

# Data Mining Techniques

CS 6220 - Section 3 - Fall 2016

## Lecture 2: Regression

Jan-Willem van de Meent

(*credit*: Yijun Zhao, Marc Toussaint, Bishop)



# Administrativa

## Instructor

Jan-Willem van de Meent

*Email:* [j.vandemeent@northeastern.edu](mailto:j.vandemeent@northeastern.edu)

*Phone:* +1 617 373-7696

*Office Hours:* 478 WVH, Wed 1.30pm - 2.30pm




## Teaching Assistants

Yuan Zhong

*E-mail:* [yzhong@ccs.neu.edu](mailto:yzhong@ccs.neu.edu)

*Office Hours:* WVH 462, Wed 3pm - 5pm



Kamlendra Kumar

*E-mail:* [kumark@zimbra.ccs.neu.edu](mailto:kumark@zimbra.ccs.neu.edu)

*Office Hours:* WVH 462, Fri 3pm - 5pm



# Administrativa

## **Course Website**

<http://www.ccs.neu.edu/course/cs6220f16/sec3/>

## **Piazza**

<https://piazza.com/northeastern/fall2016/cs622003/home>

## **Project Guidelines (Vote next week)**

<http://www.ccs.neu.edu/course/cs6220f16/sec3/project/>

# *Question*

What would ***you*** like  
to get out of this course?


# Linear Regression

# Regression Examples

**Features**

**Continuous  
Value**

**$x$**



**$y$**

- {age, major, gender, race}  $\Rightarrow$  GPA
- {income, credit score, profession}  $\Rightarrow$  Loan Amount
- {college, major, GPA}  $\Rightarrow$  Future Income

# Example: Boston Housing Data

	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	MEDV
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2
5	0.02985	0.0	2.18	0.0	0.458	6.430	58.7	6.0622	3.0	222.0	18.7	394.12	5.21	28.7
6	0.08829	12.5	7.87	0.0	0.524	6.012	66.6	5.5605	5.0	311.0	15.2	395.60	12.43	22.9
7	0.14455	12.5	7.87	0.0	0.524	6.172	96.1	5.9505	5.0	311.0	15.2	396.90	19.15	27.1
8	0.21124	12.5	7.87	0.0	0.524	5.631	100.0	6.0821	5.0	311.0	15.2	386.63	29.93	16.5
9	0.17004	12.5	7.87	0.0	0.524	6.004	85.9	6.5921	5.0	311.0	15.2	386.71	17.10	18.9
10	0.22489	12.5	7.87	0.0	0.524	6.377	94.3	6.3467	5.0	311.0	15.2	392.52	20.45	15.0

**UC Irvine Machine Learning Repository**  
(*good source for project datasets*)

<https://archive.ics.uci.edu/ml/datasets/Housing>

# Example: Boston Housing Data

1. **CRIM**: per capita crime rate by town
2. **ZN**: proportion of residential land zoned for lots over 25,000 sq.ft.
3. **INDUS**: proportion of non-retail business acres per town
4. **CHAS**: Charles River dummy variable (= 1 if tract bounds river; 0 otherwise)
5. **NOX**: nitric oxides concentration (parts per 10 million)
6. **RM**: average number of rooms per dwelling
7. **AGE**: proportion of owner-occupied units built prior to 1940
8. **DIS**: weighted distances to five Boston employment centres
9. **RAD**: index of accessibility to radial highways
10. **TAX**: full-value property-tax rate per \$10,000
11. **PTRATIO**: pupil-teacher ratio by town
12. **B**:  $1000(B_k - 0.63)^2$  where  $B_k$  is the proportion of african americans by town
13. **LSTAT**: % lower status of the population
14. **MEDV**: Median value of owner-occupied homes in \$1000's



# Example: Boston Housing Data

	<b>CRIM</b>	<b>ZN</b>	<b>INDUS</b>	<b>CHAS</b>	<b>NOX</b>	<b>RM</b>	<b>AGE</b>	<b>DIS</b>	<b>RAD</b>	<b>TAX</b>	<b>PTRATIO</b>	<b>B</b>	<b>LSTAT</b>	<b>MEDV</b>
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2
5	0.02985	0.0	2.18	0.0	0.458	6.430	58.7	6.0622	3.0	222.0	18.7	394.12	5.21	28.7
6	0.08829	12.5	7.87	0.0	0.524	6.012	66.6	5.5605	5.0	311.0	15.2	395.60	12.43	22.9
7	0.14455	12.5	7.87	0.0	0.524	6.172	96.1	5.9505	5.0	311.0	15.2	396.90	19.15	27.1
8	0.21124	12.5	7.87	0.0	0.524	5.631	100.0	6.0821	5.0	311.0	15.2	386.63	29.93	16.5
9	0.17004	12.5	7.87	0.0	0.524	6.004	85.9	6.5921	5.0	311.0	15.2	386.71	17.10	18.9
10	0.22489	12.5	7.87	0.0	0.524	6.377	94.3	6.3467	5.0	311.0	15.2	392.52	20.45	15.0

**CRIM:** per capita crime rate by town

# Example: Boston Housing Data

	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	MEDV
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2
5	0.02985	0.0	2.18	0.0	0.458	6.430	58.7	6.0622	3.0	222.0	18.7	394.12	5.21	28.7
6	0.08829	12.5	7.87	0.0	0.524	6.012	66.6	5.5605	5.0	311.0	15.2	395.60	12.43	22.9
7	0.14455	12.5	7.87	0.0	0.524	6.172	96.1	5.9505	5.0	311.0	15.2	396.90	19.15	27.1
8	0.21124	12.5	7.87	0.0	0.524	5.631	100.0	6.0821	5.0	311.0	15.2	386.63	29.93	16.5
9	0.17004	12.5	7.87	0.0	0.524	6.004	85.9	6.5921	5.0	311.0	15.2	386.71	17.10	18.9
10	0.22489	12.5	7.87	0.0	0.524	6.377	94.3	6.3467	5.0	311.0	15.2	392.52	20.45	15.0

**CHAS:** Charles River dummy variable  
(= 1 if tract bounds river; 0 otherwise)

# Example: Boston Housing Data

	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	MEDV
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2
5	0.02985	0.0	2.18	0.0	0.458	6.430	58.7	6.0622	3.0	222.0	18.7	394.12	5.21	28.7
6	0.08829	12.5	7.87	0.0	0.524	6.012	66.6	5.5605	5.0	311.0	15.2	395.60	12.43	22.9
7	0.14455	12.5	7.87	0.0	0.524	6.172	96.1	5.9505	5.0	311.0	15.2	396.90	19.15	27.1
8	0.21124	12.5	7.87	0.0	0.524	5.631	100.0	6.0821	5.0	311.0	15.2	386.63	29.93	16.5
9	0.17004	12.5	7.87	0.0	0.524	6.004	85.9	6.5921	5.0	311.0	15.2	386.71	17.10	18.9
10	0.22489	12.5	7.87	0.0	0.524	6.377	94.3	6.3467	5.0	311.0	15.2	392.52	20.45	15.0

**MEDV:** Median value of owner-occupied homes in \$1000's

# Example: Boston Housing Data

	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	MEDV
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2
5	0.02985	0.0	2.18	0.0	0.458	6.430	58.7	6.0622	3.0	222.0	18.7	394.12	5.21	28.7
6	0.08829	12.5	7.87	0.0	0.524	6.012	66.6	5.5605	5.0	311.0	15.2	395.60	12.43	22.9
7	0.14455	12.5	7.87	0.0	0.524	6.172	96.1	5.9505	5.0	311.0	15.2	396.90	19.15	27.1
8	0.21124	12.5	7.87	0.0	0.524	5.631	100.0	6.0821	5.0	311.0	15.2	386.63	29.93	16.5
9	0.17004	12.5	7.87	0.0	0.524	6.004	85.9	6.5921	5.0	311.0	15.2	386.71	17.10	18.9
10	0.22489	12.5	7.87	0.0	0.524	6.377	94.3	6.3467	5.0	311.0	15.2	392.52	20.45	15.0

⋮  $N$  data points

⋮

$D$  features

# Regression: Problem Setup

Given  $N$  observations

$$\{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_N, y_N)\}$$

learn a function

$$y_i = f(\mathbf{x}_i) \quad \forall i = 1, 2, \dots, N$$

and for a new input  $\mathbf{x}^*$  predict

$$y^* = f(\mathbf{x}^*)$$

# Linear Regression

Assume  $f$  is a linear combination of  $D$  features

$$y = w_0 + w_1x_1 + \dots + w_Dx_D = \mathbf{w}^\top \mathbf{x}$$

where  $\mathbf{x}$  and  $\mathbf{w}$  are defined as

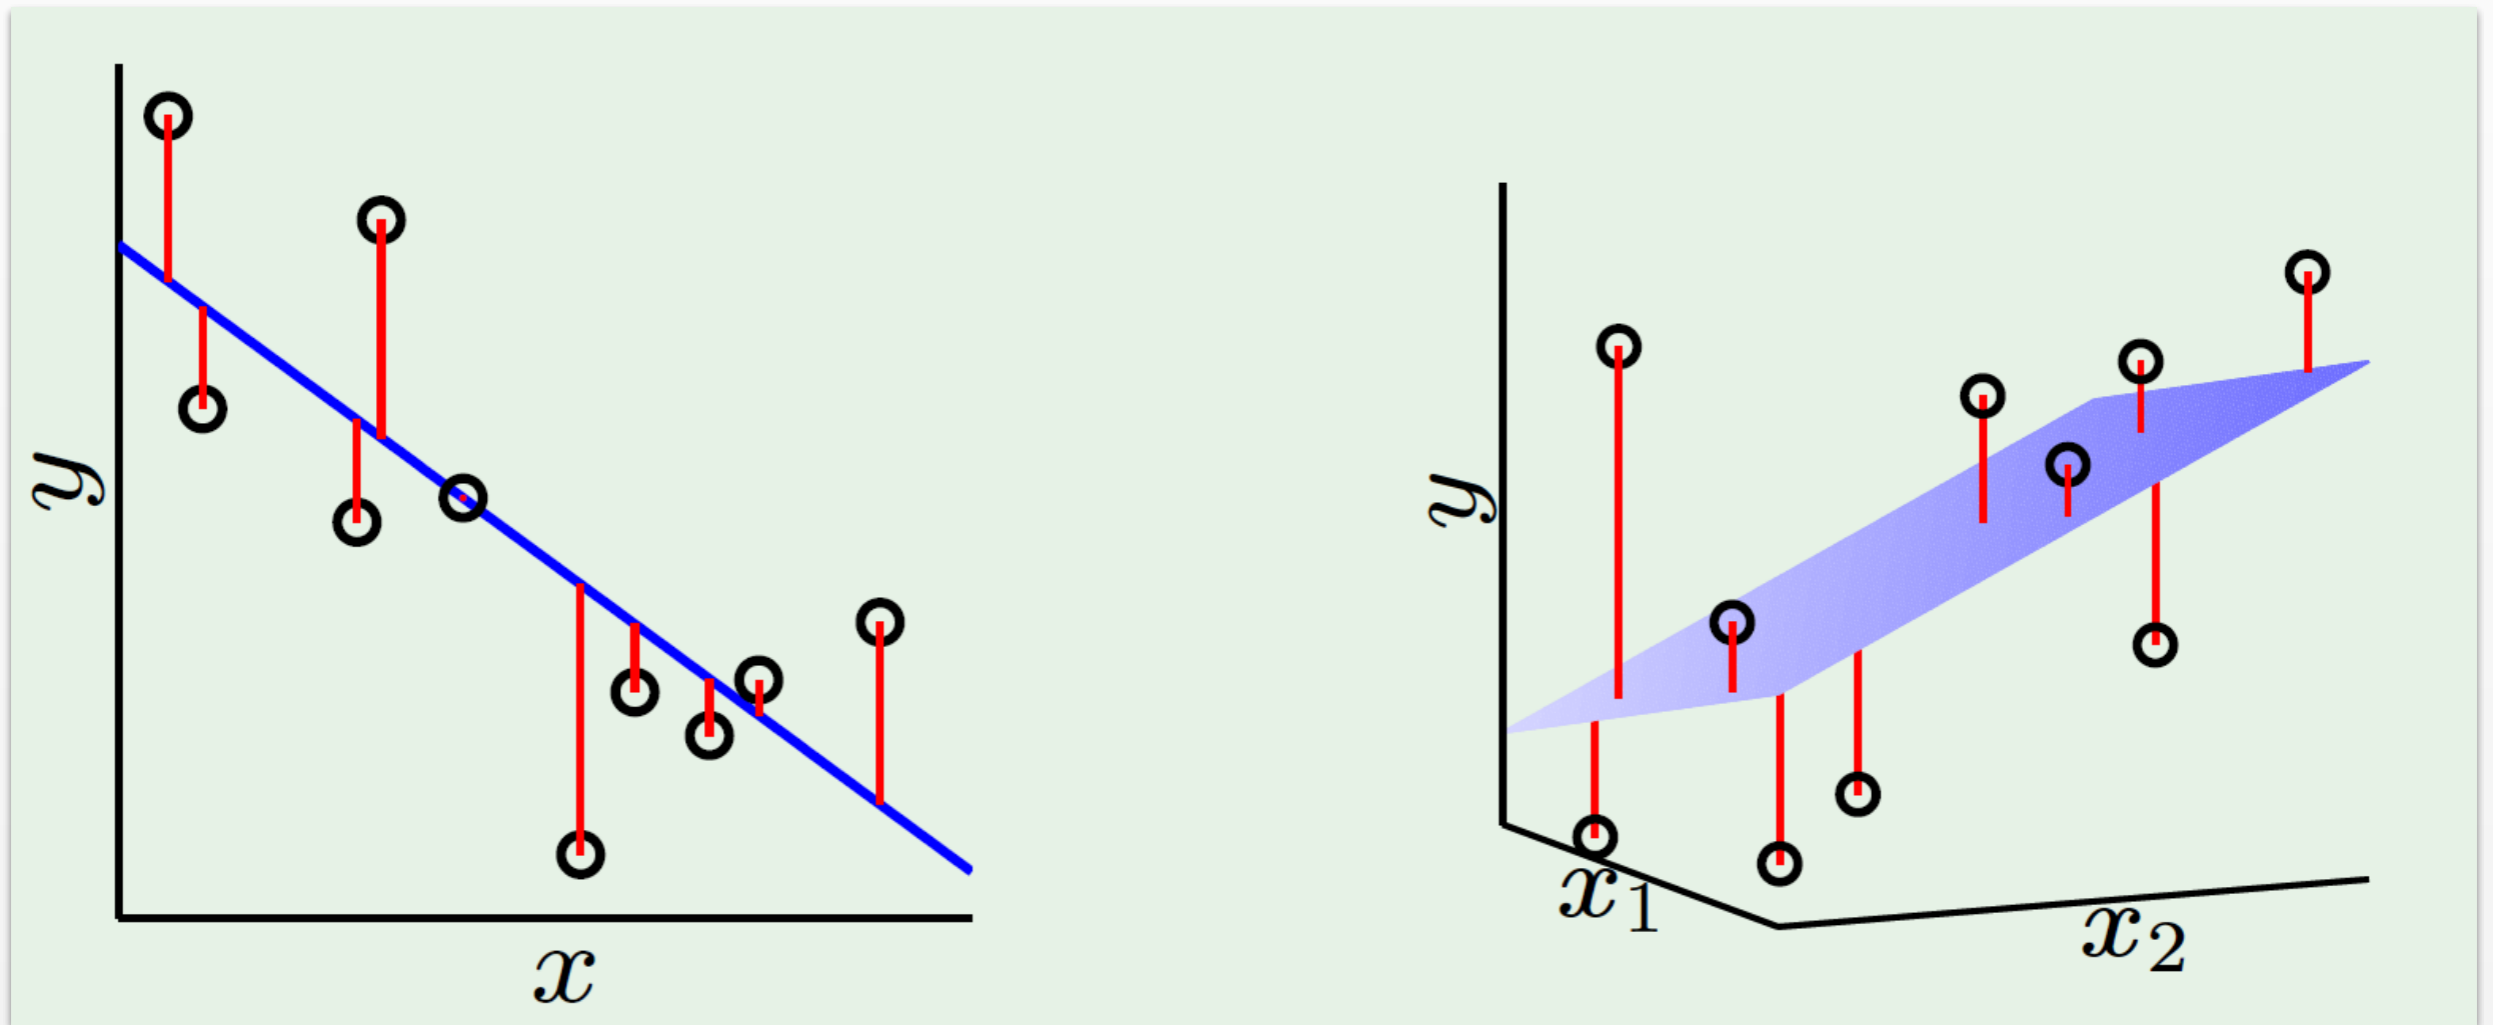
$$\mathbf{x} = (1, x_1, \dots, x_D) \quad \mathbf{w} = (w_0, w_1, \dots, w_D)$$

for  $N$  points we write

$$\mathbf{y} = \mathbf{X}\mathbf{w} \quad \mathbf{y} = (y_1, \dots, y_N) \quad \mathbf{X} = (\mathbf{x}_1^\top, \dots, \mathbf{x}_N^\top)$$

**Learning task:** Estimate  $\mathbf{w}$

# Linear Regression



# Error Measure

Mean Squared Error (MSE):

$$E(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N (\mathbf{w}^T \mathbf{x}_n - y_n)^2$$

$$= \frac{1}{N} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2$$

where

$$\mathbf{X} = \begin{bmatrix} \text{---} & \mathbf{x}_1^T & \text{---} \\ \text{---} & \mathbf{x}_2^T & \text{---} \\ & \dots & \\ \text{---} & \mathbf{x}_N^T & \text{---} \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} y_1^T \\ y_2^T \\ \dots \\ y_N^T \end{bmatrix}$$



# Minimizing the Error

$$E(\mathbf{w}) = \frac{1}{N} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2$$

$$\nabla E(\mathbf{w}) = \frac{2}{N} \mathbf{X}^T (\mathbf{X}\mathbf{w} - \mathbf{y}) = \mathbf{0}$$

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

$$\mathbf{w} = \mathbf{X}^\dagger \mathbf{y}$$

where  $\mathbf{X}^\dagger = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$  is the 'pseudo-inverse' of  $\mathbf{X}$

# Minimizing the Error

$$E(\mathbf{w}) = \frac{1}{N} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2$$

$$\nabla E(\mathbf{w}) = \frac{2}{N} \mathbf{X}^T (\mathbf{X}\mathbf{w} - \mathbf{y}) = \mathbf{0}$$

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

$$\mathbf{w} = \mathbf{X}^\dagger \mathbf{y}$$

where  $\mathbf{X}^\dagger = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$  is the 'pseudo-inverse' of  $\mathbf{X}$

**Matrix Cookbook (on course website)**

# Ordinary Least Squares


- Construct matrix  $\mathbf{X}$  and the vector  $\mathbf{y}$  from the dataset  $\{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_N, y_N)\}$  (each  $\mathbf{x}$  includes  $x_0 = 1$ ) as follows:

$$\mathbf{X} = \begin{bmatrix} \text{---} & \mathbf{x}_1^T & \text{---} \\ \text{---} & \mathbf{x}_2^T & \text{---} \\ & \dots & \\ \text{---} & \mathbf{x}_N^T & \text{---} \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} y_1^T \\ y_2^T \\ \dots \\ y_N^T \end{bmatrix}$$

- Compute  $\mathbf{X}^\dagger = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$
- Return  $\mathbf{w} = \mathbf{X}^\dagger \mathbf{y}$

# Gradient Descent

countours :  $E(w)$



# Least Mean Squares

(a.k.a. gradient descent)

- Initialize the  $\mathbf{w}(0)$  for time  $t = 0$
- for  $t = 0, 1, 2, \dots$  do
- Compute the gradient  $\mathbf{g}_t = \nabla E(\mathbf{w}(t))$
- Set the direction to move,  $\mathbf{v}_t = -\mathbf{g}_t$
- Update  $\mathbf{w}(t + 1) = \mathbf{w}(t) + \eta \mathbf{v}_t$
- Iterate until it is time to stop
- Return the final weights  $\mathbf{w}$

# *Question*

When would you want to use OLS, when LMS?

# Computational Complexity

## Ordinary least squares (OLS)

$$\mathbf{w} = (\mathbf{X}^\top \mathbf{X})^{-1} (\mathbf{X}^\top \mathbf{y})$$

$$(\mathbf{X}^\top \mathbf{y}) \quad O(DN)$$

$$(\mathbf{X}^\top \mathbf{X}) \quad O(D^2N)$$

$$(\mathbf{X}^\top \mathbf{X})^{-1} \quad O(D^3)$$

## Least Mean Squares (LMS)

$$\nabla E(\mathbf{w}) = \frac{2}{N} \mathbf{X}^\top (\mathbf{X} \mathbf{w} - \mathbf{y})$$

$$\mathbf{X} \mathbf{w} \quad O(DN)$$

$$\mathbf{X}(\mathbf{w} - \mathbf{y}) \quad O(N)$$

$$\mathbf{X}^\top (\mathbf{X} \mathbf{w} - \mathbf{y}) \quad O(DN)$$

# Computational Complexity

**Ordinary least squares (OLS)**

$$\mathbf{w} = (\mathbf{X}^\top \mathbf{X})^{-1} (\mathbf{X}^\top \mathbf{y})$$

$$(\mathbf{X}^\top \mathbf{y}) \quad O(DN)$$

$$(\mathbf{X}^\top \mathbf{X}) \quad O(D^2N)$$

$$(\mathbf{X}^\top \mathbf{X})^{-1} \quad O(D^3)$$

**Least Mean Squares (LMS)**

$$\nabla E(\mathbf{w}) = \frac{2}{N} \mathbf{X}^\top (\mathbf{X} \mathbf{w} - \mathbf{y})$$

$$\mathbf{X} \mathbf{w} \quad O(DN)$$


$$\mathbf{X} (\mathbf{w} - \mathbf{y}) \quad O(N)$$

$$\mathbf{X}^\top (\mathbf{X} \mathbf{w} - \mathbf{y}) \quad O(DN)$$

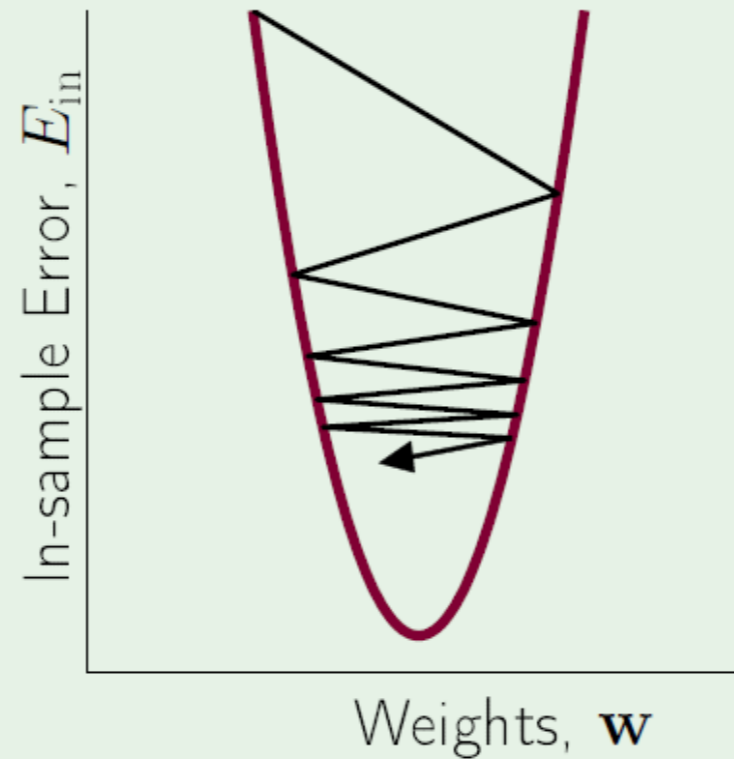
**OLS is expensive when D is large**




# Effect of step size



$\eta$  too small




$\eta$  too large



variable  $\eta$  – just right

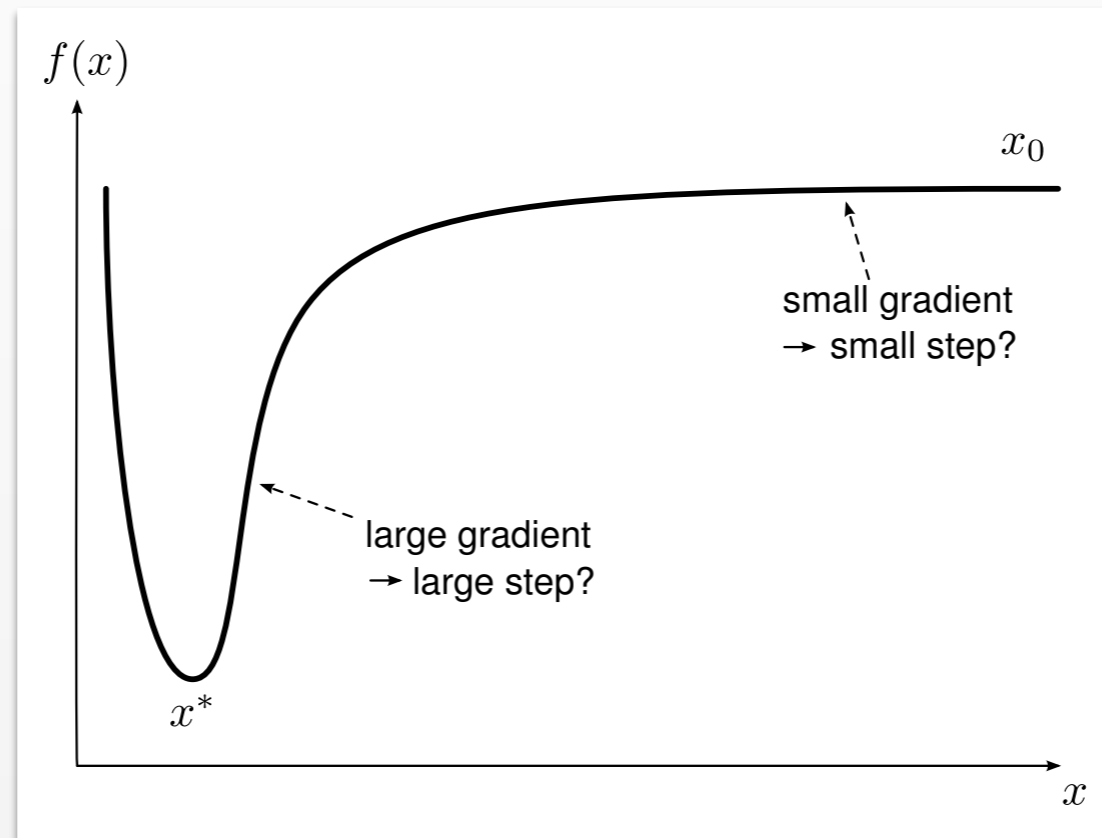
# Choosing Stepsize

Set step size proportional to  $\nabla f(x)$  ?



# Choosing Stepsize

Set step size proportional to  $\nabla f(x)$  ?



*Two commonly used techniques*

1. Stepsize adaptation
2. Line search

# Stepsize Adaptation

---

**Input:** initial  $x \in \mathbb{R}^n$ , functions  $f(x)$  and  $\nabla f(x)$ , initial stepsize  $\alpha$ , tolerance  $\theta$

**Output:**  $x$

```
1: repeat
2:    $y \leftarrow x - \alpha \frac{\nabla f(x)}{|\nabla f(x)|}$ 
3:   if [ then step is accepted]  $f(y) \leq f(x)$ 
4:      $x \leftarrow y$ 
5:      $\alpha \leftarrow 1.2\alpha$  // increase stepsize
6:   else [step is rejected]
7:      $\alpha \leftarrow 0.5\alpha$  // decrease stepsize
8:   end if
9: until  $|y - x| < \theta$  [perhaps for 10 iterations in sequence]
```

---

(“magic numbers”)

# Second Order Methods

Compute **Hessian** matrix of second derivatives

$$\mathbf{H} = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \frac{\partial^2 f}{\partial x_1 \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \frac{\partial^2 f}{\partial x_2 \partial x_1} & \frac{\partial^2 f}{\partial x_2^2} & \cdots & \frac{\partial^2 f}{\partial x_2 \partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \frac{\partial^2 f}{\partial x_n \partial x_2} & \cdots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix}$$

# Second Order Methods

Broyden-Fletcher-Goldfarb-Shanno (BFGS) method:

---

**Input:** initial  $x \in \mathbb{R}^n$ , functions  $f(x)$ ,  $\nabla f(x)$ , tolerance  $\theta$

**Output:**  $x$

- 1: initialize  $H^{-1} = \mathbf{I}_n$
  - 2: **repeat**
  - 3:     compute  $\Delta = -H^{-1}\nabla f(x)$
  - 4:     perform a line search  $\min_{\alpha} f(x + \alpha\Delta)$
  - 5:      $\Delta \leftarrow \alpha\Delta$
  - 6:      $y \leftarrow \nabla f(x + \Delta) - \nabla f(x)$
  - 7:      $x \leftarrow x + \Delta$
  - 8:     update  $H^{-1} \leftarrow \left(\mathbf{I} - \frac{y\Delta^{\top}}{\Delta^{\top}y}\right)^{\top} H^{-1} \left(\mathbf{I} - \frac{y\Delta^{\top}}{\Delta^{\top}y}\right) + \frac{\Delta\Delta^{\top}}{\Delta^{\top}y}$
  - 9: **until**  $\|\Delta\|_{\infty} < \theta$
- 

Memory-limited version: L-BFGS

# Stochastic Gradient Descent

What if  $N$  is really large?

Batch gradient descent (evaluates all data)

$$\mathbf{w}_t = \mathbf{w}_{t-1} - \alpha_t \nabla_{\mathbf{w}} E(\mathbf{y}; \mathbf{w})|_{\mathbf{w}=\mathbf{w}_{t-1}}$$

Minibatch gradient descent (evaluates subset)

$$\mathbf{w}_t = \mathbf{w}_{t-1} - \alpha_t \nabla_{\mathbf{w}} E(\mathbf{y}_t; \mathbf{w})|_{\mathbf{w}=\mathbf{w}_{t-1}} \quad \mathbf{y}_t \subset \mathbf{y}$$

Converges under Robbins-Monro conditions

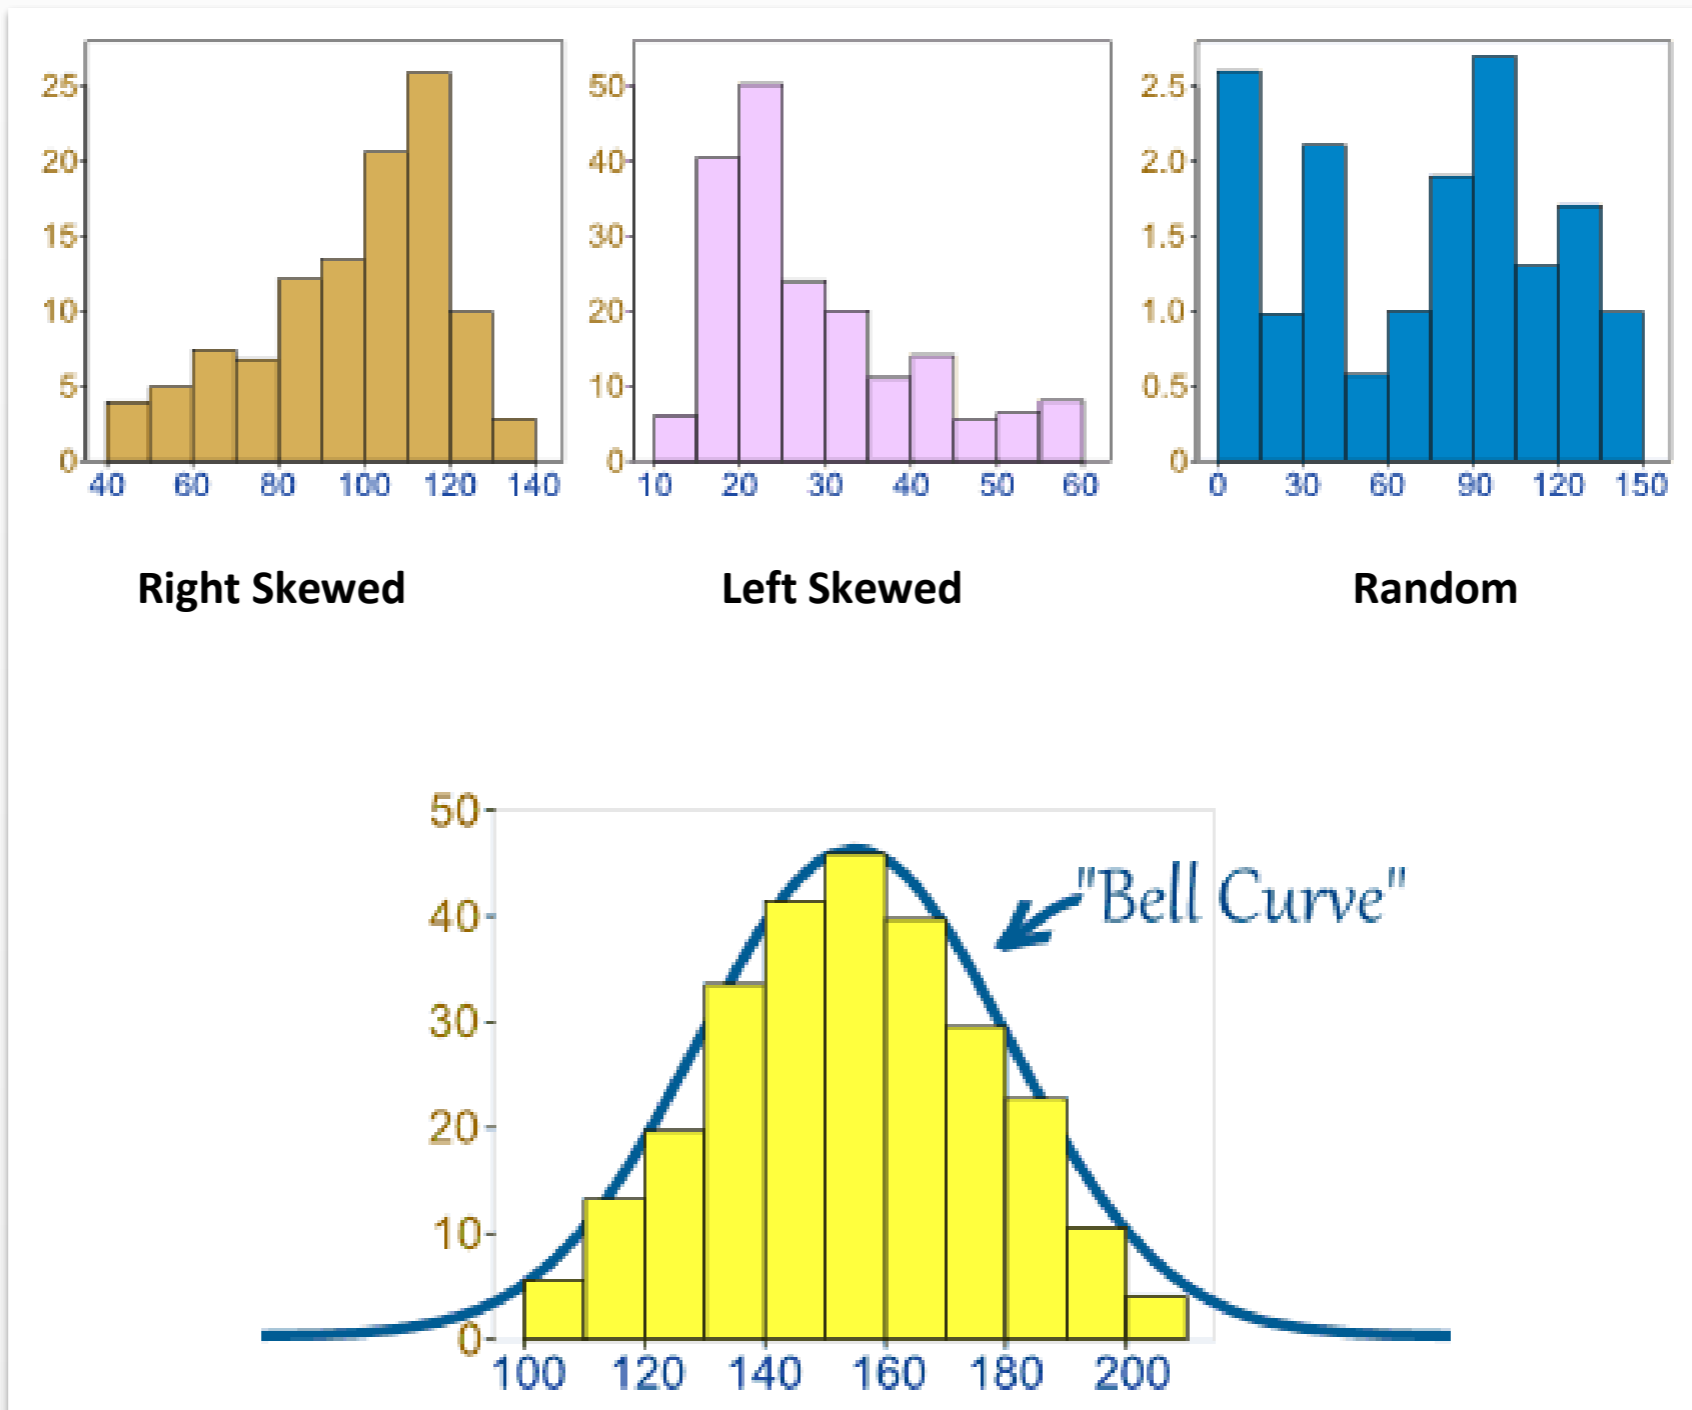
$$\sum_{t=1}^{\infty} \alpha_t = \infty \quad \sum_{t=1}^{\infty} \alpha_t^2 < \infty$$

$$\alpha_t = \frac{\alpha_0}{(\tau+t)^{\kappa}}$$


# Probabilistic Interpretation



# Normal Distribution






# Normal Distribution



Density:  $f(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp^{-\frac{1}{2}(x-\mu)^2/\sigma^2}$

# Central Limit Theorem




If  $y_1, \dots, y_n$  are

1. Independent identically distributed (i.i.d.)
2. Have finite variance  $0 < \sigma_y^2 < \infty$

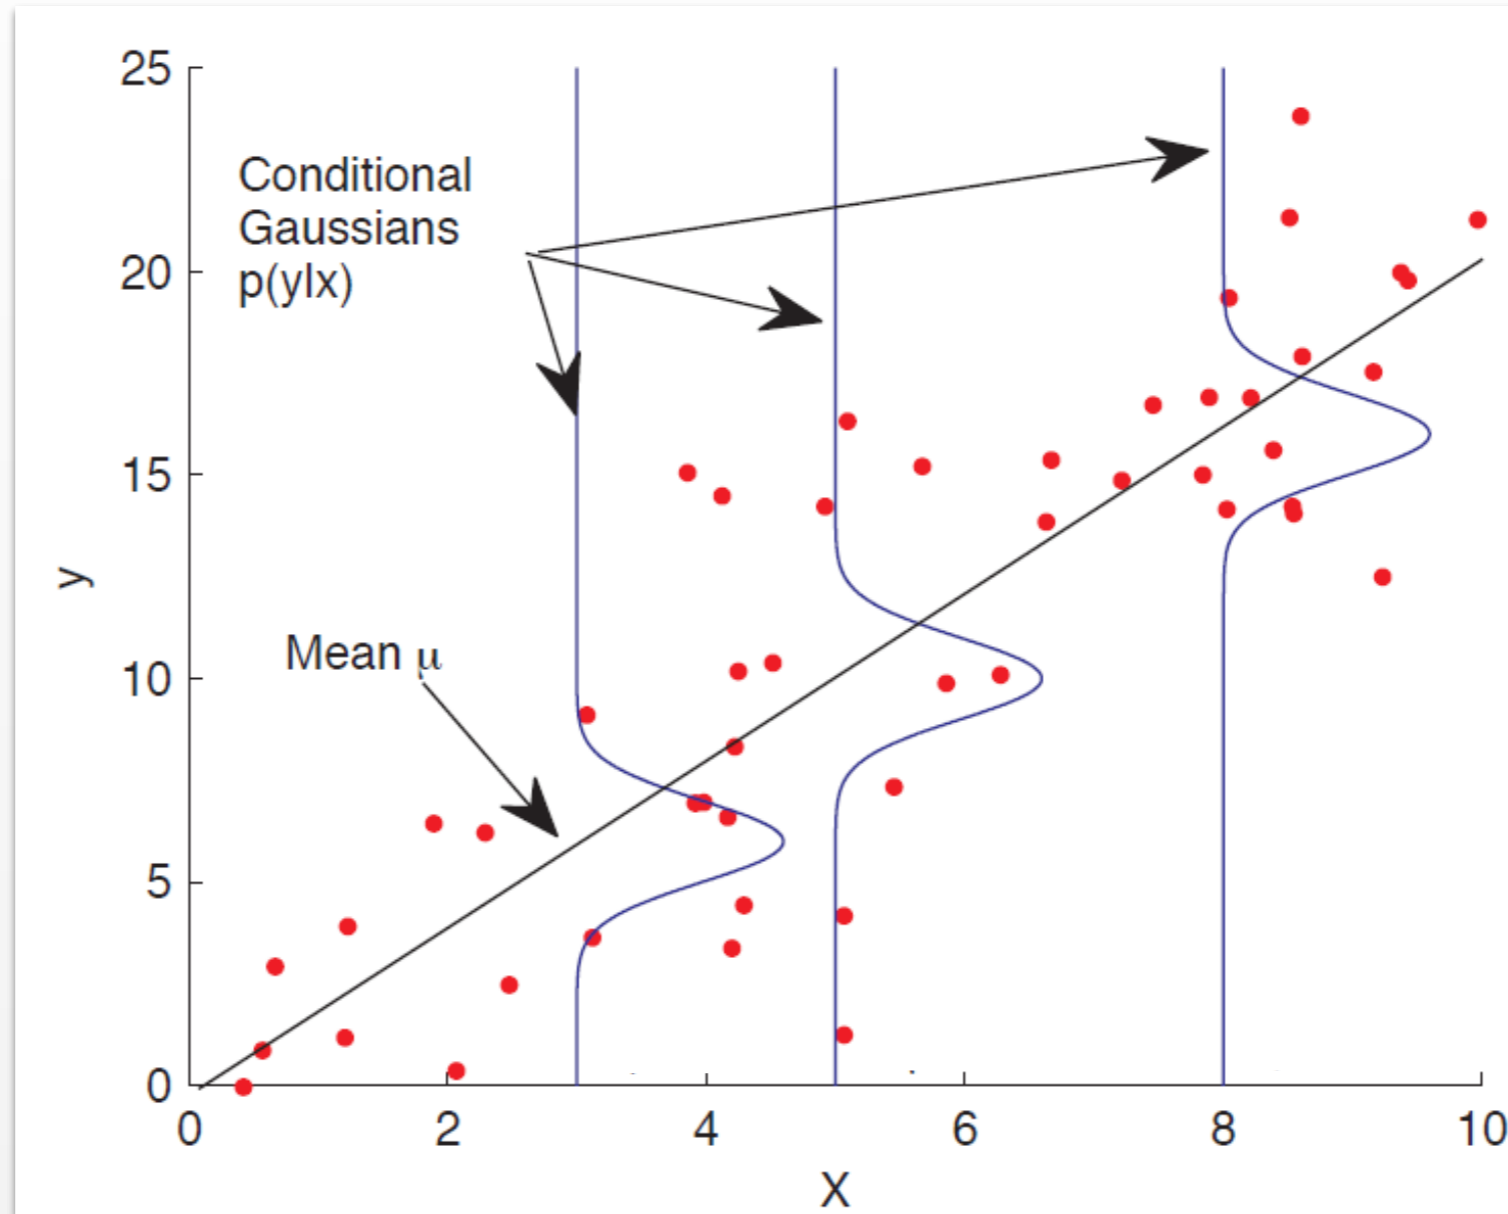
$$f(\bar{y}) = \text{Normal}(\bar{y} ; \mu_y, \sigma_y^2/N) \quad \bar{y} = \frac{1}{N} \sum_{n=1}^N y_n$$

# Multivariate Normal



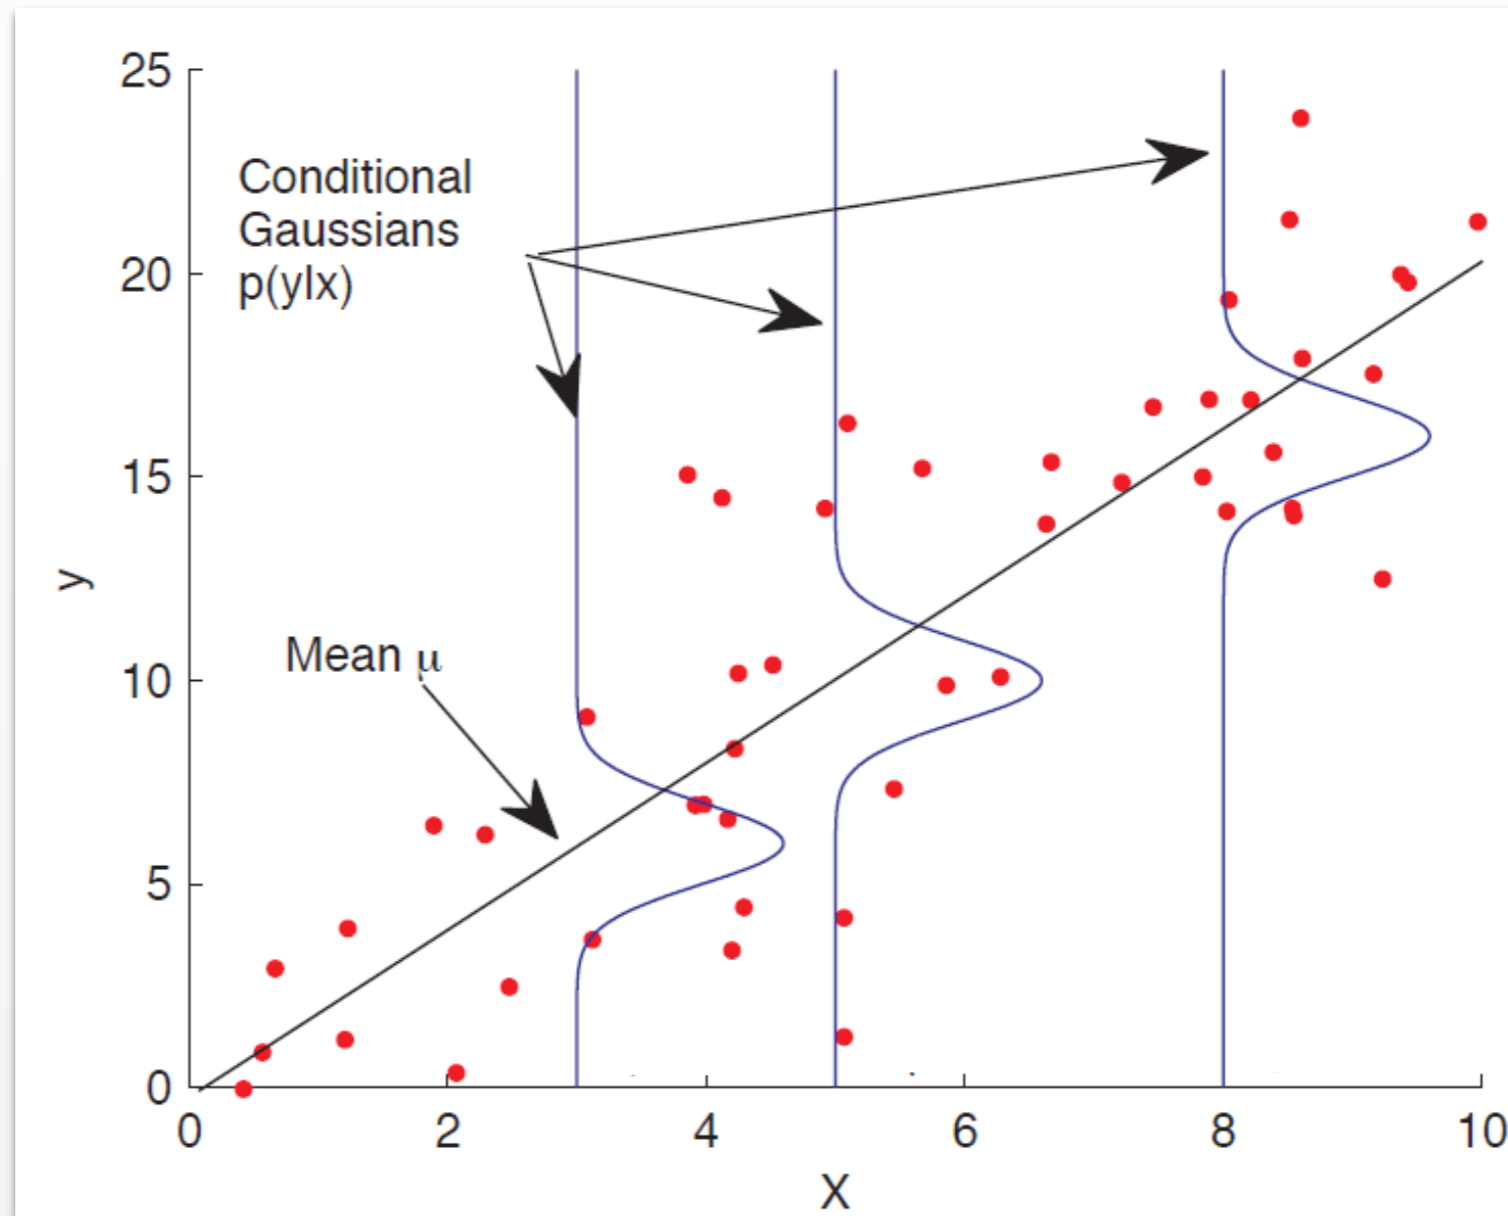
Density: 
$$f(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{(2\pi)^D |\boldsymbol{\Sigma}|}} \exp^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}$$

# Regression: Probabilistic Interpretation



$$y_n = ax_n + b + \sigma\epsilon_n \quad \epsilon \sim \text{Normal}(0, 1)$$

# Regression: Probabilistic Interpretation



$$\mu_n = \mathbf{w}^\top \mathbf{x}_n \quad y_n \sim \text{Normal}(\mu_n, \Sigma)$$

# Regression: Probabilistic Interpretation

Joint probability of  $N$  independent data points

$$\begin{aligned} p(y_1, \dots, y_N) &= \prod_{n=1}^N p(y_n) \\ &= \frac{1}{\sqrt{2\pi\sigma^2}^N} \prod_{n=1}^N \exp^{-\frac{1}{2}(x-\mu)^2/\sigma^2} \\ &= \frac{1}{\sqrt{2\pi\sigma^2}^N} \exp^{-\frac{1}{2} \sum_{n=1}^N (x-\mu)^2/\sigma^2} \end{aligned}$$

# Regression: Probabilistic Interpretation

**Log** joint probability of  $N$  independent data points

$$\begin{aligned}\log p(y_1, \dots, y_N) &= \sum_{n=1}^N \log p(y_n) \\ &= -\frac{1}{2} \left[ N \log(2\pi\sigma^2) + \sum_{n=1}^N \frac{(y_n - \mu_n)^2}{\sigma^2} \right]\end{aligned}$$



# Regression: Probabilistic Interpretation

**Log** joint probability of  $N$  independent data points

$$\begin{aligned}\log p(y_1, \dots, y_N) &= \sum_{n=1}^N \log p(y_n) \\ &= -\frac{1}{2} \left[ N \log(2\pi\sigma^2) + \sum_{n=1}^N \frac{(y_n - \mathbf{w}^\top \mathbf{x}_n)^2}{\sigma^2} \right]\end{aligned}$$

# Regression: Probabilistic Interpretation

**Log** joint probability of  $N$  independent data points

$$\begin{aligned}\log p(y_1, \dots, y_N) &= \sum_{n=1}^N \log p(y_n) \\ &= -\frac{1}{2} \left[ N \log(2\pi\sigma^2) + \sum_{n=1}^N \frac{(y_n - \mathbf{w}^\top \mathbf{x}_n)^2}{\sigma^2} \right] \\ &= -\frac{N}{2} [\text{const} + E(\mathbf{w})]\end{aligned}$$

# Regression: Probabilistic Interpretation

**Log** joint probability of  $N$  independent data points

$$\begin{aligned}\log p(y_1, \dots, y_N) &= \sum_{n=1}^N \log p(y_n) \\ &= -\frac{1}{2} \left[ N \log(2\pi\sigma^2) + \sum_{n=1}^N \frac{(y_n - \mathbf{w}^\top \mathbf{x}_n)^2}{\sigma^2} \right] \\ &= -\frac{N}{2} [\text{const} + E(\mathbf{w})]\end{aligned}$$

$$\underset{\mathbf{w}}{\operatorname{argmax}} p(y_1, \dots, y_N; \mathbf{w}) = \underset{\mathbf{w}}{\operatorname{argmin}} E(\mathbf{w})$$

*Maximum  
Likelihood*

# Basis function regression

Linear regression

$$y = w_0 + w_1 \mathbf{x}_1 + \dots + w_D \mathbf{x}_D = \mathbf{w}^T \mathbf{x}$$





Basis function regression

$$y = w_0 + w_1 \phi_1(\mathbf{x}) + \dots + w_D \phi_D(\mathbf{x})$$

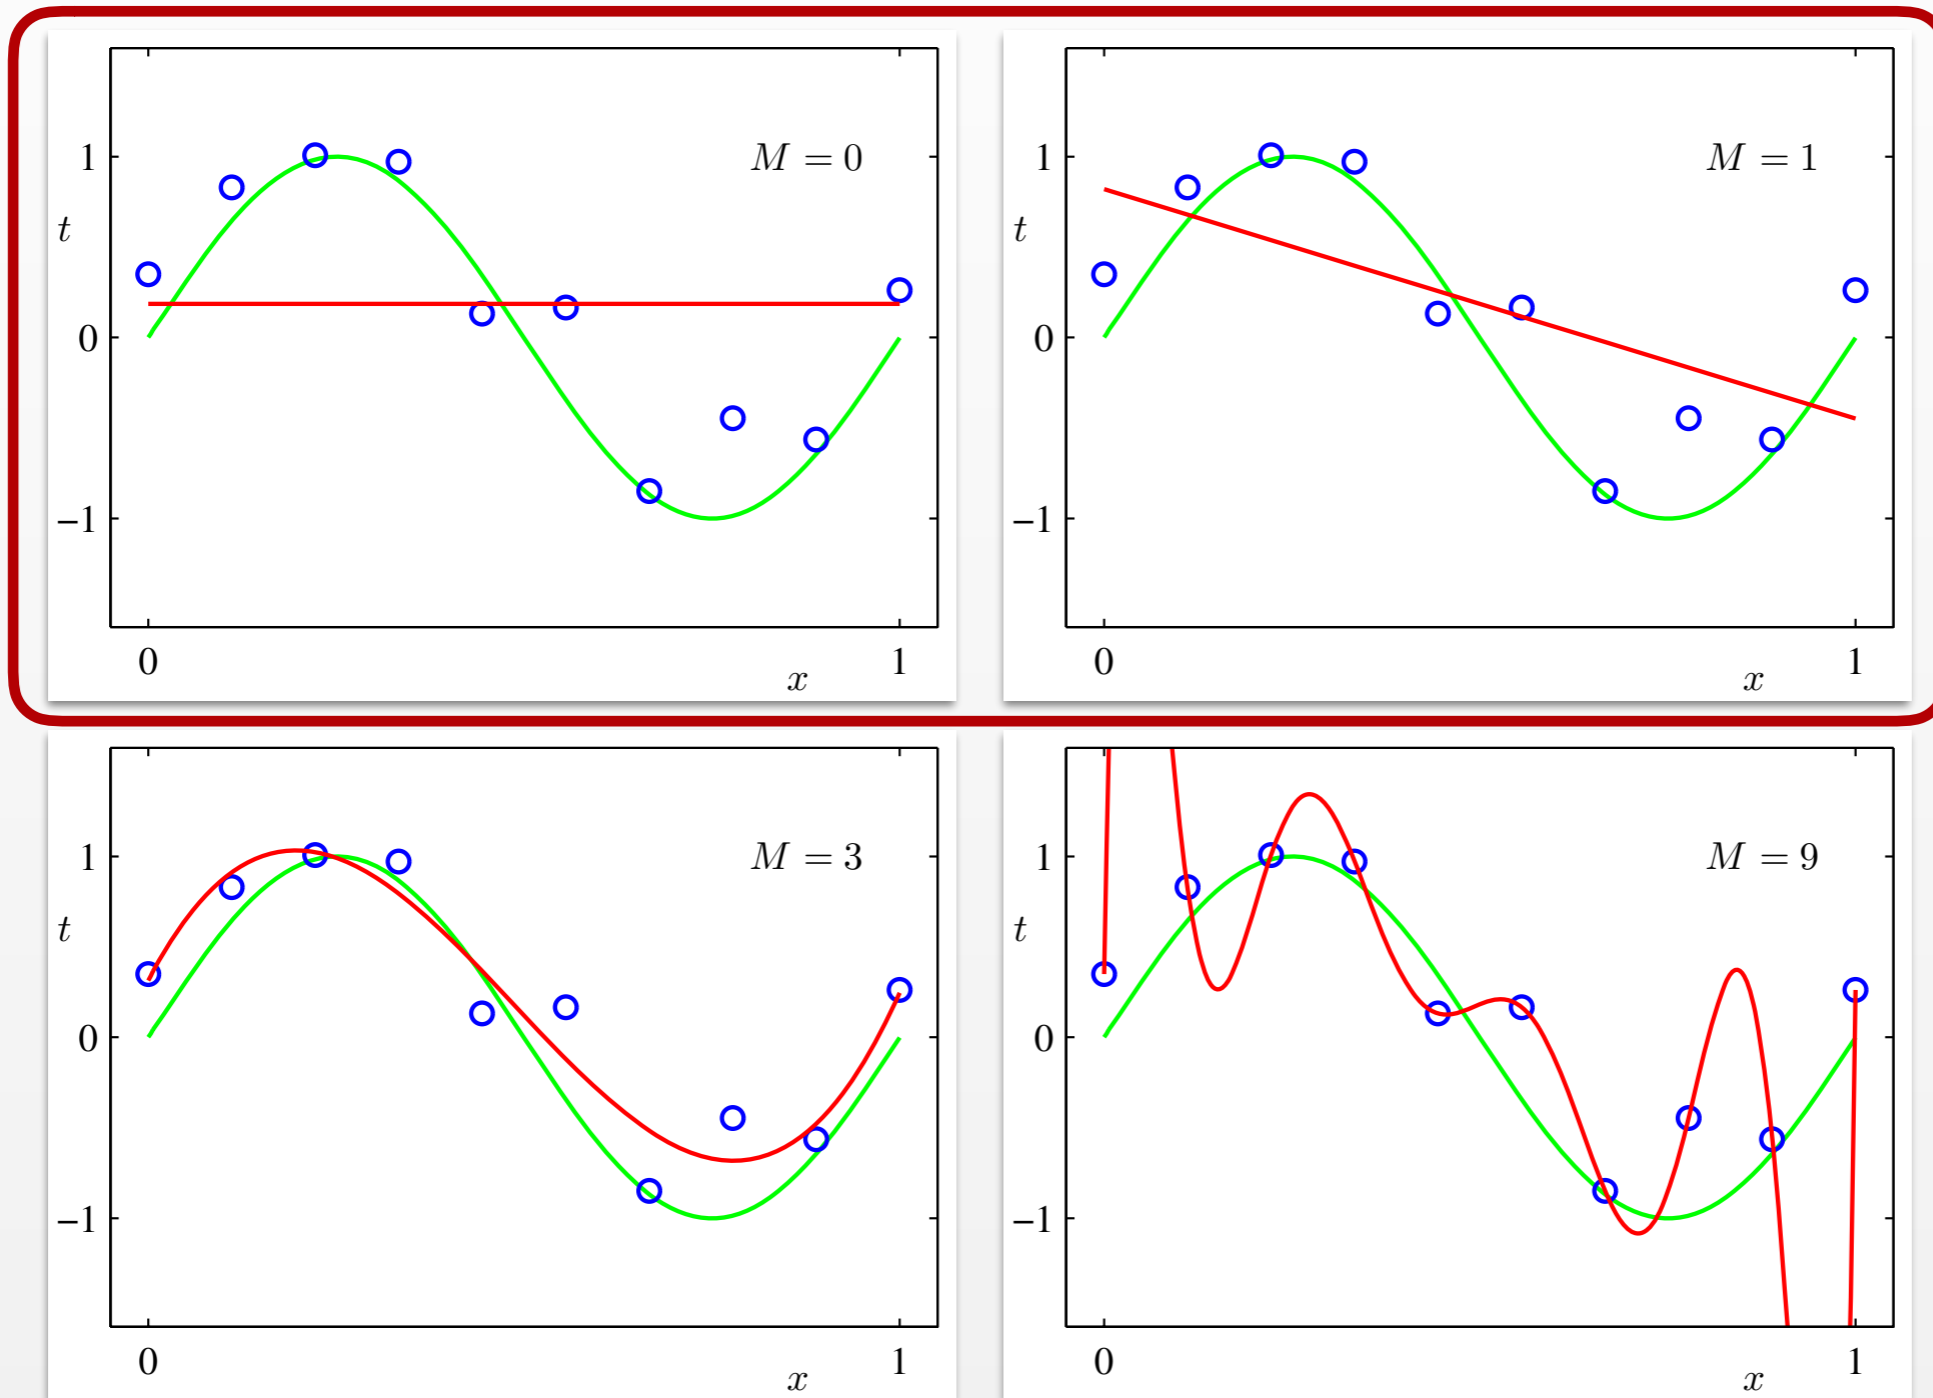
Polynomial regression

$$\mathbf{x}_d := \phi_d(\mathbf{x}) \quad \phi_d(\mathbf{x}) := \mathbf{x}^d$$

# Polynomial Regression







# Polynomial Regression



Underfit

# Polynomial Regression



Overfit

# Regularization

$L_2$  regularization (ridge regression) minimizes:

$$E(\mathbf{w}) = \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2 + \lambda \|\mathbf{w}\|^2$$

where  $\lambda \geq 0$  and  $\|\mathbf{w}\|^2 = \mathbf{w}^T \mathbf{w}$





$L_1$  regularization (LASSO) minimizes:

$$E(\mathbf{w}) = \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2 + \lambda |\mathbf{w}|_1$$

where  $\lambda \geq 0$  and  $|\mathbf{w}|_1 = \sum_{i=1}^D |\omega_i|$



# Regularization



# Regularization

$L_2$ : closed form solution

$$\mathbf{w} = (\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I})^{-1} \mathbf{X}^T \mathbf{y}$$

$L_1$ : No closed form solution. Use quadratic programming:

$$\text{minimize } \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2 \quad \text{s.t.} \quad \|\mathbf{w}\|_1 \leq s$$

# Review: Bias-Variance Trade-off

Maximum likelihood estimator

$$\hat{f} := \operatorname{argmax}_{\tilde{f}} p(\mathbf{y} | \tilde{f})$$

Bias-variance decomposition

(*expected value over possible data points*)

$$\mathbb{E}[(y - \hat{f}(\mathbf{x}))^2] = \operatorname{Bias}[\hat{f}(\mathbf{x})]^2 + \operatorname{Var}[\hat{f}(\mathbf{x})] + \sigma^2$$

$$\operatorname{Bias}[\hat{f}(\mathbf{x})] = \mathbb{E}[\hat{f}(\mathbf{x}) - f(\mathbf{x})]$$


$$\operatorname{Var}[\hat{f}(\mathbf{x})] = \mathbb{E}[\hat{f}(\mathbf{x})^2] - \mathbb{E}[\hat{f}(\mathbf{x})]^2$$

$$\sigma^2 = \mathbb{E}[y^2] - \mathbb{E}[f(\mathbf{x})]^2$$


# Bias-Variance Trade-off

Often: low bias  $\Rightarrow$  high variance  
low variance  $\Rightarrow$  high bias

Trade-off:




# K-fold Cross-Validation




1. **Divide** dataset into K “folds”
2. **Train** on all except  $k$ -th fold
3. **Test** on  $k$ -th fold
4. **Minimize** test error w.r.t.  $\lambda$

# K-fold Cross-Validation




- *Choices for K: 5, 10, N (leave-one-out)*
- Cost of computation:  $K \times$  number of  $\lambda$


# Learning Curve



# Learning Curve



Simple Model



Complex Model



# Loss Functions

squared loss:	$\frac{1}{2}(\mathbf{w}^\top \mathbf{x} - y)^2$	$y \in \mathbb{R}$
logistic loss:	$\log(1 + \exp(-y\mathbf{w}^\top \mathbf{x}))$	$y \in \{-1, +1\}$
hinge loss:	$\max\{0, 1 - y\mathbf{w}^\top \mathbf{x}\}$	$y \in \{-1, +1\}$