

CS 4774 Machine Learning

Linear Predictors

Yangfeng Ji

Information and Language Processing Lab
Department of Computer Science
University of Virginia



1. Review: Linear Functions
2. Perceptron
3. Logistic Regression
4. Linear Regression
5. ℓ_2 Regularization and Overfitting
6. Summary

Review: Linear Functions

Linear predictors discussed in this course

- ▶ halfspace predictors
- ▶ logistic regression classifiers
- ▶ linear SVMs (lecture on support vector machines)
- ▶ naive Bayes classifiers (lecture on generative models)
- ▶ linear regression predictors

Linear predictors discussed in this course

- ▶ halfspace predictors
- ▶ logistic regression classifiers
- ▶ linear SVMs (lecture on support vector machines)
- ▶ naive Bayes classifiers (lecture on generative models)
- ▶ linear regression predictors

A common core form of these linear predictors

$$h_{w,b} = \langle w, x \rangle + b = \left(\sum_{i=1}^d w_i x_i \right) + b \quad (1)$$

where w is the weights and b is the bias

Alternative Form

Given the original definition of a linear function

$$h_{w,b} = \langle w, x \rangle + b = \left(\sum_{i=1}^d w_i x_i \right) + b, \quad (2)$$

we could redefine it in a more compact form

$$\begin{aligned} w &\leftarrow (w_1, w_2, \dots, w_d, b)^\top \\ x &\leftarrow (x_1, x_2, \dots, x_d, 1)^\top \end{aligned}$$


and then

$$h_{w,b}(x) = \langle w, x \rangle \quad (3)$$

Linear Functions

Consider a two-dimensional case with $w = (1, 1, -0.5)$

$$f(\mathbf{x}) = \mathbf{w}^T \mathbf{x} = x_1 + x_2 - 0.5 \quad (4)$$




Different values of $f(\mathbf{x})$ map to different areas on this 2-D space. For example, the following equation defines the blue line L .

$$f(\mathbf{x}) = \mathbf{w}^T \mathbf{x} = 0 \quad (5)$$

Properties of Linear Functions (II)

For any two points x and x' lying in the line

$$f(x) - f(x') = w^T x - w^T x' = 0 \quad (6)$$




[Friedman et al., 2001, Section 4.5]

Properties of Linear Functions (III)

Furthermore,

$$f(x) = x_1 + x_2 - 0.5 = 0 \quad (7)$$


separates the 2-D space \mathbb{R}^2 into two **half** spaces



Properties of Linear Functions (IV)

From the perspective of linear projection, $f(\mathbf{x}) = 0$ defines the vectors on this 2-D space, whose projections onto the direction $(1, 1)$ have the same magnitude 0.5

$$x_1 + x_2 - 0.5 = 0 \Rightarrow (x_1, x_2) \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} = 0.5 \quad (8)$$




[Friedman et al., 2001, Section 4.5]

Properties of Linear Functions (IV)

From the perspective of linear projection, $f(\mathbf{x}) = 0$ defines the vectors on this 2-D space, whose projections onto the direction $(1, 1)$ have the same magnitude 0.5

$$x_1 + x_2 - 0.5 = 0 \Rightarrow (x_1, x_2) \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} = 0.5 \quad (8)$$




This idea can be generalized to compute the distance between a point and a line.

[Friedman et al., 2001, Section 4.5]

Properties of Linear Functions (V)

The distance of point x to line $L : f(x) = \langle w, x \rangle = 0$ is given by

$$\frac{f(x)}{\|w\|_2} = \frac{\langle w, x \rangle}{\|x\|_2} = \left\langle \frac{w}{\|w\|_2}, x \right\rangle \quad (9)$$



[Friedman et al., 2001, Section 4.5]

Perceptron


Halfspace Hypothesis Class

- ▶ $\mathcal{X} = \mathbb{R}^d$
- ▶ $\mathcal{Y} = \{-1, +1\}$
- ▶ Halfspace hypothesis class

$$\mathcal{H}_{\text{half}} = \{\text{sign}(\langle w, x \rangle) : w \in \mathbb{R}^d\} \quad (10)$$


which is an **infinite** hypothesis space.

The sign function $y = \text{sign}(x)$ is defined as



Linearly Separable Cases

The algorithm can find a hyperplane to separate all positive examples from negative examples





The definition of linearly separable cases is with respect to the training set S instead of \mathcal{D}

Prediction Rule

The prediction rule of a half-space predictor is based on the sign of

$$h(\mathbf{x}) = \text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle)$$

$$h(\mathbf{x}) = \begin{cases} +1 & \langle \mathbf{w}, \mathbf{x} \rangle > 0 \\ -1 & \langle \mathbf{w}, \mathbf{x} \rangle < 0 \end{cases} \quad (11)$$



Prediction Rule


The prediction rule of a half-space predictor is based on the sign of

$$h(\mathbf{x}) = \text{sign}(\langle \mathbf{w}, \mathbf{x} \rangle)$$

$$h(\mathbf{x}) = \begin{cases} +1 & \langle \mathbf{w}, \mathbf{x} \rangle > 0 \\ -1 & \langle \mathbf{w}, \mathbf{x} \rangle < 0 \end{cases} \quad (11)$$

or,

$$h(\mathbf{x}) = y' \quad \text{if } y' \in \{-1, +1\} \text{ and } y' \langle \mathbf{w}, \mathbf{x} \rangle > 0 \quad (12)$$



Perceptron Algorithm

The perceptron algorithm is defined as

- 1: **Input:** $S = \{(x_1, y_1), \dots, (x_m, y_m)\}$
- 2: Initialize $w^{(0)} = (0, \dots, 0)$

11: **Output:** $w^{(T)}$

Perceptron Algorithm

The perceptron algorithm is defined as

- 1: **Input:** $S = \{(x_1, y_1), \dots, (x_m, y_m)\}$
- 2: Initialize $w^{(0)} = (0, \dots, 0)$
- 3: **for** $t = 1, 2, \dots$ **do**
- 5: $i \leftarrow t \bmod m$
- 10: **end for**
- 11: **Output:** $w^{(T)}$

Perceptron Algorithm

The perceptron algorithm is defined as

- 1: **Input:** $S = \{(x_1, y_1), \dots, (x_m, y_m)\}$
- 2: Initialize $w^{(0)} = (0, \dots, 0)$
- 3: **for** $t = 1, 2, \dots$ **do**
- 4: *// One example at a time*
- 5: $i \leftarrow t \bmod m$
- 6: *// Only when model makes wrong prediction*
- 7: **if** $y_i \langle w^{(t)}, x_i \rangle \leq 0$ **then**
- 8: $w^{(t+1)} \leftarrow w^{(t)} + y_i x_i$ *// updating rule*
- 9: **end if**
- 10: **end for**
- 11: **Output:** $w^{(T)}$

Perceptron Algorithm

The perceptron algorithm is defined as

```
1: Input:  $S = \{(x_1, y_1), \dots, (x_m, y_m)\}$ 
2: Initialize  $w^{(0)} = (0, \dots, 0)$ 
3: for  $t = 1, 2, \dots$  do
4:   // One example at a time
5:    $i \leftarrow t \bmod m$ 
6:   // Only when model makes wrong prediction
7:   if  $y_i \langle w^{(t)}, x_i \rangle \leq 0$  then
8:      $w^{(t+1)} \leftarrow w^{(t)} + y_i x_i$  // updating rule
9:   end if
10: end for
11: Output:  $w^{(T)}$ 
```

- ▶ It will stop training when there is no prediction error
- ▶ This is an *online* learning algorithm

The updating rule can be broken down into two cases:

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} + y_i \mathbf{x}_i \quad (13)$$

- ▶ For $y_i = +1$, $\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} + \mathbf{x}_i$
- ▶ For $y_i = -1$, $\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \mathbf{x}_i$

Two Questions

The updating rule can be broken down into two cases:

$$\boldsymbol{w}^{(t+1)} \leftarrow \boldsymbol{w}^{(t)} + y_i \boldsymbol{x}_i \quad (13)$$

- ▶ For $y_i = +1$, $\boldsymbol{w}^{(t+1)} \leftarrow \boldsymbol{w}^{(t)} + \boldsymbol{x}_i$
- ▶ For $y_i = -1$, $\boldsymbol{w}^{(t+1)} \leftarrow \boldsymbol{w}^{(t)} - \boldsymbol{x}_i$

Two questions:

- ▶ How the updating rule can help?
- ▶ How many updating steps the algorithm needs?

Consider the following

- ▶ The current classification weight $w^{(t)} = (1.0, -1.0, 0.0)$
- ▶ $x_i = (0.0, 1.0, 1.0)$ with label $y_i = +1$
- ▶ Learning
 - ▶ What is the predicted label on (x_i, y_i) with $w^{(t)}$?
 - ▶ Need a update?

The Updating Rule

At time step t , given the training example (x_i, y_i) and the current weight $w^{(t)}$

$$y_i \langle w^{(t+1)}, x_i \rangle = y_i \langle w^{(t)} + y_i x_i, x_i \rangle \quad (14)$$

$$= y_i \langle w^{(t)}, x_i \rangle + \|x_i\|^2 \quad (15)$$

- ▶ $w^{(t+1)}$ gives a higher value of $y_i \langle w^{(t+1)}, x_i \rangle$ on predicting x_i than $w^{(t)}$
- ▶ the updating is affected by the norm of x_i , $\|x_i\|^2$

Assume that $\{(x_i, y_i)\}_{i=1}^m$ is linearly separable. Let

- ▶ $B = \min\{\|w\| : \forall i \in [m], y_i \langle w, x_i \rangle \geq 1\}$, and
- ▶ $R = \max_i \|x_i\|$.

Then, the Perceptron algorithm stops after at most $(RB)^2$ iterations, and when it stops it holds that $\forall i \in [m]$,

$$y_i \langle w^{(t)}, x \rangle > 0 \tag{16}$$

- ▶ A realizable case with **infinite** hypothesis space
- ▶ Finish training in **finite** steps


Logistic Regression

- ▶ The hypothesis class of logistic regression is defined as

$$\mathcal{H}_{\text{LR}} = \{\sigma(y\langle w, x \rangle) : w \in \mathbb{R}^d\} \quad (17)$$

- ▶ The sigmoid function $\sigma(a)$ with $a \in \mathbb{R}$

$$\sigma(a) = \frac{1}{1 + \exp(-a)} \quad (18)$$



- ▶ An unified form for $y \in \{-1, +1\}$

$$h(x, y) = \frac{1}{1 + \exp(-y\langle w, x \rangle)} \quad (19)$$

which is similar to the half-space predictors

- ▶ An unified form for $y \in \{-1, +1\}$

$$h(x, y) = \frac{1}{1 + \exp(-y\langle w, x \rangle)} \quad (19)$$

which is similar to the half-space predictors

- ▶ Prediction
 1. Compute the the values from Eq. 19 with $y \in \{-1, +1\}$
 2. Pick the y that has bigger value

$$y = \begin{cases} +1 & h(x, +1) > h(x, -1) \\ -1 & h(x, +1) < h(x, -1) \end{cases} \quad (20)$$

Take a close look of the uniform definition of $h(\mathbf{x}, y)$

- ▶ When $y = +1$

$$h_w(\mathbf{x}, +1) = \frac{1}{1 + \exp(-\langle \mathbf{w}, \mathbf{x} \rangle)}$$

- ▶ When $y = -1$

$$\begin{aligned} h_w(\mathbf{x}, -1) &= \frac{1}{1 + \exp(\langle \mathbf{w}, \mathbf{x} \rangle)} \\ &= \frac{\exp(-\langle \mathbf{w}, \mathbf{x} \rangle)}{1 + \exp(-\langle \mathbf{w}, \mathbf{x} \rangle)} \\ &= 1 - \frac{1}{1 + \exp(-\langle \mathbf{w}, \mathbf{x} \rangle)} \\ &= 1 - h_w(\mathbf{x}, +1) \end{aligned}$$

A Linear Classifier?

To justify this is a linear classifier, let take a look the decision boundary given by

$$h(\mathbf{x}, +1) = h(\mathbf{x}, -1) \quad (21)$$

Specifically, we have

$$\begin{aligned} \frac{1}{1 + \exp(-\langle \mathbf{w}, \mathbf{x} \rangle)} &= \frac{1}{1 + \exp(\langle \mathbf{w}, \mathbf{x} \rangle)} \\ \exp(-\langle \mathbf{w}, \mathbf{x} \rangle) &= \exp(\langle \mathbf{w}, \mathbf{x} \rangle) \\ -\langle \mathbf{w}, \mathbf{x} \rangle &= \langle \mathbf{w}, \mathbf{x} \rangle \\ 2\langle \mathbf{w}, \mathbf{x} \rangle &= 0 \end{aligned}$$

The decision boundary is a straight line (or in high-dimensional spaces, a hyperplane)

For **one** training example (x, y) , the risk/loss function is defined as the negative log of $h(x, y)$

$$\begin{aligned}L(h_w(x, y)) &= -\log \frac{1}{1 + \exp(-y\langle w, x \rangle)} \\ &= \log(1 + \exp(-y\langle w, x \rangle))\end{aligned}\tag{22}$$

Intuitively, minimizing the risk will increase the value of $h(x, y)$

The **Empirical Risk Minimization** (ERM) problem: given the training set $S = \{(x_1, y_1), \dots, (x_m, y_m)\}$, minimize the following objective function with respect to w

$$L(h_w(S)) = \frac{1}{m} \sum_{i=1}^m \log(1 + \exp(-y_i \langle w, x_i \rangle)) \quad (23)$$

- ▶ $L(h_w(S))$ is convex function with respect to w
- ▶ Estimation of w : $\hat{w} \leftarrow \operatorname{argmin}_{w'} L(h_{w'}(S))$
- ▶ Minimization can be done with gradient-based optimization¹

¹more detail will be covered in the lecture of optimization methods

- ▶ The gradient of $L(h_w(S))$ with respect to w

$$\frac{dL(h_w(S))}{dw} = \frac{1}{m} \sum_{i=1}^m \frac{\exp(-y_i \langle w, x_i \rangle)}{1 + \exp(-y_i \langle w, x_i \rangle)} \cdot (-y_i x_i) \quad (24)$$

- ▶ The gradient of $L(h_w(S))$ with respect to w

$$\frac{dL(h_w(S))}{dw} = \frac{1}{m} \sum_{i=1}^m \frac{\exp(-y_i \langle w, x_i \rangle)}{1 + \exp(-y_i \langle w, x_i \rangle)} \cdot (-y_i x_i) \quad (24)$$

- ▶ Gradient-based learning

$$w^{(\text{new})} = w^{(\text{old})} - \eta \frac{dL(h_w(S))}{dw}$$

where η is the updating step size.

- ▶ The gradient of $L(h_w(S))$ with respect to w

$$\frac{dL(h_w(S))}{dw} = \frac{1}{m} \sum_{i=1}^m \frac{\exp(-y_i \langle w, x_i \rangle)}{1 + \exp(-y_i \langle w, x_i \rangle)} \cdot (-y_i x_i) \quad (24)$$

- ▶ Gradient-based learning

$$\begin{aligned} w^{(\text{new})} &= w^{(\text{old})} - \eta \frac{dL(h_w(S))}{dw} \\ &= w^{(\text{old})} + \frac{\eta}{m} \sum_{i=1}^m \frac{\exp(-y_i \langle w, x_i \rangle)}{1 + \exp(-y_i \langle w, x_i \rangle)} \cdot (y_i x_i) \end{aligned}$$

where η is the updating step size.

Gradient-based learning

$$\mathbf{w}^{(\text{new})} = \mathbf{w}^{(\text{old})} + \frac{\eta}{m} \sum_{i=1}^m \underbrace{\frac{\exp(-y_i \langle \mathbf{w}, \mathbf{x}_i \rangle)}{1 + \exp(-y_i \langle \mathbf{w}, \mathbf{x}_i \rangle)}}_{(2)} \cdot \underbrace{(y_i \mathbf{x}_i)}_{(1)} \quad (25)$$

For each (\mathbf{x}_i, y_i) , the update is

- (1) directed by the true label y_i , as in the Perceptron algorithm
- (2) proportional to the prediction value of the opposite label, unlike the Perceptron algorithm

Updating Rules

Compare these two updating rules:

- ▶ Logistic regression

$$\boldsymbol{w}^{(\text{new})} = \boldsymbol{w}^{(\text{old})} + \frac{\eta}{m} \sum_{i=1}^m \frac{\exp(-y_i \langle \boldsymbol{w}, \boldsymbol{x}_i \rangle)}{1 + \exp(-y_i \langle \boldsymbol{w}, \boldsymbol{x}_i \rangle)} \cdot (y_i \boldsymbol{x}_i) \quad (26)$$

Updating Rules

Compare these two updating rules:

- ▶ Logistic regression

$$\mathbf{w}^{(\text{new})} = \mathbf{w}^{(\text{old})} + \frac{\eta}{m} \sum_{i=1}^m \frac{\exp(-y_i \langle \mathbf{w}, \mathbf{x}_i \rangle)}{1 + \exp(-y_i \langle \mathbf{w}, \mathbf{x}_i \rangle)} \cdot (y_i \mathbf{x}_i) \quad (26)$$

- ▶ Perceptron algorithm

$$\mathbf{w}^{(\text{new})} = \mathbf{w}^{(\text{old})} + y_i \mathbf{x}_i \quad (27)$$

only applies when the prediction is wrong

Updating Rules

Compare these two updating rules:

- ▶ Logistic regression

$$\mathbf{w}^{(\text{new})} = \mathbf{w}^{(\text{old})} + \frac{\eta}{m} \sum_{i=1}^m \frac{\exp(-y_i \langle \mathbf{w}, \mathbf{x}_i \rangle)}{1 + \exp(-y_i \langle \mathbf{w}, \mathbf{x}_i \rangle)} \cdot (y_i \mathbf{x}_i) \quad (26)$$

- ▶ Perceptron algorithm

$$\mathbf{w}^{(\text{new})} = \mathbf{w}^{(\text{old})} + y_i \mathbf{x}_i \quad (27)$$

only applies when the prediction is wrong

- ▶ Three differences
 - ▶ Online learning vs. batch learning
 - ▶ With vs. without Learning rate
 - ▶ What examples are used to update

A Probabilistic View of Logistic Regression

- ▶ From a probabilistic view, logistic regression defines the probability of a possible label y given the input x

$$p_{\mathbf{w}}(Y = y \mid \mathbf{x}) = h(\mathbf{x}, y) = \frac{1}{1 + \exp(-y\langle \mathbf{w}, \mathbf{x} \rangle)} \quad (28)$$

where Y is a random variable with $Y \in \{-1, +1\}$

A Probabilistic View of Logistic Regression

- ▶ From a probabilistic view, logistic regression defines the probability of a possible label y given the input \mathbf{x}

$$p_{\mathbf{w}}(Y = y | \mathbf{x}) = h(\mathbf{x}, y) = \frac{1}{1 + \exp(-y\langle \mathbf{w}, \mathbf{x} \rangle)} \quad (28)$$

where Y is a random variable with $Y \in \{-1, +1\}$

- ▶ The previous prediction rule is equivalent to

$$\hat{y} = \begin{cases} +1 & \text{if } p(Y = +1 | \mathbf{x}) > p(Y = -1 | \mathbf{x}) \\ -1 & \text{if } p(Y = +1 | \mathbf{x}) < p(Y = -1 | \mathbf{x}) \end{cases} \quad (29)$$

Given the training set $S = \{(x_1, y_1), \dots, (x_m, y_m)\}$, the likelihood function is defined as

$$\text{Lik}(\mathbf{x}) = \prod_{i=1}^m p_w(y_i | x_i) \quad (30)$$

Parameter Estimation: Maximum Likelihood

Given the training set S ,

- ▶ Log-likelihood function

$$\begin{aligned}\ell(\boldsymbol{w}) &= \sum_{i=1}^m \log p_{\boldsymbol{w}}(y_i | \boldsymbol{x}_i) \\ &= \sum_{i=1}^m \log \frac{1}{1 + \exp(-y_i \langle \boldsymbol{w}, \boldsymbol{x}_i \rangle)} \\ &= - \sum_{i=1}^m \log(1 + \exp(-y_i \langle \boldsymbol{w}, \boldsymbol{x}_i \rangle))\end{aligned}\tag{31}$$

Parameter Estimation: Maximum Likelihood

Given the training set S ,

- ▶ Log-likelihood function

$$\begin{aligned}\ell(\boldsymbol{w}) &= \sum_{i=1}^m \log p_{\boldsymbol{w}}(y_i | \boldsymbol{x}_i) \\ &= \sum_{i=1}^m \log \frac{1}{1 + \exp(-y_i \langle \boldsymbol{w}, \boldsymbol{x}_i \rangle)} \\ &= - \sum_{i=1}^m \log(1 + \exp(-y_i \langle \boldsymbol{w}, \boldsymbol{x}_i \rangle))\end{aligned}\tag{31}$$

- ▶ Maximize the log-likelihood function

$$\operatorname{argmax}_{\boldsymbol{w}} \ell(\boldsymbol{w}) = \operatorname{argmin}_{\boldsymbol{w}} -\ell(\boldsymbol{w}) = \operatorname{argmin}_{\boldsymbol{w}} L(h_{\boldsymbol{w}}, S)$$

learning with ERM is equivalent to the Maximum Likelihood Estimation (MLE) in Statistics

Recall the gradient-based learning on the previous slide

$$\begin{aligned} \boldsymbol{w}^{(\text{new})} &= \boldsymbol{w}^{(\text{old})} + \eta \sum_{i=1}^m \frac{\exp(-y_i \langle \boldsymbol{w}, \boldsymbol{x}_i \rangle)}{1 + \exp(-y_i \langle \boldsymbol{w}, \boldsymbol{x}_i \rangle)} \cdot (y_i \boldsymbol{x}_i) \\ &= \boldsymbol{w}^{(\text{old})} + \eta \sum_{i=1}^m (1 - p(y_i | \boldsymbol{x}_i)) \cdot y_i \boldsymbol{x}_i \end{aligned} \quad (32)$$

- ▶ If $p(y_i | \boldsymbol{x}_i) \rightarrow 0$, wrong prediction, maximal update
- ▶ If $p(y_i | \boldsymbol{x}_i) \rightarrow 1$, correct prediction, minimal update

A practical guide of building LR classifiers

Linear Regression

- ▶ The hypothesis class of linear regression predictors is defined as


$$\mathcal{H}_{\text{reg}} = \{\langle w, x \rangle : w \in \mathbb{R}^d\} \quad (33)$$

- ▶ One example hypothesis $h \in \mathcal{H}_{\text{reg}}$

$$h(x) = \langle w, x \rangle \quad (34)$$

Problem Statement

Given the training set S , in this case, $\{(x_1, y_1), \dots, (x_5, y_5)\}$, find $h \in \mathcal{H}_{\text{reg}}$ such that $h(x)$ gives the best (linear) relation between x and y



- ▶ Loss function

$$L(h, (x, y)) = (h(x) - y)^2 = (w^\top x - y)^2 \quad (35)$$

- ▶ Loss function

$$L(h, (x, y)) = (h(x) - y)^2 = (w^T x - y)^2 \quad (35)$$

- ▶ Given the training set S , the corresponding empirical risk function of linear regression is defined as

$$L(h, S) = \sum_{i=1}^m (h(x_i) - y_i)^2 \quad (36)$$


which is called **Mean Squared Error** (MSE).

Visualization

For a 1-D case, the loss function

$$L(h, S) = \sum_{i=1}^m (h(x_i) - y_i)^2 \quad (37)$$

can be visualized as



- ▶ The ERM problem

$$\operatorname{argmin}_w L_S(h_w) = \operatorname{argmin}_w \sum_{i=1}^m (\langle w, x_i \rangle - y_i)^2 \quad (38)$$

- ▶ Compute the gradient and set it to be zero

$$\begin{aligned} \frac{2}{m} \sum_{i=1}^m (\langle w, x_i \rangle - y_i) x_i &= 0 \\ \sum_{i=1}^m \langle w, x_i \rangle x_i &= \sum_{i=1}^m y_i x_i \end{aligned}$$

Empirical Risk Minimization (II)

To isolate w for solution, we have

- ▶ $\langle w, x_i \rangle x_i = (w^\top x_i) x_i = (x_i x_i^\top) w$

$$\sum_{i=1}^m (x_i x_i^\top) w = \sum_{i=1}^m y_i x_i \quad (39)$$

- ▶ then, rewrite it as

$$A w = b \quad (40)$$

with

$$A = \sum_{i=1}^m x_i x_i^\top \quad b = \sum_{i=1}^m y_i x_i \quad (41)$$

- ▶ the solution when A is invertible

$$w = A^{-1} b \quad (42)$$

The *inverse* of a square matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is denoted as \mathbf{A}^{-1} , which is the unique matrix such that

$$\mathbf{A}^{-1}\mathbf{A} = \mathbf{I} = \mathbf{A}\mathbf{A}^{-1} \quad (43)$$

- ▶ Not all matrices are invertible
 - ▶ Non-square matrices do not have inverses (by definition)
 - ▶ Not all square matrices are invertible
 - ▶ Not all symmetric matrices are invertible

Review: Symmetric Matrices

A symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is defined as

$$\mathbf{A}^T = \mathbf{A} \quad (44)$$

or, in other words,

$$a_{i,j} = a_{j,i} \quad \forall i, j \in [n] \quad (45)$$

Review: Eigen Decomposition

Every symmetric matrix \mathbf{A} can be decomposed as

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T \quad (46)$$

with

▶ $\mathbf{\Lambda} = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}$ as a diagonal matrix

▶ $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_n]$ is an orthogonal matrix

$$\langle \mathbf{u}_i, \mathbf{u}_i \rangle = \|\mathbf{u}_i\|_2^2 = 1 \text{ and } \langle \mathbf{u}_i, \mathbf{u}_j \rangle = 0$$

- ▶ If \mathbf{A} is invertible, the solution of the ERM problem is

$$w = \mathbf{A}^{-1}b \quad (47)$$

- ▶ If \mathbf{A} is invertible, the solution of the ERM problem is

$$\boldsymbol{w} = \mathbf{A}^{-1}\boldsymbol{b} \quad (47)$$

- ▶ If \mathbf{A} is not invertible, we have $\mathbf{A} = \mathbf{U}\boldsymbol{\Lambda}\mathbf{U}^\top$, and

$$\boldsymbol{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_k, \dots) \quad (48)$$

- ▶ Another common way of addressing the non-invertible issue is to add a constraint on \boldsymbol{w} as

$$L_{S,\ell_2}(h_{\boldsymbol{w}}) = \sum_{i=1}^m (h_{\boldsymbol{w}}(\boldsymbol{x}_i) - y_i)^2 + \lambda \|\boldsymbol{w}\|^2 \quad (49)$$

where λ is the regularization parameter

- ▶ Gradient of the new $L_S(h_{\boldsymbol{w}})$ as

$$\frac{dL_{S,\ell_2}(h_{\boldsymbol{w}})}{d\boldsymbol{w}} = 2 \sum_{i=1}^m (\langle \boldsymbol{w}, \boldsymbol{x}_i \rangle - y_i) \boldsymbol{x}_i + 2\lambda \boldsymbol{w} \quad (50)$$

- ▶ Solution: with the notations \mathbf{A} and b defined in Eq. (41)

$$w = (\mathbf{A} + \lambda \mathbf{I})^{-1} b \quad (51)$$

- ▶ Solution: with the notations \mathbf{A} and b defined in Eq. (41)

$$w = (\mathbf{A} + \lambda \mathbf{I})^{-1} b \quad (51)$$

- ▶ $\mathbf{A} + \lambda \mathbf{I}$ is invertible, when $d_i + \lambda \neq 0, \forall i$

$$\mathbf{A} + \lambda \mathbf{I} = \mathbf{U} \mathbf{D} \mathbf{U}^T + \lambda \mathbf{I} = \mathbf{U} (\mathbf{D} + \lambda \mathbf{I}) \mathbf{U}^T \quad (52)$$

- ▶ Solution: with the notations \mathbf{A} and b defined in Eq. (41)

$$w = (\mathbf{A} + \lambda \mathbf{I})^{-1} b \quad (51)$$


- ▶ $\mathbf{A} + \lambda \mathbf{I}$ is invertible, when $d_i + \lambda \neq 0, \forall i$

$$\mathbf{A} + \lambda \mathbf{I} = \mathbf{U} \mathbf{D} \mathbf{U}^T + \lambda \mathbf{I} = \mathbf{U} (\mathbf{D} + \lambda \mathbf{I}) \mathbf{U}^T \quad (52)$$

- ▶ Regularization will be further discussed in the next lecture on model selection

Polynomial Regression

Some learning tasks require nonlinear predictors with single variable $x \in \mathbb{R}$



$$h_w(x) = w_0 + w_1x + \cdots + w_nx^n \quad (53)$$

where $w = (w_0, w_1, \dots, w_n)$ is a vector of coefficients of size $n + 1$.

Polynomial Regression (II)

Given training examples $\{(x_i, y_i)\}_{i=1}^m$, the problem of polynomial regression

$$h_w(x) = w_0 + w_1x + \cdots + w_nx^n \quad (54)$$

can be converted to a linear regression problem

$$\begin{bmatrix} 1 & x_1 & \cdots & x_1^n \\ 1 & x_2 & \cdots & x_2^n \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_m & \cdots & x_m^n \end{bmatrix} \cdot \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix} \quad (55)$$

Polynomial Regression (II)

Given training examples $\{(x_i, y_i)\}_{i=1}^m$, the problem of polynomial regression

$$h_w(x) = w_0 + w_1x + \cdots + w_nx^n \quad (54)$$

can be converted to a linear regression problem


$$\begin{bmatrix} 1 & x_1 & \cdots & x_1^n \\ 1 & x_2 & \cdots & x_2^n \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_m & \cdots & x_m^n \end{bmatrix} \cdot \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix} \quad (55)$$

We will use polynomial regression as an example in the next section

ℓ_2 Regularization and Overfitting


Polynomial Regression: Data

Consider the following polynomial regression problem



Polynomial Regression: Data

Consider the following polynomial regression problem



Data generation process

$$y = \sin(x) + 0.3 * \epsilon \quad (56)$$

where $\epsilon \sim \mathcal{N}(0, 1)$

- ▶ We choose the hypothesis class of polynomial functions with degree 7

$$h_w(x) = w_0 + w_1x + w_2x^2 + \cdots + w_7x^7 \quad (57)$$

where $\{w_0, w_1, w_3, \dots, w_7\}$ are the parameters

- ▶ We choose the hypothesis class of polynomial functions with degree 7

$$h_w(x) = w_0 + w_1x + w_2x^2 + \dots + w_7x^7 \quad (57)$$

where $\{w_0, w_1, w_3, \dots, w_7\}$ are the parameters


- ▶ The loss function: MSE with ℓ_2 regularization

$$L_{S, \ell_2}(h_w) = \frac{1}{m} \sum_{i=1}^m (h_w(x_i) - y_i)^2 + \lambda \|\mathbf{w}\|^2 \quad (58)$$

where we can pick different values of $\lambda \in \{0, 1, 100\}$

Regression: Regularization Effects


The direct effect of regularization is to constrain the coefficient to be close to zero



Larger regularization parameter, stronger effect on reducing parameter values

Regression: Regularization for Avoiding Overfitting

By forcing the coefficient to be smaller, regularization can help avoid overfitting



Strong regularization effect will hurt the model performance.


Text classification

- ▶ Sentiment classification
- ▶ Dataset: Stanford Sentiment Treebank
- ▶ Model: Logistic regression in Sklearn
- ▶ Feature representation: Bag-of-words representation
 - ▶ Vocab size: 15K

Classification: Learning without Regularization

In the demo code, we chose $\lambda = \frac{1}{C} = 0.001$ to approximate the case without regularization.

- ▶ Training accuracy: 99.89%
- ▶ Val accuracy: 52.21%



Classification: Weights without Regularization

Here are some word features and their classification weights from the previous model without regularization. Positive weights indicate the word feature contribute to positive sentiment classification and negative weights indicate the opposite contribution

	interesting	pleasure	boring	zoe	write	workings
Without Reg	0.011	-5.63	1.80	-5.68	-8.20	14.16

Classification: Weights without Regularization

Here are some word features and their classification weights from the previous model without regularization. Positive weights indicate the word feature contribute to positive sentiment classification and negative weights indicate the opposite contribution

	interesting	pleasure	boring	zoe	write	workings
Without Reg	0.011	-5.63	1.80	-5.68	-8.20	14.16

- ▶ NEGATIVE: woody allen can **write** and deliver a one liner as well as anybody .

Classification: Weights without Regularization

Here are some word features and their classification weights from the previous model without regularization. Positive weights indicate the word feature contribute to positive sentiment classification and negative weights indicate the opposite contribution


	interesting	pleasure	boring	zoe	write	workings
Without Reg	0.011	-5.63	1.80	-5.68	-8.20	14.16

- ▶ NEGATIVE: woody allen can **write** and deliver a one liner as well as anybody .
- ▶ POSITIVE: soderbergh , like kubrick before him , may not touch the planet 's skin , but understands the **workings** of its spirit .

Classification: Learning with Regularization

We chose $\lambda = \frac{1}{C} = 10^2$

- ▶ Training accuracy: 62.54%
- ▶ Val accuracy: 63.17%



Classification: Weights with Regularization

With regularization, the classification weights make more sense to us

	interesting	pleasure	boring	zoe	write	workings
Without Reg	0.011	-5.63	1.80	-5.68	-8.20	14.16
With Reg	0.16	0.36	-0.21	-0.057	-0.066	0.040

Summary

- ▶ Perceptron
 - ▶ The hypothesis class (page 11)
 - ▶ Linearly separable cases (page 12)
 - ▶ Perceptron updating rule (page 14)
- ▶ Logistic regression
 - ▶ The hypothesis class (page 21 - 22)
 - ▶ Gradient-based updating rule (page 27 - 29)
 - ▶ Maximum likelihood estimation (page 32)
- ▶ Linear regression
 - ▶ The hypothesis class (page 35)
 - ▶ ℓ_2 regularization (page 46)
 - ▶ Maximum a posteriori (MAP) estimation (page 52)
- ▶ ℓ_2 Regularization and Overfitting



Friedman, J., Hastie, T., and Tibshirani, R. (2001).

The elements of statistical learning.

Springer.