10-11-2016

Statistical Learning II

1. Organization
2. Warm-up: decision theory, ERM
3. Interpretations of Least Squares
4. Cross-validation
5. Model Selection Intro
6. Best-subset Selection
7. Ridge Regression
8. Lasso

1. Organization

* Blackboard + Usis enrolled?
* E-mail me if entitled to more time on exam.
* Lectures notes + reading list are on website
* Book buying: go through university catalogue
  or use link on website with eduroam connection

2. Warm-up

supervised learning

classification (e.g. spam, digits)

regression (e.g. fitting polynomials)

$T = \begin{pmatrix} y_1 \\ x_1 \end{pmatrix}, \dots, \begin{pmatrix} y_N \\ x_N \end{pmatrix}$

↳ $\hat{f} \in \mathcal{F}$ and predict $\hat{y} = \hat{f}(x)$ for new X

2a Statistical Decision Theory

$EPE(f) = \underset{X,Y}{\mathbb{E}}[L(Y, f(X))]$ measures quality of f

regression: $L(Y, f(X)) = (Y - f(X))^2$ "squared error"

classification:

$$L(Y, f(X)) = \begin{cases} 0 & \text{if } f(X) = Y \\ 1 & \text{if } f(X) \neq Y \end{cases} \quad \text{"0/1-loss"}$$

Bayes optimal predictor:

$$f_B = \underset{f}{\operatorname{argmin}} \; EPE(f)$$

regression: $f_B(x) = \mathbb{E}[Y|X]$

classification: $f_B(x) = \underset{\hat{Y}}{\operatorname{argmax}} \Pr(Y = \hat{Y}|X)$

Estimators:

$\hat{f}$ depends on $\mathcal{T}$ $\Rightarrow$ $EPE(\hat{f})$ depends on $\mathcal{T}$

$\Rightarrow$ evaluate $\underset{\mathcal{T}}{\mathbb{E}}[EPE(\hat{f})]$

2b) Empirical Risk Minimization

$$\hat{f} = \underset{f \in \mathcal{F}}{\operatorname{argmin}} \; \frac{1}{N} \sum_{i=1}^{N} L(y_i, f(x_i))$$

classification: - minimize #mistakes

- usually cannot compute efficiently for 0/1-loss

regression: - minimize sum of squared errors

- can compute if $\mathcal{F}$ is linear model

3. Regression I: least squares

Linear model: $\hat{Y} = X^T\hat{\beta} = \hat{f}(x)$

$$RSS(\beta) = \sum_{i=1}^{N} (y_i - x_i^T\beta)^2$$

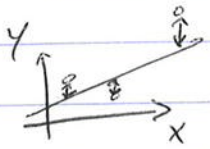

$\hat{\beta} = \underset{\beta}{\operatorname{argmin}}\ RSS(\beta)$ ← may not have unique solution, e.g. if two features are the same. In general: if features are not linearly independent.

a) Interpretations

I. ERM for squared error

$$\mathcal{F} = \{ f_\beta(x) = x^T\beta \mid \beta \in \mathbb{R}^p \}$$

$$\hat{f} = f_{\hat{\beta}} = \underset{f \in \mathcal{F}}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^{N} L(y_i, f(x_i))$$

II. Maximum Likelihood for Gaussian Errors

Model: $Y = x^T\beta + \varepsilon \qquad \varepsilon \sim \mathcal{N}(0, \sigma^2)$

$$p_\beta(y_i|x_i) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y_i - x_i^T\beta)^2}{2\sigma^2}}$$

$\log(a \cdot b) = \log a + \log b$

$$\hat{\beta} = \underset{\beta}{\operatorname{argmax}} \prod_{i=1}^{N} p_\beta(y_i|x_i)$$

$$= \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^{N} -\log\left(\frac{1}{\sqrt{2\pi\sigma^2}} \cdot e^{-\frac{(y_i - x_i^T\beta)^2}{2\sigma^2}} \right)$$

$$= \underset{\beta}{\operatorname{argmin}}\ N \cdot \left(-\log \frac{1}{\sqrt{2\pi\sigma^2}}\right) + \sum_{i=1}^{N} -\log e^{-\frac{(y_i - x_i^T\beta)^2}{2\sigma^2}}$$

$$= \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^{N} \frac{1}{2\sigma^2} (y_i - x_i^T\beta)^2$$

$$= \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^{N} (y_i - x_i^T\beta)^2$$

← noise magnitude σ^2 does not depend on x_i

So least squares is often associated with assumption of Gaussian errors, but it can be useful more widely! (see I or III)

4. k-Fold Cross-Validation

Hyperparameter $m \in \{1, \ldots, M\}$

E.g. * m is degree of polynomial in linear regression
* m is k in k-nearest neighbour

Training set → $\hat{f}_m$ for each m.

~~What happens if I look at the fit on training data~~
~~1) Construct least squares estimate $\hat{f}_m$ for each degree polynomial~~

What happens if I look at the fit on the training data to choose m (i.e. I use ERM)?

- for degree of polynomial.
- for k in k-NN.

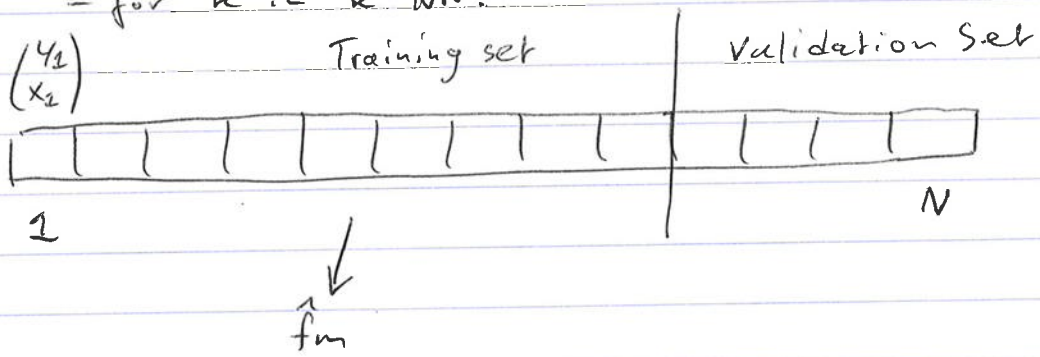

Hold-out estimate:

- Randomly set aside validation set V
- Construct $\hat{f}_m$ on rest of training data for all m
- Use ERM on validation set:

$$\hat{m} = \underset{m}{\operatorname{argmin}} \frac{1}{|V|} \sum_{i \in V} L(Y_i, \hat{f}_m(X_i))$$

- Works well if V is large enough, but that reduces training set a lot

K-Fold Cross validation:

Try to get away with smaller validation set by averaging multiple choices.

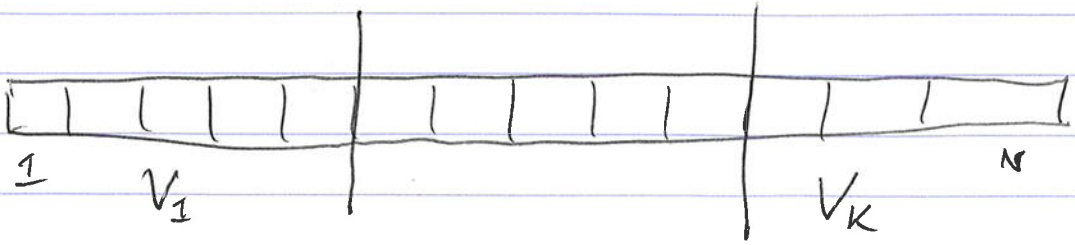

For $k = 1, \ldots, K$:

- train $\hat{f}_m$ on all data except V_k (for each m)
- evaluate on V_k: (use V_k as validation set)

$$\hat{L}_m^k = \frac{1}{|V_k|} \sum_{i \in V_k} L(y_i, \hat{f}_m(x_i))$$

$$\hat{m} = \underset{m}{\operatorname{argmin}} \frac{1}{K} \sum_{k=1}^{K} \hat{L}_m^k$$

(average over all choices of validation set)

Often train final $\hat{f}_{\hat{m}}$ on all data, or take the ~~mean~~ average of the ~~estimates~~ K estimators found for $\hat{m}$ during cross-validation.

$K = N$: "leave-one-out cross-validation"

(story about Mervyn Stone and LOO.)

5. Model Selection for Regression Intro

Suppose many features p. Maybe even $p \gg N$!

often not satisfied with least squares, because we want:

1. better prediction accuracy:

   least squares often has low bias, but high variance

2. better interpretability:

   If many features, which ones are most important?

   E.g. If we have 20 000 genes, want to know which ones are (most) relevant for prediction.

2. Best-subset Selection

Want to use only m most important features.
Select $m \leq p$ features as follows:

1. Run least squares for each of the $\binom{p}{m} \approx \left(\frac{p}{m}\right)^m$ subsets of m features.
2. Choose the subset such that $RSS(\hat{\beta})$ is smallest.

Different ways to choose m. E.g. use cross-validation.

Problems:

* Computation: exponentially many subsets in m
* Variance: discrete choice of subset; if multiple subsets approximately equally good, then which one of those is chosen varies a lot if we would sample a training set multiple times.

Computable approximations:

* Forward Stepwise: start with no variables,
  greedily choose one variable to add
  until we have m variables
* Backward Stepwise: start with all variables,
  greedily choose one variable to remove
  until we have m left.

3. Ridge regression

a) Definition

We set the first feature to 1: $X = \begin{pmatrix} 1 \\ x_2 \\ x_3 \\ \vdots \\ x_p \end{pmatrix}$

ridge regression: $\hat{\beta}_{ridge} = \underset{\beta}{\operatorname{argmin}} \; RSS(\beta) + \lambda \sum_{j=2}^{p} \beta_j^2$

$$= \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^{N} \left(y_i - \beta_1 - \sum_{j=2}^{p} x_{ij} \beta_j \right)^2 + \lambda \sum_{j=2}^{p} \beta_j^2$$

alternative formulation: ~~If~~ If $\hat{\beta}$ is unique:

$$\hat{\beta}_{ridge} = \underset{\beta}{\operatorname{argmin}} \; RSS(\beta)$$

$$\text{subject to } \sum_{j=2}^{p} \beta_j^2 \leq t$$

~~If β is unique then~~

One-to-one relation between λ and t ~~that is unique~~

b) Motivation

* More stable / less variance than least squares for highly correlated features

Example: two features the same: $x_2 = x_3$

Then $RSS(\beta) = \sum_{i=1}^{N} \left(y_i - \beta_1 - \sum_{j=2}^{p} x_{ij} \beta_j \right)^2$ the same for

1.] $\beta = (\beta_1, \beta_2, \beta_3, \ldots, \beta_p)$ and

2.] $\beta = (\beta_1, \beta_2 + a, \beta_3 - a, \ldots, \beta_p)$ for any a.

so least squares does not distinguish

ridge regression prefers equal weights (situation 2)

For highly correlated features it makes the weights approximately equal

* $\hat{\beta}^{ridge}$ is always uniquely defined for $\lambda > 0$

Because $RSS(\beta) + \lambda \sum_{j=2}^{p} \beta_j^2$ is strictly convex.

Warning: need to normalize features
- If we make feature x_j 10 times as small, then least squares makes $\hat{\beta}_j$ 10 times as big and its predictions don't change.
- This does not hold for ridge regression, so need to center and rescale features to standard range

c) Computation

- Can interpret as least squares with fake training data.

$$\begin{pmatrix} y_{N+1} \\ x_{N+1} \end{pmatrix}, \dots, \begin{pmatrix} y_{N+p-1} \\ x_{N+p-1} \end{pmatrix}$$

$$y_{N+j} = 0 \qquad x_{N+j} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ \sqrt{\lambda} \\ 0 \\ \vdots \\ 0 \end{pmatrix} \leftarrow (j+1)\text{-th position} \qquad \text{for } j = 1, \dots, p-1$$

$$(y_{N+j} - x_{N+j}^T \beta)^2 = \lambda \beta_{j+1}^2$$

$$X = \begin{pmatrix} x_{11} & \cdots & x_{1p} \\ \vdots & & \\ x_{N1} & & x_{Np} \end{pmatrix} \qquad y = \begin{pmatrix} y_1 \\ \vdots \\ y_N \end{pmatrix}$$

with fake training data:

$$\bar{X} = \begin{pmatrix} x_{11} & \cdots & x_{1p} \\ \vdots & & \\ x_{N} & & x_{Np} \\ 0 & \sqrt{\lambda} & 0 & 0 & \cdots \\ 0 & 0 & \sqrt{\lambda} & 0 & \cdots \\ 0 & 0 & 0 & \sqrt{\lambda} & \\ 0 & 0 & 0 & 0 & \sqrt{\lambda} \end{pmatrix} \qquad \bar{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_N \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

$$\hat{\beta}_{LS} = (X^T X)^{-1} X^T y$$

$$\hat{\beta}_{ridge} = (\bar{X}^T \bar{X})^{-1} \bar{X}^T \bar{y}$$

$$= \left(X^T X + \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & \lambda & 0 & 0 & 0 \\ 0 & 0 & \lambda & 0 & 0 \\ 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & 0 & \lambda \end{pmatrix}\right)^{-1} Xy$$

$$= (X^T X + \lambda A)^{-1} Xy \qquad A = \begin{pmatrix} 0 & & & \\ & 1 & & 0 \\ & & 1 & \\ & 0 & & \ddots \\ & & & & 1 \end{pmatrix}$$

Different from book, because deals with intercept by pre-processing the data, but is equivalent.

4. Lasso

a) Definition

$$\hat{\beta}_{Lasso} = \underset{\beta}{\operatorname{argmin}} \; RSS(\beta) + \lambda \sum_{j=2}^{p} |\beta_j|$$

If $\hat{\beta}_{LS}$ is unique, then

$$\hat{\beta}_{Lasso} = \underset{\beta}{\operatorname{argmin}} \; RSS(\beta)$$

$$\text{subject to } \sum_{j=2}^{p} |\beta_j| \leq t$$

One-to-one relation: $t \leftrightarrow \lambda$

b) Motivation

* Sets many weights β_j to 0 to find only most important features: good for prediction and interpretation

* $p \gg N$ allowed

NB Need to normalize features

c) Computation

- Exist multiple efficient algorithms
- LARS computes solutions for all values of t in one pass