

Overfitting

Machine Learning – CSE446

Carlos Guestrin

University of Washington

April 8, 2013

©2005-2013 Carlos Guestrin


1

Bias-Variance Tradeoff

- Choice of hypothesis class introduces learning bias

- ☐ More complex class → less bias


- ☐ More complex class → more variance



Select points by clicking on the graph or press [Example](#)

Degree of polynomial: 1 ☐ Fit Y to X
☐ Fit X to Y


[Calculate](#) [View Polynomial](#) [Reset](#)



Select points by clicking on the graph or press [Example](#)

Degree of polynomial: 13 ☐ Fit Y to X
☐ Fit X to Y

[Calculate](#) [View Polynomial](#) [Reset](#)



Select points by clicking on the graph or press [Example](#)

Degree of polynomial: 13 ☐ Fit Y to X
☐ Fit X to Y

[Calculate](#) [View Polynomial](#) [Reset](#)

©2005-2013 Carlos Guestrin

2

Training set error

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \sum_j \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

- Given a dataset (Training data)
- Choose a loss function
 - e.g., squared error (L_2) for regression
- **Training set error:** For a particular set of parameters, loss function on training data:


$$\underset{\text{params}}{\overset{\text{min}}{\text{error}_{\text{train}}(\mathbf{w})}} = \frac{1}{N_{\text{train}}} \sum_{j=1}^{N_{\text{train}}} \left(\underset{\text{train data}}{t(\mathbf{x}_j)} - \sum_i \underset{\text{truth}}{w_i} \underset{\text{prediction}}{h_i(\mathbf{x}_j)} \right)^2$$

Handwritten notes: "Squared" with an arrow pointing to the squared term in the equation.

©2005-2013 Carlos Guestrin

3

Training set error as a function of model complexity



©2005-2013 Carlos Guestrin

4

Prediction error

- Training set error can be poor measure of “quality” of solution
- Prediction error:** We really care about error over all possible input points, not just training data:


$$\begin{aligned}
 error_{true}(\mathbf{w}) &= E_{\mathbf{x}} \left[\left(t(\mathbf{x}) - \sum_i w_i h_i(\mathbf{x}) \right)^2 \right] \\
 &= \int_{\mathbf{x}} \left(t(\mathbf{x}) - \sum_i w_i h_i(\mathbf{x}) \right)^2 p(\mathbf{x}) d\mathbf{x}
 \end{aligned}$$

Handwritten notes:
 - "all possible examples brain sends" points to $E_{\mathbf{x}}$
 - "expected error over all possible points" points to the integral
 - "squared error" points to the square in the integrand
 - "approximate over all possible points" points to the integral

©2005-2013 Carlos Guestrin

5


Prediction error as a function of model complexity: Bias/Variance tradeoff



©2005-2013 Carlos Guestrin


6

Prediction error as a function of model complexity: train v. true error



$$error_{train}(\mathbf{w}) = \frac{1}{N_{train}} \sum_{j=1}^{N_{train}} \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

$$error_{true}(\mathbf{w}) = \int_{\mathbf{x}} \left(t(\mathbf{x}) - \sum_i w_i h_i(\mathbf{x}) \right)^2 p(\mathbf{x}) d\mathbf{x}$$



©2005-2013 Carlos Guestrin

7

Computing prediction error

Computing prediction

- ☐ Hard integral
- ☐ May not know $t(\mathbf{x})$ for every \mathbf{x}

$$error_{true}(\mathbf{w}) = \int_{\mathbf{x}} \left(t(\mathbf{x}) - \sum_i w_i h_i(\mathbf{x}) \right)^2 p(\mathbf{x}) d\mathbf{x}$$

Monte Carlo integration (sampling approximation)

- ☐ Sample a set of i.i.d. points $\{\mathbf{x}_1, \dots, \mathbf{x}_M\}$ from $p(\mathbf{x})$
- ☐ Approximate integral with sample average

$$error_{true}(\mathbf{w}) \approx \frac{1}{M} \sum_{j=1}^M \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

Handwritten red annotations:

- approximated by "samples"* (pointing to the sum)
- train data or test data* (pointing to the $t(\mathbf{x}_j)$ term)

©2005-2013 Carlos Guestrin

8

Why training set error doesn't approximate prediction error?

- Sampling approximation of prediction error:

$$error_{true}(\mathbf{w}) \approx \frac{1}{M} \sum_{j=1}^M \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

- Training error :

$$error_{train}(\mathbf{w}) = \frac{1}{N_{train}} \sum_{j=1}^{N_{train}} \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

- Very similar equations!!!

- ☐ Why is training set a bad measure of prediction error???

©2005-2013 Carlos Guestrin

9

Why training set error doesn't approximate prediction error?

Because you cheated!!!

Training error good estimate for a single \mathbf{w} ,
But you optimized \mathbf{w} with respect to the training error,
and found \mathbf{w} ^{only} that is good for this set of samples

**Training error is a (optimistically) biased
estimate of prediction error**

- Very similar equations!!!

- ☐ Why is training set a bad measure of prediction error???

©2005-2013 Carlos Guestrin

10

Test set error

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \sum_j \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

- Given a dataset, randomly split it into two parts:
 - Training data – $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_{\text{train}}}\}$ ←
 - Test data – $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_{\text{test}}}\}$ ←
- Use training data to optimize parameters \mathbf{w}
- **Test set error:** For the final output $\hat{\mathbf{w}}$, evaluate the error using:

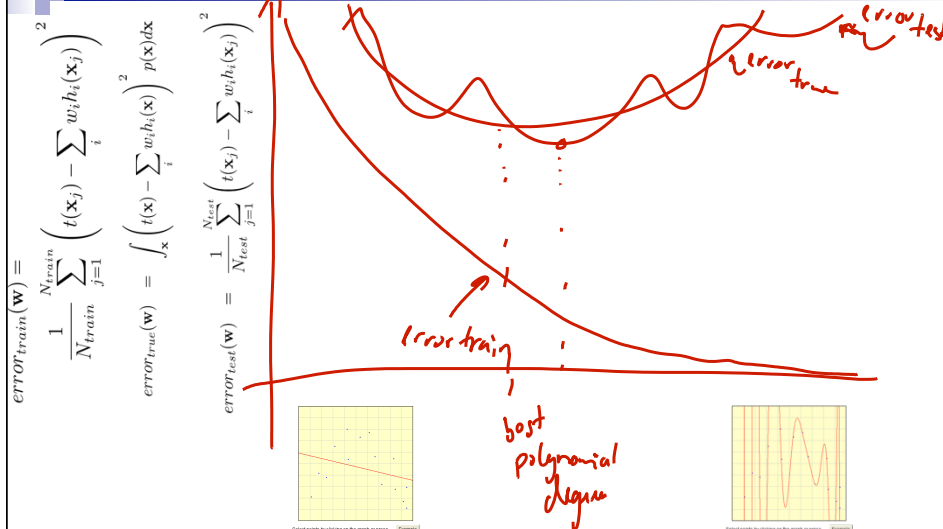
$$error_{test}(\mathbf{w}) = \frac{1}{N_{test}} \sum_{j=1}^{N_{test}} \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

©2005-2013 Carlos Guestrin

11

Test set error as a function of model complexity

has extreme
in expectation
variability
error_{true} = error_{test}




©2005-2013 Carlos Guestrin

12

Overfitting

- **Overfitting:** a learning algorithm overfits the training data if it outputs a solution $\hat{\mathbf{w}}$ when there exists another solution \mathbf{w}^* such that:

$$[error_{train}(\hat{\mathbf{w}}) < error_{train}(\mathbf{w}^*)] \wedge [error_{true}(\mathbf{w}^*) < error_{true}(\hat{\mathbf{w}})]$$



©2005-2013 Carlos Guestrin

13

How many points to I use for training/testing?

- Very hard question to answer!
 - Too few training points, learned \mathbf{w} is bad
 - Too few test points, you never know if you reached a good solution
- Bounds, such as Hoeffding's inequality can help:

$$P(|\hat{\theta} - \theta^*| \geq \epsilon) \leq 2e^{-2N\epsilon^2}$$

- More on this later this quarter, but still hard to answer
- Typically:
 - If you have a reasonable amount of data, pick test set "large enough" for a "reasonable" estimate of error, and use the rest for learning
 - If you have little data, then you need to pull out the big guns...
 - e.g., bootstrapping

(cross validation)

©2005-2013 Carlos Guestrin

14

Error estimators

$$error_{true}(\mathbf{w}) = \int_{\mathbf{x}} \left(t(\mathbf{x}) - \sum_i w_i h_i(\mathbf{x}) \right)^2 p(\mathbf{x}) d\mathbf{x}$$

← gold standard


$$error_{train}(\mathbf{w}) = \frac{1}{N_{train}} \sum_{j=1}^{N_{train}} \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

← data used to learn, optimistically biased

$$error_{test}(\mathbf{w}) = \frac{1}{N_{test}} \sum_{j=1}^{N_{test}} \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$


← test our final result

Error as a function of number of training examples for a fixed model complexity



$$error_{train}(\mathbf{w}) = \frac{1}{N_{train}} \sum_{j=1}^{N_{train}} \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

$$error_{true}(\mathbf{w}) = \int_{\mathbf{x}} \left(t(\mathbf{x}) - \sum_i w_i h_i(\mathbf{x}) \right)^2 p(\mathbf{x}) d\mathbf{x}$$



Error estimators

Be careful!!!

Test set only unbiased if you never never ever ever
do any any any any learning on the test data

For example, if you use the test set to select
the degree of the polynomial... no longer unbiased!!!
(We will address this problem later in the quarter)

$$error_{test}(\mathbf{w}) = \frac{1}{N_{test}} \sum_{j=1}^{N_{test}} \left(t(\mathbf{x}_j) - \sum_i w_i h_i(\mathbf{x}_j) \right)^2$$

What you need to know

- True error, training error, test error

- ☐ Never learn on the test data
- ☐ Never learn on the test data
- ☐ Never learn on the test data
- ☐ Never learn on the test data
- ☐ Never learn on the test data

- Overfitting

Regularization

Machine Learning – CSE446

Carlos Guestrin

University of Washington

April 8, 2013

©2005-2013 Carlos Guestrin

19


Regularization in Linear Regression

- Overfitting usually leads to very large parameter choices, e.g.:

$$-2.2 + 3.1 X - 0.30 X^2$$



$$-1.1 + 4,700,910.7 X - 8,585,638.4 X^2 + \dots$$



- Regularized** or **penalized** regression aims to impose a “complexity” penalty by penalizing large weights

- “Shrinkage” method

L_2 regularization \rightarrow penalizes large weights, smoother functions

©2005-2013 Carlos Guestrin

20

Ridge Regression

$$\lambda > 0 \quad \| \cdot \|_2^2 = \sum_i \cdot^2$$

- Ameliorating issues with overfitting:

penalizing large weights
"regularization"

- New objective:

$$\min_w \underbrace{\sum_{j=1}^N \left(t(x_j) - \left(w_0 + \sum_{i=1}^k w_i h_i(x_j) \right) \right)^2}_{\text{error train}} + \lambda \underbrace{\sum_{i=1}^k w_i^2}_{\|w_{1:k}\|_2^2}$$

trade off error train
with magnitude of coeffs.

L_2 regularization

importantly, don't
penalize w_0
let me free