$$\text{So } E[\text{Loss}] = \text{bias}^2 + \text{Var}(\hat{\theta})$$

2 predictors (one with high bias, one with high variance) can have the same L_2 loss



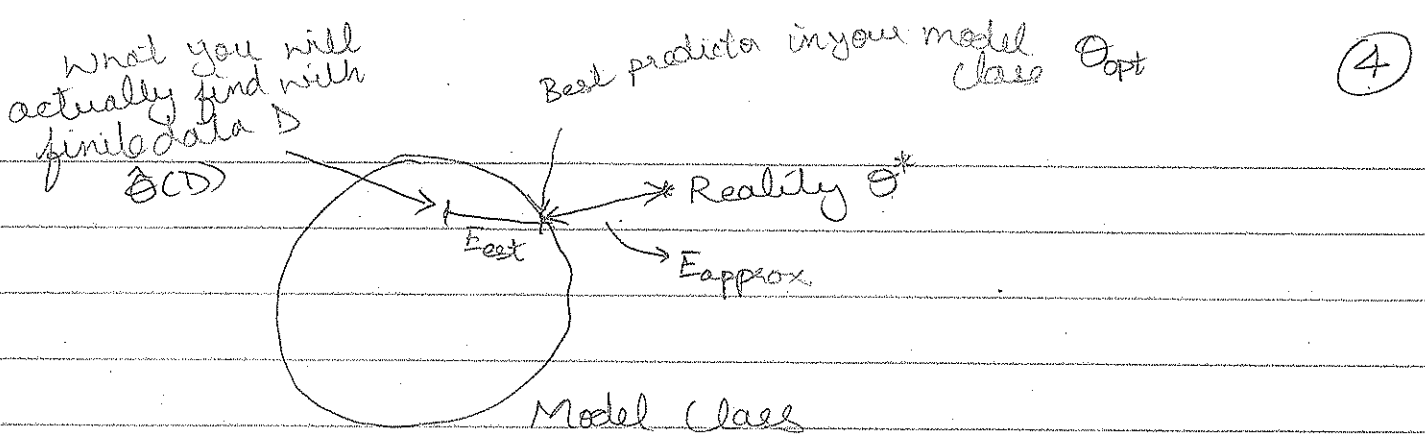
⑥ General Error Decomposition

Expected Error = "Approximation Error" (or \approx Bias)
→ you approximated reality with your model

+ "Estimation Error" (or \approx Variance)
→ you estimated with finite data

+ "Optimization Error"
→ you didn't / couldn't maximize MLE exactly

+ Bayes Error
→ lower bound on error





⑦ Model Selection via Regularization
 (let's put in a preference for simpler models)

$$\min_w \sum_{train} \mathcal{L}(y_i, \hat{y}(w; x_i)) + \lambda \|w\|_2^2$$

Can also be viewed as a Gaussian prior on w
 (MAP estimator)

⑧ Effect of ^{More} Data / Learning Curves



Simple Model Class

→ More data won't help

Too

Expressive Model Class

→ More data will help, but very slowly.

