

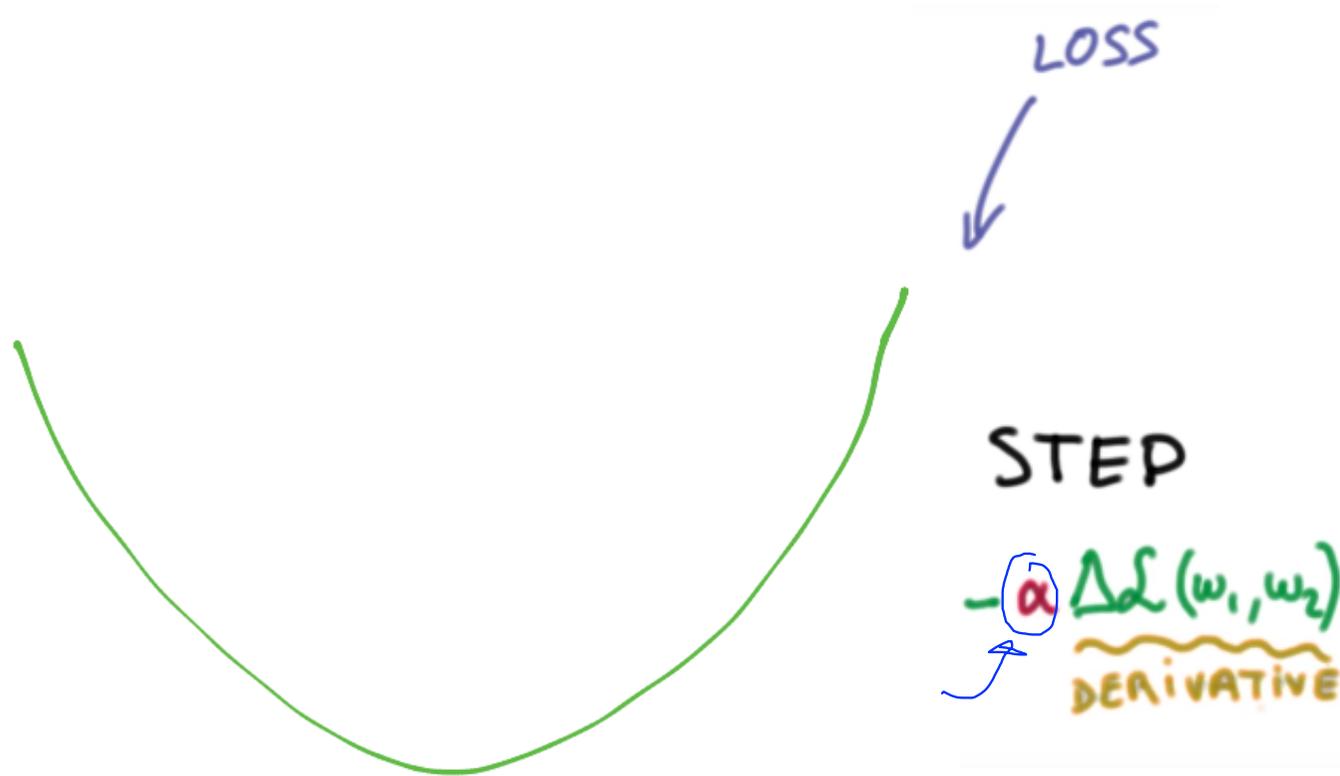
# Lecture 7-1

Application & Tips:

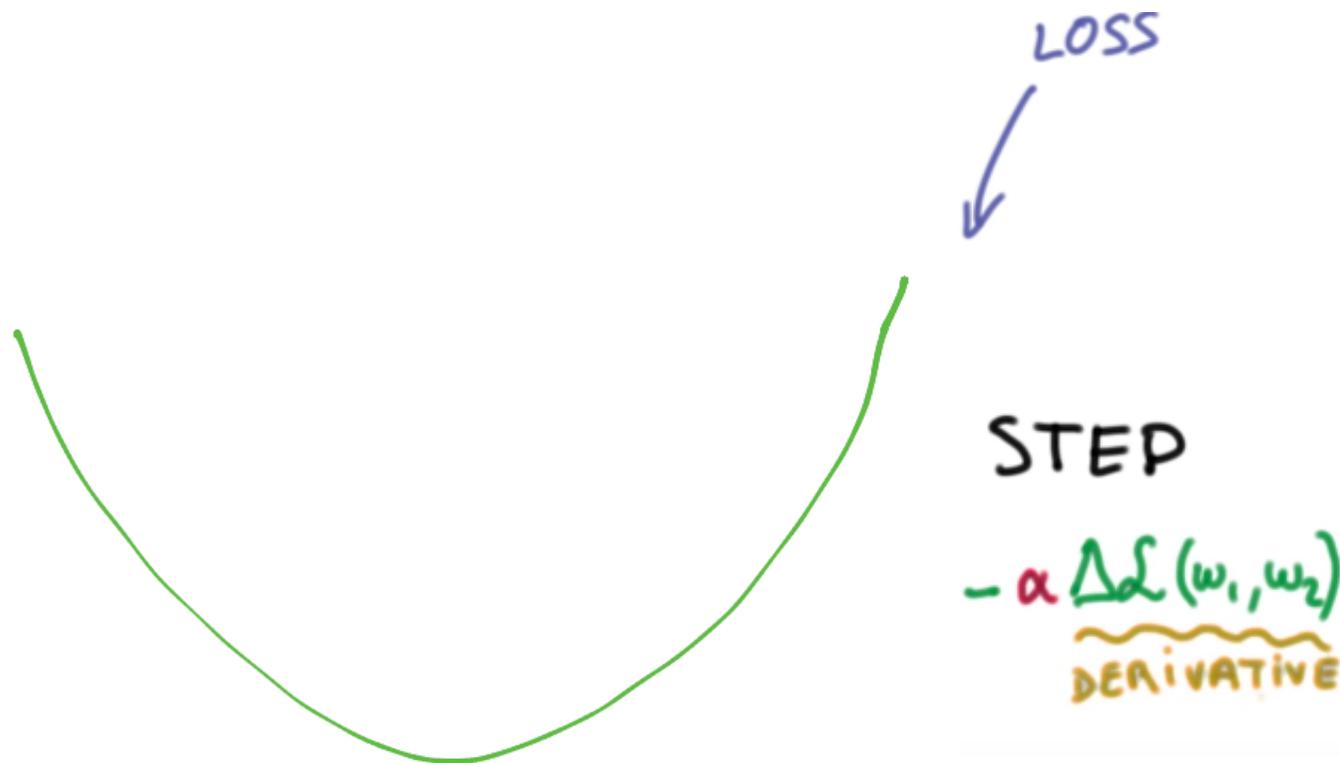
Learning rate, data preprocessing, <sup>★</sup>overfitting

Sung Kim <hunkim+mr@gmail.com>

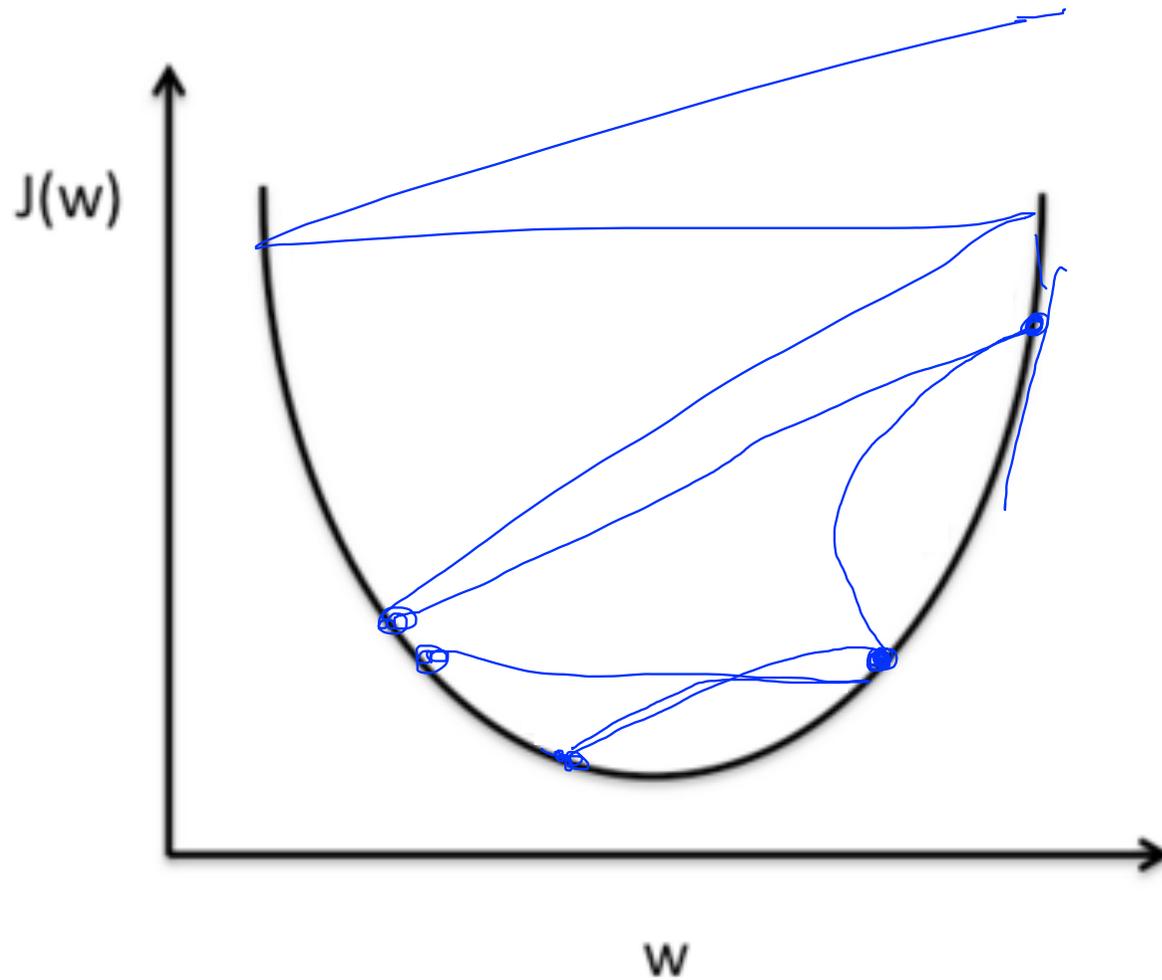
# Gradient descent



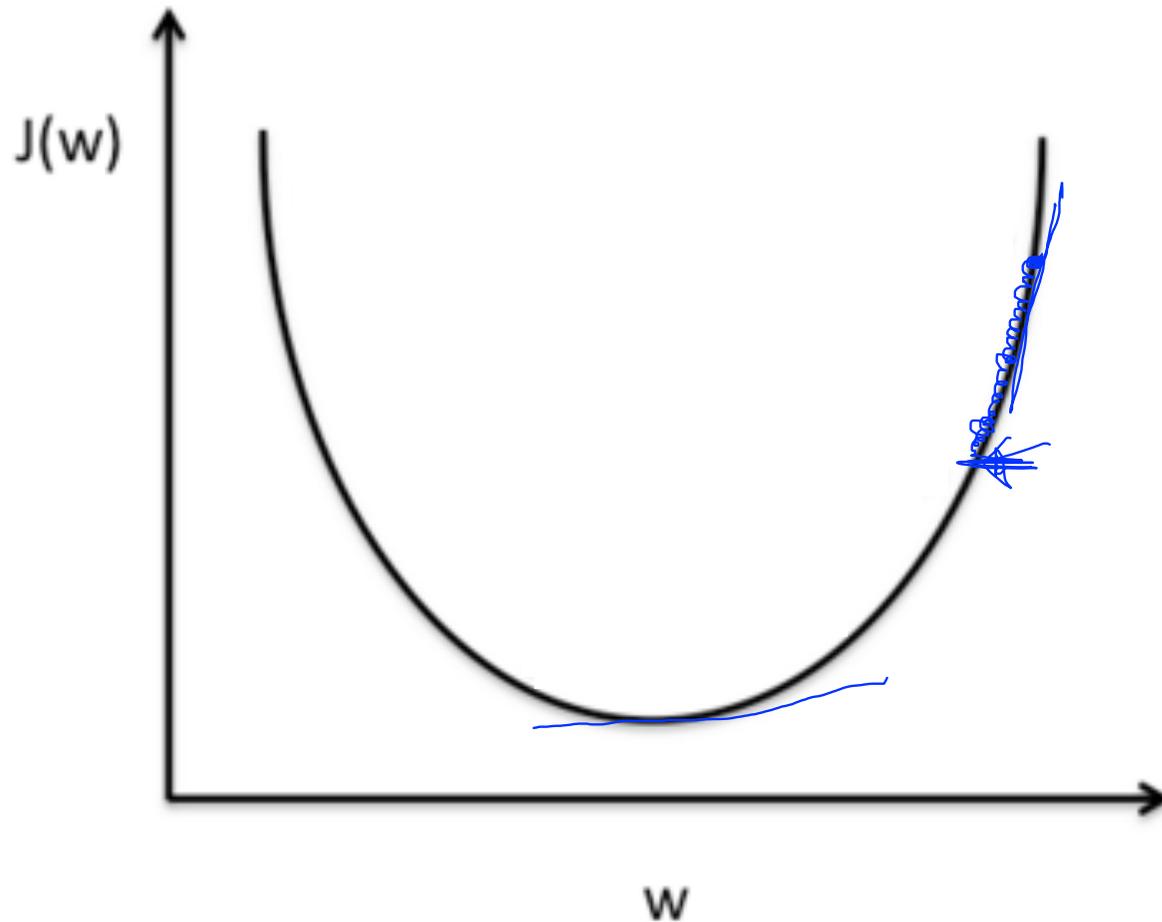
```
# Minimize error using cross entropy
learning_rate = 0.001
cost = tf.reduce_mean(-tf.reduce_sum(Y*tf.log(hypothesis), reduction_indices=1)) # Cross entropy
optimizer = tf.train.GradientDescentOptimizer(learning_rate).minimize(cost) # Gradient Descent
```



Large learning rate: overshooting



Small learning rate:  
takes too long, stops at local minimum

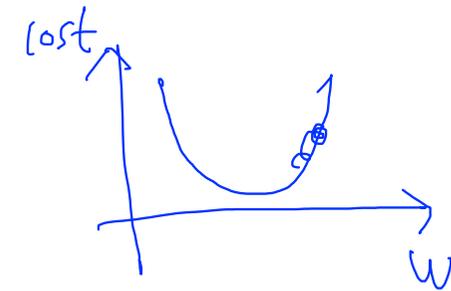
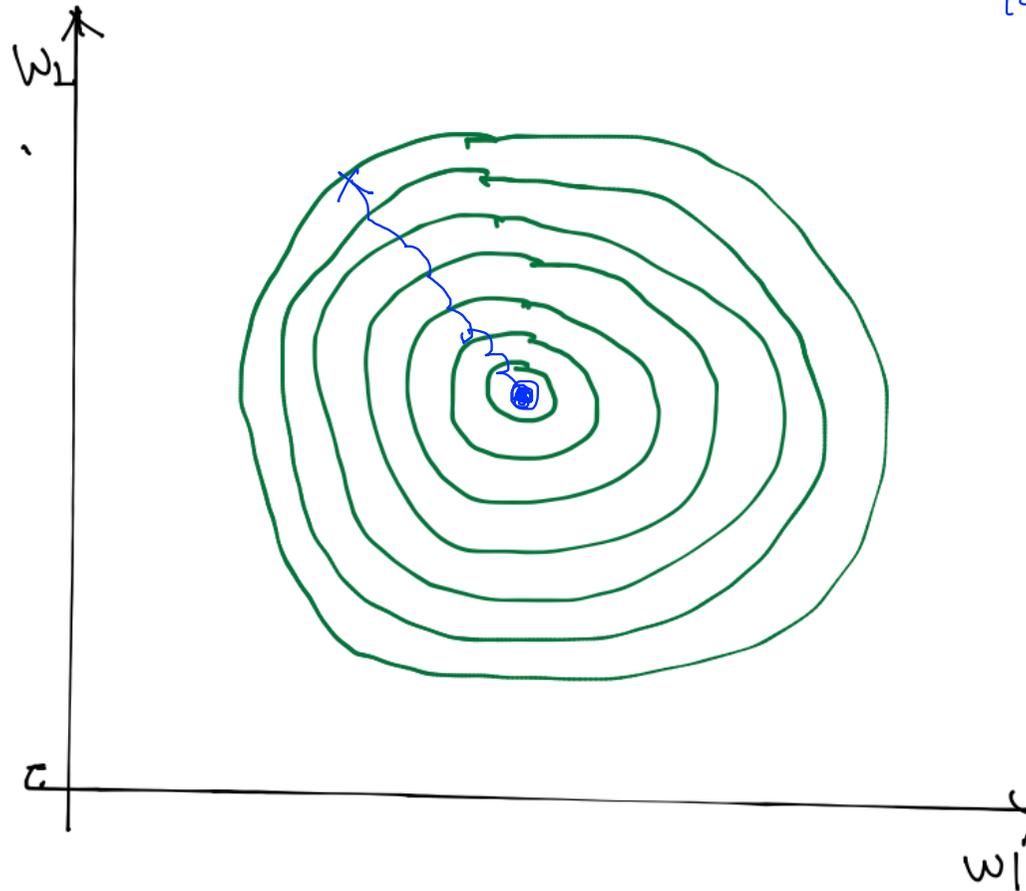


# Try several learning rates

0.01

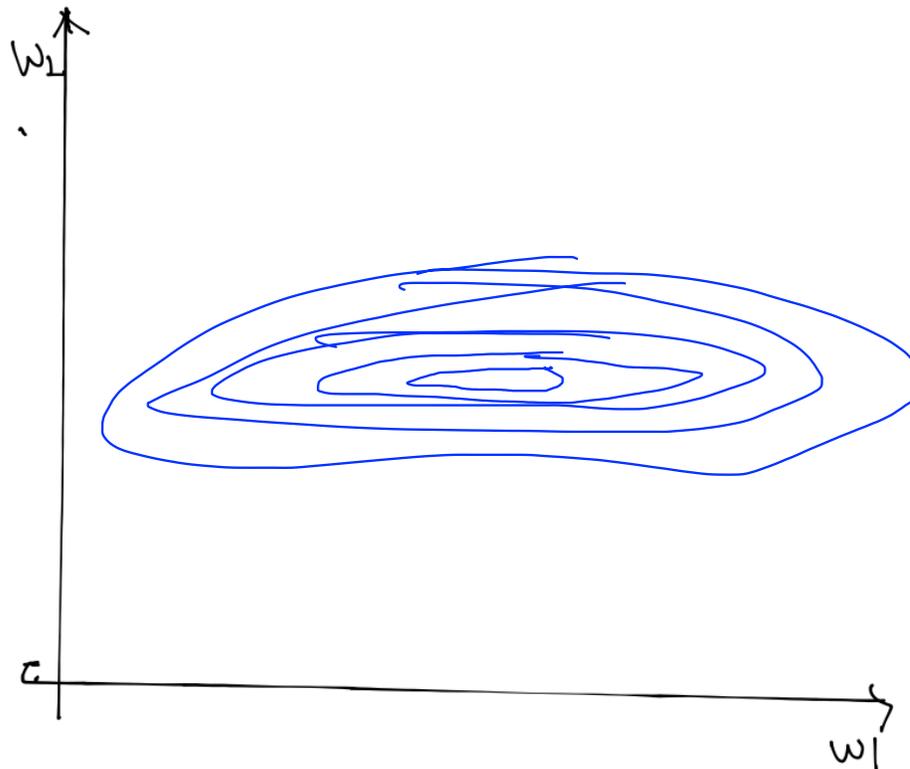
- Observe the cost function
- Check it goes down in a reasonable rate

# Data (X) preprocessing for gradient descent



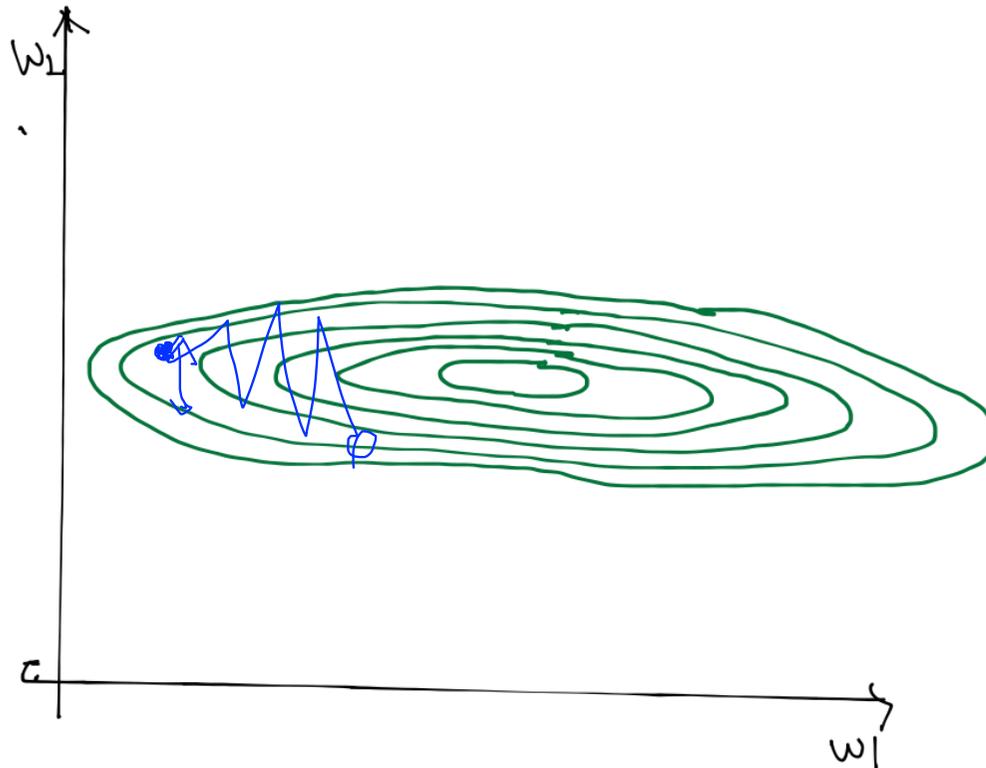
# Data (X) preprocessing for gradient descent

x1	x2	y
1	9000	A
2	-5000	A
4	-2000	B
6	8000	B
9	9000	C

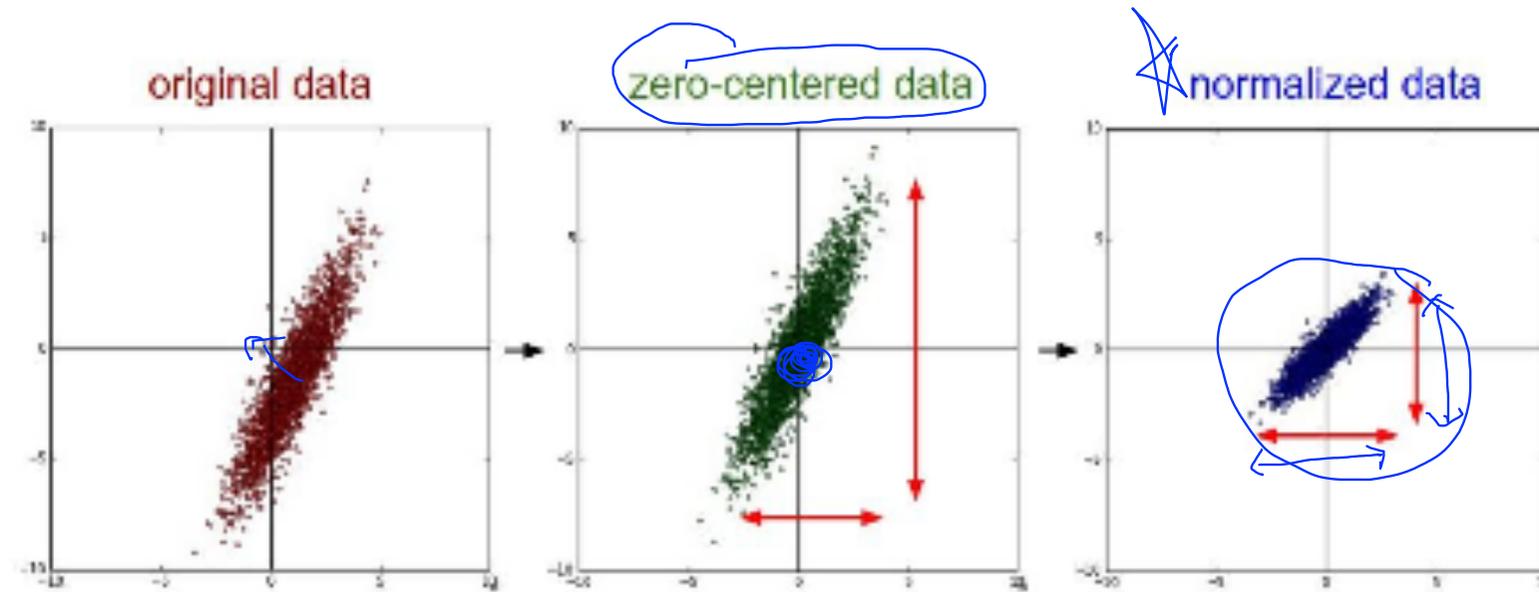


# Data (X) preprocessing for gradient descent

x1	x2	y
1	9000	A
2	-5000	A
4	-2000	B
6	8000	B
9	9000	C



# Data (X) preprocessing for gradient descent



# Standardization

$$\mathbf{x}'_j = \frac{\mathbf{x}_j - \mu_j}{\sigma_j}$$

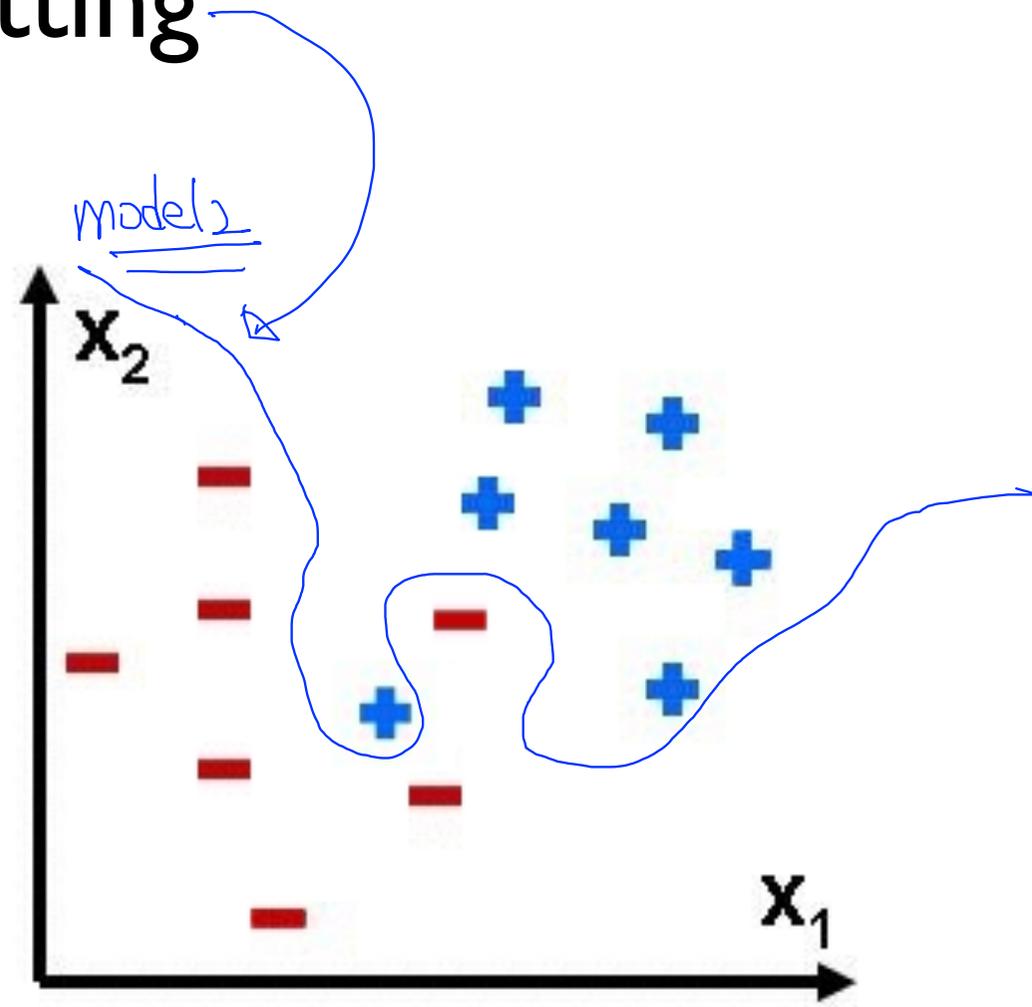
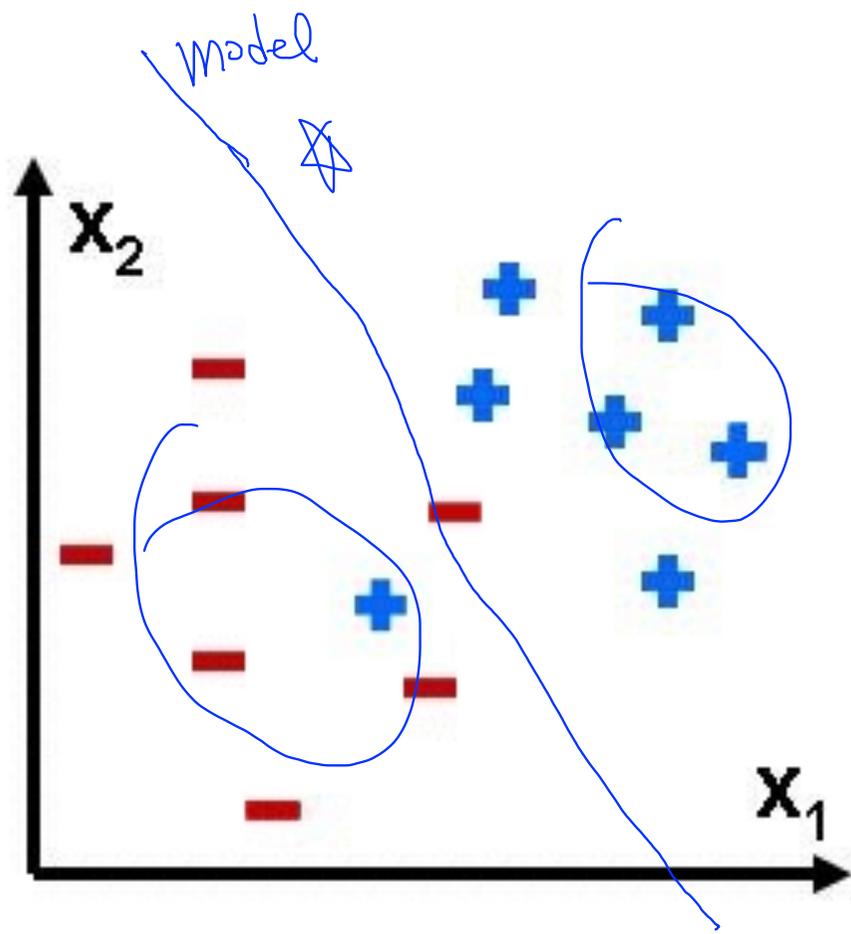
```
x_std[:,0] = (x[:,0] - x[:,0].mean()) / x[:,0].std()
```

[http://sebastianraschka.com/Articles/2015\\_singlelayer\\_neurons.html](http://sebastianraschka.com/Articles/2015_singlelayer_neurons.html)

# Overfitting

- Our model is very good with training data set (with memorization)
- Not good at test dataset or in real use

# Overfitting



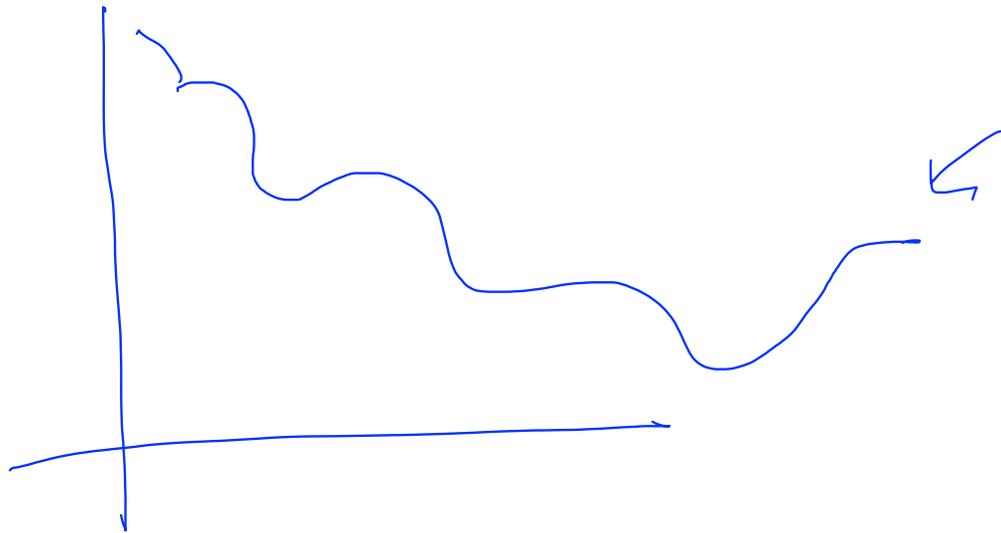
# Solutions for overfitting

- More training data!
- Reduce the number of features
- Regularization



# Regularization

- Let's not have too big numbers in the weight



# Regularization

- Let's not have too big numbers in the weight

A handwritten diagram illustrating the loss function. The equation is  $\mathcal{L} = \frac{1}{N} \sum_i \mathcal{D}(s(wx_i + b), L_i)$ . The word "LOSS" is written in blue above the equation, with a blue arrow pointing to the  $\mathcal{L}$  term. The word "TRAINING SET" is written in blue below the equation, with a blue arrow pointing to the summation index  $i$ . Another blue arrow points from "TRAINING SET" to the  $x_i$  term in the function. A third blue arrow points from "TRAINING SET" to the  $L_i$  term in the function. The terms  $\mathcal{L}$ ,  $\frac{1}{N}$ ,  $\sum$ ,  $\mathcal{D}$ ,  $s$ ,  $w$ ,  $x_i$ , and  $L_i$  are written in red, green, orange, and blue respectively.

$$\mathcal{L} = \frac{1}{N} \sum_i \mathcal{D}(s(wx_i + b), L_i)$$

# Regularization

- Let's not have too big numbers in the weight

LOSS

$$\mathcal{L} = \frac{1}{N} \sum_i \mathcal{D}(s(wX_i + b), L_i) + \lambda \sum W^2$$

TRAINING SET

regularization strength

0 X  
↓ ↑  
0.001

# Regularization

- Let's not have too big numbers in the weight

LOSS

`l2reg = 0.001 * tf.reduce_sum(tf.square(W))` ← regularization strength

$$\mathcal{L} = \frac{1}{N} \sum_i \mathcal{D}(s(wX_i + b), L_i) + \lambda \sum W^2$$

TRAINING SET

# Summary

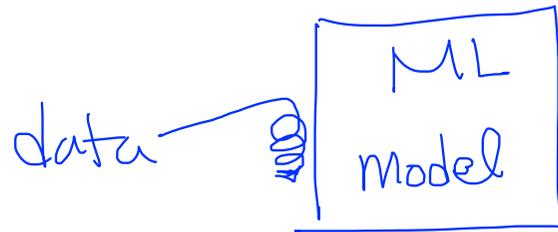
- Learning rate ✓
- Data preprocessing
- Overfitting ✱
  - More training data
  - Regularization

# Lecture 7-2

## Application & Tips: Learning and test data sets

Sung Kim <[hunkim+mr@gmail.com](mailto:hunkim+mr@gmail.com)>

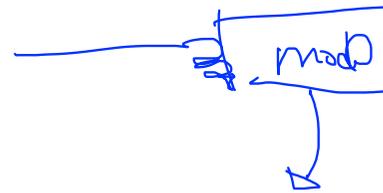
# Performance evaluation: is this good?



# Evaluation using training set?

Size	Price
2104	400
1600	330
2400	369
1416	232
3000	540
1985	300
1534	315
1427	199
1380	212
1494	243

training set

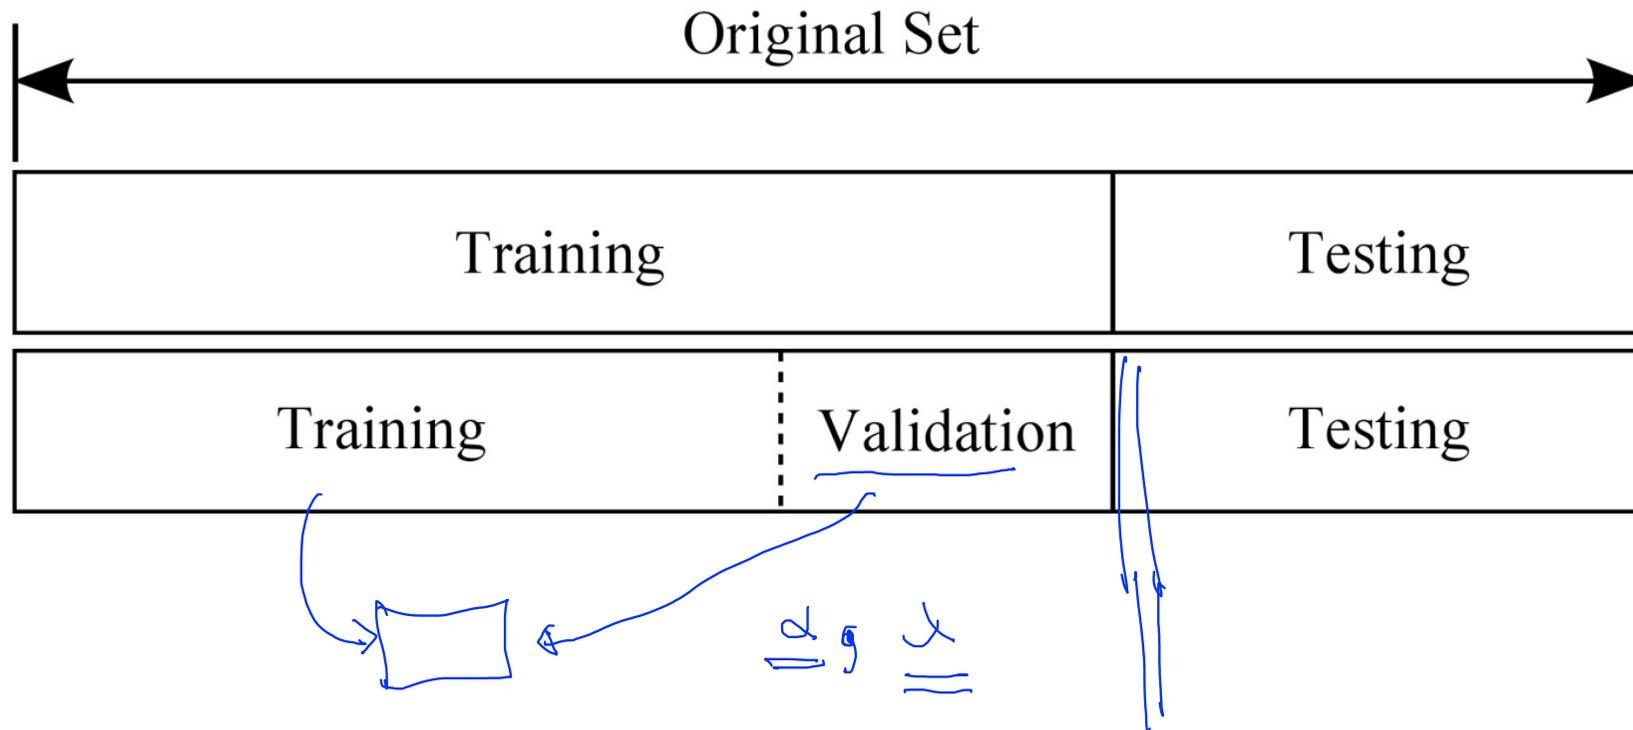


- 100% correct (accuracy)
- Can memorize

# Training and test sets

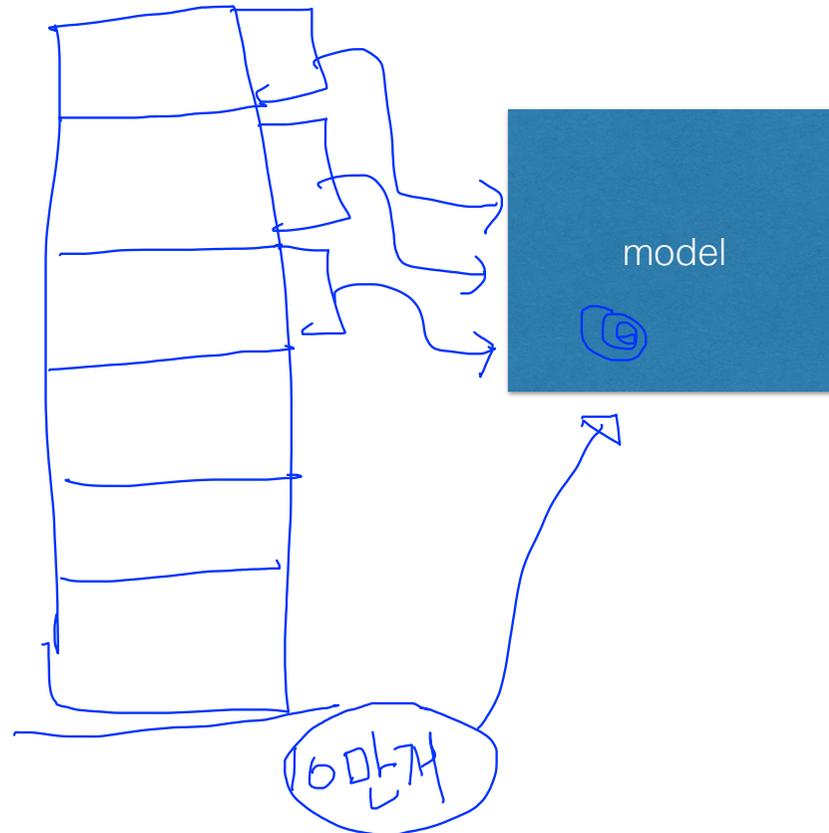
	Size	Price
training (교과서)	2104	400
	1600	330
	2400	369
	1416	232
	3000	540
	1985	300
	1534	315
test set (시험)	1427	199
	1380	212
	1494	243

# Training, validation and test sets



# Online learning

100만개

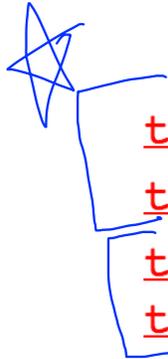


# M NIST Dataset

Zip: 0363



0  
1  
2



- [train-images-idx3-ubyte.gz](#): training set images (9912422 bytes)
- [train-labels-idx1-ubyte.gz](#): training set labels (28881 bytes)
- [t10k-images-idx3-ubyte.gz](#): test set images (1648877 bytes)
- [t10k-labels-idx1-ubyte.gz](#): test set labels (4542 bytes)

<http://yann.lecun.com/exdb/mnist/>

# Accuracy

- How many of your predictions are correct?
- 95% ~ 99%?
- Check out the lab video

