




CPSC540



Linear Predictive Models



Nando de Freitas
September, 2011
University of British Columbia

Outline

Linear regression and prediction are simple **supervised learning** tasks.
But,

- Many real processes can be approximated with linear models.
- Linear regression often appears as a module of larger systems.
- Linear problems can be solved analytically.
- Linear prediction provides an introduction to many of the core concepts of machine learning, while still allowing for analytical tractability.

We are given a training dataset of n instances of input-output pairs $\{\mathbf{x}_{1:n}, \mathbf{y}_{1:n}\}$. Each input $\mathbf{x}_i \in \mathbb{R}^{1 \times d}$ is a vector with d attributes. The inputs are also known as predictors or covariates. The output, often referred to as the target, will be assumed to be univariate, $\mathbf{y}_i \in \mathbb{R}$, for now.



A typical dataset with $n = 4$ instances and 2 attributes would look like the following table:

Wind speed	People inside building	Energy requirement
100	2	5
50	42	25
45	31	22
60	35	18

Given the training set $\{\mathbf{x}_{1:n}, \mathbf{y}_{1:n}\}$, we would like to learn a model of how the inputs affect the outputs. Given this model and a new value of the input \mathbf{x}_{n+1} , we can use the model to make a prediction $\hat{y}(\mathbf{x}_{n+1})$.

Prostate cancer example



□ **Goal:** Predict a prostate-specific antigen (log of **lpsa**) from a number of clinical measures in men who are about to receive a radical prostatectomy.

□ The **inputs** are:

- Log cancer volume (**lcavol**)
- Log prostate weight (**lweight**)
- Age
- Log of the amount of benign prostatic hyperplasia (**lbph**)
- Seminal vesicle invasion (**svi**) - *binary*
- Log of capsular penetration (**lcp**)
- Gleason score (**gleason**) – *ordered categorical*
- Percent of Gleason scores 4 or 5 (**pgg45**)




[Hastie, Tibshirani & Friedman book]

Linear regression in 1D

$$\hat{y}(\mathbf{x}_i) = \theta_1 - x_i\theta_2$$

$$J(\theta) = \sum_{i=1}^n (y_i - \hat{y}_i)^2 = \sum_{i=1}^n (y_i - \theta_1 - x_i\theta_2)^2$$



Linear prediction

In general, the linear model is expressed as follows:

$$\hat{y}_i = \sum_{j=1}^d x_{ij}\theta_j,$$

where we have assumed that $x_{i1} = 1$ so that θ_1 corresponds to the intercept of the line with the vertical axis. θ_1 is known as the bias or offset.

In matrix form, the expression for the linear model is:

$$\hat{\mathbf{y}} = \mathbf{X}\boldsymbol{\theta},$$

with $\hat{\mathbf{y}} \in \mathbb{R}^{n \times 1}$, $\mathbf{X} \in \mathbb{R}^{n \times d}$ and $\boldsymbol{\theta} \in \mathbb{R}^{d \times 1}$. That is,

$$\begin{bmatrix} \hat{y}_1 \\ \vdots \\ \hat{y}_n \end{bmatrix} = \begin{bmatrix} x_{11} & \cdots & x_{1d} \\ \vdots & \vdots & \vdots \\ x_{n1} & \cdots & x_{nd} \end{bmatrix} \begin{bmatrix} \theta_1 \\ \vdots \\ \theta_d \end{bmatrix}.$$

Linear prediction

Wind speed	People inside building	Energy requirement
100	2	5
50	42	25
45	31	22
60	35	18

For our energy prediction example, we would form the following matrices with $n = 4$ and $d = 3$:

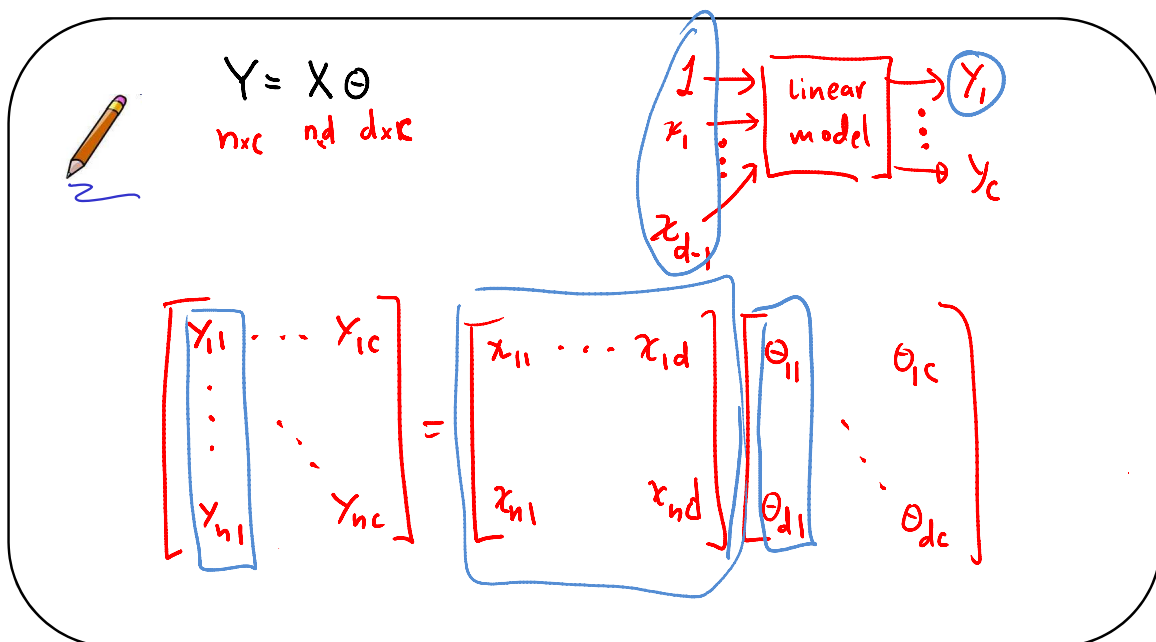
$$\mathbf{y} = \begin{bmatrix} 5 \\ 25 \\ 22 \\ 18 \end{bmatrix}, \quad \mathbf{X} = \begin{bmatrix} 1 & 100 & 2 \\ 1 & 50 & 42 \\ 1 & 45 & 31 \\ 1 & 60 & 35 \end{bmatrix}, \quad \boldsymbol{\theta} = \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{bmatrix}.$$

Suppose that $\boldsymbol{\theta} = [1 \ 0 \ 0.5]^T$. Then, by multiplying \mathbf{X} times $\boldsymbol{\theta}$, we would get the following predictions on the training set:

$$\hat{\mathbf{y}} = \begin{bmatrix} 2 \\ 22 \\ 16.5 \\ 18.5 \end{bmatrix} = \begin{bmatrix} 1 & 100 & 2 \\ 1 & 50 & 42 \\ 1 & 45 & 31 \\ 1 & 60 & 35 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0.5 \end{bmatrix}.$$

Multiple outputs

If we have several outputs $\mathbf{y}_i \in \mathbb{R}^c$, our linear regression expression becomes:



Linear classification with indicators

$Y = \begin{bmatrix} c_1 & c_2 & c_3 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \\ \vdots & \vdots & \vdots \\ 0 & 0 & 1 \end{bmatrix}$

$\hat{Y}_c = X\Theta$

$\hat{Y}_c = \begin{bmatrix} 1.1 & 0.1 & 0.2 \\ \vdots & \vdots & \vdots \\ 0 & 0.1 & 0.9 \end{bmatrix} = \begin{bmatrix} x_{n1} & & \\ & \ddots & \\ & & x_{nd} \end{bmatrix} \begin{bmatrix} \theta_{11} & \theta_{12} & \theta_{13} \\ \vdots & \vdots & \vdots \\ \theta_{d1} & \theta_{d2} & \theta_{d3} \end{bmatrix} C_1$

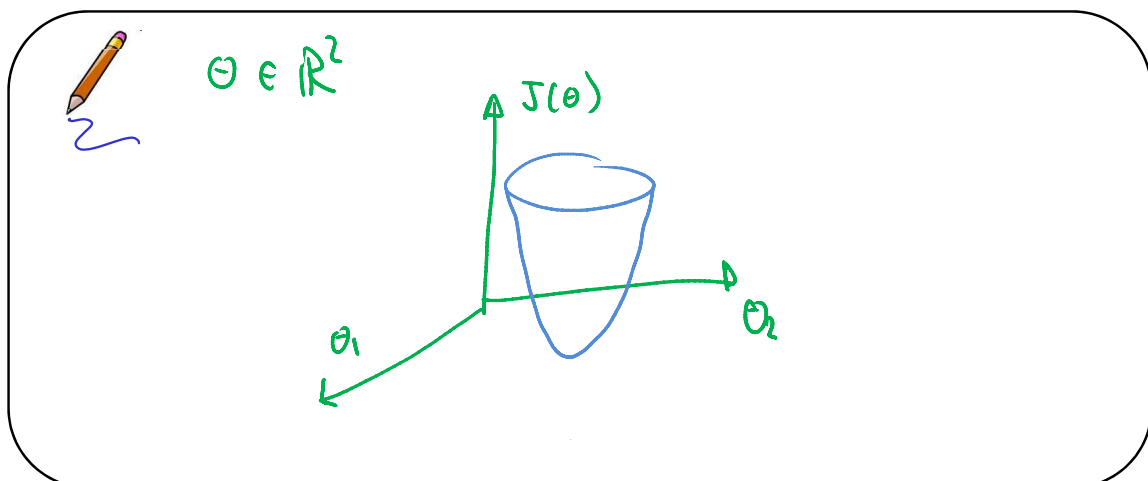
$\hat{y}_i = \operatorname{argmax}_c \hat{Y}_{ci}$

As discussed in Ch 4 of [Hastie, Tibshirani & Friedman book], this approach is easy, but not recommended. To describe better approaches, we will have to introduce notions of probability and optimization.

Optimization approach

Our aim is to minimise the quadratic cost between the output labels and the model predictions

$$J(\theta) = (\mathbf{y} - \mathbf{X}\theta)^T (\mathbf{y} - \mathbf{X}\theta) = \sum_{i=1}^n (y_i - \mathbf{x}_i^T \theta)^2$$



Optimization approach

We will need the following results from matrix differentiation:

$$\frac{\partial \mathbf{A}\theta}{\partial \theta} = \mathbf{A}^T \text{ and } \frac{\partial \theta^T \mathbf{A}\theta}{\partial \theta} = 2\mathbf{A}^T \theta$$



$$\begin{aligned} \frac{\partial J(\theta)}{\partial \theta} &= \frac{\partial}{\partial \theta} (y - X\theta)^T (y - X\theta) \\ &= \frac{\partial}{\partial \theta} (y^T y + \theta^T X^T X \theta - 2y^T X \theta) \\ &= 0 + 2X^T X \theta - 2X^T y \\ X^T X \theta &= X^T y \end{aligned}$$

Optimization approach



$$\begin{matrix} X^T & X & \theta & = & X^T & y \\ d \times n & n \times d & d \times 1 & & d \times n & n \times 1 \end{matrix}$$

These are the **normal equations**. The solution (estimate) is:

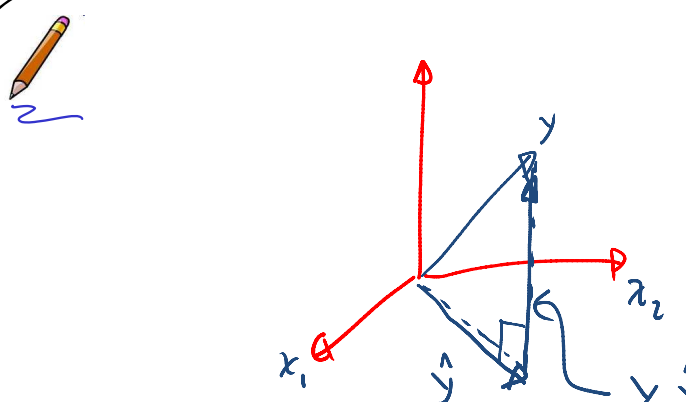
$$\hat{\theta} = (X^T X)^{-1} X^T y$$

The corresponding predictions are

$$\hat{y} = \mathbf{H}y = X^* \hat{\theta} = X^* (X^T X)^{-1} X^T y$$

where \mathbf{H} is the “hat” matrix.

Geometric approach



$$\mathbf{X}^T = \begin{bmatrix} x_1^T \\ x_2^T \end{bmatrix}$$

$$\mathbf{X}^T(\mathbf{y} - \hat{\mathbf{y}}) = \begin{bmatrix} 0 \\ 0 \end{bmatrix} = \underline{0}$$

$$x_1^T y = x_1^T \hat{y}$$


$$x_2^T y = x_2^T \hat{y} \Rightarrow \Theta = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T y$$

Probability approach: Univariate Gaussian distribution

The probability density function (pdf) of a Gaussian distribution is given by

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(x-\mu)^2}$$

where μ is the mean or center of mass and σ^2 is the variance.



Multivariate Gaussian distribution

Let $\mathbf{y} \in \mathbb{R}^{n \times 1}$, then pdf of an n-dimensional Gaussian is given by


$$p(\mathbf{x}) = |2\pi\Sigma|^{-1/2} e^{-\frac{1}{2}(\mathbf{y}-\boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{y}-\boldsymbol{\mu})},$$

where

$$\boldsymbol{\mu} = \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix} = \begin{pmatrix} \mathbb{E}(y_1) \\ \vdots \\ \mathbb{E}(y_n) \end{pmatrix}$$

and

$$\Sigma = \begin{pmatrix} \sigma_{11} & \cdots & \sigma_{1n} \\ \cdots & \cdots & \cdots \\ \sigma_{n1} & \cdots & \sigma_{nn} \end{pmatrix} = \mathbb{E}[(\mathbf{y} - \boldsymbol{\mu})(\mathbf{y} - \boldsymbol{\mu})^T]$$



Energy (“distance”) and probability

$d(\mathbf{y})$

The exponent $\frac{1}{2}(\mathbf{y} - \boldsymbol{\mu})^T \Sigma^{-1}(\mathbf{y} - \boldsymbol{\mu})$ is called the Mahalanobis distance. Conceptually, it measures the weighted Euclidean distance between \mathbf{y} and $\boldsymbol{\mu}$.



$$P(\mathbf{y}) = |2\pi\Sigma|^{-1/2} e^{-\frac{1}{2}d(\mathbf{y})}$$

$$\frac{d}{2} = E$$

$$P(\mathbf{y}) = Z e^{-E(\mathbf{y})}$$

$$-\text{Log} P(\mathbf{y}) = E(\mathbf{y}) - \text{Log} Z$$

Maximum likelihood

We will assume that our errors in the target predictions are Gaussian distributed as follows:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\theta} + \mathcal{N}(0, \sigma^2 I_n)$$

$n \times 1$ $n \times d$ $d \times 1$


$I_n = \begin{bmatrix} 1 & & & \\ & \ddots & & \\ & & 1 & \\ & & & \ddots & \\ 0 & & & & 1 \end{bmatrix}$

$$\begin{aligned}
 p(\mathbf{y} | \mathbf{X}, \boldsymbol{\theta}, \sigma) &= (2\pi\sigma^2)^{-n/2} e^{-\frac{1}{2\sigma^2}(\mathbf{y} - \mathbf{X}\boldsymbol{\theta})^T(\mathbf{y} - \mathbf{X}\boldsymbol{\theta})} \\
 &= (2\pi\sigma^2)^{-n/2} e^{-\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - \mathbf{x}_i^T \boldsymbol{\theta})^2} \\
 P(y_1, y_2) &= P(y_1)P(y_2) \quad \text{+ independent} \\
 &= \prod_{i=1}^n (2\pi\sigma^2)^{-1/2} e^{-\frac{1}{2\sigma^2} (y_i - \mathbf{x}_i^T \boldsymbol{\theta})^2} \\
 &= \prod_{i=1}^n p(y_i | \mathbf{x}_i, \boldsymbol{\theta}, \sigma).
 \end{aligned}$$

Maximum likelihood

$$J(\boldsymbol{\theta}) = \sum_{i=1}^n (y_i - \hat{y}_i)^2 = \sum_{i=1}^n (y_i - \theta_1 - x_i \theta_2)^2$$

$$\begin{aligned}
 p(\mathbf{y} | \mathbf{X}, \boldsymbol{\theta}, \sigma) &= (2\pi\sigma^2)^{-n/2} e^{-\frac{1}{2\sigma^2}(\mathbf{y} - \mathbf{X}\boldsymbol{\theta})^T(\mathbf{y} - \mathbf{X}\boldsymbol{\theta})} \\
 &= (2\pi\sigma^2)^{-n/2} e^{-\frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - \mathbf{x}_i^T \boldsymbol{\theta})^2}
 \end{aligned}$$



Maximum likelihood

The maximum likelihood estimate (MLE) of θ is obtained by taking the derivative of the log-likelihood, $\log p(\mathbf{y}|\mathbf{X}, \theta, \sigma)$. The goal is to maximize the likelihood of seeing the training data \mathbf{y} by modifying the parameters (θ, σ) .



$$P(\mathbf{y}|\mathbf{X}, \theta, \sigma^2) = (2\pi\sigma^2)^{-n/2} e^{-\frac{1}{2\sigma^2}(\mathbf{y}-\mathbf{X}\theta)^T(\mathbf{y}-\mathbf{X}\theta)}$$

$$+ \log P(\mathbf{y}|\mathbf{X}, \theta, \sigma^2) = -\frac{n}{2} \log(2\pi\sigma^2)$$

$$-\frac{1}{2\sigma^2} (\mathbf{y}-\mathbf{X}\theta)^T (\mathbf{y}-\mathbf{X}\theta)$$

Maximum likelihood



The ML estimate of θ is:

$$\frac{\partial \log P(\mathbf{y}|\mathbf{X}, \theta, \sigma^2)}{\partial \theta} = 0 \quad -\frac{1}{2\sigma^2} \frac{\partial}{\partial \theta} (\mathbf{y}-\mathbf{X}\theta)^T (\mathbf{y}-\mathbf{X}\theta)$$

$$= -\frac{1}{2\sigma^2} (2\mathbf{X}^T\mathbf{X}\theta - 2\mathbf{X}^T\mathbf{y}) \rightarrow 0$$

$$\mathbf{X}^T\mathbf{X}\theta = \mathbf{X}^T\mathbf{y}$$

$$\hat{\theta}_{ML} = (\mathbf{X}^T\mathbf{X})^{-1} \mathbf{X}^T\mathbf{y}$$

Maximum likelihood



Proceeding in the same way, the ML estimate of σ is:


$$\frac{\partial}{\partial \sigma^2} \log P(y|x, \theta, \sigma^2) \rightarrow 0$$

[e.]

Probabilistic graphical models (DAGs, Bayesian networks)




All r.v.s are binary



nodes \equiv r.v.s

edges \equiv dependencies (statistical)


shaded nodes \equiv observation



$M \perp\!\!\!\perp E \mid A$
 \uparrow
 independent

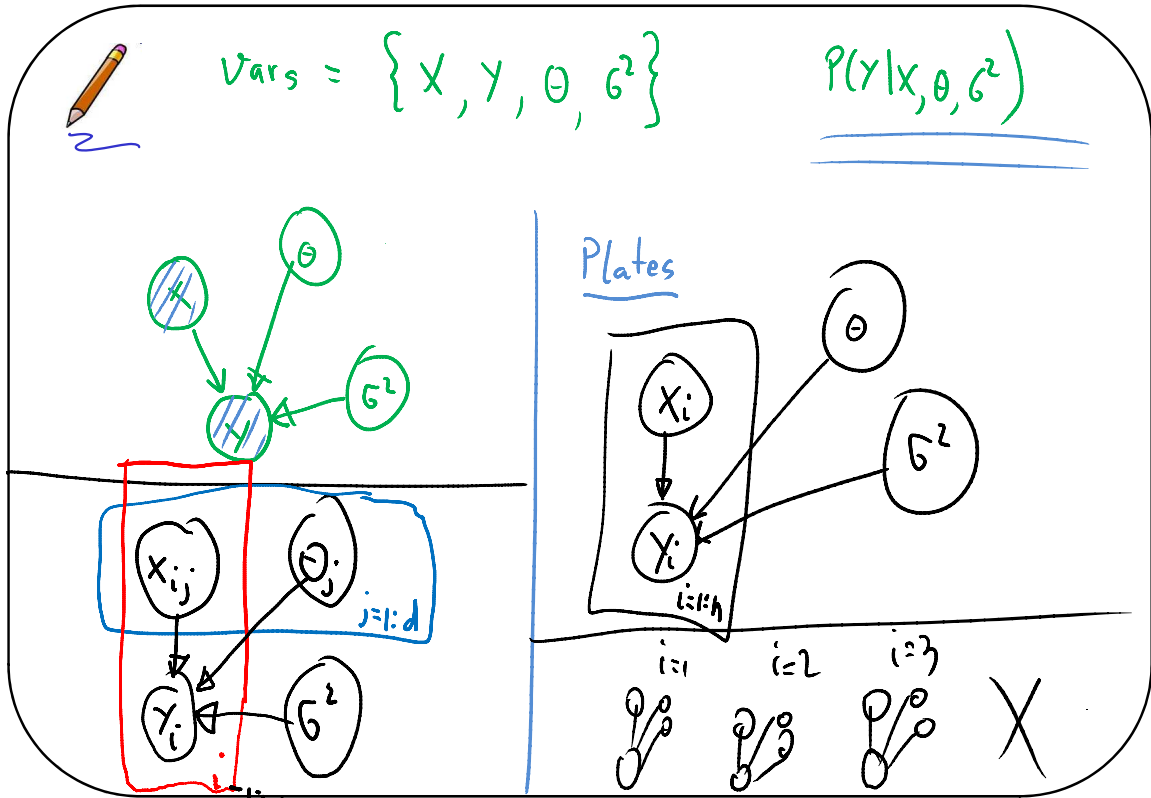
given J|M|A

B|E|A



Markov Blanket of A

Linear regression DAG



Making predictions

Predictions for September 2010 from June

Observed and Predicted Ice Extent from the Sea Ice Index

Area (10^6 km²)

Predictor: G1.0m
Prediction: 3.96 +/- 0.34
R² of Fit: 0.84

Correlation

Sep IE & Jun G1.0m G1.0m Anomaly in June 2010 G1.0m Weighted Anomaly (R²)

$P(Y|X, \theta, \sigma^2)$

Given new x^*


Use estimated $\hat{\theta}$ and $\hat{\sigma}_2$

Prediction:
 $P(Y|x^*, \hat{\theta}, \hat{\sigma}_2)$

$E(Y|x^*, \hat{\theta}, \hat{\sigma}_2) = x^* \hat{\theta}$

Why the MLE: Hallucination!

The MLE assumes that the data has been generated by a distribution $p(\mathbf{Y}|\mathbf{X}, \theta_0)$ for some *true parameter* θ_0 .



$$\begin{aligned}
 \hat{\theta} &= \arg \max_{\theta} \prod_{i=1}^n p(y_i | \mathbf{x}_i, \theta) \\
 &= \arg \max_{\theta} \sum_{i=1}^n \log p(y_i | \mathbf{x}_i, \theta) \\
 &= \arg \max_{\theta} \frac{1}{n} \sum_{i=1}^n \log p(y_i | \mathbf{x}_i, \theta) - \frac{1}{n} \sum_{i=1}^n \log p(y_i | \mathbf{x}_i, \theta_0) \\
 &= \arg \max_{\theta} \frac{1}{n} \sum_{i=1}^n \log \frac{p(y_i | \mathbf{x}_i, \theta)}{p(y_i | \mathbf{x}_i, \theta_0)} \\
 &\rightarrow \arg \min_{\theta} \int \log \frac{p(y | \mathbf{x}, \theta)}{p(y | \mathbf{x}, \theta_0)} p(y | \mathbf{x}, \theta_0) dy \\
 &= \arg \min_{\theta} \int \log p(y | \mathbf{x}, \theta) p(y | \mathbf{x}, \theta_0) dy - \int \log p(y | \mathbf{x}, \theta_0) p(y | \mathbf{x}, \theta_0) dy
 \end{aligned}$$

$\nearrow \hat{\theta}$
 \nearrow KL divergence

Why the MLE (frequentist view)

Under smoothness and identifiability assumptions, the MLE is *consistent*:

$$\hat{\theta} \xrightarrow{P} \theta_0$$

or equivalently,

$$\lim_{N \rightarrow \infty} P(|\hat{\theta} - \theta_0| > \alpha) \rightarrow 0$$

The MLE is *asymptotically normal*. That is, as $N \rightarrow \infty$, we have:

$$\hat{\theta} - \theta_0 \implies N(0, I^{-1})$$

where I is the *Fisher Information matrix*.

It is asymptotically optimal or *efficient*.

Regularization

All the answers so far are of the form

$$\hat{\theta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$$

They require the inversion of $\mathbf{X}^T \mathbf{X}$. This can lead to problems if the system of equations is poorly conditioned. A solution is to add a small element to the diagonal:

$$\hat{\theta} = (\mathbf{X}^T \mathbf{X} + \delta^2 I_d)^{-1} \mathbf{X}^T \mathbf{y}$$

This is the ridge regression estimate. It is the solution to the following **regularised quadratic cost function**

$$J(\theta) = (\mathbf{y} - \mathbf{X}\theta)^T (\mathbf{y} - \mathbf{X}\theta) + \delta^2 \theta^T \theta$$

Proof



$$\begin{aligned} J(\theta) &= (\mathbf{y} - \mathbf{X}\theta)^T (\mathbf{y} - \mathbf{X}\theta) + \delta^2 \theta^T \theta \\ &= \mathbf{y}^T \mathbf{y} - 2\mathbf{y}^T \mathbf{X}\theta + \theta^T \mathbf{X}^T \mathbf{X}\theta + \delta^2 \theta^T \theta \end{aligned}$$

$$\frac{\partial J(\theta)}{\partial \theta} = 0 - 2\mathbf{X}^T \mathbf{y} + 2\mathbf{X}^T \mathbf{X}\theta + 2\delta^2 \theta$$

$$= -2\mathbf{X}^T \mathbf{y} + 2(\mathbf{X}^T \mathbf{X} + \delta^2 \mathbf{I})\theta$$

equated to zero

$$\hat{\theta}_{\text{ridge}} = (\mathbf{X}^T \mathbf{X} + \delta^2 \mathbf{I})^{-1} \mathbf{X}^T \mathbf{y}$$

Ridge regression as constrained optimization

$J(\theta) = (y - X\theta)^T (y - X\theta) + \delta^2 \theta^T \theta$ Lagrangian

$\theta : \theta^T \theta \leq t(\delta)$!!! $\{ (y - X\theta)^T (y - X\theta) \}$

$\delta \text{ large} \Rightarrow t \text{ small}$

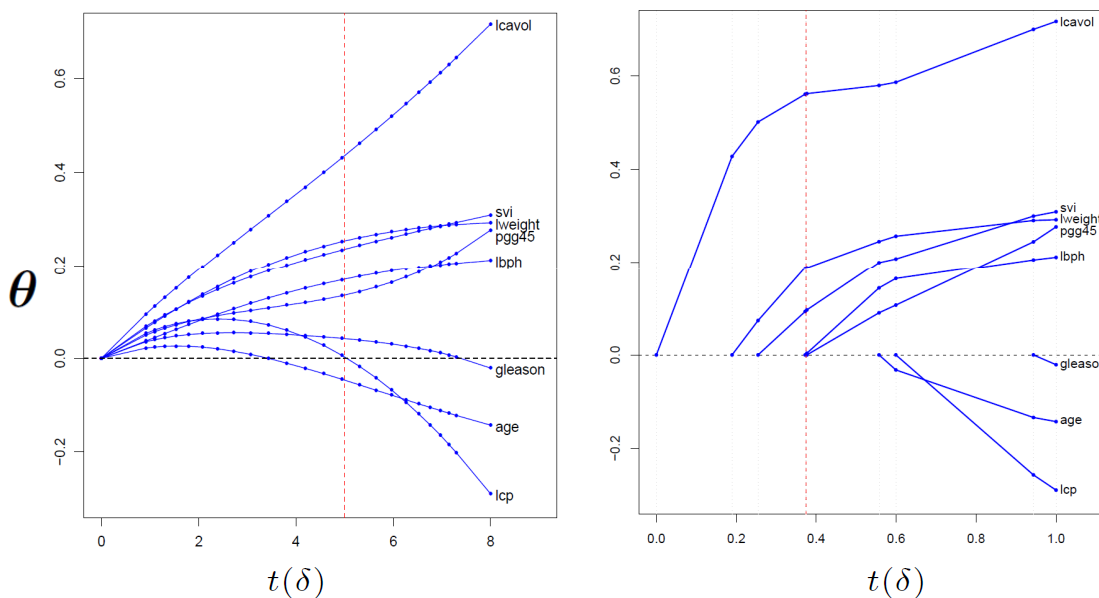
$d=2$

$\hat{\theta}_{ML} = \hat{\theta}_{ridge} (S^2 = 0)$

$\tilde{\theta}_{ridge} (\delta^2 \rightarrow \infty)$

$\theta^T \theta = [\theta_1 \ \theta_2] \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} = t = \theta_1^2 + \theta_2^2$

Ridge as constrained optimization




Ridge, feature selection, shrinkage and weight decay

Large values of θ are penalised. We are *shrinking* θ towards zero. This can be used to carry out *feature weighting*. An input $x_{i,d}$ weighted by a small θ_d will have less influence on the output y_i . This penalization with a regularizer is also known as weight decay in the neural networks literature.

Note that shrinking the bias term θ_1 is undesirable. To keep the notation simple, we will assume that the mean of \mathbf{y} has been subtracted from \mathbf{y} . This mean is indeed our estimate $\widehat{\theta}_1$.

The Lasso: least absolute selection and shrinkage operator



Spectral view of ridge regression

$$\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^T = \sum_{i=1}^d \mathbf{u}_i \sigma_i \mathbf{v}_i^T \quad \left| \begin{array}{l} \hat{\mathbf{y}} = \mathbf{U}\mathbf{U}^T \mathbf{y} \\ \hat{\mathbf{y}} = \sum_i u_i u_i^T \mathbf{y} \end{array} \right.$$



$$\begin{aligned} \mathbf{X}^T \mathbf{X} \hat{\boldsymbol{\theta}} &= \mathbf{X}^T \mathbf{y} \\ \hat{\boldsymbol{\theta}} &= (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y} \end{aligned}$$

$$\hat{\mathbf{y}} = \mathbf{X} \hat{\boldsymbol{\theta}}$$

$$\begin{aligned} (\mathbf{V}\Sigma\mathbf{U}^T)(\mathbf{U}\Sigma\mathbf{V}^T) \hat{\boldsymbol{\theta}} &= \mathbf{V}\Sigma\mathbf{U}^T \mathbf{y} \\ \mathbf{V}^T \mathbf{V} \Sigma^2 \mathbf{V}^T \hat{\boldsymbol{\theta}} &= \mathbf{V}^T \mathbf{V} \Sigma \mathbf{U}^T \mathbf{y} \\ \Sigma \mathbf{V}^T \hat{\boldsymbol{\theta}} &= \mathbf{U}^T \mathbf{y} \end{aligned}$$

$$\hat{\boldsymbol{\theta}} = \mathbf{V} \Sigma^{-1} \mathbf{U}^T \mathbf{y}$$

$$\hat{\mathbf{y}} = \mathbf{U} \Sigma \mathbf{V}^T \mathbf{V} \Sigma^{-1} \mathbf{U}^T \mathbf{y}$$

$$\hat{\mathbf{y}} = \mathbf{U} \mathbf{U}^T \mathbf{y}$$


$$\hat{\mathbf{y}} = \sum_k u_k u_k^T \mathbf{y}$$



$$\hat{\mathbf{y}}_{\text{ridge}} = \sum_{i=1}^d \left(\frac{\sigma_i^2}{\sigma_i^2 + \delta^2} \right) \mathbf{u}_i \mathbf{u}_i^T \mathbf{y}$$

$$\hat{\mathbf{y}}_{\text{ML}} = \sum_{i=1}^d \mathbf{u}_i \mathbf{u}_i^T \mathbf{y}$$

$$(\mathbf{X}^T \mathbf{X} + \delta^2 \mathbf{I}) \boldsymbol{\theta} = \mathbf{X}^T \mathbf{y} \quad [e.]$$



Say $\delta^2 = 10$, $\sigma_i^2 = 0.1$ $f_i \approx 0.01$

$\delta^2 = 10$, $\sigma_i^2 = 100$ $f_i \approx 1$

Regularization and noise filtering

The filter factor

$$f_i = \frac{\sigma_i^2}{\sigma_i^2 + \delta^2}$$

penalises small values of σ^2 (they go to zero at a faster rate).

Minimax and cross-validation

Cross-validation is a widely used technique for choosing δ . Here's an example:

δ^2	Train error $\sum_{i \in \text{train}} (y_i - \hat{y}_i)^2$	Test error $\sum_{i \in \text{test}} (y_i - \hat{y}_i)^2$	max worst	min-max e_1 (worst case)	avg e_2
0.1	100	2	100		
1	10	11	11	X	
10	1	19	19		X 10
50	20	0	20		X 10
100	100	1000	1000		

demo

Least Angle Regression

Algorithm 3.2 *Least Angle Regression.*

1. Standardize the predictors to have mean zero and unit norm. Start with the residual $\mathbf{r} = \mathbf{y} - \bar{y}$, $\beta_1, \beta_2, \dots, \beta_p = 0$.
2. Find the predictor \mathbf{x}_j most correlated with \mathbf{r} .
3. Move β_j from 0 towards its least-squares coefficient $\langle \mathbf{x}_j, \mathbf{r} \rangle$, until some other competitor \mathbf{x}_k has as much correlation with the current residual as does \mathbf{x}_j .
4. Move β_j and β_k in the direction defined by their joint least squares coefficient of the current residual on $(\mathbf{x}_j, \mathbf{x}_k)$, until some other competitor \mathbf{x}_l has as much correlation with the current residual.
5. Continue in this way until all p predictors have been entered. After $\min(N - 1, p)$ steps, we arrive at the full least-squares solution.

$$\hat{\mathbf{y}} = \mathbf{X} \boldsymbol{\beta}$$

$n \times 1$ $n \times p$ $p \times 1$

$$\beta_{\mathcal{A}_k}(\alpha) = \beta_{\mathcal{A}_k} + \alpha \cdot \delta_k$$

$$\delta_k = (\mathbf{X}_{\mathcal{A}_k}^T \mathbf{X}_{\mathcal{A}_k})^{-1} \mathbf{X}_{\mathcal{A}_k}^T \mathbf{r}_k$$


$$\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_j \dots \mathbf{x}_p]$$

Suppose \mathcal{A}_k is the active set of variables


$\mathbf{r}_k = \mathbf{y} - \mathbf{X}_{\mathcal{A}_k} \boldsymbol{\beta}_{\mathcal{A}_k}$ is the current residual.

[Hastie, Tibshirani & Friedman book]

Fitting the residuals



The Lasso & LAR



Algorithm 3.2 Least Angle Regression.

1. Standardize the predictors to have mean zero and unit norm. Start with the residual $\mathbf{r} = \mathbf{y} - \bar{\mathbf{y}}$, $\beta_1, \beta_2, \dots, \beta_p = 0$.
 2. Find the predictor \mathbf{x}_j most correlated with \mathbf{r} .
 3. Move β_j from 0 towards its least-squares coefficient $\langle \mathbf{x}_j, \mathbf{r} \rangle$, until some other competitor \mathbf{x}_k has as much correlation with the current residual as does \mathbf{x}_j .
 4. Move β_j and β_k in the direction defined by their joint least squares coefficient of the current residual on $(\mathbf{x}_j, \mathbf{x}_k)$, until some other competitor \mathbf{x}_l has as much correlation with the current residual.
 5. Continue in this way until all p predictors have been entered. After $\min(N - 1, p)$ steps, we arrive at the full least-squares solution.
-

Algorithm 3.2a Least Angle Regression: Lasso Modification.

- 4a. If a non-zero coefficient hits zero, drop its variable from the active set of variables and recompute the current joint least squares direction.
-

[Hastie, Tibshirani & Friedman book]

The Lasso as LAR intuition


When the active attributes are tied, in the sense that the absolute value of their inner-product with the residuals is the same (having standardized the inputs so that instead of using correlations, we are now using inner-products), we have:

$$\mathbf{x}_j^T (\mathbf{y} - \mathbf{X}\boldsymbol{\theta}) = \gamma \text{sign}(\mathbf{x}_j^T (\mathbf{y} - \mathbf{X}\boldsymbol{\theta})), \quad \forall j \in \mathcal{A}$$

where $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_d]$, with $\mathbf{x}_j \in \mathbb{R}^{n \times 1}$, γ is the common absolute value and the sign function $\text{sign}(\mathbf{x}_j^T (\mathbf{y} - \mathbf{X}\boldsymbol{\theta}))$ is in $\{-1, 1\}$. The set \mathcal{A} is the active set.

“Differentiating” the Lasso’s Lagrangian

$$J(\boldsymbol{\theta}) = (\mathbf{y} - \mathbf{X}\boldsymbol{\theta})^T (\mathbf{y} - \mathbf{X}\boldsymbol{\theta}) + \delta^2 \sum_{j=1}^d |\theta_j|$$



with respect to θ_j and equating to zero, we get

$$-\mathbf{x}_j^T (\mathbf{y} - \mathbf{X}\boldsymbol{\theta}) + \delta^2 \text{sign}(\theta_j) = 0$$

$$\mathbf{x}_j^T (\mathbf{y} - \mathbf{X}\boldsymbol{\theta}) = \delta^2 \text{sign}(\theta_j), \quad \forall j \in \mathcal{B}$$

[Hastie, Tibshirani & Friedman book]

Lasso using optimization


$$\frac{\partial}{\partial \theta_j} J(\boldsymbol{\theta}) = a_j \theta_j - c_j + \delta^2 \frac{\partial}{\partial \theta_j} |\theta_j|$$

$$a_j = 2 \sum_{i=1}^n x_{ij}^2$$

$$c_j = 2 \sum_{i=1}^n x_{ij} (y_i - \boldsymbol{\theta}_{-j}^T \mathbf{x}_{i,-j})$$

where $\boldsymbol{\theta}_{-j} = \boldsymbol{\theta}$ without component j and likewise for $\mathbf{x}_{i,-j}$. We see that c_j is (proportional to) the correlation between the j 'th feature \mathbf{x}_j and the residual due to the other features, $\mathbf{r}_{-j} = \mathbf{y} - \mathbf{X}_{-j} \boldsymbol{\theta}_{-j}$.


Subdifferentials



$$\partial f(x) = \begin{cases} \{-1\} & \text{if } x < 0 \\ [-1, 1] & \text{if } x = 0 \\ \{+1\} & \text{if } x > 0 \end{cases}$$

Lasso using optimization

$$\partial_{\theta_j} J(\boldsymbol{\theta}) = a_j \theta_j - c_j + \delta^2 \partial_{\theta_j} |\theta_j| = \begin{cases} \{a_j \theta_j - c_j - \delta^2\} & \text{if } \theta_j < 0 \\ [-c_j - \delta^2, -c_j + \delta^2] & \text{if } \theta_j = 0 \\ \{a_j \theta_j - c_j + \delta^2\} & \text{if } \theta_j > 0 \end{cases}$$



$$\hat{\theta}_j = \begin{cases} (c_j + \delta^2)/a_j & \text{if } c_j < -\delta^2 = -\lambda \\ 0 & \text{if } c_j \in [-\delta^2, \delta^2] = [-\lambda, \lambda] \\ (c_j - \delta^2)/a_j & \text{if } c_j > \delta^2 = \lambda \end{cases}$$

[Kevin Murphy's book]

Coordinate descent for lasso

Algorithm 11.1: Coordinate descent for lasso (aka shooting algorithm)

- 1 Initialize $\mathbf{w} = (\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I})^{-1} \mathbf{X}^T \mathbf{y}$;
- 2 **repeat**
- 3 **for** $j = 1, \dots, D$ **do**
- 4 $a_j = 2 \sum_{i=1}^n x_{ij}^2$;
- 5 $c_j = 2 \sum_{i=1}^n x_{ij} (y_i - \mathbf{w}^T \mathbf{x}_i + w_j x_{ij})$;
- 6 **if** $c_j < -\lambda$ **then**
- 7 $w_j = \frac{c_j + \lambda}{a_j}$
- 8 **else if** $c_j > \lambda$ **then**
- 9 $w_j = \frac{c_j - \lambda}{a_j}$
- 10 **else**
- 11 $w_j = 0$
- 12 **end for**
- 13 **until** converged;

$W = \Theta$

$\lambda = \delta^2$

predictor $D = p$

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

[Kevin Murphy's book]