

What is Machine Learning?

- **Machine Learning (ML)**: branch of artificial intelligence (AI) with goals to extract structures from large data sets
- Applications: identifying patterns and making decisions,
image recognition, fraud detection,
extracting meaning from text, ...
- Specifically, Face ID, Google Translate, self-driving cars,
cancer diagnosis, product recommendations,
ChatGPT, ...

- Recently, Large Language Models and Artificial General Intelligence have become hot issues.
- Clearly, AI and Machine Learning are making revolutions in many areas.
- Machine Learning brings together ideas from Computer Science, Statistics, and **Mathematics**.
- In this course, we will discuss some of the essential mathematical ideas in Machine Learning and practice computational techniques using Python.

Types of Machine Learning

- Regression:
predict numerical values
example: the value of a house
- Classification:
predict categorical labels
example: hand-written digits

Types of Machine Learning

- Supervised learning:
train a machine with known classification labels
- Unsupervised learning:
recognize patterns or clusters from an unlabelled dataset
- Reinforcement learning:
enable an agent to learn in an environment by trial and error

Typical ML problems

- Estimate the value of a house

3 bd | 2 ba | 1,996 sqft

● Off market

Zestimate®: **\$412,200** | Rent Zestimate®: **\$2,500**

Est. refi payment: \$2,496/mo [Refinance your loan](#)

[Home value](#) | [Owner tools](#) | [Home details](#) | [Neighborhood details](#)

Home value


Zestimate
\$412,200



Zestimate range

\$392,000 - \$433,000



Last 30-day change

+ \$5,239 (+1.3%)




Zestimate per sqft

\$207

\mathcal{F} : (# of bedrooms, square footage, year built, location, ...)
→ (the value of a house)

Typical ML problems


- Recognizing hand-written digits



$$\mathcal{C} : \{ \text{pictures} \} \longrightarrow \{0, 1, \dots, 9\}$$

Typical ML problems

- Clustering the data points of hand-written digits




TSNE visualization of MNIST data set encodings (VAE encoder). 10000 examples are shown.

Linear Regression – part (1)

- Input: $x \in \mathbb{R}$ Output: $t \in \mathbb{R}$
Observations: $(x_1, t_1), (x_2, t_2), \dots, (x_N, t_N)$
- Use these observations as **training examples**.
Task: Given a new input \tilde{x} , predict the output \tilde{t} .
- If t is a continuous variable, this task is called **regression**.
It is an example of **supervised learning**.

- For example, $x \in [0, 1)$, $N = 10$

x	t
0.884644066199	-0.864791215635069
0.793349886821	-1.32738612014193
0.735440841558	-1.18222466237236
0.421871764847	0.304255805886633
0.0118832729931	0.101594120287724
0.226770188973	1.13377458999431
0.978530671629	-0.147028527196347
0.0431076970157	0.247622971933151
0.890003286931	-0.605625802202937
0.888362799625	-0.649537521948140



- Fit the data using a polynomial

$$y(x, \mathbf{w}) = w_0 + w_1 x + w_2 x^2 + \cdots + w_K x^K,$$

where $\mathbf{w} = [w_0, w_1, \dots, w_K]^\top$.

- How to obtain the best approximation?
- Want: Minimize the distance between

$$(t_1, t_2, \dots, t_N) \quad \text{and} \quad (y_1, y_2, \dots, y_N),$$

where $y_n = y(x_n, \mathbf{w})$.

- Minimize the error function

$$\begin{aligned} E(\mathbf{w}) &= \sum_{n=1}^N \{y(x_n, \mathbf{w}) - t_n\}^2 \\ &= \sum_{n=1}^N \{w_0 + w_1 x_n + w_2 x_n^2 + \cdots + w_K x_n^K - t_n\}^2. \end{aligned}$$

- Introduce the following matrices

$$X = \begin{bmatrix} 1 & x_1 & \cdots & x_1^K \\ 1 & x_2 & \cdots & x_2^K \\ 1 & x_3 & \cdots & x_3^K \\ \vdots & \vdots & & \vdots \\ 1 & x_N & \cdots & x_N^K \end{bmatrix}, \quad \mathbf{w} = \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_K \end{bmatrix}, \quad \text{and } \mathbf{t} = \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_N \end{bmatrix}.$$

Then

$$E(\mathbf{w}) = \|\mathbf{X}\mathbf{w} - \mathbf{t}\|^2.$$

- The matrix X can be considered as a **data matrix**.

Tidy Convention:

- Each row is an **observation**.
- Each column is a **feature**.

- How can we find a minimum of a function?
- We take the “derivative”.
- $E(\mathbf{w})$ is a multi-variable function.

K features, and let $M = K + 1$.

- Let $f : \mathbb{R}^M \rightarrow \mathbb{R}$ be a differentiable function. The **gradient** of f at $\mathbf{x} \in \mathbb{R}^M$ is defined by

$$\nabla f(\mathbf{x}) = \begin{bmatrix} \frac{\partial f}{\partial x_1}(\mathbf{x}) \\ \frac{\partial f}{\partial x_2}(\mathbf{x}) \\ \vdots \\ \frac{\partial f}{\partial x_M}(\mathbf{x}) \end{bmatrix}.$$

Lemma

Assume that A is an $N \times M$ matrix and $\mathbf{x} \in \mathbb{R}^M$, $\mathbf{b} \in \mathbb{R}^N$. Then we have

$$\nabla \|A\mathbf{x} + \mathbf{b}\|^2 = 2A^\top(A\mathbf{x} + \mathbf{b}).$$

Proof. $A = [a_{ij}]$, $\mathbf{x} = [x_j]_{j=1,\dots,M}$, $\mathbf{b} = [b_i]_{i=1,\dots,N}$

$$\|\mathbf{Ax} + \mathbf{b}\|^2 = \|\left[\sum_j a_{ij}x_j + b_i\right]_{i=1,\dots,N}\|^2 = \sum_i \left(\sum_j a_{ij}x_j + b_i\right)^2$$

$$\left[\frac{\partial \|\mathbf{Ax} + \mathbf{b}\|^2}{\partial x_k}\right]_{k=1,\dots,N} = \left[\sum_i 2\left(\sum_j a_{ij}x_j + b_i\right)a_{ik}\right]_{k=1,\dots,N}$$

$$\begin{aligned} 2A^\top(\mathbf{Ax} + \mathbf{b}) &= 2[a_{ij}]^\top \left[\sum_j a_{ij}x_j + b_i\right]_{i=1,\dots,N} \\ &= 2\left[\sum_i a_{ik}\left(\sum_j a_{ij}x_j + b_i\right)\right]_{k=1,\dots,N} \end{aligned}$$



Theorem (Fermat's Theorem)

Let $f : \mathbb{R}^M \rightarrow \mathbb{R}$ be differentiable. If f has a local extremum at \mathbf{x} , then $\nabla f(\mathbf{x}) = 0$.

- Is the converse of Fermat's Theorem true?
- No, in general.

- A set $\mathcal{X} \in \mathbb{R}^M$ is said to be **convex** if for any $\mathbf{x}, \mathbf{y} \in \mathcal{X}$,

$$\{\alpha \mathbf{x} + (1 - \alpha) \mathbf{y} : 0 \leq \alpha \leq 1\} \subseteq \mathcal{X}.$$

- Let \mathcal{X} be a convex set. A function $f : \mathcal{X} \rightarrow \mathbb{R}$ is said to be **convex** if for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ and $\alpha \in [0, 1]$,

$$f(\alpha \mathbf{x} + (1 - \alpha) \mathbf{y}) \leq \alpha f(\mathbf{x}) + (1 - \alpha) f(\mathbf{y}).$$

Theorem

Assume \mathcal{X} is convex. Then a differentiable function $f : \mathcal{X} \rightarrow \mathbb{R}$ is convex if and only if for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$,

$$f(\mathbf{y}) - f(\mathbf{x}) \geq \nabla f(\mathbf{x}) \cdot (\mathbf{y} - \mathbf{x}) = \nabla f(\mathbf{x})^\top (\mathbf{y} - \mathbf{x}).$$

Corollary

Assume \mathcal{X} is convex. Then a convex differentiable function $f : \mathcal{X} \rightarrow \mathbb{R}$ has a global minimum at \mathbf{x} if and only if $\nabla f(\mathbf{x}) = 0$.

Proof.

(\Rightarrow) Fermat's Theorem

(\Leftarrow) Theorem in the above



X : $N \times M$ data matrix $\mathbf{t} \in \mathbb{R}^N$

Proposition

The function $E(\mathbf{w}) = \|X\mathbf{w} - \mathbf{t}\|^2$ is convex and differentiable.

Proof.

- $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$ (triangle inequality)
- x^2 is convex.

Now we have

$$\begin{aligned} E(\alpha\mathbf{w} + (1 - \alpha)\mathbf{v}) &= \|X(\alpha\mathbf{w} + (1 - \alpha)\mathbf{v}) - \mathbf{t}\|^2 \\ &= \|\alpha(X\mathbf{w} - \mathbf{t}) + (1 - \alpha)(X\mathbf{v} - \mathbf{t})\|^2 \\ &\leq (\alpha\|X\mathbf{w} - \mathbf{t}\| + (1 - \alpha)\|X\mathbf{v} - \mathbf{t}\|)^2 \\ &\leq \alpha E(\mathbf{w}) + (1 - \alpha)E(\mathbf{v}). \end{aligned}$$



- Since $\nabla E(\mathbf{w}) = 2X^\top(X\mathbf{w} - \mathbf{t})$, the convex function $E(\mathbf{w})$ has a global minimum if and only if

$$X^\top X\mathbf{w} = X^\top \mathbf{t}.$$

- When $S := X^\top X$ is invertible, there is a unique solution

$$\boxed{\mathbf{w} = S^{-1}X^\top \mathbf{t}}.$$


- S is not invertible if and only if $\det S = 0$.

When $\det S \approx 0$, the formula for \mathbf{w} is unstable.

- If $\det S = 0$, what can we say about the features?

- Back to the example, $x \in [0, 1)$, $N = 10$


x	t
0.884644066199	-0.864791215635069
0.793349886821	-1.32738612014193
0.735440841558	-1.18222466237236
0.421871764847	0.304255805886633
0.0118832729931	0.101594120287724
0.226770188973	1.13377458999431
0.978530671629	-0.147028527196347
0.0431076970157	0.247622971933151
0.890003286931	-0.605625802202937
0.888362799625	-0.649537521948140




Using the formula, we obtain the following.

w	$K = 1$	$K = 3$	$K = 6$	$K = 9$
w_0	0.53	-0.10	0.03	-1.61
w_1	-1.41	12.21	4.49	200.18
w_2		-37.18	34.15	-5191.34
w_3		25.37	-211.33	41628.04
w_4			339.78	-141666.70
w_5			-210.56	210688.85
w_6			43.45	-61808.88
w_7				-189181.01
w_8				214844.27
w_9				-69519.92


$$K = 1$$




$K = 3$



$$K = 6$$



$$K = 9$$



- If we choose $K = 6$ as our model, then we get

$$y(x) = 0.03 + 4.49x + 34.15x^2 - 211.33x^3 \\ + 339.78x^4 - 201.78x^5 + 43.45x^6.$$

- For a new input $x = 0.741234$, the **prediction** is

$$y(0.741234) = -1.28212.$$

- This is how machine learning works in this example.

- The case $K = 9$ is **over-fitting**.

It is a perfect fit. However, it is useless in machine learning.

- In order to avoid over-fitting, we can use **regularization**.

Observe that when $K = 9$ the coordinates of \mathbf{w} tend to be large.

- Ridge regression

$$\tilde{E}(\mathbf{w}) = \sum_{n=1}^N \{y(x_n, \mathbf{w}) - t_n\}^2 + \lambda \|\mathbf{w}\|^2.$$

This can be considered as a result of Bayesian Learning.

- Lasso regression

$$\hat{E}(\mathbf{w}) = \sum_{n=1}^N \{y(x_n, \mathbf{w}) - t_n\}^2 + \lambda \sum_{n=0}^K |w_n|.$$

- For other situations, we can take different basis functions $\phi_n(x)$ instead of power functions x^n , and consider

$$y(x, \mathbf{w}) = w_0 + w_1\phi_1(x) + w_2\phi_2(x) + \cdots + w_k\phi_K(x).$$

- For example,

$$\phi_n(x) = \exp\left(-\frac{(x-\mu_n)^2}{2\sigma^2}\right).$$

- Minimize the error function

$$E(\mathbf{w}) = \sum_{n=1}^N \{y(x_n, \mathbf{w}) - t_n\}^2.$$

- We have

$$X = \begin{bmatrix} 1 & \phi_1(x_1) & \cdots & \phi_K(x_1) \\ 1 & \phi_1(x_2) & \cdots & \phi_K(x_2) \\ 1 & \phi_1(x_3) & \cdots & \phi_K(x_3) \\ \vdots & \vdots & & \vdots \\ 1 & \phi_1(x_N) & \cdots & \phi_K(x_N) \end{bmatrix}.$$

- We obtain the solution

$$\boxed{\mathbf{w} = S^{-1} X^\top \mathbf{t}},$$

where $S = X^\top X$.

Linear Regression – part (2)

- Input: $(x_1, x_2, \dots, x_K) \in \mathbb{R}^K$ Output: $t \in \mathbb{R}$
Observations: $(x_{11}, x_{12}, \dots, x_{1K}; t_1), \dots, (x_{N1}, x_{N2}, \dots, x_{NK}; t_N)$
- Use these observations as **training examples**.
Task: Given a new input $(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_K)$, predict the output \tilde{t} .

- Fit the data using a linear function

$$y(\mathbf{x}, \mathbf{w}) = w_0 + w_1 x_1 + w_2 x_2 + \cdots + w_K x_K,$$

where $\mathbf{w} = [w_0, w_1, \dots, w_K]^\top$.

- w_0 is called a **bias**.

- Want: Minimize the distance between

$$(t_1, t_2, \dots, t_N) \quad \text{and} \quad (y_1, y_2, \dots, y_N),$$

where $y_n = y(\mathbf{x}_n, \mathbf{w})$, $n = 1, 2, \dots, N$.

- Minimize the error function

$$\begin{aligned} E(\mathbf{w}) &= \sum_{n=1}^N \{y_n - t_n\}^2 = \sum_{n=1}^N \{y(\mathbf{x}_n, \mathbf{w}) - t_n\}^2 \\ &= \sum_{n=1}^N \{w_0 + w_1 x_{n1} + w_2 x_{n2} + \dots + w_K x_{nK} - t_n\}^2. \end{aligned}$$

- Introduce the following matrices

$$X = \begin{bmatrix} 1 & x_{11} & \cdots & x_{1K} \\ 1 & x_{21} & \cdots & x_{2K} \\ 1 & x_{31} & \cdots & x_{3K} \\ \vdots & \vdots & & \vdots \\ 1 & x_{N1} & \cdots & x_{NK} \end{bmatrix}, \quad \mathbf{w} = \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_K \end{bmatrix}, \quad \text{and } \mathbf{t} = \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_N \end{bmatrix}.$$

Then

$$E(\mathbf{w}) = \|\mathbf{X}\mathbf{w} - \mathbf{t}\|^2.$$

- The matrix X is the **data matrix**.

- Since $\nabla E(\mathbf{w}) = 2X^\top(X\mathbf{w} - \mathbf{t})$, the convex function $E(\mathbf{w})$ has a global minimum if and only if

$$X^\top X\mathbf{w} = X^\top \mathbf{t}.$$

- When $S := X^\top X$ is invertible, there is a unique solution

$$\boxed{\mathbf{w} = S^{-1}X^\top \mathbf{t}}.$$

- When $(\tilde{x}_1, \dots, \tilde{x}_K)$ is a new input, the prediction is given by

$$\tilde{t} = w_0 + w_1\tilde{x}_1 + \dots + w_K\tilde{x}_K = [1, \tilde{x}_1, \dots, \tilde{x}_K]\mathbf{w}.$$

Geometry of Linear Regression

- $\tilde{\mathbf{t}} := X\mathbf{w}$ is the best approximation (or the closest point) to $\mathbf{t} \in \mathbb{R}^N$.
- $\tilde{\mathbf{t}}$ is in the column space of X .
- $\tilde{\mathbf{t}} - \mathbf{t}$ is orthogonal to the column space of X .
($\because \nabla E(\mathbf{w}) = 2X^\top(X\mathbf{w} - \mathbf{t}) = 0$ if and only if $X^\top(\tilde{\mathbf{t}} - \mathbf{t}) = 0$)

if and only if $\mathbf{c}_n \cdot (\tilde{\mathbf{t}} - \mathbf{t}) = 0$ for all $n = 0, 1, \dots, K$,

where \mathbf{c}_n are the column vectors of X .)

- $\tilde{\mathbf{t}}$ is the orthogonal projection of \mathbf{t} on to the column space of X .
- column space = feature space

Proposition

The matrix $S := X^\top X$ is not invertible if and only if the columns of X are linearly dependent.

Proof.

- \Rightarrow) There exists $\mathbf{w} \neq 0$ such that $S\mathbf{w} = X^\top X\mathbf{w} = 0$. Then $X\mathbf{w}$ is orthogonal to the column space of X . However, $X\mathbf{w}$ is actually in the column space of X . Therefore $X\mathbf{w} = 0$, and the columns of X are linearly dependent.
- \Leftarrow) There exists $\mathbf{w} \neq 0$ such that $X\mathbf{w} = 0$. Then $S\mathbf{w} = X^\top X\mathbf{w} = 0$, implying that S is not invertible.

Binary Classification

- Given a dataset \mathcal{D} , we like to construct a function

$$f : \mathcal{D} \rightarrow \{0, 1\},$$

where 0 = (category 0) and 1 = (category 1).

- Examples

- 1 $\mathcal{D} = \{\text{SAT scores}\}$, 0 = rejection, 1 = admission
- 2 $\mathcal{D} = \{\text{emails}\}$, 0 = non-spam, 1 = spam

- More precisely, we will construct a function

$$f : \mathcal{D} \rightarrow [0, 1]$$

so that $f(d)$, $d \in \mathcal{D}$, represents the **probability**
that d belongs to (category 1).

- In the SAT scores example,

$$f(1350) = 0.732$$

will mean “A student with SAT score 1350 is accepted with probability 0.732.”

Q: How can we construct such a function? **Logistic Regression**

- In linear regression, we use a linear model and minimize the mean square error (MSE).
- In **logistic regression**, our strategy will be similar to that of linear regression, and the method is called **maximum likelihood estimation** (MLE).

A dataset \mathcal{D} is given with known classification.

- Step 1: Using a **probabilistic model**, write a function out of \mathcal{T} with unknown *parameters* or *weights*.
- Step 2: Determine the parameters so that the known classification may have the **maximum likelihood**.

It is customary to take the negative log of the likelihood function, and the resulting function is called the **cross-entropy**. Then we need to *minimize* the cross-entropy.

Learning \iff Maximizing Certainty
 \iff Minimizing Randomness (Entropy)

- $\mathcal{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ with known classification $\{t_1, \dots, t_N\}$
- $y_n = f(\mathbf{x}_n, \mathbf{w})$: probability that \mathbf{x}_n belongs to class 1
Then $1 - y_n$ is the probability that \mathbf{x}_n belongs to class 0.
- Apply this to $\{t_1, \dots, t_N\}$ to compute the total probability.

- For example,

t_n	1	0	1	1	0
p_n	y_1	$1 - y_2$	y_3	y_4	$1 - y_5$

Assuming that data points are independent, we obtain

$$\text{likelihood} = y_1(1 - y_2)y_3y_4(1 - y_5).$$

- We can write

$$\text{likelihood} = \prod_{n=1}^N y_n^{t_n} (1 - y_n)^{1-t_n},$$

where $y_n = f(\mathbf{x}_n, \mathbf{w})$.

- Want: Determine \mathbf{w} that maximizes the likelihood.
- Equivalently, minimize the cross-entropy

$$E(\mathbf{w}) := -\log(\text{likelihood}).$$


- In logistic regression, we take

$$y_n = f(\mathbf{x}_n, \mathbf{w}) = \sigma(w_0 + w_1 x_{n1} + w_2 x_{n2} + \cdots + w_K x_{nK}),$$

where $\mathbf{x}_n = (x_{n1}, \dots, x_{nK})$, $\mathbf{w} = (w_0, w_1, \dots, w_K)$ and

$$\sigma(x) := \frac{1}{1 + e^{-x}}.$$

- The function $\sigma(x)$ is called the **sigmoid** function.



The cross-entropy function is given by

$$E(\mathbf{w}) = - \sum_{n=1}^N \{t_n \ln y_n + (1 - t_n) \ln(1 - y_n)\},$$

where

$$y_n = \sigma(w_0 + w_1 x_{n1} + w_2 x_{n2} + \cdots + w_K x_{nK})$$

and

$$\sigma(x) = \frac{1}{1 + e^{-x}}.$$

Goal: Determine $\mathbf{w} = (w_0, w_1, w_2, \dots, w_K)$ which minimizes $E(\mathbf{w})$.

- Compute the gradient $\nabla E(\mathbf{w})$.
- Crucial identity:

$$\sigma'(x) = \sigma(x)(1 - \sigma(x))$$

Thus it is a solution to

$$y' = y(1 - y),$$

which is called a **logistic equation**.

$$(\ln y)' = \frac{1}{y}y' = 1 - y.$$

$$(\ln(1 - y))' = \frac{1}{1 - y}(-y') = -y.$$

$$E(\mathbf{w}) = - \sum_{n=1}^N \{t_n \ln y_n + (1 - t_n) \ln(1 - y_n)\},$$

$$y_n = \sigma(w_0 + w_1 x_{n1} + w_2 x_{n2} + \cdots + w_K x_{nK}).$$

Write $x_n = w_0 + w_1 x_{n1} + w_2 x_{n2} + \cdots + w_K x_{nK}$ and $y_n = \sigma(x_n)$.

$$\begin{aligned} \frac{\partial E}{\partial w_j} &= - \sum \frac{\partial}{\partial w_j} \{t_n \ln y_n + (1 - t_n) \ln(1 - y_n)\} \\ &= - \sum \left\{ t_n \frac{\partial}{\partial x_n} \ln y_n \cdot \frac{\partial x_n}{\partial w_j} + (1 - t_n) \frac{\partial}{\partial x_n} \ln(1 - y_n) \cdot \frac{\partial x_n}{\partial w_j} \right\} \\ &= - \sum \{t_n(1 - y_n)x_{nj} + (1 - t_n)(-y_n)x_{nj}\} \\ &= \sum (y_n - t_n)x_{nj} \end{aligned}$$

$$\boxed{\frac{\partial E}{\partial w_j} = \sum_{n=1}^N (y_n - t_n)x_{nj} .}$$

- We have

$$\nabla E(\mathbf{w}) = \left[\sum_{n=1}^N (y_n - t_n) x_{nj} \right] = X^{\top}(\mathbf{y} - \mathbf{t}),$$

where

$$X = \begin{bmatrix} 1 & x_{11} & \cdots & x_{1K} \\ 1 & x_{21} & \cdots & x_{2K} \\ \vdots & & \vdots & \vdots \\ 1 & x_{N1} & \cdots & x_{NK} \end{bmatrix}, \quad \mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}, \quad \text{and } \mathbf{t} = \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_N \end{bmatrix}.$$

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

- Assume $\nabla E = 0$.

Can we determine \mathbf{w} explicitly?

When $j = 0$, we have $x_{n,0} = 1$ for all n , and thus

$$\sum_{n=1}^N y_n = \sum_{n=1}^N t_n.$$

Therefore, we have

$$\begin{aligned}\sum_{n=1}^N y_n &= \sum_{n=1}^N \frac{1}{1 + \exp(-(w_0 + w_1 x_{n1} + w_2 x_{n2} + \cdots + w_K x_{nK}))} \\ &= \sum_{n=1}^N t_n.\end{aligned}$$

We have more complicated equations for $j \neq 0$. Clearly, it is not feasible to write w_0, w_1, \dots, w_K explicitly.

- The cross-entropy $E(\mathbf{w})$ is not linear, and it is **not possible** to calculate a **closed-form formula** for \mathbf{w}_* which minimizes $E(\mathbf{w})$.
- On the other hand, it can be shown that $E(\mathbf{w})$ is convex, and a global minimum exists.
- We will find an approximate value for \mathbf{w}_* using **gradient descent** and **Newton's method**.
- The method of gradient descent is widely used in many other parts of machine learning.

Gradient Descent

Consider a function $E : \mathbb{R}^M \rightarrow \mathbb{R}$, $\mathbf{w} = (w_1, w_2, \dots, w_M) \mapsto E(\mathbf{w})$. The **gradient** ∇E of E is defined by

$$\nabla E := \left(\frac{\partial E}{\partial w_1}, \frac{\partial E}{\partial w_2}, \dots, \frac{\partial E}{\partial w_M} \right).$$

Proposition

$E(\mathbf{w})$: differentiable in a nbhd of \mathbf{w}

The function $E(\mathbf{w})$ decreases fastest in the direction of $-\nabla E(\mathbf{w})$.

Proof: For a unit vector \mathbf{u} , the directional derivative $D_{\mathbf{u}}E$ is given by

$$D_{\mathbf{u}}E = \lim_{t \rightarrow 0} \frac{E(\mathbf{w}_0 + t\mathbf{u}) - E(\mathbf{w}_0)}{t} = g'(0),$$

where $g(t) = E(\mathbf{w}_0 + t\mathbf{u})$. Let $\mathbf{w} = \mathbf{w}_0 + t\mathbf{u}$. Using the chain rule,

$$g'(0) = \sum_{i=1}^m \frac{\partial E}{\partial w_i} \cdot \frac{dw_i}{dt} = \sum_{i=1}^m \frac{\partial E}{\partial w_i} \cdot \frac{du_i}{dt} = \nabla E \cdot \mathbf{u}.$$

Furthermore,

$$D_{\mathbf{u}}E = \nabla E \cdot \mathbf{u} = |\nabla E| |\mathbf{u}| \cos \theta = |\nabla E| \cos \theta,$$

where θ is the angle between ∇E and \mathbf{u} . The minimum value of $D_{\mathbf{u}}E$ occurs when $\cos \theta$ is -1 . \square

- Choose an initial point \mathbf{w}_0 .
- Set

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta_k \nabla E(\mathbf{w}_k)$$

where η_k is the step size or **learning rate**.

Usually, for some $\eta > 0$, we set $\eta_k = \eta$ for all k .

- From Proposition, we have

$$E(\mathbf{w}_k) \geq E(\mathbf{w}_{k+1}).$$

- Under some moderate conditions,

$$E(\mathbf{w}_k) \rightarrow \text{local minimum} \quad \text{as } k \rightarrow \infty.$$

In particular, this is true when E is **convex**.

- E is convex \iff

for all $0 \leq \alpha \leq 1$ and all $x_1, x_2 \in \mathbb{R}^n$, we have

$$E(\alpha x_1 + (1 - \alpha)x_2) \leq \alpha E(x_1) + (1 - \alpha)E(x_2).$$

- $\mathbf{w}_* = \lim_{k \rightarrow \infty} \mathbf{w}_k$

We use \mathbf{w}_k for a sufficiently large k as an approximation of \mathbf{w}_* .

This method is called **gradient descent**.

Caveat: Making a right choice of η is crucial.

Example

- Consider $E(\mathbf{w}) = E(w_1, w_2) = w_1^4 + w_2^4 - 16w_1 w_2$.

Then $\nabla E(\mathbf{w}) = [4w_1^3 - 16w_2, 4w_2^3 - 16w_1]$.

Choose $\mathbf{w}_0 = (1, 1)$ and $\eta = 0.01$.

$$\mathbf{w}_{30} = (1.99995558586289, 1.99995558586289)$$

$$E(\mathbf{w}_{30}) = -31.99999999368777$$

- We see that $\mathbf{w}_k \rightarrow (2, 2)$ and $E(2, 2) = -32$.
- Indeed, when $\mathbf{w} = (2, 2)$, a local minimum of $E(\mathbf{w})$ is -32 .

k	w1	w2	E(w1,w2)
1	1.120000000000000	1.120000000000000	-16.9233612800000
2	1.243002880000000	1.243002880000000	-19.9465014818312
3	1.36506297054983	1.36506297054983	-22.8698545020842
4	1.48172688079195	1.48172688079195	-25.4876645161458
5	1.58867706472624	1.58867706472624	-27.6422269714610
6	1.68247924276483	1.68247924276483	-29.2656452783487
7	1.76116971206054	1.76116971206054	-30.3861816086105
8	1.82445074094736	1.82445074094736	-31.0984991504577
9	1.87344669354831	1.87344669354831	-31.5194128485897
10	1.91018104795404	1.91018104795404	-31.7533053700606
11	1.93701591038872	1.93701591038872	-31.8770223901250
12	1.95622873443784	1.95622873443784	-31.9400248989010
13	1.96977907222858	1.96977907222858	-31.9712142030755
14	1.97923168007769	1.97923168007769	-31.9863406140263
15	1.98577438322011	1.98577438322011	-31.9935701975589
16	1.99027812738069	1.99027812738069	-31.9969902100773
17	1.99336647981957	1.99336647981957	-31.9985965516271
18	1.99547865709166	1.99547865709166	-31.9993473166738
19	1.99692058430943	1.99692058430943	-31.9996970174121
20	1.99790372262623	1.99790372262623	-31.9998595272283
21	1.99857347710339	1.99857347710339	-31.9999349274762
22	1.99902947615421	1.99902947615421	-31.9999698732955
23	1.99933981776146	1.99933981776146	-31.9999860577045
24	1.99955097148756	1.99955097148756	-31.9999935493971
25	1.99969461222478	1.99969461222478	-31.9999970160815
26	1.99979231393118	1.99979231393118	-31.9999986198712
27	1.99985876312152	1.99985876312152	-31.9999993617137
28	1.99990395413526	1.99990395413526	-31.9999997048203
29	1.99993468659806	1.99993468659806	-31.9999998634976
30	1.99995558586289	1.99995558586289	-31.9999999368777

Back to Binary Classification

$$E(\mathbf{w}) = - \sum_{n=1}^N \{t_n \ln y_n + (1 - t_n) \ln(1 - y_n)\},$$

$$y_n = \sigma(\mathbf{w}_0 + \mathbf{w}_1 x_{n1} + \mathbf{w}_2 x_{n2} + \cdots + \mathbf{w}_K x_{nK}).$$

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

- We will perform gradient descent using this.

Newton's Method

- Second-order approximation

Much faster in convergence, more expensive (and more subtle)

- $f(x)$: single-variable, convex, differentiable function

Find a local minimum

$$\iff \text{Find } x_* \text{ such that } f'(x_*) = 0$$

Choose an initial point x_0 and set $x = x_0 + h$.

- Using Taylor's expansion,

$$\begin{aligned}f(x) &= f(x_0 + h) \approx f(x_0) + f'(x_0)h + \frac{1}{2}f''(x_0)h^2 \\f'(x) &\approx \frac{d}{dh} \left(f(x_0) + f'(x_0)h + \frac{1}{2}f''(x_0)h^2 \right) \frac{dh}{dx} \\&= f'(x_0) + f''(x_0)h\end{aligned}$$

From $f'(x) = 0$, we approximately obtain

$$0 = f'(x_0) + f''(x_0)h, \quad h = -f'(x_0)/f''(x_0).$$

- We have shown that

$$x_1 = x_0 - f'(x_0)/f''(x_0)$$

is an approximation of x_* .

- Repeat the process to obtain

$$x_{k+1} = x_k - f'(x_k)/f''(x_k),$$

and $x_k \rightarrow x_*$ as $k \rightarrow \infty$.

- This is Newton's method for a single-variable function, and we generalize it to a multi-variable function.

- $F(\mathbf{x})$: multi-variable, convex, differentiable function

The **Hessian matrix** is defined by

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

In short, $\mathbf{H}F = [\frac{\partial^2 F}{\partial x_i \partial x_j}]$.

- Recall

$$x_{k+1} = x_k - f'(x_k)/f''(x_k).$$

- Generalizing the single-variable case,

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \mathbf{H}F(\mathbf{x}_k)^{-1} \nabla F(\mathbf{x}_k).$$

Proof: Using Taylor's expansion, we have for $\mathbf{h} \in \mathbb{R}^m$,

$$F(\mathbf{x}) = \boxed{F(\mathbf{x}_0 + \mathbf{h}) \approx F(\mathbf{x}_0) + \mathbf{h}^\top \nabla F(\mathbf{x}_0) + \frac{1}{2} \mathbf{h}^\top \mathbf{H} F(\mathbf{x}_0) \mathbf{h}}.$$

Write $\mathbf{H} F(\mathbf{x}_0) = [H_{k\ell}]$, and

$$\begin{aligned} \mathbf{h}^\top \nabla F(\mathbf{x}_0) &= \sum_k \frac{\partial F}{\partial x_k} h_k, \\ \mathbf{h}^\top \mathbf{H} F(\mathbf{x}_0) \mathbf{h} &= \sum_{k,\ell} H_{k\ell} h_k h_\ell. \end{aligned}$$

We obtain

$$\begin{aligned} \frac{\partial F}{\partial x_i}(\mathbf{x}) &\approx \frac{\partial}{\partial h_i} \left(F(\mathbf{x}_0) + \mathbf{h}^\top \nabla F(\mathbf{x}_0) + \frac{1}{2} \mathbf{h}^\top \mathbf{H} F(\mathbf{x}_0) \mathbf{h} \right) \\ &= \frac{\partial F}{\partial x_i} + \sum_{k=1}^M H_{ik} h_k. \end{aligned}$$

Thus

$$\nabla F(\mathbf{x}) \approx \nabla F(\mathbf{x}_0) + \mathbf{H}F(\mathbf{x}_0)\mathbf{h}.$$

From $\nabla F(\mathbf{x}) = 0$, we approximately obtain

$$\mathbf{h} = -\mathbf{H}F(\mathbf{x}_0)^{-1}\nabla F(\mathbf{x}_0).$$



- Using a step size η_k , the formula may be modified to be

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \eta_k \mathbf{H}F(\mathbf{x}_k)^{-1} \nabla F(\mathbf{x}_k).$$

- Newton's method is much faster than gradient descent. However, it may be expensive to compute $\mathbf{H}F(\mathbf{x}_k)^{-1}$. Sometimes, $\mathbf{H}F(\mathbf{x}_k)$ is close to a singular matrix.

Example

- Consider $E(\mathbf{w}) = E(w_1, w_2) = w_1^4 + w_2^4 - 16w_1 w_2$.

Then $\nabla E(\mathbf{w}) = [4w_1^3 - 16w_2, 4w_2^3 - 16w_1]^\top$.

$$\mathbf{H}E(\mathbf{w}) = \begin{bmatrix} 12w_1^2 & -16 \\ -16 & 12w_2^2 \end{bmatrix}$$

$$\mathbf{H}E(\mathbf{w})^{-1} = \frac{1}{9w_1^2 w_2^2 - 16} \begin{bmatrix} \frac{3}{4}w_2^2 & 1 \\ 1 & \frac{3}{4}w_1^2 \end{bmatrix}$$

$$\mathbf{H}E^{-1} \nabla E = \frac{1}{9w_1^2 w_2^2 - 16} \begin{bmatrix} 3w_1^3 w_2^2 - 8w_2^3 - 16w_1 \\ 3w_1^2 w_2^3 - 8w_1^3 - 16w_2 \end{bmatrix}$$

- Choose $\mathbf{w}_0 = (1, 1)$ and $\eta = 1$.
Then $\mathbf{w}_1 = (2, 2)$ and $E(\mathbf{w}_1) = -32$.
- Choose $\mathbf{w}_0 = (1.2, 1.2)$ and $\eta = 1$.
Then $\mathbf{w}_9 = (2.00000004189571, 2.00000004189571)$,
 $E(\mathbf{w}_9) = -31.99999999999999$.

k	w1	w2	E(w1,w2)
0	1.200000000000000	1.200000000000000	-18.8928000000000
1	10.800000000000000	10.800000000000000	25343.5392000001
2	7.28325624421832	7.28325624421832	4778.98521693644
3	4.98069646698406	4.98069646698406	833.890570717962
4	3.50906808575457	3.50906808575457	106.230520855080
5	2.62345045192591	2.62345045192591	-15.3824765840014
6	2.16920289601164	2.16920289601164	-31.0047054152139
7	2.01793795417254	2.01793795417254	-31.9896107961456
8	2.00023638179330	2.00023638179330	-31.9999982117454
9	2.00000004189571	2.00000004189571	-31.9999999999999
10	2.000000000000000	2.000000000000000	-32.0000000000000

- Apply Newton's method to our main example:

$$E(\mathbf{w}) = - \sum_{n=1}^N \{t_n \ln y_n + (1 - t_n) \ln(1 - y_n)\},$$

where $y_n = \sigma(w_0 + w_1 x_{n1} + w_2 x_{n2} + \cdots + w_K x_{nK})$.

- Recall that $\sigma'(x) = \sigma(x)(1 - \sigma(x))$.

- We have

$$\nabla E(\mathbf{w}) = \left[\sum_{n=1}^N (y_n - t_n) x_{nj} \right] = X^T (\mathbf{y} - \mathbf{t}),$$

where

$$X = \begin{bmatrix} 1 & x_{11} & \cdots & x_{1K} \\ 1 & x_{21} & \cdots & x_{2K} \\ \vdots & & \vdots & \vdots \\ 1 & x_{N1} & \cdots & x_{NK} \end{bmatrix}, \quad \mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}, \quad \text{and } \mathbf{t} = \begin{bmatrix} t_1 \\ t_2 \\ \vdots \\ t_N \end{bmatrix}.$$

- Calculate $\frac{\partial^2 E}{\partial w_i \partial w_j}$.

- Recall $\frac{\partial E}{\partial w_j} = \sum_{n=1}^N (y_n - t_n) x_{nj}$.

- Here $y_n = \sigma(w_0 + w_1 x_{n1} + w_2 x_{n2} + \cdots + w_K x_{nK})$.

Write $x_n = w_0 + w_1 x_{n1} + w_2 x_{n2} + \cdots + w_K x_{nK}$ and $y_n = \sigma(x_n)$.

$$\begin{aligned} \frac{\partial^2 E}{\partial w_i \partial w_j} &= \frac{\partial}{\partial w_i} \sum_{n=1}^N (y_n - t_n) x_{nj} = \sum_{n=1}^N \frac{\partial (y_n - t_n) x_{nj}}{\partial x_n} \frac{\partial x_n}{\partial w_i} \\ &= \sum_{n=1}^N y_n (1 - y_n) x_{nj} x_{ni} \end{aligned}$$

Thus we have

$$\boxed{\frac{\partial^2 E}{\partial w_i \partial w_j} = \sum_{n=1}^N y_n (1 - y_n) x_{ni} x_{nj} .}$$



$$\frac{\partial^2 E}{\partial \mathbf{w}_i \partial \mathbf{w}_j} = \sum_{n=1}^N y_n(1 - y_n) x_{ni} x_{nj}$$

- We get

$$\mathbf{H}E = \left[\sum_{n=1}^N y_n(1 - y_n) x_{ni} x_{nj} \right] = \mathbf{X}^\top \mathbf{R} \mathbf{X},$$

where $\mathbf{R} = \text{diag}(y_n(1 - y_n))$.

- Then we have

$$\boxed{\mathbf{w}_{k+1} = \mathbf{w}_k - (\mathbf{X}^\top \mathbf{R} \mathbf{X})^{-1} \mathbf{X}^\top (\mathbf{y} - \mathbf{t})},$$

where \mathbf{R} and \mathbf{y} are determined by \mathbf{w}_k in each step.

Stochastic Gradient Descent (SGD)

- Typically in Machine Learning, we minimize a function $E(\mathbf{w})$ given by a sum of the form

$$E(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N E_n(\mathbf{w}),$$

where N is the number of elements in the training set.

- When N is large, computation of the gradient ∇E may be expensive.

- The SGD selects a **sample** from the training set in each iteration step instead of using the whole batch of the training set, and use

$$\frac{1}{L} \sum_{i=1}^L \nabla E_{n_i}(\mathbf{w}),$$

where L is the size of the sample and

$$\{n_1, n_2, \dots, n_L\} \subset \{1, 2, \dots, N\}.$$

- The SGD is commonly used in implementations of many Machine Learning algorithms.

Suppose that we are performing an experiment.

- S : set of all outcomes, called the **sample space**
- An **event** is a subset of S .

Example: Roll a die. Then the sample space is given by

$$S = \{1, 2, 3, 4, 5, 6\}.$$

The subset $A = \{1, 3, 5\}$ is an event, which corresponds to the statement “The outcome is an odd number”.

Assume that S is a sample space.

- A **probability** is a function $P : \{\text{events}\} \rightarrow [0, 1]$ satisfying
 - 1 $P(S) = 1$;
 - 2 if E_1, E_2, \dots are events such that $E_i \cap E_j = \emptyset$ for $i \neq j$, then

$$P\left(\bigcup_i E_i\right) = \sum_i P(E_i).$$

Example: Roll a die. Then we have the usual probability defined by

$$P(\{1\}) = P(\{2\}) = P(\{3\}) = P(\{4\}) = P(\{5\}) = P(\{6\}) = \frac{1}{6}.$$


We obtain

$$P(\{1, 3, 5\}) = P(\{1\} \cup \{3\} \cup \{5\}) = \frac{1}{6} + \frac{1}{6} + \frac{1}{6} = \frac{3}{6} = \frac{1}{2}.$$

Example: Let $S = \mathbb{R}$. Define a function $f : S \rightarrow \mathbb{R}$ by

$$f(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}.$$

The graph of f is



For an event $A = [a, b] \subset S = \mathbb{R}$, define

$$P(A) = \int_A f(x)dx = \int_a^b f(x)dx.$$

Then P is a probability.

Proof: We have

$$P(S) = \int_S f(x) dx = \int_{-\infty}^{\infty} f(x) dx.$$

$$\begin{aligned} P(S)^2 &= \int_{-\infty}^{\infty} f(x) dx \int_{-\infty}^{\infty} f(y) dy = \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-(x^2+y^2)/2} dx dy \\ &= \frac{1}{2\pi} \int_0^{2\pi} \int_0^{\infty} e^{-r^2/2} r dr d\theta = \int_0^{\infty} e^{-u} du = 1 \end{aligned}$$

Thus $P(S) = 1$. The second condition is satisfied by the properties of an integral. □

Conditional Probability

The **conditional probability** of B given that A has occurred is defined to be

$$P(B|A) = \frac{P(A \cap B)}{P(A)}.$$

Example: Roll a die.

A = odd number, B = prime number

$$P(B|A) = \frac{P(A \cap B)}{P(A)} = \frac{2/6}{3/6} = \frac{2}{3}$$

- From the definition $P(B|A) = \frac{P(A \cap B)}{P(A)}$, we obtain

$$P(A \cap B) = P(A)P(B|A).$$

- Two events A and B are said to be **independent** if

$$P(A \cap B) = P(A)P(B).$$

In this case, $P(B|A) = \frac{P(A \cap B)}{P(A)} = \frac{P(A)P(B)}{P(A)} = P(B)$, and similarly, $P(A|B) = P(A)$.

Exmaples:

- ① Draw a card from a deck of 52.

A = the card is an ace, B = the card is a spade

$$P(A) = \frac{1}{13}, \quad P(B) = \frac{1}{4}, \quad P(A \cap B) = \frac{1}{52} = P(A)P(B)$$

- ② Draw two cards from a deck of 52.

A = the first card is a spade, B = the second card is a spade

$$P(A) = \frac{1}{4}, \quad P(B) = \frac{1}{4}, \quad P(A \cap B) = \frac{13}{52} \times \frac{12}{51} \neq P(A)P(B)$$

Bayes' Theorem

$$P(H|E) = \frac{P(E \cap H)}{P(E)} = \frac{P(E|H)P(H)}{P(E)}$$

$$P(H|E) = \frac{P(E|H)}{P(E)} P(H)$$

- $P(H)$: **prior** degree of belief in hypothesis H

$P(H|E)$: **posterior** degree of belief after evidence E

$\frac{P(E|H)}{P(E)}$: support that E provides for H

Previously,

Learning \iff Maximizing Certainty

\iff Minimizing Randomness (Entropy)

The Bayesian viewpoint is

Learning \iff Updating Belief through Evidences

Since $P(E) = P(E|H)P(H) + P(E|H^c)P(H^c)$, we have

$$\begin{aligned} P(H|E) &= \frac{P(E|H)P(H)}{P(E)} \\ &= \frac{P(E|H)P(H)}{P(E|H)P(H) + P(E|H^c)P(H^c)}. \end{aligned}$$

Example: A new test for cancer is developed. When a person has cancer, this test is positive with probability 0.9. When a person does not have cancer, it is positive with probability 0.1. It is known that a person has cancer with probability 0.01. Bob receives the test and the result is positive. What is the probability that Bob has cancer?

$$P(C|+) = \frac{P(C \cap +)}{P(+)} = \frac{0.01 \times 0.9}{0.01 \times 0.9 + 0.99 \times 0.1} \approx 0.083$$

$$P(NC|+) = 1 - P(C|+) \approx 0.917$$

$$P(C|-) = \frac{P(C \cap -)}{P(-)} = \frac{0.01 \times 0.1}{0.01 \times 0.1 + 0.99 \times 0.9} \approx 0.001$$

$$P(NC|-) = 1 - P(C|-) \approx 0.999$$

- Bayesian interpretation

$$P(C|+) = \frac{P(+|C)}{P(+)} P(C)$$

- prior $P(C) = 0.01$, posterior $P(C|+) \approx 0.083$

$$\text{update } \frac{P(+|C)}{P(+)} \approx 8.3$$

Random Variables

Let S be a sample space with probability P .

- A function $X : S \rightarrow \mathbb{R}$ is said to be a **random variable**.

Examples:

- 1 Roll a die, $X =$ the number that appears
- 2 Roll two dice, $X =$ the sum of two numbers that appear
- 3 Drive from Storrs to Boston, $X =$ the driving time in minutes

When the values of X is discrete, the **probability mass function (pmf)** of X is defined by

$$p(x) := P(X = x) = P(\{\omega \in \mathcal{S} : X(\omega) = x\}).$$

Example: Consider an experiment whose outcome is a success or a failure.

- Let $Y = \begin{cases} 1 & \text{if a success,} \\ 0 & \text{if a failure.} \end{cases}$
- This random variable Y is called a **Bernoulli** random variable. Such an experiment is called a Bernoulli trial.
- If we set $p(1) = P(Y = 1) = p$, then $p(0) = P(Y = 0) = 1 - p$, and it defines a pmf.

Example: Repeat a Bernoulli trial n times, and assume each trial is independent.

- Let X = number of successes out of n trials.
- This random variable X is called a **binomial** random variable with parameters (n, p) .
- The pmf is given by

$$p(k) = P(X = k) = \binom{n}{k} p^k (1 - p)^{n-k}, \quad k = 0, 1, \dots, n,$$

where $\binom{n}{k} := \frac{n!}{k!(n-k)!}$ is “ n choose k ”.

When the values of X is continuous, the **probability density function (pdf)** of X is defined to be a function $f(x) \geq 0$ such that

$$P(a \leq X \leq b) = \int_a^b f(x) dx \quad (a \leq b).$$

Then we have

$$f(x) = \frac{d}{dx} P(-\infty \leq X \leq x).$$

Examples:

- ❶ The random variable X with density function

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

is called the **standard normal** random variable.

- ❷ For $\alpha, \lambda > 0$, the random variable X with density function

$$f(x) = \frac{\lambda}{\Gamma(\alpha)} e^{-\lambda x} (\lambda x)^{\alpha-1} \quad (x \geq 0)$$

is called a **gamma** random variable, where $\Gamma(\alpha) = \int_0^\infty e^{-x} x^{\alpha-1} dx$ is the gamma function. We have

$$\Gamma(n) = \int_0^\infty e^{-x} x^{n-1} dx = (n-1)!, \quad n = 1, 2, 3, \dots$$

- X : random variable
- The **expected value** or **mean** of X is defined to be

$$E(X) = \sum x p(x) \text{ or } \int x f(x) dx.$$

- The **variance** of X is defined by

$$\text{Var}(X) = E[(X - E(X))^2] = E(X^2) - E(X)^2.$$

- The **standard deviation** of X is defined by

$$\sigma(X) = \sqrt{\text{Var}(X)}.$$

- $E(aX + b) = aE(X) + b$
- $\text{Var}(aX + b) = a^2 \text{Var}(X)$

Example: Roll a die, X = the number that appears

$$p(x) = P(X = x) = \frac{1}{6} \quad \text{for } x = 1, 2, 3, 4, 5, 6$$

$$\begin{aligned} E(X) &= \sum x p(x) = 1 \times \frac{1}{6} + 2 \times \frac{1}{6} + \cdots + 6 \times \frac{1}{6} \\ &= \frac{1 + 2 + \cdots + 6}{6} = \frac{21}{6} = \frac{7}{2} \end{aligned}$$

$$\begin{aligned} \text{Var}(X) &= \sum (x - E(X))^2 p(x) = \sum x^2 p(x) - E(X)^2 \\ &= \frac{91}{6} - \left(\frac{7}{2}\right)^2 = \frac{35}{12} \approx 2.9167 \end{aligned}$$

$$\sigma(X) = \sqrt{\text{Var}(X)} = \sqrt{\frac{35}{12}} \approx 1.708$$

Example: $X = \text{binom}(n, p)$

$$E(X) = np \quad \text{and} \quad \text{Var}(X) = np(1 - p)$$

Example

- X : standard normal, $f(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$
- Since $xf(x)$ is odd, we have

$$E(X) = \int_{-\infty}^{\infty} xf(x)dx = 0.$$

$$\begin{aligned} E(X^2) &= \int_{-\infty}^{\infty} x^2 \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx = \int_0^{\infty} \frac{2}{\sqrt{2\pi}} x^2 e^{-\frac{x^2}{2}} dx \\ &= \frac{2}{\sqrt{2\pi}} (-xe^{-\frac{x^2}{2}}) \Big|_0^{\infty} + \int_0^{\infty} \frac{2}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx = 1 \end{aligned}$$

$$\text{Var}(X) = E(X^2) - E(X)^2 = 1$$

- Let Z be the standard normal random variable.

$$E(Z) = 0 \quad \text{and} \quad \text{Var}(Z) = 1$$

- Define $X = \sigma Z + \mu$. Then X is a **normal** random variable.
- $E(X) = E(\sigma Z + \mu) = \sigma E(Z) + \mu = \mu$
- $\text{Var}(X) = \text{Var}(\sigma Z + \mu) = \sigma^2 \text{Var}(Z) = \sigma^2$

$$E(X) = \mu \quad \text{and} \quad \text{Var}(X) = \sigma^2$$

- The density function of f is given by

$$\mathcal{N}(x|\mu, \sigma) := f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}.$$

- Conversely, we can standardize a normal random variable by

$$Z = \frac{X - \mu}{\sigma}.$$

Multivariate Distributions

Consider two random variables X and Y at the same time.

- When X and Y are discrete, the **joint mass function** is defined by

$$p(x, y) = P(X = x, Y = y).$$

- Set $p_X(x) = \sum_y p(x, y)$ and $p_Y(y) = \sum_x p(x, y)$. They are pmfs of X and Y , respectively, and are called the **marginal** mass functions.

- When X and Y are continuous, the **joint density function** of X and Y is defined to be a function $f(x, y) \geq 0$ such that, for any $A \subseteq \mathbb{R}^2$,

$$P((X, Y) \in A) = \iint_A f(x, y) dx dy.$$

- Define $f_X(x) = \int f(x, y) dy$ and $f_Y(y) = \int f(x, y) dx$. They are pdfs of X and Y , respectively, and are called the **marginal density functions**.
- In particular, we have

$$\iint_{\mathbb{R}^2} f(x, y) dx dy = 1.$$

Example

- X, Y : two standard normal variables
- Define

$$f(x, y) = \frac{1}{2\pi} e^{-\frac{x^2+y^2}{2}}.$$

Then $f(x, y)$ is a joint density function.

- $P(X^2 + Y^2 \leq 1) = ??$

$$\begin{aligned} P(X^2 + Y^2 \leq 1) &= \iint_{x^2+y^2 \leq 1} f(x, y) dx dy = \int_0^{2\pi} \int_0^1 \frac{1}{2\pi} e^{-\frac{r^2}{2}} r dr d\theta \\ &= \int_0^1 e^{-\frac{r^2}{2}} r dr = -e^{-u} \Big|_0^1 = 1 - e^{-1} \approx 0.6321 \end{aligned}$$

- Two random variables X and Y are said to be **independent** if

$$P(X \leq a, Y \leq b) = P(X \leq a)P(Y \leq b) \quad \text{for all } a \text{ and } b.$$

- discrete case: X and Y are independent

$$\iff P(X = a, Y = b) = P(X = a)P(Y = b) \quad \text{for all } a \text{ and } b.$$

- continuous case: X and Y are independent

$$\iff f(x, y) = f_X(x)f_Y(y).$$

- X, Y : random variables

Then we have

$$E(X + Y) = E(X) + E(Y).$$

In general, for random variables X_1, X_1, \dots, X_n , we have

$$E\left(\sum X_i\right) = \sum E(X_i).$$

- X, Y : independent random variables

Then we have

$$E(XY) = E(X)E(Y).$$

- The covariance of X and Y is defined by

$$\text{Cov}(X, Y) = E[(X - E(X))(Y - E(Y))].$$

We have

$$\text{Cov}(X, Y) = E(XY) - E(X)E(Y).$$

- If X and Y are independent, we have

$$\text{Cov}(X, Y) = 0.$$

- We have

$$\text{Var}(X + Y) = V(X) + V(Y) + 2 \text{Cov}(X, Y).$$

In general,

$$\text{Var}\left(\sum X_i\right) = \sum \text{Var}(X_i) + 2 \sum_{i < j} \text{Cov}(X_i, X_j).$$

- If X_1, X_2, \dots, X_n are independent, then

$$\text{Var}\left(\sum X_i\right) = \sum \text{Var}(X_i).$$

Maximum Likelihood Estimate (MLE)

A student practices three-point shots in basketball. He tried 100 shots and made 67 shots successful. What is the probability of success for the next shot?

Choose your model and compute the likelihood.

$$L(p) = \binom{100}{67} p^{67} (1 - p)^{33}$$

Determine the parameter p so that the likelihood function $L(p)$ is maximized.

It is more convenient to consider $\log L(p)$.

$$\log L(p) = 67 \log p + 33 \log(1 - p) + C$$

$$(\log L(p))' = \frac{67}{p} - \frac{33}{1-p} = \frac{67(1-p) - 33p}{p(1-p)} = \frac{67 - 100p}{p(1-p)}$$

When $p = \frac{67}{100}$, the likelihood $L(p)$ is maximized.

Prediction: the next shot will be successful with probability $\frac{67}{100}$.

This method is called the **maximum likelihood estimate (MLE)**.

Linear Regression as MLE

- Input: $\mathbf{x} = (x_1, x_2, \dots, x_K) \in \mathbb{R}^K$ Output: $t \in \mathbb{R}$
Observations: $(x_{11}, x_{12}, \dots, x_{1K}; t_1), \dots, (x_{N1}, x_{N2}, \dots, x_{NK}; t_N)$
- Use these observations as **training examples**.
Task: Given a new input $(\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_K)$, predict the output \tilde{t} .

- Choose a model:

Given \mathbf{x} , the corresponding value of t has a normal distribution with mean

$$y(\mathbf{x}, \mathbf{w}) = w_0 + w_1 x_1 + w_2 x_2 + \cdots + w_K x_K,$$

where $\mathbf{w} = [w_0, w_1, \dots, w_K]^\top$.

- Then we have

$$t = y(\mathbf{x}, \mathbf{w}) + \epsilon,$$

where ϵ is the Gaussian noise with mean 0.

- The density function is given by

$$p(t|\mathbf{x}, \mathbf{w}, \beta) = \mathcal{N}(t|y(\mathbf{x}, \mathbf{w}), \beta^{-1}),$$

where β is a parameter corresponding to the inverse variance, called the **precision**.

- Assume that each observation is independent, and that the variance β^{-1} is all the same.
- Then the likelihood density is

$$\begin{aligned} L(\mathbf{w}) &= p(\mathbf{t}|\mathbf{X}, \mathbf{w}, \beta) \\ &= \prod_{n=1}^N \mathcal{N}(t_n|y_n(\mathbf{x}_n, \mathbf{w}), \beta^{-1}) \\ &= \left(\sqrt{\frac{\beta}{2\pi}} \right)^N \exp \left[-\frac{\beta}{2} \sum (t_n - y_n)^2 \right]. \end{aligned}$$

- This is our probabilistic model.

Taking log, we obtain

$$\begin{aligned}\log L(\mathbf{w}) &= -\frac{\beta}{2} \sum (t_n - y_n)^2 + (\text{constant}) \\ &= -\frac{\beta}{2} \|\mathbf{t} - X\mathbf{w}\|^2 + (\text{constant}).\end{aligned}$$

Thus MLE is equivalent to minimizing the error function

$$E(\mathbf{w}) = \sum (t_n - y_n)^2 = \|\mathbf{t} - X\mathbf{w}\|^2.$$

- Maximum likelihood

\iff Minimizing $\|\mathbf{y} - X\mathbf{w}\|^2$

\iff Linear regression

Bayesian Linear Regression

- **Bayesian linear regression** avoids the over-fitting problem of maximum likelihood.
- ① Consider \mathbf{w} as a random variable.
- ② Choose **prior** information on \mathbf{w} in the form of probability distribution.
- ③ Given observations $\{(\mathbf{x}_n, t_n)\}$, update \mathbf{w} using **Bayes' formula** to obtain the **posterior** distribution of \mathbf{w} .
- ④ Determine \mathbf{w} so that the posterior probability is maximized; that is, take the mode of the posterior.

- M -dimensional Gaussian distribution:

$$\mathcal{N}(\mathbf{w}|\boldsymbol{\mu}, \Sigma) := \frac{1}{(2\pi)^{M/2}} \frac{1}{|\Sigma|^{1/2}} \exp \left\{ -\frac{1}{2}(\mathbf{w} - \boldsymbol{\mu})^\top \Sigma^{-1}(\mathbf{w} - \boldsymbol{\mu}) \right\},$$

where the M -dimensional vector $\boldsymbol{\mu}$ is the mean, the $M \times M$ matrix Σ is the covariance, and $|\Sigma|$ is the determinant of Σ .

- Choose a **prior** distribution for \mathbf{w} :

$$p(\mathbf{w}|\alpha) = \mathcal{N}(\mathbf{w}|0, \alpha^{-1}I).$$

- We have

$$-\ln p(\mathbf{w}|\alpha) = \frac{\alpha}{2} \mathbf{w}^\top \mathbf{w} + (\text{constant}) = \frac{\alpha}{2} \|\mathbf{w}\|^2 + (\text{constant}).$$

- Bayes' Theorem gives the **posterior**

$$p(\mathbf{w}|X, \mathbf{t}, \alpha, \beta) \propto p(\mathbf{t}|X, \mathbf{w}, \beta) p(\mathbf{w}|\alpha).$$

Recall $P(A|B) = \frac{P(B|A)P(A)}{P(B)}$.

Task: Determine \mathbf{w} so that the posterior is maximized.

This process is called a **maximum a posteriori (MAP)** estimation.

- Take the negative logarithm of the posterior

$$\begin{aligned} -\log p(\mathbf{w}|X, \mathbf{t}, \alpha, \beta) &= -\log [p(\mathbf{t}|X, \mathbf{w}, \beta) p(\mathbf{w}|\alpha)] + (\text{constant}) \\ &= \frac{\beta}{2} \|\mathbf{t} - X\mathbf{w}\|^2 + \frac{\alpha}{2} \|\mathbf{w}\|^2 + (\text{constant}) \end{aligned}$$

- The maximum of the posterior is given by the minimum of

$$\tilde{E}(\mathbf{w}) = \frac{\beta}{2} \|\mathbf{t} - X\mathbf{w}\|^2 + \frac{\alpha}{2} \|\mathbf{w}\|^2.$$

Thus, maximizing the posterior distribution is equivalent to minimizing the **regularized** sum-of-square error function, i.e., **ridge regression**.

- We can compute \mathbf{w} explicitly:

$$\tilde{E}(\mathbf{w}) = \frac{\beta}{2} \|X\mathbf{w} - \mathbf{t}\|^2 + \frac{\alpha}{2} \|\mathbf{w}\|^2,$$


and

$$\nabla \tilde{E}(\mathbf{w}) = \beta X^\top (X\mathbf{w} - \mathbf{t}) + \alpha \mathbf{w} = \mathbf{0}.$$


Thus

$$\mathbf{w} = \beta S X^\top \mathbf{t} \quad \text{with} \quad S^{-1} = \alpha I + \beta X^\top X.$$

$N = 9$, $\alpha = 0.01$ and $\beta = 1000$



Recall the maximum likelihood gave us



- The posterior can be computed explicitly, since the prior and the likelihood are all Gaussian.
- Indeed, we obtain

$$p(\mathbf{w}|X, \mathbf{t}, \alpha, \beta) = \mathcal{N}(\mathbf{w}|\mathbf{m}, S),$$

where

$$\mathbf{m} = \beta S X^T \mathbf{t} \quad \text{with} \quad S^{-1} = \alpha I + \beta X^T X.$$

Naive Bayes Classifier

Consider the binary classification

$$\{\text{sentences}\} \longrightarrow \{\text{positive, negative}\}.$$

For example, take a dataset of sentences from Amazon reviews on products.

Task: Given a new sentence, determine whether it is positive or negative.

Question: How can we convert a sentence into a vector?

- Fix a vocabulary of K words.
- Make an ordered list w_1, w_2, \dots, w_K of the words in the vocabulary.
- Replace a sentence with a binary vector:

$$\text{sentence} \longrightarrow (x_1, x_2, \dots, x_K),$$

where $x_i = 1$ if w_i occurs and $x_i = 0$ otherwise.

- For simplicity, we don't record how many times a word occurs. Similarly, we ignore the order in which different words appear.

Analyze a dataset of 1000 sentences.

Denote by G an occurrence of “great” in a sentence and by G^c no occurrence of “great”.

	+	-	total
G	92	5	97
G^c	408	495	903
total	500	500	1000

For the classification, we are interested in $P(+|G)$.

How can we compute $P(+|G)$??

- Conditional probability

$$P(+|G) = \frac{P(+ \cap G)}{P(G)} = \frac{92/1000}{97/1000} = \frac{92}{97} \approx 0.9485$$

- Bayes' Theorem

$$P(+|G) = \frac{P(G|+)}{P(G)} P(+) = \frac{92/500}{97/1000} 0.5 \approx 0.9485$$

We have

$$P(-|G) = 1 - 0.9485 = 0.0515,$$

$$P(+|G^c) = \frac{P(+ \cap G^c)}{P(G^c)} = \frac{408}{903} = 0.4518,$$

$$P(-|G^c) = 1 - 0.4518 = 0.5482.$$

Denote by W an occurrence of “waste” in a sentence and by W^c no occurrence of “waste”.

	+	−	total
W	0	14	14
W^c	500	486	986
total	500	500	1000

We have

$$P(+|W) = \frac{P(+ \cap W)}{P(W)} = 0, \quad P(-|W) = 1,$$

$$P(+|W^c) = \frac{P(+ \cap W^c)}{P(W^c)} = \frac{500}{986} = 0.5071,$$

$$P(-|W^c) = 1 - 0.5071 = 0.4929.$$

We need to consider many different words all together.

Assumption: Given a collection of words, their presence and absence in a (postive or negative) review are independent.

Then we have

$$P(G, W|\pm) = P(G|\pm)P(W|\pm), \quad P(G, W^c|\pm) = P(G|\pm)P(W^c|\pm),$$

$$P(G^c, W|\pm) = P(G^c|\pm)P(W|\pm), \quad P(G^c, W^c|\pm) = P(G^c|\pm)P(W^c|\pm).$$

The above ssumption is naive, but it enables us to compute relavant probabilities.

- Consider a sentence with G and W^c .
- Is this sentence positive or negative?
- To be precise, we compare $P(+|G, W^c)$ with $P(-|G, W^c)$.

Using Bayes' Theorem, we have

$$P(+|G, W^c) = \frac{P(G, W^c|+)P(+)}{P(G, W^c)} = \frac{P(G|+)P(W^c|+)P(+)}{P(G, W^c)},$$

$$P(-|G, W^c) = \frac{P(G, W^c|-)P(-)}{P(G, W^c)} = \frac{P(G|-)P(W^c|-)P(-)}{P(G, W^c)}.$$

Since the denominators are the same, it is enough to consider the numerators for comparison.

$$L(+|G, W^c) := P(G|+)P(W^c|+)P(+)=\frac{92}{500}\frac{500}{500}\frac{1}{2}=0.092$$

$$L(-|G, W^c) := P(G|-)P(W^c|-)P(-)=\frac{5}{500}\frac{486}{500}\frac{1}{2}=0.00486$$

- Assume that we have extracted key words w_1, w_2, \dots, w_K from the dataset.
- Replace a sentence with a binary vector:

$$\text{sentence} \longrightarrow (x_1, x_2, \dots, x_K),$$

where $x_i = 1$ if w_i occurs and $x_i = 0$ otherwise.

- By Assumption, x_i 's correspond to independent Bernoulli random variables.
- If there are N sentences in the dataset, we obtain a data matrix X of size $N \times K$.

- Let $\mathbf{t} = [t_1, t_2, \dots, t_N]^\top$ be the column vector, where $t_n = 1$ if the n^{th} review is positive and $t_n = 0$ otherwise.
- The number of positive reviews is equal to $N_+ := \mathbf{t}^\top \mathbf{t}$, and the number of negative reviews is equal to $N_- := N - N_+$.
- Notice that

$$\mathbf{t}^\top X = [\#(+ \cap w_1), \#(+ \cap w_2), \dots, \#(+ \cap w_K)],$$

where $\#(+ \cap w_i)$ is the number of positive reviews that contain w_i .

- We have

$$P_+ := \frac{1}{N_+} \mathbf{t}^\top X = [P(w_1|+), P(w_2|+), \dots, P(w_K|+)].$$

- Let $\mathbb{1}$ be the matrix with all entries equal to 1, whose size is to be determined by the context.
- We have

$$(\mathbb{1} - \mathbf{t})^\top X = [\#(- \cap \mathbf{w}_1), \#(- \cap \mathbf{w}_2), \dots, \#(- \cap \mathbf{w}_K)].$$

- It follows that

$$P_- := \frac{1}{N_-} (\mathbb{1} - \mathbf{t})^\top X = [P(\mathbf{w}_1|-), P(\mathbf{w}_2|-), \dots, P(\mathbf{w}_K|-)].$$

- Let $\mathbf{x} = (x_1, x_2, \dots, x_K)$ be a data vector.

For example, $\mathbf{x} = (1, 1, 0, 0, 1)$. Then

$$P(\mathbf{x}|+) = P(w_1|+)P(w_2|+)(1 - P(w_3|+))(1 - P(w_4|+))P(w_5|+).$$

- We see that

$$P(\mathbf{x}|\pm) = \prod_i P(w_i|\pm)^{x_i} (1 - P(w_i|\pm))^{(1-x_i)}.$$

- Taking log of both sides, we get

$$\log P(\mathbf{x}|\pm) = \sum_i x_i \log P(w_i|\pm) + (1 - x_i) \log(1 - P(w_i|\pm)).$$

- Bayes' Theorem

$$P(\pm|\mathbf{x}) = \frac{P(\mathbf{x}|\pm)P(\pm)}{P(\mathbf{x})}$$

- By taking log, we have

$$\log P(\pm|\mathbf{x}) = \log P(\mathbf{x}|\pm) + \log P(\pm) - \log P(\mathbf{x}),$$

where

$$P(\pm) = \frac{N_{\pm}}{N}.$$

- A review is positive if $\log P(+|\mathbf{x}) > \log P(-|\mathbf{x})$ and negative otherwise.
- It follows from

$$\log P(\pm|\mathbf{x}) = \log P(\mathbf{x}|\pm) + \log P(\pm) - \log P(\mathbf{x})$$

that a review is **positive** if

$$\log P(\mathbf{x}|+) + \log P(+)> \log P(\mathbf{x}|-) + \log P(-)$$

and negative otherwise.

- In a previous example, we had $P(+|W) = 0$.
Then $\log P(+|W)$ is not defined!
- Introduce a “fake sentence” into both classes in which every vocabulary word appears.
This guarantees that no relevant probability is zero.

- Recall

$$\log P(\mathbf{x}|\pm) = \sum_i x_i \log P(w_i|\pm) + (1 - x_i) \log(1 - P(w_i|\pm)).$$

- Let X be the data matrix with each row \mathbf{x} corresponding to a sentence.
- Define $\log P(X|\pm)$ to be the column vector consisting of $\log P(\mathbf{x}|\pm)$.
- We have

$$\log P(X|\pm) = X(\log P_{\pm})^{\top} + (\mathbb{1} - X) \log(\mathbb{1} - P_{\pm})^{\top}.$$

Multi-class Logistic Regression

- Assume that there are s different classes.

Data set: $\{(x_{n,1}, \dots, x_{n,K}; t_n)\}$, $t_n = 1, 2, \dots, s$, $n = 1, \dots, N$

- How to generalize logistic regression?
- $t_n = 1 \Leftrightarrow [1, 0, \dots, 0]$, $t_n = 2 \Leftrightarrow [0, 1, 0, \dots, 0]$, \dots ,
 $t_n = s \Leftrightarrow [0, \dots, 0, 1]$
- Obtain an $N \times s$ matrix $\mathbf{t} = [t_{n,m}]$ such that

$$t_{n,m} = \begin{cases} 1 & \text{if } t_n = m, \\ 0 & \text{if } t_n \neq m. \end{cases}$$

- Define $\sigma : \mathbb{R}^s \rightarrow (0, 1)^s$ by

$$\sigma(\mathbf{a}) = \left(\frac{e^{a_1}}{\sum_{i=1}^s e^{a_i}}, \dots, \frac{e^{a_s}}{\sum_{i=1}^s e^{a_i}} \right),$$

where $\mathbf{a} = (a_1, a_2, \dots, a_s)$.

The function σ is called the **softmax** function.

- When $s = 2$, we have

$$\sigma(\mathbf{a}) = \left(\frac{e^{a_1}}{e^{a_1} + e^{a_2}}, \frac{e^{a_2}}{e^{a_1} + e^{a_2}} \right) = \left(\frac{1}{1 + e^{a_2 - a_1}}, \frac{e^{a_2 - a_1}}{1 + e^{a_2 - a_1}} \right).$$

- The softmax function is a generalization of the sigmoid function.

- Define $\mathbf{y} = [y_1, \dots, y_s] = \boldsymbol{\sigma}(\mathbf{a})$.
- For $m, j = 1, \dots, s$,

$$\frac{\partial y_m}{\partial a_j} = y_m(\delta_{j,m} - y_j),$$

where $\delta_{j,m}$ is the Kronecker's delta, i.e.

$$\delta_{j,m} = \begin{cases} 1 & \text{if } j = m, \\ 0 & \text{otherwise.} \end{cases}$$

Proof: Recall $y_m = \frac{e^{a_m}}{\sum_i e^{a_i}}$. Then we have

$$\begin{aligned}\frac{\partial y_m}{\partial a_j} &= \frac{\delta_{j,m} e^{a_m} (\sum e^{a_i}) - e^{a_m} e^{a_j}}{(\sum e^{a_i})^2} \\ &= \frac{e^{a_m} (\delta_{j,m} \sum e^{a_i} - e^{a_j})}{(\sum e^{a_i})^2} \\ &= \frac{e^{a_m}}{\sum e^{a_i}} \frac{\delta_{j,m} \sum e^{a_i} - e^{a_j}}{\sum e^{a_i}} = y_m (\delta_{j,m} - y_j).\end{aligned}$$

□

- Consider a $(K + 1) \times s$ matrix $\mathbf{w} = [w_{p,q}]$.

Define $\mathbf{y} = \sigma(X\mathbf{w}) = [y_{n,m}]$,

where X is as before and σ is applied to the rows of $X\mathbf{w}$.

Each row of \mathbf{y} consists of probabilities for classes 1 through s .

- The likelihood function is given by

$$L(\mathbf{w}) = \prod_{n=1}^N \prod_{m=1}^s y_{n,m}^{t_{n,m}}.$$

- The cross-entropy is

$$E(\mathbf{w}) = - \sum_{n=1}^N \sum_{m=1}^s t_{n,m} \ln y_{n,m}.$$

- We have

$$\nabla E(\mathbf{w}) = \left[\frac{\partial E}{\partial \mathbf{w}_{p,q}} \right] = \left[\sum_{n=1}^N (y_{n,q} - t_{n,q}) x_{n,p} \right] = X^T (\mathbf{y} - \mathbf{t}).$$

Proof: Notice

$$\sum_{m=1}^s t_{n,m} (\delta_{q,m} - y_{n,q}) = \sum_{m=1}^s (t_{n,m} \delta_{q,m} - t_{n,m} y_{n,q}) = t_{n,q} - y_{n,q}.$$

Write $X\mathbf{w} = [\mathbf{a}_{n,j}]$ and $\mathbf{a}_n = (a_{n,1}, \dots, a_{n,s})$. Then $a_{n,j} = \sum_k x_{n,k} w_{k,j}$ and $\mathbf{y}_n = \sigma(\mathbf{a}_n) = (y_{n,1}, \dots, y_{n,s})$. We have

$$\frac{\partial y_{n,m}}{\partial a_{n,j}} = y_{n,m} (\delta_{j,m} - y_{n,j}).$$

$$E(\mathbf{w}) = - \sum_{n=1}^N \sum_{m=1}^s t_{n,m} \ln y_{n,m}$$

$$\begin{aligned} \frac{\partial E}{\partial w_{p,q}} &= - \sum_n \sum_m t_{n,m} \frac{1}{y_{n,m}} \frac{\partial y_{n,m}}{\partial w_{p,q}} \\ &= - \sum_n \sum_m t_{n,m} \frac{1}{y_{n,m}} \sum_j \frac{\partial y_{n,m}}{\partial a_{n,j}} \frac{\partial a_{n,j}}{\partial w_{p,q}} \\ &= - \sum_n \sum_m t_{n,m} \frac{1}{y_{n,m}} \sum_j y_{n,m} (\delta_{j,m} - y_{n,j}) \delta_{j,q} x_{n,p} \\ &= - \sum_n \sum_m t_{n,m} (\delta_{q,m} - y_{n,q}) x_{n,p} = \sum_n (y_{n,q} - t_{n,q}) x_{n,p} \end{aligned}$$

□

- Gradient Descent

$$\mathbf{w}_{i+1} = \mathbf{w}_i - \eta \mathbf{X}^\top (\mathbf{y} - \mathbf{t}).$$

- Let $\mathbf{w}_i \rightarrow \mathbf{w}_*$ as $i \rightarrow \infty$.

Given $\mathbf{x} = [1, x_1, \dots, x_K]$, the coordinates of the vector

$$\mathbf{y} = \sigma(\mathbf{x}\mathbf{w}_*)$$

represent the probabilities for the classes.

- The (multi-class) logistic regression is the simplest **neural network**.

Dimensionality Reduction


- Typically, our data points live in a high dimensional space.
- In a high dimensional space,
 - the data set becomes very sparse;
 - our geometric intuitions can fail.
- Severe difficulties that can arise in spaces of high dimensions are usually called the **curse of dimensionality** (Bellman, 1961).

- Frequently, dimensionality reduction is very helpful and Principal Component Analysis (PCA) is a common method.
- PCA is an unsupervised learning method.
- Other methods for dimensionality reduction:
 - Linear Discriminative Analysis (LDA),
 - t-distributed Stochastic Neighbor Embedding (t-SNE),
 - Uniform Manifold Approximation and Projection (UMAP)

Principal Component Analysis (PCA)

- $\{\mathbf{x}_n \in \mathbb{R}^K\}$: dataset, $n = 1, \dots, N$

For example, when $K = 2$,



Which direction represents the dataset best?

- Goal: Project the dataset onto a subspace of dimension $D < K$ such that the variance is maximized.
Typically, we take $D = 2$ or $D = 3$.
- It is called the **Principal Component Analysis** (PCA).

- Consider \mathbf{x}_n as a row vector.
- Write $\bar{\mathbf{x}} = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n$.
- Let \mathbf{u} be a column vector with $\|\mathbf{u}\| = 1$.

The **mean** of the \mathbf{u} -coordinates is

$$\frac{1}{N} \sum_{n=1}^N \mathbf{x}_n \mathbf{u} = \bar{\mathbf{x}} \mathbf{u}.$$

- The **variance** of the \mathbf{u} -coordinates is

$$\begin{aligned}\frac{1}{N} \sum_{n=1}^N \{\mathbf{x}_n \mathbf{u} - \bar{\mathbf{x}} \mathbf{u}\}^2 &= \frac{1}{N} \sum_{n=1}^N \mathbf{u}^\top (\mathbf{x}_n - \bar{\mathbf{x}})^\top (\mathbf{x}_n - \bar{\mathbf{x}}) \mathbf{u} \\ &= \boxed{\mathbf{u}^\top \mathbf{S} \mathbf{u}},\end{aligned}$$

where \mathbf{S} is a $K \times K$ symmetric matrix given by

$$\boxed{\mathbf{S} = \frac{1}{N} \sum_{n=1}^N (\mathbf{x}_n - \bar{\mathbf{x}})^\top (\mathbf{x}_n - \bar{\mathbf{x}})}.$$

- Write $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_K)$ and $\mathbf{x}_n = (x_{n1}, \dots, x_{nK})$.
- Let S_{ij} be the (i, j) -entry of S . Then we have

$$S_{ij} = \frac{1}{N} \sum_{n=1}^N (x_{ni} - \bar{x}_i)(x_{nj} - \bar{x}_j).$$

- This is the covariance of the i -th feature and the j -th feature. In particular, the diagonal entry S_{ii} is the variance of the i -th feature.
- The symmetric matrix S is called the **covariance matrix**.

- Write X_0 be the matrix whose rows are $\mathbf{x}_n - \bar{\mathbf{x}}$, $n = 1, 2, \dots, N$.

Then we have

$$S = \frac{1}{N} X_0^\top X_0.$$

- The covariance S_{ij} measures the dependency between the i -feature and the j -th feature.

However, S_{ij} may be big simply because x_{ni} and x_{nj} are big.

- Define the **correlation** of X and Y by

$$\text{Corr}(X, Y) = \frac{\text{Cov}(X, Y)}{\sigma(X)\sigma(Y)}.$$

- We have $-1 \leq \text{Corr}(X, Y) \leq 1$. In particular, $\text{Corr}(X, X) = 1$.
- In our case, the correlation matrix $C = [C_{ij}]$ is given by

$$C_{ij} = \frac{S_{ij}}{\sqrt{S_{ii}S_{jj}}}.$$

- Task: Maximize the variance $\mathbf{u}^\top S \mathbf{u}$ for \mathbf{u} such that

$$\|\mathbf{u}\|^2 = \mathbf{u}^\top \mathbf{u} = 1.$$

- It is an optimization problem with constraints:
maximize $f(\mathbf{x})$ for \mathbf{x} subject to $g(\mathbf{x}) = c$?
- We can use Lagrange Multiplier.

Theorem (Lagrange Multiplier)

Assume that $g(\mathbf{x}) = c$ and $\nabla g \neq 0$. If \mathbf{x}_* is a maximum or a minimum of $f(\mathbf{x})$, then

$$\nabla f(\mathbf{x}_*) = \lambda \nabla g(\mathbf{x}_*)$$

for some $\lambda \in \mathbb{R}$.

- Let $f(\mathbf{u}) = \mathbf{u}^\top S \mathbf{u}$ and $g(\mathbf{u}) = \mathbf{u}^\top \mathbf{u}$.

Then $\nabla f(\mathbf{u}) = 2S\mathbf{u}$ and $\nabla g(\mathbf{u}) = 2\mathbf{u}$.

From Lagrange multiplier, we obtain

$$S\mathbf{u} = \lambda \mathbf{u}.$$

- \mathbf{u} is an eigenvector with eigenvalue λ .

- In this case, the variance is

$$\mathbf{u}^\top \mathbf{S} \mathbf{u} = \lambda \mathbf{u}^\top \mathbf{u} = \lambda,$$

and it is maximized when λ is the **largest eigenvalue** of \mathbf{S} .

- The corresponding eigenvector is the **first** principal direction.
- Actually, we can always find K principal directions, or K linearly independent eigenvectors.
- Recall that \mathbf{S} is symmetric.

Theorem (Spectral Theorem)

Let A be a $K \times K$ *symmetric* matrix. Then the following is true.

- 1 The matrix A has K *real* eigenvalues, counting multiplicities.
- 2 The eigenvectors corresponding to different eigenvalues are orthogonal.
- 3 The matrix A is orthogonally diagonalizable.

- In particular, the covariance matrix S is diagonalizable with real eigenvalues.

- Let

$$\lambda_1 \geq \cdots \geq \lambda_K$$

be the eigenvalues of S , and $\mathbf{u}_1, \dots, \mathbf{u}_K$ are corresponding eigenvectors with $\|\mathbf{u}_i\| = 1$.

- The variance in the \mathbf{u}_i direction is

$$\mathbf{u}_i^\top S \mathbf{u}_i = \lambda_i \quad \text{for } i = 1, \dots, K.$$

- For $D \leq K$, the D -dimensional projection onto the subspace spanned by $\mathbf{u}_1, \dots, \mathbf{u}_D$ produces the desired dimensionality reduction. (Typically, $D = 2$ or 3 .)

How to compute the projection:

- In order to obtain a centered picture, use X_0 instead of X .
- The \mathbf{u}_i -coordinate of $\mathbf{x}_n - \bar{\mathbf{x}}$ is simply

$$(\mathbf{x}_n - \bar{\mathbf{x}})\mathbf{u}_i.$$

- Let U be the $K \times D$ matrix whose columns are $\mathbf{u}_1, \dots, \mathbf{u}_D$. Then the coordinates of the projection is given by the $N \times D$ matrix

$$X_0 U = X_0[\mathbf{u}_1 \cdots \mathbf{u}_D].$$

Question: Can we see the original features in the projection?

- Consider

$$\mathbf{e}_1 = (1, 0, \dots, 0), \mathbf{e}_2 = (0, 1, 0, \dots, 0), \dots, \mathbf{e}_K = (0, \dots, 0, 1).$$

Then \mathbf{e}_i is supported only by the i -th feature.

- The projection $\mathbf{e}_i U$ is nothing but the i -th row of U .
- The i -th row vector of U is the direction of the i -th feature, and it is called the **loading** of the i -th feature.

PCA Approximation

- We may use PCA for approximation.
- Let $\{\mathbf{u}_1, \dots, \mathbf{u}_K\}$ be an orthonormal basis consisting of eigenvectors with $\lambda_1 \geq \dots \geq \lambda_K$.
- Then a PCA approximation $\tilde{\mathbf{x}}_n$ to a data vector \mathbf{x}_n is given by


$$\tilde{\mathbf{x}}_n = \bar{\mathbf{x}} + \sum_{i=1}^D (\mathbf{x}_n - \bar{\mathbf{x}}) \mathbf{u}_i \mathbf{u}_i^{\top}.$$

Example: Hand-written digits


Original Data: $K = 28 \times 28 = 784$

Taking $N = 10$ and $D = 1$, we obtain the following:

Original



PCA approximation with $D = 1$



Singular Value Decomposition (SVD)

- PCA can be understood as SVD.
- SVD also generalizes the notion of eigenvectors from square matrices to matrices of arbitrary sizes.
- Recall that a matrix A is **orthogonal** if the columns form an orthonormal set, or equivalently,

$$A^T A = I.$$

Theorem

Any (real) $N \times K$ matrix X can be decomposed as follows:

$$X = U\Lambda V^T,$$

where U is an $N \times N$ orthogonal matrix, V is a $K \times K$ orthogonal matrix, and Λ is an $N \times K$ matrix containing $r := \min(N, K)$ values, called *singular values*, on the main diagonal with 0's elsewhere.

- Write $\Lambda = \text{diag}(\sigma_1, \dots, \sigma_r)$. The *pseudoinverse* X^\dagger is defined by

$$X^\dagger = V\Lambda^\dagger U^T,$$

where $\Lambda^\dagger = \text{diag}(1/\sigma_1, \dots, 1/\sigma_r)$.

Example

- Let $X = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$. Then the SVD gives

$$U \approx \begin{bmatrix} -0.3863 & -0.9224 \\ -0.9224 & 0.3863 \end{bmatrix}, \quad \Lambda \approx \begin{bmatrix} 9.508 & 0.0 & 0.0 \\ 0.0 & 0.7729 & 0.0 \end{bmatrix},$$

$$V \approx \begin{bmatrix} -0.4287 & 0.806 & 0.4082 \\ -0.5663 & 0.1124 & -0.8165 \\ -0.7039 & -0.5812 & 0.4082 \end{bmatrix}.$$

- The pseudoinverse is given by

$$X^\dagger = V\Lambda^\dagger U^\top = \begin{bmatrix} -17/18 & 4/9 \\ -1/9 & 1/9 \\ 13/18 & -2/9 \end{bmatrix}, \quad \text{where } \Lambda^\dagger \approx \begin{bmatrix} 1/9.508 & 0.0 \\ 0.0 & 1/0.7729 \\ 0.0 & 0.0 \end{bmatrix}.$$

- We have $XX^\dagger = I$, but $X^\dagger X \neq I$.

- Let $X = \begin{bmatrix} 3 & 1 \\ 2 & 7 \\ 1 & 5 \end{bmatrix}$. Then the SVD gives

$$U \approx \begin{bmatrix} -0.2068 & 0.9702 & 0.1261 \\ -0.803 & -0.0947 & -0.5885 \\ -0.559 & -0.2229 & 0.7986 \end{bmatrix}, \quad \Lambda \approx \begin{bmatrix} 9.0613 & 0.0 \\ 0.0 & 2.6255 \\ 0.0 & 0.0 \end{bmatrix},$$

$$V \approx \begin{bmatrix} -0.3074 & 0.9516 \\ -0.9516 & -0.3074 \end{bmatrix}.$$

Eigenvalues and Eigenvectors

- Suppose that $X = U\Lambda V^T$ is a SVD.

Then

$$X^T X = V\Lambda^T U^T U \Lambda V^T = V(\Lambda^T \Lambda) V^T,$$

where $\Lambda^T \Lambda$ is a diagonal matrix.

Hence

$$(X^T X)V = V(\Lambda^T \Lambda).$$

- eigenvectors of $X^T X$ = columns of V
eigenvalues of $X^T X$ = diagonal entries of $\Lambda^T \Lambda$
- Recall that we used an eigenvector of $X^T X$ in the PCA example.

- Similarly,

$$XX^T = U\Lambda V^T V\Lambda^T U^T = U(\Lambda\Lambda^T)U^T$$

$$(XX^T)U = U(\Lambda\Lambda^T)$$

- eigenvectors of XX^T = columns of U
eigenvalues of XX^T = diagonal entries of $\Lambda\Lambda^T$

Example

- Let $X = \begin{bmatrix} 3 & 1 \\ 2 & 7 \\ 1 & 5 \end{bmatrix}$. Then $X^\top X = \begin{bmatrix} 14 & 22 \\ 22 & 75 \end{bmatrix}$.

Recall that $\Lambda \approx \begin{bmatrix} 9.0613 & 0.0 \\ 0.0 & 2.6255 \\ 0.0 & 0.0 \end{bmatrix}$.

The eigenvalues are given by $\Lambda^\top \Lambda \approx \begin{bmatrix} 82.1065 & 0.0 \\ 0.0 & 6.8935 \end{bmatrix}$;

the eigenvectors are given by $V \approx \begin{bmatrix} -0.3074 & 0.9516 \\ -0.9516 & -0.3074 \end{bmatrix}$.

Linear Discriminant Analysis

Multi-class classification

$$\mathbf{x} = (x_1, \dots, x_k) \rightsquigarrow P(t|\mathbf{x}), \quad t = 1, 2, \dots, s$$

- We studied logistic regression.
- The **Linear Discriminant Analysis (LDA)** is based on Bayesian inference.
- LDA can be used for both supervised and unsupervised learning.

- Recall the density function f of the normal distribution is given by

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}.$$

- K -dimensional Gaussian distribution:

$$\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \Sigma) = \frac{1}{(2\pi)^{K/2}} \frac{1}{|\Sigma|^{1/2}} \exp \left\{ -\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})\Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu})^\top \right\}$$

where the K -dimensional vector $\boldsymbol{\mu}$ is the mean, the $K \times K$ matrix Σ is the covariance, and $|\Sigma|$ is the determinant of Σ .

- $\pi_t := P(t)$ prior probability that an observation belongs to class t
 $f_t(\mathbf{x}) := P(\mathbf{x}|t)$ likelihood

Bayes' Theorem:

$$P(t|\mathbf{x}) \propto P(t)P(\mathbf{x}|t) = \pi_t f_t(\mathbf{x})$$

- LDA assumes

(1) $f_t(\mathbf{x})$ is normal, i.e., $f_t(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_t, \Sigma_t)$

(2) all the covariances are the same,

$$\text{i.e., } \Sigma := \Sigma_1 = \cdots = \Sigma_s.$$

- Recall

$$\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_t, \Sigma) = \frac{1}{(2\pi)^{K/2}} \frac{1}{|\Sigma|^{1/2}} \exp \left\{ -\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_t)\Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu}_t)^\top \right\}.$$

- We have

$$\begin{aligned}\ln P(t|\mathbf{x}) &= \ln \pi_t + \ln f_t(\mathbf{x}) + (\text{terms without } t) \\ &= \ln \pi_t - \frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_t)\Sigma^{-1}(\mathbf{x} - \boldsymbol{\mu}_t)^\top + (\text{terms without } t) \\ &= \ln \pi_t - \frac{1}{2}\boldsymbol{\mu}_t\Sigma^{-1}\boldsymbol{\mu}_t^\top + \mathbf{x}\Sigma^{-1}\boldsymbol{\mu}_t^\top + (\text{terms without } t).\end{aligned}$$

(Note that $\mathbf{x}\Sigma^{-1}\boldsymbol{\mu}_t^\top$ is scalar and $\mathbf{x}\Sigma^{-1}\boldsymbol{\mu}_t^\top = (\mathbf{x}\Sigma^{-1}\boldsymbol{\mu}_t^\top)^\top = \boldsymbol{\mu}_t\Sigma^{-1}\mathbf{x}^\top$.)

- The part (terms without t) does not contribute when we compare $\ln P(t|\mathbf{x})$ for $t = 1, \dots, s$.

- Define the **discriminant function** by

$$\delta_t(\mathbf{x}) := \ln \pi_t - \frac{1}{2} \boldsymbol{\mu}_t^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_t + \mathbf{x}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_t$$

for $t = 1, \dots, s$.

- Given \mathbf{x} , if $\delta_{t_*}(\mathbf{x})$ is the largest,
observation \mathbf{x} belongs to class t_* with largest probability.

~~~~~> Classification

- In LDA, we use the training data to approximate  $\delta_t(\mathbf{x})$ .

Given  $\mathcal{D} = \mathcal{D}_1 \sqcup \mathcal{D}_2 \sqcup \cdots \sqcup \mathcal{D}_s$  (disjoint union),  
set

$$N_t := \#(\mathcal{D}_t), \quad t = 1, 2, \dots, s, \quad N := N_1 + \cdots + N_s.$$

We make the following estimates:

$$\hat{\pi}_t = N_t/N,$$

$$\hat{\boldsymbol{\mu}}_t = \frac{1}{N_t} \sum_{\mathbf{x} \in \mathcal{D}_t} \mathbf{x},$$

$$\hat{\boldsymbol{\Sigma}} = \frac{1}{N-s} \sum_{t=1}^s \sum_{\mathbf{x} \in \mathcal{D}_t} (\mathbf{x} - \hat{\boldsymbol{\mu}}_t)^\top (\mathbf{x} - \hat{\boldsymbol{\mu}}_t).$$

The use of  $N - s$  in  $\hat{\boldsymbol{\Sigma}}$  is called Bessel's correction.

Let  $\hat{X}_0$  be the  $N \times K$  matrix whose rows are  $\mathbf{x} - \hat{\boldsymbol{\mu}}_t$ .

Then we have

$$\hat{\Sigma} = \frac{1}{N-s} \hat{X}_0^\top \hat{X}_0.$$

- An approximation of the discriminant function is given by

$$\hat{\delta}_t(\mathbf{x}) := \ln \hat{\pi}_t - \frac{1}{2} \hat{\boldsymbol{\mu}}_t^{\top} \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\mu}}_t + \mathbf{x}^{\top} \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\mu}}_t.$$

- When  $\pi_1 = \pi_2 = \cdots = \pi_s$ , the **decision boundaries** are given by

$$-\frac{1}{2} \hat{\boldsymbol{\mu}}_i^{\top} \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\mu}}_i + \mathbf{x}^{\top} \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\mu}}_i = -\frac{1}{2} \hat{\boldsymbol{\mu}}_j^{\top} \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\mu}}_j + \mathbf{x}^{\top} \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\mu}}_j$$

for  $i \neq j$ .

# Dimensionality Reduction through LDA

- We studied **dimensionality reduction** through PCA.
- LDA can be used for dimensionality reduction.

It can be considered as a “supervised PCA”.

- Given  $\mathcal{D} = \mathcal{D}_1 \sqcup \mathcal{D}_2 \sqcup \cdots \sqcup \mathcal{D}_s$  (disjoint union),  
set

$$N_t := \#(\mathcal{D}_t), \quad t = 1, 2, \dots, s, \quad N := N_1 + \cdots + N_s.$$



- Main Idea: Shrink class  $\mathcal{D}_t$  into a single point  $\mu_t$  and do PCA.  
~~~~~> Principal directions for **classes**.

Set


$$\mu_t = \frac{1}{N_t} \sum_{\mathbf{x} \in \mathcal{D}_t} \mathbf{x} \quad \text{and} \quad \mu = \frac{1}{N} \sum_{\mathbf{x} \in \mathcal{D}} \mathbf{x}.$$

However, the point μ_t must have multiplicity N_t .

- Find a vector \mathbf{a} which maximizes

$$\mathbf{a}^\top \left(\sum_{t=1}^s N_t (\mu_t - \mu)^\top (\mu_t - \mu) \right) \mathbf{a}.$$

- Constraint: Overlap between classes needs to be controlled to achieve better class separation.



[Hastie, Tibshirani and Friedman, *The elements of statistical learning*]

Although the line joining the centroids defines the direction of greatest centroid spread, there is considerable overlap between the projected classes.

- A vector \mathbf{a} may capture more variability from one class than from the other classes.
- To control this, we fix the (weighted) sum of variances. That is,

$$\sum_{t=1}^s \mathbf{a}^\top \left(\sum_{\mathbf{x} \in \mathcal{D}_t} (\mathbf{x} - \boldsymbol{\mu}_t)^\top (\mathbf{x} - \boldsymbol{\mu}_t) \right) \mathbf{a} = 1.$$

- In short, our goal is to maximize the between-class variance relative to the within-class variance.

- Define

$$B = \sum_{t=1}^s N_t (\boldsymbol{\mu}_t - \boldsymbol{\mu})^\top (\boldsymbol{\mu}_t - \boldsymbol{\mu})$$
$$W = \sum_{t=1}^s \sum_{\mathbf{x} \in \mathcal{D}_t} (\mathbf{x} - \boldsymbol{\mu}_t)^\top (\mathbf{x} - \boldsymbol{\mu}_t).$$

Let X_0 be the $N \times K$ matrix whose rows are $\mathbf{x} - \boldsymbol{\mu}_t$.

Then we have

$$W = X_0^\top X_0.$$

- Task: Maximize $\mathbf{a}^\top \mathbf{B}\mathbf{a}$, subject to $\mathbf{a}^\top \mathbf{W}\mathbf{a} = 1$.

- Lagrange Multiplier

$$f(\mathbf{a}) = \mathbf{a}^\top \mathbf{B}\mathbf{a}, \quad g(\mathbf{a}) = \mathbf{a}^\top \mathbf{W}\mathbf{a}$$

- We obtain $\nabla f = 2\mathbf{B}\mathbf{a}$, $\nabla g = 2\mathbf{W}\mathbf{a}$, and

$$\mathbf{B}\mathbf{a} = \lambda \mathbf{W}\mathbf{a} \quad \text{for some } \lambda \quad \Longleftrightarrow \quad \boxed{W^{-1}B\mathbf{a} = \lambda \mathbf{a}}.$$

Critical points are eigenvectors of $W^{-1}B$.

- Note that $W^{-1}B$ is symmetric.
Thus the Spectral Theorem applies.

- If \mathbf{a} is an eigenvector of $W^{-1}B$ such that $\mathbf{a}^\top W\mathbf{a} = 1$, then

$$\mathbf{a}^\top B\mathbf{a} = \mathbf{a}^\top WW^{-1}B\mathbf{a} = \lambda \mathbf{a}^\top W\mathbf{a} = \lambda.$$

Recall that we are maximizing $\mathbf{a}^\top B\mathbf{a}$.

Thus an eigenvector of the **largest eigenvalue** is the first principal direction.

- Take D -many principal directions for $D < K$.

Project data points onto the subspace of the principal directions.

~~~~~> Dimensionality Reduction

# Support Vector Machines

- Assume that there are two sets  $\mathcal{A}_+$  and  $\mathcal{A}_-$  of datapoints.
- Suppose that  $\mathcal{A}_+$  and  $\mathcal{A}_-$  can be separated by a hyperplane.
- That is, there is a hyperplane  $H$  given by

$$f(\mathbf{x}) = f(x_1, \dots, x_K) = 0,$$

where  $f(\mathbf{x}) = w_0 + w_1 x_1 + \dots + w_K x_K = w_0 + \mathbf{x}\mathbf{w}$  and

$$\mathbf{w} = [w_1, \dots, w_K]^T,$$

such that the points  $\mathbf{x}$  in  $\mathcal{A}_+$  satisfy  $f(\mathbf{x}) > 0$  and  
those  $\mathbf{x}$  in  $\mathcal{A}_-$  satisfy  $f(\mathbf{x}) < 0$ .

- The vector  $\mathbf{w}$  is orthogonal (or normal) to the hyperplane  $H$ .

Proof. Assume  $f(\mathbf{a}) = f(\mathbf{b}) = 0$ . Then

$$(\mathbf{a} - \mathbf{b})\mathbf{w} = (w_0 + \mathbf{a}\mathbf{w}) - (w_0 + \mathbf{b}\mathbf{w}) = f(\mathbf{a}) - f(\mathbf{b}) = 0. \quad \square$$

- For  $\mathbf{x} \in \mathbb{R}^K$ , the distance from  $\mathbf{x}$  to  $H$  is

$$\frac{|f(\mathbf{x})|}{\|\mathbf{w}\|}.$$

Proof. Assume  $\mathbf{a} \in H$ . The distance between  $\mathbf{x}$  and  $H$  is the length of  $\mathbf{x} - \mathbf{a}$  in the  $\mathbf{w}$ -direction, i.e.

$$\left| (\mathbf{x} - \mathbf{a}) \frac{\mathbf{w}}{\|\mathbf{w}\|} \right| = \frac{|w_0 + \mathbf{x}\mathbf{w} - (w_0 + \mathbf{a}\mathbf{w})|}{\|\mathbf{w}\|} = \frac{|f(\mathbf{x})|}{\|\mathbf{w}\|}. \quad \square$$



- We have  $w_0 + \mathbf{p}w > 0 > w_0 + \mathbf{q}w$  for  $\mathbf{p} \in \mathcal{A}_+$  and  $\mathbf{q} \in \mathcal{A}_-$ .  
Thus  $(\mathbf{p} - \mathbf{q})w > 0$  for  $\mathbf{p} \in \mathcal{A}_+$  and  $\mathbf{q} \in \mathcal{A}_-$ .

- Find  $\mathbf{a} \in \mathcal{A}_+$  such that  $\mathbf{a}w \leq \mathbf{p}w$  for all  $\mathbf{p} \in \mathcal{A}_+$ .  
Define  $w_0^+ = -\mathbf{a}w$  and  $f_{w,+}(\mathbf{x}) = w_0^+ + \mathbf{x}w$ .

- Similarly, find  $\mathbf{b} \in \mathcal{A}_-$  such that  $\mathbf{b}w \geq \mathbf{q}w$  for all  $\mathbf{q} \in \mathcal{A}_-$ .  
Define  $w_0^- = -\mathbf{b}w$  and  $f_{w,-}(\mathbf{x}) = w_0^- + \mathbf{x}w$ .

- We have

$$(\mathbf{p} - \mathbf{q})w \geq (\mathbf{a} - \mathbf{b})w = w_0^- - w_0^+ > 0$$

for  $\mathbf{p} \in \mathcal{A}_+$  and  $\mathbf{q} \in \mathcal{A}_-$ .

- $H_{\mathbf{w},+}$ : hyperplane defined by  $f_{\mathbf{w},+}(\mathbf{x}) := w_0^+ + \mathbf{x}\mathbf{w} = 0$   
We have  $f_{\mathbf{w},+}(\mathbf{x}) \geq 0$  for  $\mathbf{x} \in \mathcal{A}_+$  and  $\exists \mathbf{x} \in \mathcal{A}_+ \cap H_{\mathbf{w},+}$ .
- $H_{\mathbf{w},-}$ : hyperplane defined by  $f_{\mathbf{w},-}(\mathbf{x}) := w_0^- + \mathbf{x}\mathbf{w} = 0$   
We have  $f_{\mathbf{w},-}(\mathbf{x}) \leq 0$  for  $\mathbf{x} \in \mathcal{A}_-$  and  $\exists \mathbf{x} \in \mathcal{A}_- \cap H_{\mathbf{w},-}$ .

## Definition

Let  $\mathcal{A} \subset \mathbb{R}^K$ . A hyperplane  $H$  defined by  $f(\mathbf{x}) = 0$  is called a **supporting hyperplane** of  $\mathcal{A}$  if  $f(\mathbf{x}) \geq 0$  for all  $\mathbf{x} \in \mathcal{A}$  or  $f(\mathbf{x}) \leq 0$  for all  $\mathbf{x} \in \mathcal{A}$  and  $\exists \mathbf{x} \in \mathcal{A} \cap H$ .

- By this definition,  $H_{\mathbf{w},+}$  and  $H_{\mathbf{w},-}$  are supporting hyperplanes of  $\mathcal{A}_+$  and  $\mathcal{A}_-$ , respectively.

- Define  $d(\mathbf{w})$  to be the distance between  $H_{\mathbf{w},+}$  and  $H_{\mathbf{w},-}$ . Then we have

$$d(\mathbf{w}) = \frac{|w_0^+ - w_0^-|}{\|\mathbf{w}\|}.$$

Proof. Assume  $\mathbf{q} \in H_{\mathbf{w},-}$ . Then  $d(\mathbf{w})$  is equal to the distance between  $\mathbf{q}$  and  $H_{\mathbf{w},+}$ . Thus we have

$$\begin{aligned} d(\mathbf{w}) &= \frac{|f_{\mathbf{w},+}(\mathbf{q})|}{\|\mathbf{w}\|} = \frac{|w_0^+ + \mathbf{q}\mathbf{w}|}{\|\mathbf{w}\|} = \frac{|w_0^- + \mathbf{q}\mathbf{w} + (w_0^+ - w_0^-)|}{\|\mathbf{w}\|} \\ &= \frac{|w_0^+ - w_0^-|}{\|\mathbf{w}\|}. \quad \square \end{aligned}$$

- The hyperplane  $H_{\mathbf{w}}$  defined by

$$\frac{w_0^+ + w_0^-}{2} + \mathbf{x}\mathbf{w} = 0$$

is at the same distance from  $H_{\mathbf{w},+}$  and  $H_{\mathbf{w},-}$ .

Proof. Assume  $\mathbf{p} \in H_{\mathbf{w}}$ .

$$\begin{aligned} \frac{|w_0^+ + \mathbf{p}\mathbf{w}|}{\|\mathbf{w}\|} &= \frac{\left| \frac{w_0^+ + w_0^-}{2} + \frac{w_0^+ - w_0^-}{2} + \mathbf{p}\mathbf{w} \right|}{\|\mathbf{w}\|} = \frac{\left| \frac{w_0^+ - w_0^-}{2} \right|}{\|\mathbf{w}\|} \\ \frac{|w_0^- + \mathbf{p}\mathbf{w}|}{\|\mathbf{w}\|} &= \frac{\left| \frac{w_0^+ + w_0^-}{2} + \frac{w_0^- - w_0^+}{2} + \mathbf{p}\mathbf{w} \right|}{\|\mathbf{w}\|} = \frac{\left| \frac{w_0^- - w_0^+}{2} \right|}{\|\mathbf{w}\|} \quad \square \end{aligned}$$

- Now let  $\mathbf{w}$  vary.
- What is the best  $\mathbf{w}$  for separating  $\mathcal{A}_+$  from  $\mathcal{A}_-$ ?
- Define the **optimal margin**  $d(\mathcal{A}_+, \mathcal{A}_-)$  by

$$d(\mathcal{A}_+, \mathcal{A}_-) := \max_{\mathbf{w}} d(\mathbf{w}).$$

- When  $d(\mathbf{w}_*) = d(\mathcal{A}_+, \mathcal{A}_-)$ , the hyperplane  $H_*$  defined by

$$H_* : \frac{w_0^+ + w_0^-}{2} + \mathbf{x}\mathbf{w}_* = 0$$

is called the **classifying hyperplane**.



- The hyperplane  $H_*$  is at the same distance from  $H_{\mathbf{w}_*,+}$  and  $H_{\mathbf{w}_*,-}$ .

Question: How to find the classifying hyperplane?

Naive Idea: Find  $\mathbf{a} \in \mathcal{A}_+$  and  $\mathbf{b} \in \mathcal{A}_-$  be such that

$$\|\mathbf{a} - \mathbf{b}\| \leq \|\mathbf{p} - \mathbf{q}\| \quad \text{for all } \mathbf{p} \in \mathcal{A}_+, \mathbf{q} \in \mathcal{A}_-.$$

Set  $\mathbf{w} = (\mathbf{a} - \mathbf{b})^\top$  and use  $\mathbf{w}$  to construct the classifying hyperplane.



This idea works if we consider the line!

- Recall a set  $\mathcal{X} \subset \mathbb{R}^K$  is said to be **convex** if for any  $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ ,

$$\{\lambda \mathbf{x} + (1 - \lambda) \mathbf{y} : 0 \leq \lambda \leq 1\} \subseteq \mathcal{X},$$

or equivalently,

$$\{\lambda_1 \mathbf{x} + \lambda_2 \mathbf{y} : \lambda_1 + \lambda_2 = 1, \lambda_i \in \mathbb{R}_{\geq 0}\} \subseteq \mathcal{X}.$$

- Let  $\mathcal{A} \subset \mathbb{R}^K$ . The **convex hull**  $C(\mathcal{A})$  of  $\mathcal{A}$  is defined to be the smallest convex set that contains  $\mathcal{A}$ .
- When  $\mathcal{A} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ , the convex hull  $C(\mathcal{A})$  is given by

$$C(\mathcal{A}) = \left\{ \sum_{i=1}^N \lambda_i \mathbf{x}_i : \sum_{i=1}^N \lambda_i = 1, \lambda_i \in \mathbb{R}_{\geq 0} \right\}.$$

Proof. Exercise





## Lemma

*Let  $\mathcal{A} \subset \mathbb{R}^k$  be a finite set. Then a hyperplane  $H$  is a supporting hyperplane of  $\mathcal{A}$  if and only if it is a supporting hyperplane of  $C(\mathcal{A})$ .*

Proof. Let  $\mathcal{A} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ .

$\Rightarrow$ ) Assume  $H$  is a supporting hyperplane of  $\mathcal{A}$  defined by  $f(\mathbf{x}) = w_0 + \mathbf{x}\mathbf{w} = 0$ . WLOG, suppose  $f(\mathbf{x}) \geq 0$  for  $\mathbf{x} \in \mathcal{A}$ . Let  $\sum \lambda_i \mathbf{x}_i \in C(\mathcal{A})$  with  $\sum \lambda_i = 1$  and  $\lambda_i \geq 0$ . Then

$$f\left(\sum \lambda_i \mathbf{x}_i\right) = w_0 + \sum \lambda_i \mathbf{x}_i \mathbf{w} = \sum \lambda_i (w_0 + \sum \mathbf{x}_i \mathbf{w}) \geq 0.$$

There exists  $\mathbf{x} \in \mathcal{A} \cap H \subset C(\mathcal{A}) \cap H$ . Thus  $H$  is a supporting hyperplane of  $C(\mathcal{A})$ .

$\Leftarrow$ ) Assume  $H$  is a supporting hyperplane of  $C(\mathcal{A})$  defined by  $f(\mathbf{x}) = w_0 + \mathbf{x}\mathbf{w} = 0$ . WLOG, suppose  $f(\mathbf{x}) \geq 0$  for  $\mathbf{x} \in C(\mathcal{A})$ . Then  $f(\mathbf{x}) \geq 0$  for  $\mathbf{x} \in \mathcal{A} \subset C(\mathcal{A})$ . Suppose  $f(\sum \lambda_i \mathbf{x}_i) = 0$  with  $\sum \lambda_i = 1$  and  $\lambda_i \geq 0$ . Then

$$f\left(\sum \lambda_i \mathbf{x}_i\right) = \sum \lambda_i (w_0 + \mathbf{x}_i \mathbf{w}) = 0.$$

We must have  $\lambda_{i_*} > 0$  for some  $i_*$ , and  $w_0 + \mathbf{x}_{i_*} \mathbf{w} = 0$ . Thus  $H$  is a supporting hyperplane of  $\mathcal{A}$ . □

## Proposition

Let  $\mathbf{p} \in C(\mathcal{A}_+)$  and  $\mathbf{q} \in C(\mathcal{A}_-)$ . Then

$$\|\mathbf{p} - \mathbf{q}\| \geq d(\mathcal{A}_+, \mathcal{A}_-).$$

Proof. For any  $\mathbf{w}$ , we have

$$\|\mathbf{p} - \mathbf{q}\| \|\mathbf{w}\| \geq |(\mathbf{p} - \mathbf{q})\mathbf{w}| \geq |w_0^+ - w_0^-|$$

(the proof the second inequality is almost the same as the one for  $\mathcal{A}_+$  and  $\mathcal{A}_-$  since  $C(\mathcal{A}_+)$  and  $C(\mathcal{A}_-)$  are compact) and hence

$$\|\mathbf{p} - \mathbf{q}\| \geq \frac{|w_0^+ - w_0^-|}{\|\mathbf{w}\|} = d(\mathbf{w}).$$

In particular,  $\|\mathbf{p} - \mathbf{q}\| \geq d(\mathcal{A}_+, \mathcal{A}_-)$ . □

## Theorem

Let  $\mathbf{a} \in C(\mathcal{A}_+)$  and  $\mathbf{b} \in C(\mathcal{A}_-)$  be such that

$$\|\mathbf{a} - \mathbf{b}\| \leq \|\mathbf{p} - \mathbf{q}\| \quad \text{for all } \mathbf{p} \in C(\mathcal{A}_+), \mathbf{q} \in C(\mathcal{A}_-).$$

Set  $\mathbf{w} = (\mathbf{a} - \mathbf{b})^\top$ ,  $w_0^+ = -\mathbf{a}\mathbf{w}$  and  $w_0^- = -\mathbf{b}\mathbf{w}$ . Then the supporting hyperplanes  $H_{\mathbf{w},+}$  and  $H_{\mathbf{w},-}$  of  $\mathcal{A}_+$  and  $\mathcal{A}_-$  are respectively defined by

$$w_0^+ + \mathbf{x}\mathbf{w} = 0 \quad \text{and} \quad w_0^- + \mathbf{x}\mathbf{w} = 0,$$

and the optimal margin is equal to

$$d(\mathcal{A}_+, \mathcal{A}_-) = \|\mathbf{a} - \mathbf{b}\| = \frac{|w_0^+ - w_0^-|}{\|\mathbf{w}\|}.$$

Proof. For the first part, we need to show  $\mathbf{aw} \leq \mathbf{pw}$  for all  $\mathbf{p} \in C(\mathcal{A}_+)$  and  $\mathbf{bw} \geq \mathbf{qw}$  for all  $\mathbf{q} \in C(\mathcal{A}_-)$ . Then, from the previous discussion, the hyperplanes  $H_{\mathbf{w},+}$  and  $H_{\mathbf{w},-}$  are the supporting hyperplanes of  $\mathcal{A}_+$  and  $\mathcal{A}_-$ , respectively. Suppose  $\exists \mathbf{p} \in C(\mathcal{A}_+)$  such that  $\mathbf{aw} > \mathbf{pw}$ . Consider  $(1 - \lambda)\mathbf{a} + \lambda\mathbf{p} \in C(\mathcal{A}_+)$  for  $0 \leq \lambda \leq 1$ . Define

$$g(\lambda) = \|(1 - \lambda)\mathbf{a} + \lambda\mathbf{p} - \mathbf{b}\|^2 = \|\lambda(\mathbf{p} - \mathbf{a}) + \mathbf{a} - \mathbf{b}\|^2.$$

Then

$$g'(0) = 2(\mathbf{p} - \mathbf{a})(\mathbf{a} - \mathbf{b})^\top = 2(\mathbf{p} - \mathbf{a})\mathbf{w} < 0.$$

This is a contradiction because  $g(\lambda)$  has a minimum when  $\lambda = 0$ . Thus  $\mathbf{aw} \leq \mathbf{pw}$  for all  $\mathbf{p} \in C(\mathcal{A}_+)$ . Similarly,  $\mathbf{bw} \geq \mathbf{qw}$  for all  $\mathbf{q} \in C(\mathcal{A}_-)$ .

For the second part, it follows from Proposition that

$$\begin{aligned} d(\mathcal{A}_+, \mathcal{A}_-) &\geq \frac{|w_0^+ - w_0^-|}{\|\mathbf{w}\|} = \frac{|(\mathbf{a} - \mathbf{b})\mathbf{w}|}{\|\mathbf{w}\|} \\ &= \frac{\|\mathbf{w}\|^2}{\|\mathbf{w}\|} = \|\mathbf{w}\| = \|\mathbf{a} - \mathbf{b}\| \geq d(\mathcal{A}_+, \mathcal{A}_-). \end{aligned}$$



Task: Find  $\mathbf{a} \in C(\mathcal{A}_+)$  and  $\mathbf{b} \in C(\mathcal{A}_-)$  such that

$$\|\mathbf{a} - \mathbf{b}\| \leq \|\mathbf{p} - \mathbf{q}\| \quad \text{for all } \mathbf{p} \in C(\mathcal{A}_+), \mathbf{q} \in C(\mathcal{A}_-).$$

- There is an algorithm for this task called **Sequential Minimal Optimization** invented by John Platt in 1998.
- We will use the implementation in scikit-learn which in turn uses the LIBSVM library.

# Inseparable sets

Frequently, two sets  $\mathcal{A}_+$  and  $\mathcal{A}_-$  are not separable by a hyperplane.

- Option 1: Shrink the convex hulls by considering points  $\sum \lambda_i \mathbf{x}_i$  where  $\sum_i \lambda_i = 1$  and  $0 \leq \lambda_i \leq \alpha$  for some  $\alpha < 1$ . This excludes outliers.
- Option 2: Use a nonlinear kernel. That is, use nonlinear hypersurfaces.



# Multi-class classification


Question: How can we generalize SVM to multi-class cases?

Assume that there are  $s$  classes.

- One-versus-One: Compute all  $\binom{s}{2}$  pairwise classifiers. For each test point, the predicted class is the one that wins the most pairwise contests.
- One-versus-All: Each class is compared to all the others in  $s$  two-class comparisons. For each test point, compute the value of the separating hyperplane function  $f(\mathbf{x})$  for each of the  $s$  classifiers. The predicted class is the one with the largest value.

# Neural Networks

- **Neural Networks:** mimicking the operation of neurons in the brain.





- Form a graph with **input nodes** and **output nodes**,  
and with **hidden layers** of extra nodes between them.
- The output of each layer is the input to the next layer.  
The forwarded input goes through an **activation** at each hidden layer.
- Each node is meant to play the role of a **neuron** in the brain.
- The neural network is also called the **multi-layer perceptron** or **MLP**, in short.

- In each layer of a typical neural network, the process is similar to a logistic regression.

Neural networks  $\approx$  multiple layers of logistic regressions

- Activation functions are given  
by a sigmoid or a rectified linear unit (relu).

$$r(x) = \max(x, 0), \quad r'(x) = u_0(x) = \begin{cases} 1 & \text{if } x > 0, \\ 0 & \text{if } x < 0. \end{cases}$$



- Multi-class Classification:

$$\mathcal{D} = \{(x_1, \dots, x_k; t)\}, \quad t \in \{1, 2, \dots, s\}$$

- Build a network with  $k$  input nodes and  $s$  output nodes.  
Put one hidden layer with  $m$  nodes.
- Each output node will produce probability for the input to be in the corresponding class.  
(Multi-class logistic regression = one-layer neural network)

For simplicity, we take  $N = 1$  in what follows.

input:  $\mathbf{x} = [1, x_1, \dots, x_k]$

hidden layer:  $\mathbf{xw}^{(1)} = [z_1, \dots, z_m]$

with  $\mathbf{w}^{(1)}$  of size  $(k + 1) \times m$

activation:  $\sigma(z_i)$  ( $\sigma$  : sigmoid)

$$\mathbf{h} = [1, h_1, \dots, h_m]$$

$$= [1, \sigma(z_1), \dots, \sigma(z_m)]$$

output:  $\mathbf{y} = [y_1, \dots, y_s] = \sigma(\mathbf{hw}^{(2)})$  ( $\sigma$  : softmax)

with  $\mathbf{w}^{(2)}$  of size  $(m + 1) \times s$

As in logistic regression, we want to learn

the best values of  $\mathbf{w}^{(\ell)}$  ( $\ell = 1, 2$ ) using the training data.

- We have

$$\boxed{x \rightsquigarrow h \rightsquigarrow y}.$$

This process is considered as **forward propagation**.

- We can develop **more general** neural networks by considering more complex directed graphs with many layers and by adopting activation functions different from  $\sigma$ .
- However, there should be no oriented cycles in the directed graph to ensure that the outputs are deterministic functions of the inputs. In other words, a network must be **feed-forward**.

- The likelihood function is given by

$$L(\mathbf{w}^{(1)}, \mathbf{w}^{(2)}) = \prod_{i=1}^s y_i^{t_i},$$

where  $t$  is identified with a binary vector  $\mathbf{t} = [t_1, \dots, t_n]$ .

- Cross-entropy

$$E(\mathbf{w}^{(1)}, \mathbf{w}^{(2)}) = - \sum_{i=1}^s t_i \ln y_i$$

- We will use gradient descent and need to take derivatives.  
In particular, we need to use the chain rule inductively.

$$\nabla_{\mathbf{w}^{(2)}} E \rightsquigarrow \nabla_{\mathbf{w}^{(1)}} E$$

- This process is called **backpropagation**.



- Recall that the softmax  $\sigma : \mathbb{R}^s \rightarrow (0, 1)^s$  is defined by

$$\sigma(\mathbf{a}) = \left( \frac{e^{a_1}}{\sum_{i=1}^s e^{a_i}}, \dots, \frac{e^{a_s}}{\sum_{i=1}^s e^{a_i}} \right),$$

where  $\mathbf{a} = (a_1, a_2, \dots, a_s)$ .

- Write  $\mathbf{y} = [y_1, \dots, y_s] = \sigma(\mathbf{a})$ , and  
for  $i, j = 1, \dots, s$ ,

$$\frac{\partial y_j}{\partial a_i} = y_j(\delta_{i,j} - y_i).$$

- Write  $\mathbf{xw}^{(1)} = [z_1^{(1)}, \dots, z_m^{(1)}]$  and  $\mathbf{hw}^{(2)} = [z_1^{(2)}, \dots, z_s^{(2)}]$ .

In particular,

$$z_i^{(2)} = \sum_r h_r w_{r,i}^{(2)}.$$

- As in multi-class logistic regression,

$$\frac{\partial E}{\partial z_i^{(2)}} = - \sum_{j=1}^s t_j \frac{1}{y_j} \frac{\partial y_j}{\partial z_i^{(2)}} = - \sum_{j=1}^s t_j (\delta_{i,j} - y_i) = y_i - t_i,$$

$$\frac{\partial E}{\partial w_{p,q}^{(2)}} = \sum_{i=1}^s \frac{\partial E}{\partial z_i^{(2)}} \frac{\partial z_i^{(2)}}{\partial w_{p,q}^{(2)}} = h_p (y_q - t_q),$$

and

$$\nabla_{\mathbf{w}^{(2)}} E = \mathbf{h}^\top (\mathbf{y} - \mathbf{t}).$$

Next we have

$$\begin{aligned}
 \frac{\partial E}{\partial z_i^{(1)}} &= \sum_{j=1}^s \frac{\partial E}{\partial z_j^{(2)}} \frac{\partial z_j^{(2)}}{\partial z_i^{(1)}} = \sum_{j=1}^s (y_j - t_j) \sum_{\ell=1}^m \frac{\partial z_j^{(2)}}{\partial h_\ell} \frac{\partial h_\ell}{\partial z_i^{(1)}} \\
 &= \sum_{j=1}^s (y_j - t_j) \sum_{\ell=1}^m w_{\ell j}^{(2)} \delta_{i,\ell} h_i (1 - h_i) = h_i (1 - h_i) \sum_{j=1}^s (y_j - t_j) w_{ij}^{(2)}, \\
 \frac{\partial E}{\partial w_{p,q}^{(1)}} &= \sum_{i=1}^m \frac{\partial E}{\partial z_i^{(1)}} \frac{\partial z_i^{(1)}}{\partial w_{p,q}^{(1)}} = x_p h_q (1 - h_q) \sum_{j=1}^s (y_j - t_j) w_{qj}^{(2)},
 \end{aligned}$$

$$\text{(Note that } z_i^{(1)} = \sum_r x_r w_{r,i}^{(1)} \text{.)}$$

and

$$\nabla E_{\mathbf{w}^{(1)}} = \mathbf{x}^T \left[ h_q (1 - h_q) \sum_{j=1}^s (y_j - t_j) w_{qj}^{(2)} \right]_{q=1, \dots, m}.$$

- Write

$$\delta_i^{(a)} = \frac{\partial E}{\partial z_i^{(a)}} \quad \text{for } a = 1, 2.$$

Then

$$\delta_i^{(2)} = \frac{\partial E}{\partial z_i^{(2)}} = y_i - t_i,$$

$$\delta_i^{(1)} = \frac{\partial E}{\partial z_i^{(1)}} = h_i(1 - h_i) \sum_{j=1}^s (y_j - t_j) w_{ij}^{(2)}.$$

Notice that  $h_i = \sigma(z_i^{(1)})$  and  $h_i(1 - h_i) = \sigma'(z_i^{(1)})$ . The formula

$$\delta_i^{(1)} = \sigma'(z_i^{(1)}) \sum_{j=1}^s \delta_j^{(2)} w_{ij}^{(2)}$$

is called the **backpropagation formula**.

$\nabla E_{\mathbf{w}^{(1)}}$  and  $\nabla E_{\mathbf{w}^{(2)}}$   $\rightsquigarrow$  Gradient Decent

### Caveats:

- The error function  $E$  is **non-convex** (and non-concave).
- Initialization – don't take the zero matrix for  $\mathbf{w}^{(1)}$
- More susceptible to **over-fitting** – a lot more of parameters

# Remarks

- Complex neural networks with multiple layers are usually called **deep** neural networks.
- Most **deep learning** models are based on **convolutional** neural networks.
- **TensorFlow** is an open-source software library for machine learning with a particular focus on deep neural networks, and **Keras** provides a Python interface for the TensorFlow library.
- **PyTorch** is an open-source machine learning library designed to provide flexibility and speed for deep neural network implementation.

# Convolution Neural Networks (CNN)

- Introduced by Yann LeCun in 1989
- Convolution

$$(f * g)(x) = \int_{-\infty}^{\infty} f(x - t)g(t) dt$$

$$(f * g)(x) = \sum_{t=-\infty}^{\infty} f(x - t)g(t)$$

- In Math 3160, when  $X$  and  $Y$  are independent,

$$f_{X+Y} = f_X * f_Y.$$

- Convolution is [translation-invariant](#), i.e., when  $\tau_a f(x) = f(x - a)$ ,

$$\tau_a(f * g) = \tau_a(f) * g.$$

- In Machine Learning, the convolution of 2-D arrays is defined by

input:  $X = [X_{i,j}]$       kernel (or filter):  $\mathbf{w} = [w_{i,j}]$

$$\text{output: } (X * \mathbf{w})_{i,j} = \sum_{c,d} X_{i+c,j+d} w_{c,d}$$

- For example,

|   |   |   |   |
|---|---|---|---|
| 1 | 2 | 0 | 3 |
| 0 | 2 | 1 | 2 |
| 1 | 0 | 0 | 2 |
| 3 | 0 | 2 | 0 |

 \* 

|   |   |
|---|---|
| 1 | 0 |
| 1 | 2 |


 = 

|   |   |   |
|---|---|---|
| 5 | 6 | 5 |
| 1 | 2 | 5 |
| 4 | 4 | 2 |

- Clearly, this operation is translation-invariant.



- Convolution brings about **parameter sharing** and **sparse connectivity**.



- After convolution, **pooling** is performed. Two typical methods are the max pooling and the average pooling.
- Pooling is approximately translation-invariant and helps reduce noise.

For example,

|   |   |   |
|---|---|---|
| 1 | 2 | 0 |
| 3 | 1 | 0 |
| 0 | 0 | 0 |

|   |   |   |
|---|---|---|
| 0 | 0 | 0 |
| 0 | 1 | 2 |
| 0 | 3 | 1 |

Notice that the maximum or the average does not change.

- The result of convolution and pooling will go through **activation** and then move forward to the next layer.

# Ideas of CNN

- Sparse connectivity


Nearby pixels on an image are more strongly correlated than more distant pixels. Recognizing local features can enhance the performance of a classifier.

- Parameter sharing

Any useful features that are detected in some portion of the input may be valid in other parts. The weights are shared across the layer, and the number of parameters is reduced.

- Translation invariance

Features of images are invariant under translation.




# Some specifics for code

- We use the max pooling and the relu function for activation.
- We do not perform padding for simplicity. In general, padding brings several benefits.
- Multiple for-loops make computations very slow.

We need to **vectorize** our computations.

For a convolution with a  $2 \times 2$  kernel, we convert



and use a matrix multiplication once.

- We need to record where the maxima occur in the max pooling.

- We may use multiple kernels (or filters).  
Different kernels learn different features.
- In order to boost speed, we use stochastic gradient descent.

input:  $\mathbf{x} = [x_1, \dots, x_k]$  (no bias here)

hidden layer: convolution & pooling with filter  $\mathbf{w}^{(1)}$

$$\rightsquigarrow [z_1^{(1)}, \dots, z_m^{(1)}]$$

activation:  $\sigma(z_i^{(1)}) = r(w_0^{(1)} + z_i^{(1)})$

$r(x)$  is the relu and  $w_0^{(1)}$  is a bias.

$$\mathbf{h} = [1, h_1, \dots, h_m] = [1, \sigma(z_1^{(1)}), \dots, \sigma(z_m^{(1)})]$$

output:  $\mathbf{y} = [y_1, \dots, y_s] = \sigma(\mathbf{h}\mathbf{w}^{(2)})$  ( $\sigma$  : softmax)

with  $\mathbf{w}^{(2)}$  of size  $(m+1) \times s$

$$\text{and } \mathbf{h}\mathbf{w}^{(2)} = [z_1^{(2)}, \dots, z_s^{(2)}]$$

For  $\mathbf{w}^{(2)}$ , we have the same formula

$$\nabla_{\mathbf{w}^{(2)}} E = \mathbf{h}^\top (\mathbf{y} - \mathbf{t}).$$

Recall the backpropagation:

$$\delta_i^{(1)} = \sigma'(z_i^{(1)}) \sum_{j=1}^s \delta_j^{(2)} w_{ij}^{(2)}.$$

$$\delta_i^{(2)} = \frac{\partial E}{\partial z_i^{(2)}} = y_i - t_i,$$

$$\delta_i^{(1)} = \frac{\partial E}{\partial z_i^{(1)}} = u_0(w_0^{(1)} + z_i^{(1)}) \sum_{j=1}^s (y_j - t_j) w_{ij}^{(2)}.$$



Since  $\frac{\partial h_\ell}{\partial w_0^{(1)}} = u_0(w_0^{(1)} + z_\ell^{(1)})$  for any  $\ell$ , we see that

$$\begin{aligned}\frac{\partial E}{\partial w_0^{(1)}} &= \sum_{j=1}^s \frac{\partial E}{\partial z_j^{(2)}} \frac{\partial z_j^{(2)}}{\partial w_0^{(1)}} = \sum_{j=1}^s (y_j - t_j) \sum_{\ell=1}^m \frac{\partial z_j^{(2)}}{\partial h_\ell} \frac{\partial h_\ell}{\partial w_0^{(1)}} \\ &= \sum_{j=1}^s (y_j - t_j) \sum_{\ell=1}^m w_{\ell j}^{(2)} u_0(z_\ell^{(1)} + w_0^{(1)}) = \sum_{i=1}^m \delta_i^{(1)}, \\ \frac{\partial E}{\partial w_{p,q}^{(1)}} &= \sum_{i=1}^m \frac{\partial E}{\partial z_i^{(1)}} \frac{\partial z_i^{(1)}}{\partial w_{p,q}^{(1)}} = \sum_{i=1}^m \delta_i^{(1)} x_{i,p,q},\end{aligned}$$

where  $x_{i,p,q}$  are determined by pooling.