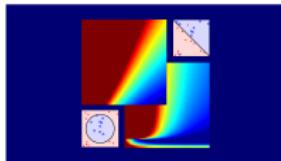


Machine Learning Foundations (機器學習基石)



Lecture 9: Linear Regression

Hsuan-Tien Lin (林軒田)

htlin@csie.ntu.edu.tw

Department of Computer Science
& Information Engineering

National Taiwan University
(國立台灣大學資訊工程系)



Roadmap

- ① When Can Machines Learn?
- ② Why Can Machines Learn?

Lecture 8: Noise and Error

learning can happen
with **target distribution $P(y|x)$** and **low E_{in} w.r.t. err**

- ③ **How** Can Machines Learn?

Lecture 9: Linear Regression

- Linear Regression Problem
- Linear Regression Algorithm
- Generalization Issue

- ④ How Can Machines Learn Better?

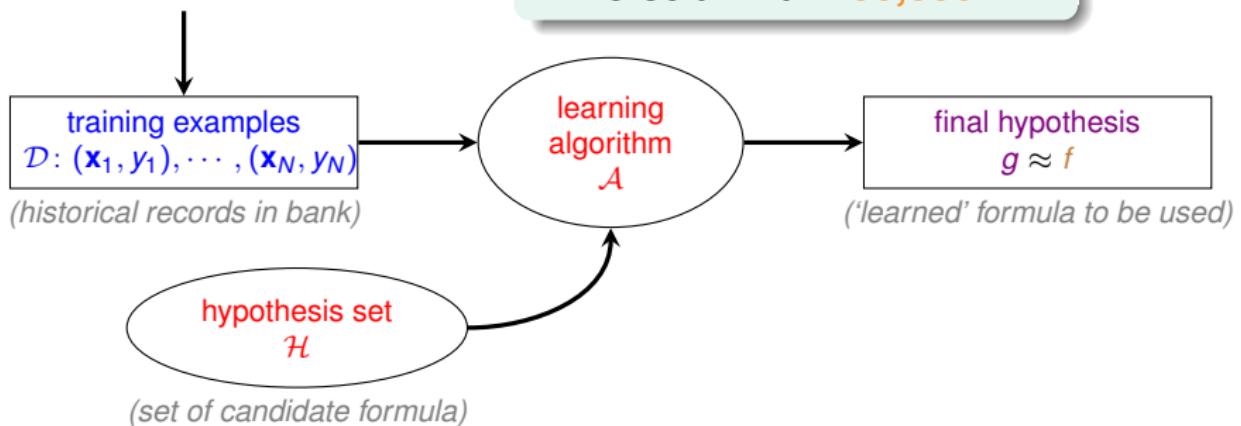
Credit Limit Problem

unknown target function
 $f: \mathcal{X} \rightarrow \mathcal{Y}$

(ideal credit **limit** formula)

age	23 years
gender	female
annual salary	NTD 1,000,000
year in residence	1 year
year in job	0.5 year
current debt	200,000

credit limit? **100,000**



$\mathcal{Y} = \mathbb{R}$: **regression**

Linear Regression Hypothesis

age	23 years
annual salary	NTD 1,000,000
year in job	0.5 year
current debt	200,000

- For $\mathbf{x} = (x_0, x_1, x_2, \dots, x_d)$ ‘features of customer’, approximate the **desired credit limit** with a **weighted sum**:

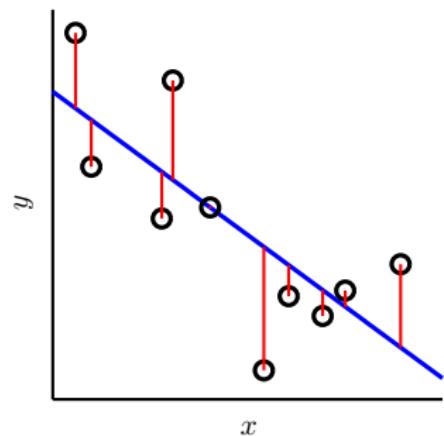
$$y \approx \sum_{i=0}^d w_i x_i$$

- linear regression hypothesis: $h(\mathbf{x}) = \mathbf{w}^T \mathbf{x}$

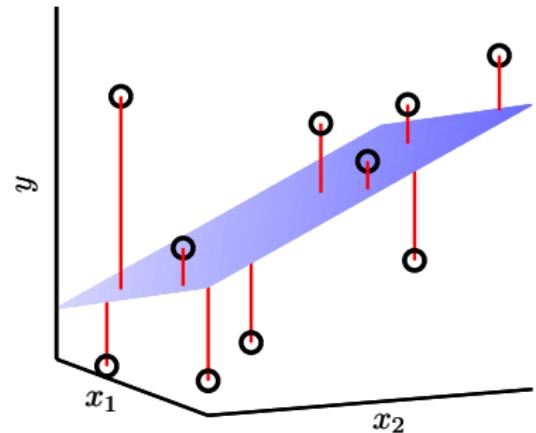
$h(\mathbf{x})$: like **perceptron**, but without the **sign**

Illustration of Linear Regression

$$\mathbf{x} = (x) \in \mathbb{R}$$



$$\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$$



linear regression:
find **lines/hyperplanes** with small **residuals**

The Error Measure

popular/historical error measure:

$$\text{squared error } \text{err}(\hat{y}, y) = (\hat{y} - y)^2$$

in-sample

$$E_{\text{in}}(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N (\underbrace{h(\mathbf{x}_n)}_{\mathbf{w}^T \mathbf{x}_n} - y_n)^2$$

out-of-sample

$$E_{\text{out}}(\mathbf{w}) = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P} (\mathbf{w}^T \mathbf{x} - y)^2$$

next: how to minimize $E_{\text{in}}(\mathbf{w})$?

Fun Time

Consider using linear regression hypothesis $h(\mathbf{x}) = \mathbf{w}^T \mathbf{x}$ to predict the credit limit of customers \mathbf{x} . Which feature below shall have a positive weight in a **good hypothesis** for the task?

- ① birth month
- ② monthly income
- ③ current debt
- ④ number of credit cards owned

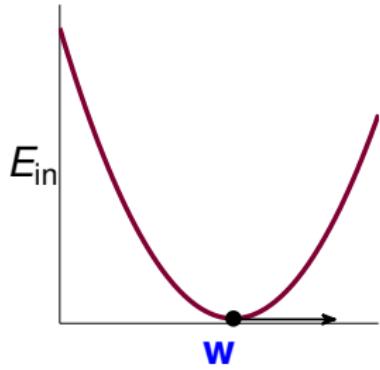
Reference Answer: ②

Customers with higher monthly income should naturally be given a higher credit limit, which is captured by the positive weight on the ‘monthly income’ feature.

Matrix Form of $E_{\text{in}}(\mathbf{w})$

$$\begin{aligned}
 E_{\text{in}}(\mathbf{w}) &= \frac{1}{N} \sum_{n=1}^N (\mathbf{w}^T \mathbf{x}_n - y_n)^2 = \frac{1}{N} \sum_{n=1}^N (\mathbf{x}_n^T \mathbf{w} - y_n)^2 \\
 &= \frac{1}{N} \left\| \begin{array}{c} \mathbf{x}_1^T \mathbf{w} - y_1 \\ \mathbf{x}_2^T \mathbf{w} - y_2 \\ \vdots \\ \mathbf{x}_N^T \mathbf{w} - y_N \end{array} \right\|^2 \\
 &= \frac{1}{N} \left\| \begin{bmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_N^T \end{bmatrix} \mathbf{w} - \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix} \right\|^2 \\
 &= \frac{1}{N} \| \underbrace{\mathbf{X}}_{N \times d+1} \underbrace{\mathbf{w}}_{d+1 \times 1} - \underbrace{\mathbf{y}}_{N \times 1} \|^2
 \end{aligned}$$

$$\min_{\mathbf{w}} E_{\text{in}}(\mathbf{w}) = \frac{1}{N} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2$$



- $E_{\text{in}}(\mathbf{w})$: continuous, differentiable, **convex**
- necessary condition of ‘best’ \mathbf{w}

$$\nabla E_{\text{in}}(\mathbf{w}) \equiv \begin{bmatrix} \frac{\partial E_{\text{in}}}{\partial w_0}(\mathbf{w}) \\ \frac{\partial E_{\text{in}}}{\partial w_1}(\mathbf{w}) \\ \vdots \\ \frac{\partial E_{\text{in}}}{\partial w_d}(\mathbf{w}) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

—not possible to ‘roll down’

task: find \mathbf{w}_{LIN} such that $\nabla E_{\text{in}}(\mathbf{w}_{\text{LIN}}) = \mathbf{0}$

The Gradient $\nabla E_{\text{in}}(\mathbf{w})$

$$E_{\text{in}}(\mathbf{w}) = \frac{1}{N} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2 = \frac{1}{N} \left(\mathbf{w}^T \mathbf{X}^T \mathbf{X} \mathbf{w} - 2\mathbf{w}^T \mathbf{X}^T \mathbf{y} + \mathbf{y}^T \mathbf{y} \right)$$

A b c

one \mathbf{w} only

$$E_{\text{in}}(\mathbf{w}) = \frac{1}{N} (a\mathbf{w}^2 - 2b\mathbf{w} + c)$$

$$\nabla E_{\text{in}}(\mathbf{w}) = \frac{1}{N} (2a\mathbf{w} - 2b)$$

simple! :-)

vector \mathbf{w}

$$E_{\text{in}}(\mathbf{w}) = \frac{1}{N} (\mathbf{w}^T \mathbf{A} \mathbf{w} - 2\mathbf{w}^T \mathbf{b} + c)$$

$$\nabla E_{\text{in}}(\mathbf{w}) = \frac{1}{N} (2\mathbf{A}\mathbf{w} - 2\mathbf{b})$$

similar (derived by definition)

$$\nabla E_{\text{in}}(\mathbf{w}) = \frac{2}{N} (\mathbf{X}^T \mathbf{X} \mathbf{w} - \mathbf{X}^T \mathbf{y})$$

Optimal Linear Regression Weights

task: find \mathbf{w}_{LIN} such that $\frac{2}{N} (\mathbf{X}^T \mathbf{X} \mathbf{w} - \mathbf{X}^T \mathbf{y}) = \nabla E_{\text{in}}(\mathbf{w}) = \mathbf{0}$

invertible $\mathbf{X}^T \mathbf{X}$

- **easy!** unique solution

$$\mathbf{w}_{\text{LIN}} = \underbrace{(\mathbf{X}^T \mathbf{X})^{-1}}_{\text{pseudo-inverse } \mathbf{x}^\dagger} \mathbf{X}^T \mathbf{y}$$

- often the case because $N \gg d + 1$

singular $\mathbf{X}^T \mathbf{X}$

- **many** optimal solutions
- one of the solutions

$$\mathbf{w}_{\text{LIN}} = \mathbf{X}^\dagger \mathbf{y}$$

by defining \mathbf{X}^\dagger in other ways

practical suggestion:

use **well-implemented \dagger routine**

instead of $(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$

for numerical stability when **almost-singular**

Linear Regression Algorithm

- ① from \mathcal{D} , construct input matrix \mathbf{X} and output vector \mathbf{y} by

$$\mathbf{X} = \underbrace{\begin{bmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_N^T \end{bmatrix}}_{N \times (d+1)} \quad \mathbf{y} = \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}}_{N \times 1}$$

- ② calculate pseudo-inverse $\underbrace{\mathbf{X}^\dagger}_{(d+1) \times N}$
- ③ return $\underbrace{\mathbf{w}_{\text{LIN}}}_{(d+1) \times 1} = \mathbf{X}^\dagger \mathbf{y}$

simple and efficient
with good † routine

Fun Time

After getting \mathbf{w}_{LIN} , we can calculate the predictions $\hat{\mathbf{y}}_n = \mathbf{w}_{\text{LIN}}^T \mathbf{x}_n$. If all $\hat{\mathbf{y}}_n$ are collected in a vector $\hat{\mathbf{y}}$ similar to how we form \mathbf{y} , what is the matrix formula of $\hat{\mathbf{y}}$?

- ① \mathbf{y}
- ② $\mathbf{X}\mathbf{X}^T \mathbf{y}$
- ③ $\mathbf{X}\mathbf{X}^\dagger \mathbf{y}$
- ④ $\mathbf{X}\mathbf{X}^\dagger \mathbf{X}\mathbf{X}^T \mathbf{y}$

Reference Answer: ③

Note that $\hat{\mathbf{y}} = \mathbf{X}\mathbf{w}_{\text{LIN}}$. Then, a simple substitution of \mathbf{w}_{LIN} reveals the answer.

Is Linear Regression a ‘Learning Algorithm’?

$$\mathbf{w}_{\text{LIN}} = \mathbf{X}^\dagger \mathbf{y}$$

No!

- analytic (**closed-form**) solution, ‘instantaneous’
- not improving E_{in} nor E_{out} iteratively

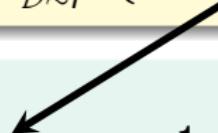
Yes!

- good E_{in} ?
yes, optimal!
- good E_{out} ?
yes, finite d_{VC} like perceptrons
- improving iteratively?
somewhat, within an iterative pseudo-inverse routine

if $E_{\text{out}}(\mathbf{w}_{\text{LIN}})$ is good, **learning ‘happened’!**

Benefit of Analytic Solution: 'Simpler-than-VC' Guarantee

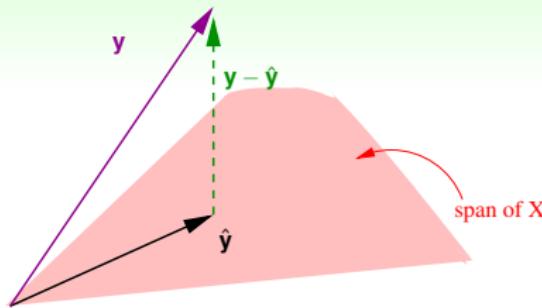
$$\overline{E_{\text{in}}} = \mathbb{E}_{\mathcal{D} \sim P_N} \left\{ E_{\text{in}}(\mathbf{w}_{\text{LIN}} \text{ w.r.t. } \mathcal{D}) \right\} \text{ to be shown} = \text{noise level} \cdot \left(1 - \frac{d+1}{N}\right)$$



$$\begin{aligned} E_{\text{in}}(\mathbf{w}_{\text{LIN}}) &= \frac{1}{N} \|\mathbf{y} - \underbrace{\hat{\mathbf{y}}}_{\text{predictions}}\|^2 = \frac{1}{N} \|\mathbf{y} - \underbrace{\mathbf{X} \mathbf{X}^\dagger \mathbf{y}}_{\mathbf{w}_{\text{LIN}}}\|^2 \\ &= \frac{1}{N} \|(\underbrace{\mathbf{I}}_{\text{identity}} - \mathbf{X} \mathbf{X}^\dagger) \mathbf{y}\|^2 \end{aligned}$$

call $\mathbf{X} \mathbf{X}^\dagger$ the **hat matrix H**
because it **puts \wedge on \mathbf{y}**

Geometric View of Hat Matrix

in \mathbb{R}^N

- $\hat{\mathbf{y}} = \mathbf{X}\mathbf{w}_{\text{LIN}}$ within the **span of X columns**
- $\mathbf{y} - \hat{\mathbf{y}}$ smallest: $\mathbf{y} - \hat{\mathbf{y}} \perp \text{span}$
- \mathbf{H} : project \mathbf{y} to $\hat{\mathbf{y}} \in \text{span}$
- $\mathbf{I} - \mathbf{H}$: transform \mathbf{y} to $\mathbf{y} - \hat{\mathbf{y}} \perp \text{span}$

claim: $\text{trace}(\mathbf{I} - \mathbf{H}) = N - (d + 1)$. **Why? :-)**

The Hat Matrix

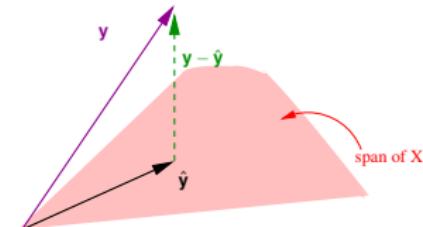
when $X^T X$ invertible, hat matrix $H = \mathbf{X} \mathbf{X}^\dagger = \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$

Claim: $H^{1126} = H$

proof (when $X^T X$ invertible):

$$\begin{aligned} H^{1126} &= HHH^{1124} \\ &= X(X^T X)^{-1} X^T X (X^T X)^{-1} X^T H^{1124} \\ &= X(X^T X)^{-1} (X^T X) (X^T X)^{-1} X^T H^{1124} \\ &= X(X^T X)^{-1} X^T H^{1124} \\ &= H^{1125} \end{aligned}$$

... and you know the rest



geometrically, **projecting 1126 times**
 \equiv projecting once

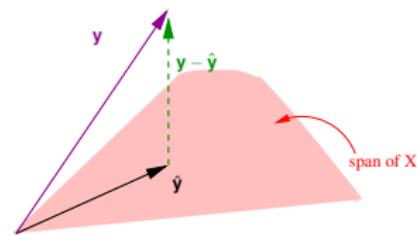
Trace of The Hat Matrix

when $X^T X$ invertible, hat matrix $H = \mathbf{X}\mathbf{X}^\dagger = \mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T$

Claim: $\text{trace}(H) = d + 1$
when $X^T X$ invertible

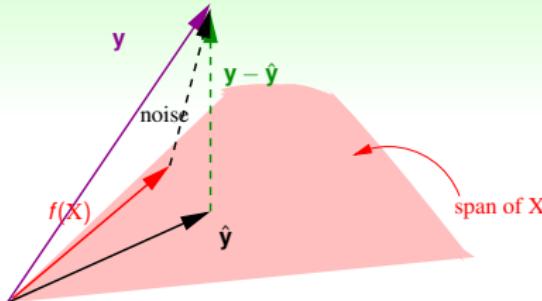
proof:

$$\begin{aligned}\text{trace}(H) &= \text{trace}(\mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T) \\ &= \text{trace}(\mathbf{X}^T\mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1}) \\ &= \text{trace}(\mathbf{I}_{d+1}) \\ &= d + 1\end{aligned}$$



geometrically, H projects to
a $(d + 1)$ -dimensional subspace

An Illustrative ‘Proof’, Corrected



- if \mathbf{y} comes from some ideal $f(\mathbf{X}) \in \text{span}$ plus **noise**
- **noise** with per-dimension ‘noise level’ σ^2 transformed by $\mathbf{I} - \mathbf{H}$ to be $\mathbf{y} - \hat{\mathbf{y}}$

$$\begin{aligned} E_{\text{in}}(\mathbf{w}_{\text{LIN}}) &= \frac{1}{N} \|\mathbf{y} - \hat{\mathbf{y}}\|^2 = \frac{1}{N} \|(\mathbf{I} - \mathbf{H})\mathbf{noise}\|^2 \\ &= \frac{1}{N} (N - (d + 1)) \sigma^2 \end{aligned}$$

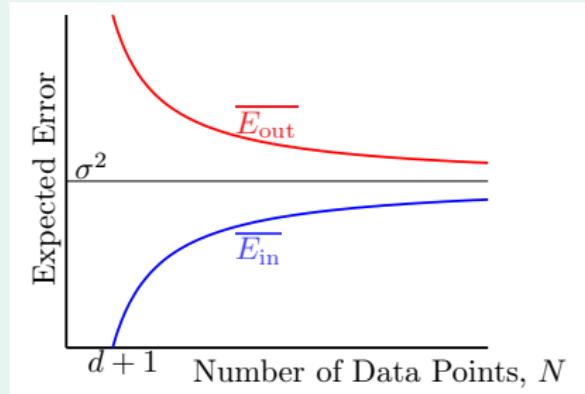
$$\overline{E_{\text{in}}} = \sigma^2 \cdot \left(1 - \frac{d+1}{N}\right)$$

$$\overline{E_{\text{out}}} = \sigma^2 \cdot \left(1 + \frac{d+1}{N}\right) \text{(complicated!)}$$

The Learning Curve

$$\overline{E}_{\text{out}} = \text{noise level} \cdot \left(1 + \frac{d+1}{N}\right)$$

$$\overline{E}_{\text{in}} = \text{noise level} \cdot \left(1 - \frac{d+1}{N}\right)$$



- both converge to σ^2 (**noise** level) for $N \rightarrow \infty$
- expected generalization error: $\frac{2(d+1)}{N}$
—similar to worst-case guarantee from VC

linear regression (LinReg):
learning 'happened'!

Fun Time

Which of the following property about H is not true?

- ① H is symmetric
- ② $H^2 = H$ (double projection = single one)
- ③ $(I - H)^2 = I - H$ (double residual transform = single one)
- ④ none of the above

Reference Answer: ④

You can conclude that ② and ③ are true by their physical meanings! :-)

Summary

① When Can Machines Learn?

② Why Can Machines Learn?

Lecture 8: Noise and Error

③ How Can Machines Learn?

Lecture 9: Linear Regression

- Linear Regression Problem
use hyperplanes to approximate real values
- Linear Regression Algorithm
analytic solution with pseudo-inverse
- Generalization Issue
 $E_{\text{out}} - E_{\text{in}} \approx \frac{2(d+1)}{N}$ on average

• next: binary classification, regression, and then?

④ How Can Machines Learn Better?