

L19: Feature Selection

Nov 5, 2023


Data Mining

Jeff M. Phillips



Data $X = \{x_1, x_2, \dots, x_n\} \subset \mathbb{R}^d$

$$X \in \mathbb{R}^{n \times d}$$



x_j
column
a feature

Goal $\Rightarrow$




x_1, x_2, x_3

$$X' \in \mathbb{R}^{n \times k}$$


$$k \ll d$$

Recall Matrix sketching

$\rightarrow$ row sampling.




$\Rightarrow$



$$Q = \frac{d}{k^2}$$

Transpose
 $\rightarrow$ column
sampling
L2 score

Reduce


$$x \in \mathbb{R}^{n \times d} \rightarrow x' \in \mathbb{R}^{n \times d'}$$

but for what goal?

Input $(X, y) \in \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times 1}$

point (x_i, y_i)

$\uparrow$ label

goal $f(x_i) \approx y_i$

Regression

Linear Least Squares Reg.

$$S_f(X, y) = \sum_{i=1}^n (f(x_i) - y_i)^2$$

$$f_\alpha(x_i) = \langle \alpha, x_i \rangle$$

solution

$$\alpha = (X^T X)^{-1} X^T y$$

Regularized Lin Regression

$$S(\alpha; X, y) = \sum_{i=1}^n (\langle x_i, \alpha \rangle - y_i)^2 = \|X\alpha - y\|_2^2$$


argmin $\alpha \in \mathbb{R}^d$ $\|X\alpha - y\|_2 + \underbrace{s \|\alpha\|_2}_{\text{regularization term}}$

parameter

Ridge Regression

Tikhonov Regularization

$$\alpha_s = (X^T X + s^2 I)^{-1} X^T y$$



Lasso

$$\alpha_s^* = \operatorname{argmin}_{\alpha \in \mathbb{R}^d} \|X\alpha - y\|_2 + \underbrace{\lambda}_{\text{parameter}} \underbrace{\|\alpha\|_1}_{\text{Lasso}}$$

— in high-dimensions $\Rightarrow$ least $\&$ sparse solution

many $\alpha_j = 0 \Rightarrow f_x$ ignores those features

features not ignored $\alpha_j \neq 0$

$\rightarrow$ selected

larger λ fewer features selected


Why does Lasso give sparse solution?

Rewrite as

$$\alpha_t^* \rightarrow \arg \min_{\alpha \in \mathbb{R}^d} \|X\alpha - g\|_2 \quad \text{s.t.} \quad \|\alpha\|_1 \leq t$$

for all s in lasso, $\exists t$ s.t. $\alpha_s^* = \alpha_t^*$

$$t = \|\alpha_s^*\|_1$$



$$\text{Lasso } \alpha_1^* = (1, 0)$$

Orthogonal Matching Pursuit

Set residual $r = y$ $\alpha_j = 0 \quad \forall j$

for $i = 1$ to bz

Set $x_j = \underset{x_j \in X}{\operatorname{argmax}} |\langle r, x_j \rangle|$

Set $\alpha_j = \underset{\alpha}{\operatorname{argmin}} \|r - x_j \alpha\|^2 + s|\alpha|$

Set $r = r - x_j \alpha_j$

Return α

Units

$$x_i = (\text{height}_i, \text{weight}_i, \text{age}_i)$$

inches lbs years

$$f_\alpha(x_i) = \langle x_i, \alpha \rangle = x_{i1}\alpha_1 + x_{i2}\alpha_2 + x_{i3}\alpha_3$$

not a problem

unit of $\alpha_1 = \frac{1}{\text{inches}}$

$$\alpha_2 = \frac{1}{\text{lbs}}$$


$$\alpha_3 = \frac{1}{\text{years}}$$

$$\alpha = \arg \min_{\alpha \in \mathbb{R}^d} \|X\alpha - y\| + \lambda \|\alpha\|_1$$


$$\hookrightarrow \|\alpha\|_1 = |\alpha_1| + |\alpha_2| + |\alpha_3|$$

Normalizer first

• standardize



• 0-1 Normalization



for each column / feature

$$X_{ij} \in X_j$$

① subtract mean $\frac{1}{n} \sum_i X_{ij}$

② divide variance $\sigma^2 = \|X_j'\|^2$


$$X_{ij} \in X_j \Rightarrow \in [0, 1]$$

① subtract $\min_i X_{ij}$ from all X_{ij}

② divide all X_{ij} by $\max_i X_{ij}$

Problems w/ Normalization

- affected by outliers



- change 1 data point
↳ changes normalization.

- features may be colinear

X_1 & X_2 have similar correlation w/ y .
"fixed" using principal components

What can we learn from sparse

$$\alpha_s^* = \arg \min_{\alpha} \|X\alpha - y\|_2 \text{ s.t. } \|\alpha\|_1 \leq s$$

$$\alpha_s^* = (0.3, 0, 0, 0, 0.2, 0, 0, 7)$$

$j \quad 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \quad 8$

$$\max_j \alpha_j^*$$

weight

h₁

most

important

feature

$\alpha_8^* \leftarrow \text{largest}$

$\hookrightarrow$ maybe not!