

Ch 5.1.1-2: Leave One Out Cross-validation

Lecture 12 - CMSE 381

Prof. Elizabeth Munch

Michigan State University

::

Dept of Computational Mathematics, Science & Engineering

Mon, Sep 30, 2024

Announcements

Last time:

- Exam

Announcements:

- Exam 1 grades.... hopefully soon
- HW #4 Posted.
 - ▶ Changed Deadline! Due Wednesday Oct 9.

Lec #	Date			Reading	HW
12	Mon	9/30	Leave one out CV	5.1.1, 5.1.2	
13	Wed	10/2	k-fold CV	5.1.3	
14	Fri	10/4	More k-fold CV,	5.1.4-5	
15	Mon	10/7	k-fold CV for classification	5.1.5	
16	Wed	10/9	Resampling methods: Bootstrap	5.2	HW #4 Due Weds 10/9
17	Fri	10/11	Subset selection	6.1	
18	Mon	10/14	Shrinkage: Ridge	6.2.1	
19	Wed	10/16	Shrinkage: Lasso	6.2.2	
20	Fri	10/18	Dimension Reduction	6.3	HW #5 Due Fri 10/18
	Mon	10/21	No class - Fall break		
	Wed	10/23	Review		
	Fri	10/25	Midterm #2		

Covered in this lecture

- Validation Set
- LOO CV

Section 1

Validation set

What's the problem?



- How well is my ML method doing? *Model Assessment*
- Which method is best for our data?
- How many features should I use? Which ones? *Model selection*
- What is the uncertainty in the learned parameters?

Training Error vs Testing Error


Training Error

Testing Error


Throw-back Monday



Model tradeoffs




Validation set approach



- Divide randomly into two parts:
 - ▶ Training set
 - ▶ Validation/Hold-out/Testing set
- Fit model on training set
- Use fitted model to predict response for observations in the test set
- Evaluate quality (e.g. MSE)


Example with the auto data




Predicting mpg using horsepower:

$$\text{mpg} = \beta_0 + \beta_1 \text{hp} + \beta_2 \text{hp}^2 + \cdots + \beta_p \text{hp}^p$$

Rinse and repeat



Again example with auto data




Coding example in jupyter notebook

Section 2

Leave-One-Out Cross-Validation (LOOCV)

The idea



The idea in mathy words

- Remove (x_1, y_1) for testing.
- Train the model on $n - 1$ points:
 $\{(x_2, y_2), \dots, (x_n, y_n)\}$
- Calculate $\text{MSE}_1 = (y_1 - \hat{y}_1)^2$


- Remove (x_2, y_2) for testing.
- Train the model on $n - 1$ points:
 $\{(x_1, y_1), (x_3, y_3), \dots, (x_n, y_n)\}$
- Calculate $\text{MSE}_2 = (y_2 - \hat{y}_2)^2$

- Rinse and repeat

Return the score:

$$CV_{(n)} = \frac{1}{n} \sum_{i=1}^n \text{MSE}_i$$

Again example with auto data




Do the LOOCV coding section

LOOCV Pros and Cons

Advantages:

Disadvantages:

Validation set



LOO-CV



LOO-CV Score

$$CV_{(n)} = \frac{1}{n} \sum_{i=1}^n \text{MSE}_i$$

Next time

Lec #	Date		Reading	HW	
12	Mon	9/30	Leave one out CV	5.1.1, 5.1.2	
13	Wed	10/2	k-fold CV	5.1.3	
14	Fri	10/4	More k-fold CV,	5.1.4-5	
15	Mon	10/7	k-fold CV for classification	5.1.5	
16	Wed	10/9	Resampling methods: Bootstrap	5.2	HW #4 Due Weds 10/9
17	Fri	10/11	Subset selection	6.1	
18	Mon	10/14	Shrinkage: Ridge	6.2.1	
19	Wed	10/16	Shrinkage: Lasso	6.2.2	
20	Fri	10/18	Dimension Reduction	6.3	HW #5 Due Fri 10/18
	Mon	10/21	No class - Fall break		
	Wed	10/23	Review		
	Fri	10/25	Midterm #2		