


Introduction



Course Details and Topics

Orientation

- This course addresses three cohorts of students:
 - ★ Third year Artificial Intelligence Students
 - ★ MSc in Machine Learning Students
 - ★ Volunteers
- It is dual coded (AICE3001/AICE6002), with subtle difference in expectations reflected in slightly different evaluation
- Notes are on Moodle and <https://ecs-vlc.github.io/aice3001/>


Why Mathematics?

- Learning machines are notoriously difficult to understand
- Mathematics allows us to reason rigorously
- Posing question mathematically and seeking mathematical solutions provide allow us to build a strong intuition—we can't cheat that much
- Mathematics also allow us to describe the world precisely
- It also provides an arsenal of tools that turn out to be hugely useful

Starting from Different Places

- You all start from different places
- Some of you may have a degree in mathematics
- For others it may have been your least favourite subject
- I invite strong mathematicians to become even stronger—the best way to learn is through teaching others
- For students starting from a weak base, judge yourselves by your progress, not by comparing yourself to others

1. Why are you doing this course?
2. What we assume that you know



Why are We Teaching You This?

- In a world where LLMs can do everything, why bother with learning mathematics?
- You need to learn to trust your own judgement
- LLMs can help us a lot, but if you don't trust yourself you are disempowered
- This course is aimed at building up trust in your understanding of how machine learning works

Cracking the Code

- Mathematics is a language that takes time to master
- For some it is easy, for others it will be a struggle
- But given determination everyone can crack the code and understand what the mathematics is saying
- This comes from engaging and working with maths—this course has tutorials that should be considered compulsory—not because there are marks attached, but because without doing the work you aren't going to crack the code

The Golden Rule of Mathematics

- Understand every step
- When you don't understand a step, stop, break the problem down, ask someone
- Mathematics uses a lot of short-cuts, which speeds-up calculations, but you need to be able to justify every step
- It takes bravery to admit you don't understand and insist that you will
- Asking "stupid" questions isn't stupid, it is how we all learn

- We are going to visit lots of aspects of mathematics that have proved fruitful in machine learning
- Some of this will be very concrete, while in other cases by look at a higher level of abstraction we can generalise methods across different application areas
- We will cover, calculus, vector spaces, convexity, optimisation, entropy, probability, statistics and much else besides
- The journey will pay you off, these are the foundations of most of machine learning and will give you super-powers that are hard to acquire without mathematics

What we Assume you Know

- You should all have A-level mathematics and have done some mathematics at Uni
- You should know what a function is, e.g. $f(x) = x^2 + 3$
- You should be able to do simple algebraic manipulations
- And some elementary calculus

$$f'(x) = \frac{d(x^2 + 3)}{dx} = 2x$$

Matrices

- Matrices are blocks of number that define linear mappings between vectors, e.g. $M \in \mathbb{R}^{2 \times 3}$
- An $m \times n$ matrix (i.e. $M \in \mathbb{R}^{m \times n}$) has m rows and n columns
- An $m \times n$ matrix acting on a vector in $\mathbb{R}^n$ transforms it to $\mathbb{R}^m$ vector, e.g.


$$Mx = \begin{pmatrix} 1.1 & -2.1 & 6.2 \\ 9.3 & 3.1 & -4.4 \end{pmatrix} \begin{pmatrix} 1.2 \\ 2.3 \\ 4.7 \end{pmatrix} = \begin{pmatrix} 28.63 \\ -2.39 \end{pmatrix}$$

Matrix Multiplication

- The combined mapping AB can be represented by an $\ell \times n$ matrix C which we can compute using matrix multiplication (rows times columns)

- Matrix multiplication is associative: $(AB)C = A(BC)$
- But in general not commutative $AB \neq BA$

1. Why are you doing this course?
2. What we assume that you know



Vectors

- Ordinary vectors (what mathematicians call elements of $\mathbb{R}^n$) are much used in machine learning and should be familiar
- We will use bold symbols to represent column vectors, e.g.

$$x = \begin{pmatrix} 1.2 \\ 2.3 \\ 4.7 \end{pmatrix}$$

- We denote row vectors as the transpose of column vectors $x^T = (1.2, 2.3, 4.7)$
- We care whether vectors are row or column vectors as it allows us to write equation involving vectors succinctly and unambiguously

Composition of Mappings

- Matrices, A , B , etc. represent mappings between spaces
- Provided the matrices are the correct size we can compose these mappings, e.g. if $A \in \mathbb{R}^{\ell \times m}$ and $B \in \mathbb{R}^{m \times n}$ then AB represents the effect of
 - ★ first applying B to vectors in $\mathbb{R}^n$ to obtain vectors in $\mathbb{R}^m$
 - ★ and then applying A to get vectors in $\mathbb{R}^\ell$

$$ABx = A(Bx)$$

Transpose

- We often encounter transposes of a matrix M^T where we have exchanged the rows and columns: $[M^T]_{ij} = [M]_{ji}$
- Clear $(M^T)^T = M$
- It is simple question of book keeping to show that

$$(AB)^T = B^T A^T$$
- For a symmetric matrix $M = M^T$
- Since 1×1 matrices are always symmetric

$$a^T b = (a^T b)^T = b^T a$$

$$a^T M b = b^T M^T a$$

Matrix-Vector Multiplication

- We can treat a column vector as a $n \times 1$ matrix
- Then matrix-vector multiplication can be treated the same as matrix multiplication

$$\begin{pmatrix} \blacksquare & \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare & \blacksquare & \blacksquare \end{pmatrix} \begin{pmatrix} \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \end{pmatrix} = \begin{pmatrix} \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \end{pmatrix}$$

Inner Products

- The dot or inner-product between two vectors with the same number of elements, can be written $a^T b$

$$\begin{pmatrix} \blacksquare & \blacksquare & \blacksquare & \blacksquare & \blacksquare \end{pmatrix} \begin{pmatrix} \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \end{pmatrix} = \blacksquare$$

- Inner products encode angles between vectors

$$a^T b = \|a\| \|b\| \cos(\theta)$$

Outer Products

- We can also define the outer-product between any size vectors, ab^T

$$\begin{pmatrix} \blacksquare \\ \blacksquare \\ \blacksquare \\ \blacksquare \end{pmatrix} \begin{pmatrix} \blacksquare & \blacksquare & \blacksquare \end{pmatrix} = \begin{pmatrix} \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare \\ \blacksquare & \blacksquare & \blacksquare \end{pmatrix}$$

- Notice that we need to know whether we are dealing with row or column vectors

Norms

- Norms are non-negative numbers associated with the "size" of a vector
- The norm you should be familiar with is the **Euclidean Norm** or 2-norm

$$\|x\|_2 = \sqrt{\sum_{i=1}^n x_i^2} = \sqrt{x_1^2 + x_2^2 + \dots + x_n^2} = \sqrt{x^T x}$$

- We will see there exist other norms (using different ways of measuring size)
- The Euclidean norm is intimately related to the inner-product, we will see other examples of this for different vector spaces

Matrices as Linear Mappings

- Matrices represent mappings between vectors

$$b = Ma$$

- The elements of b are **linear sums** of the elements of a


$$b_i = \sum_j M_{ij} a_j = M_{i1} a_1 + M_{i2} a_2 + \dots + M_{in} a_n$$

- Two important consequences of this linearity are for any real c

$$\begin{aligned} M(ca) &= c(Ma) \\ M(a+b) &= (Ma) + (Mb) \end{aligned}$$

Action Of Matrices on Volumes

- For a closed set of points corresponding to a volume, V in $\mathbb{R}^n$, treating each point as a vector, we can ask what happens to the volume after mapping the vector by a matrix



- For square matrices, M , the determinant is equal to $\det(M) = \pm V'/V$
- The sign tells us if there has been a reflection

Determinants

- A consequence of the fact that AB denotes the composition, apply mapping B then apply mapping A

$$\det(AB) = \det(A) \times \det(B)$$

- Determinants are only defined for square matrices, as non-square matrices project into different size spaces, where the volumes cannot be compared

- Singular matrices** projects all points onto a lower dimensional space, so have determinant 0

Multivariate Functions

- In machine learning we heavily use multivariate functions $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$
- Often our functions will be scalar value $f: \mathbb{R}^n \rightarrow \mathbb{R}$
- E.g. a learning machine for regression might be $\hat{f}(x, w)$ or we might have a loss function $L(w)$
- We say a function f is linear if

$$f(cx) = cf(x) \quad f(x+y) = f(x) + f(y)$$

- The gradient of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ at a point $\mathbf{x}^*$ is defined as the column vector

$$\nabla f(\mathbf{x}^*) = \begin{pmatrix} \frac{\partial f(\mathbf{x}^*)}{\partial x_1} \\ \frac{\partial f(\mathbf{x}^*)}{\partial x_2} \\ \vdots \\ \frac{\partial f(\mathbf{x}^*)}{\partial x_n} \end{pmatrix}$$

- The gradient vector is orthogonal to the contour plane and points in the direction in which the function increases fastest

Probability Distributions

- For discrete events, we can define a **probability mass function** $\mathbb{P}(X = x)$ that the random variable takes the value x
- This is a number between 0 and 1 with $\sum_{x \in \mathcal{X}} \mathbb{P}(X = x) = 1$
- For continuous events the probability of X taking a particular value x is zero
- Instead we consider the probability density $f_X(x)$ which we can define as

$$f_X(x) = \lim_{\delta x \rightarrow 0} \frac{\mathbb{P}(x < X \leq x + \delta x)}{\delta x}$$

Lesson

- I've covered material that we assume you know
- If you don't, you are going to struggle following any machine learning course, so grab a maths book and revise this
- We will revisit most of these topics to gain a better understanding, but you will be lost in the other machine learning modules if you are not familiar with everything I've covered
- There are exercises we will go through in the tutorial session where you can test your proficiency

- We will have a lot to say about probability later in this course
- To reason about events with uncertain outcomes we introduce **random variables**
- Random variables corresponds to objects that we assign numbers to depending on the outcome of a random event
- E.g. we might assign 0 to the outcome of tossing a tail and 1 for a head—this is an example of a discrete random variable
- A continuous random variable might be the temperature at 12 noon tomorrow in Hyde park

Expectations


- One of the most useful quantities in dealing with random numbers is to look at the average value, aka as the **expectation**, for some function $g(X)$ of the random variable

$$\mathbb{E}[g(X)] = \begin{cases} \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) g(x) \\ \int f_X(x) g(x) dx \end{cases}$$

- Expectations are linear operators


$$\begin{aligned} \mathbb{E}[cg(X)] &= c\mathbb{E}[g(x)] \\ \mathbb{E}[g(X) + h(X)] &= \mathbb{E}[g(X)] + \mathbb{E}[h(X)] \end{aligned}$$

- The expectation of a constant is itself $\mathbb{E}[c] = c$

Differential Calculus*Differentiation, product and chain rules, vectors and matrices***Why Calculus?**

- Calculus is a fundamental tool of mathematical analysis
- In machine learning differentiation is fundamental tool in optimisation
- Integration is an essential tool in taking expectations over continuous distributions
- Both differentiation and integration crop up elsewhere


1. Why Calculus?
2. Differentiation
3. Vector Calculus
4. Matrix Calculus

**Back to Basics**

- You have all done A-level maths so should be familiar with the rules of calculus
- But, it is easy to forget the rules and sometimes we use quite sophisticated tricks
- Although the sophisticated tricks really speed up calculations, it pays to be able to understand where these tricks come from

Outline


1. Why Calculus?
2. Differentiation
3. Vector Calculus
4. Matrix Calculus

**Polynomials**

- $f(x) = x^2$

$$\frac{dx^2}{dx} = \lim_{\epsilon \rightarrow 0} \frac{(x+\epsilon)^2 - x^2}{\epsilon} = \lim_{\epsilon \rightarrow 0} \frac{(x^2 + 2\epsilon x + \epsilon^2) - x^2}{\epsilon} = \lim_{\epsilon \rightarrow 0} 2x + \epsilon = 2x$$
- $(x+\epsilon)^n = (x+\epsilon)(x+\epsilon)\cdots(x+\epsilon) = x^n + n\epsilon x^{n-1} + O(\epsilon^2)$


$$\frac{dx^n}{dx} = \lim_{\epsilon \rightarrow 0} \frac{(x+\epsilon)^n - x^n}{\epsilon} = \lim_{\epsilon \rightarrow 0} n x^{n-1} + O(\epsilon) = n x^{n-1}$$

Differentiation**Linearity of derivatives**

- Consider the derivative of $af(x) + bg(x)$

$$\begin{aligned} \frac{d(af(x) + bg(x))}{dx} &= \lim_{\epsilon \rightarrow 0} \frac{(af(x+\epsilon) + bg(x+\epsilon)) - (af(x) + bg(x))}{\epsilon} \\ &= \lim_{\epsilon \rightarrow 0} \frac{a(f(x+\epsilon) - f(x)) + b(g(x+\epsilon) - g(x))}{\epsilon} \\ &= \lim_{\epsilon \rightarrow 0} a \left(\frac{f(x+\epsilon) - f(x)}{\epsilon} \right) + b \left(\frac{g(x+\epsilon) - g(x)}{\epsilon} \right) \\ &= af'(x) + bg'(x) \end{aligned}$$

- Differentiation is a linear operation



Chain Rule

- Recall $f(x + \epsilon) = f(x) + \epsilon f'(x) + O(\epsilon^2)$
- Let $h(x) = f(g(x))$
- Then

$$\begin{aligned} h(x + \epsilon) &= f(g(x + \epsilon)) = f(g(x) + \epsilon g'(x) + O(\epsilon^2)) \\ &= f(g(x)) + \epsilon g'(x) f'(g(x)) + O(\epsilon^2) \end{aligned}$$

- Thus

$$h'(x) = \lim_{\epsilon \rightarrow 0} \frac{h(x + \epsilon) - h(x)}{\epsilon} = g'(x) f'(g(x))$$

- This is the famous **chain rule**. Together with the product rule it means you can differentiate almost everything

Inverse functions

- Suppose $g(y) = f^{-1}(y)$ is the inverse of $f(x)$ in the sense that $g(f(x)) = f^{-1}(f(x)) = x$
- Using the chain rule

$$\frac{dg(f(x))}{dx} = f'(x) g'(f(x)) = 1$$

since $g(f(x)) = x$

- So $g'(f(x)) = 1/f'(x)$
- Writing $y = f(x)$ so that $x = f^{-1}(y) = g(y)$ we find $g'(y) = 1/f'(g(y))$ that is

$$\frac{dg(y)}{dy} = \frac{1}{f'(g(y))} \quad \frac{df^{-1}(y)}{dy} = \frac{1}{f'(f^{-1}(y))}$$

Functions of Exponentials

- What about $f(x) = e^{cx}$

$$\frac{de^{cx}}{dx} = \frac{de^{cx}}{dcx} \frac{dcx}{dx} = ce^{cx}$$

- More generally using the chain rule

$$\frac{de^{g(x)}}{dx} = g'(x) e^{g(x)}$$

- Also $a^{bc} = (a^b)^c$ (that is we multiply a together $b \times c$ times)

$$\frac{da^x}{dx} = \frac{d(e^{\ln(a)x})}{dx} = \frac{de^{\ln(a)x}}{dx} = \ln(a) e^{\ln(a)x} = \ln(a) a^x$$

- Recall $f(x + \epsilon) = f(x) + \epsilon f'(x) + O(\epsilon^2)$
- If $h(x) = f(x)g(x)$

$$\begin{aligned} h'(x) &= \lim_{\epsilon \rightarrow 0} \frac{f(x + \epsilon)g(x + \epsilon) - f(x)g(x)}{\epsilon} \\ &= \lim_{\epsilon \rightarrow 0} \frac{(f(x) + \epsilon f'(x) + O(\epsilon^2))(g(x) + \epsilon g'(x) + O(\epsilon^2)) - f(x)g(x)}{\epsilon} \\ &= \lim_{\epsilon \rightarrow 0} \frac{\epsilon(f'(x)g(x) + f(x)g'(x)) + O(\epsilon^2)}{\epsilon} = f'(x)g(x) + f(x)g'(x) \end{aligned}$$

- This is the **product rule**

More on chain rules

- We can also write the chain rule as

$$\frac{df(g(x))}{dx} = \frac{df(g)}{dg} \frac{dg(x)}{dx}$$


- Sometimes this is neater or easier to remember

$$\begin{aligned} \frac{de^{\cos(x^2)}}{dx} &= \frac{de^{\cos(x^2)}}{d\cos(x^2)} \frac{d\cos(x^2)}{dx^2} \frac{dx^2}{dx} \\ &= e^{\cos(x^2)} (-\sin(x^2)) 2x \\ &= -2x \sin(x^2) e^{\cos(x^2)} \end{aligned}$$

Exponentials

- Note that $a^{b+c} = a^b a^c$ (that is we multiply a together $b + c$ times)

- Now $e^\epsilon \approx (1 + \epsilon)$



- But $e^{x+\epsilon} = e^x e^\epsilon = e^x (1 + \epsilon + O(\epsilon^2)) = e^x + \epsilon e^x + O(\epsilon^2)$

$$\frac{de^x}{dx} = \lim_{\epsilon \rightarrow 0} \frac{e^{x+\epsilon} - e^x}{\epsilon} = \lim_{\epsilon \rightarrow 0} \frac{\epsilon e^x + O(\epsilon^2)}{\epsilon} = e^x$$

Natural Logarithms


- The natural logarithm is defined as the inverse of e^x

$$\ln(e^x) = x \quad e^{\ln(y)} = y$$


- Recall that if $g(y) = f^{-1}(y)$ then $g'(y) = 1/f'(g(y))$

- Consider $g(y) = \ln(y)$ and $f(x) = e^x$ (with $f'(x) = e^x$)

$$\frac{d\ln(y)}{dy} = \frac{1}{e^{\ln(y)}} = \frac{1}{y}$$



1. Why Calculus?
2. Differentiation
3. **Vector Calculus**
4. Matrix Calculus



Derivatives in High Dimensions

- When working with functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$ in many dimensions then there will typically be different derivative in different directions
- To compute the derivative in a direction $\mathbf{u} \in \mathbb{R}^n$ (where $\|\mathbf{u}\| = 1$) at a point $\mathbf{x} \in \mathbb{R}^n$ we use

$$\partial_{\mathbf{u}} F(\mathbf{x}) = \lim_{\epsilon \rightarrow 0} \frac{f(\mathbf{x} + \epsilon \mathbf{u}) - f(\mathbf{x})}{\epsilon}$$

- If $\mathbf{u} = \delta_i = (0, \dots, 0, 1, 0, \dots, 0)^T$ (i.e. $u_i = 1$) then

$$\frac{\partial f(\mathbf{x})}{\partial x_i} = \lim_{\epsilon \rightarrow 0} \frac{f(\mathbf{x} + \epsilon \delta_i) - f(\mathbf{x})}{\epsilon}$$

Computing Gradients 1

- We can compute the gradient by writing out $f(\mathbf{x})$ componentwise and performing the partial derivative with respect to x_i

$$\begin{aligned} \nabla \mathbf{w}^T \mathbf{M} \mathbf{w} &= \begin{pmatrix} \frac{\partial}{\partial w_1} \\ \frac{\partial}{\partial w_2} \\ \frac{\partial}{\partial w_3} \\ \vdots \end{pmatrix} \sum_{i,j} w_i M_{ij} w_j = \begin{pmatrix} \sum_j M_{1j} w_j + \sum_i w_i M_{i1} \\ \sum_j M_{2j} w_j + \sum_i w_i M_{i2} \\ \sum_j M_{3j} w_j + \sum_i w_i M_{i3} \\ \vdots \end{pmatrix} \\ &= \mathbf{M} \mathbf{w} + \mathbf{M}^T \mathbf{w} = (\mathbf{M} + \mathbf{M}^T) \mathbf{w} \end{aligned}$$

- It is tedious to compute these things component-wise, but when you need to understand what is going on then go back to the basics

Function Composition

- The gradient tells us about the change in a function value when we perturb an input $\mathbf{x}$ by a small amount in the directions $\mathbf{u}$

$$f(\mathbf{x} + \epsilon \mathbf{u}) = f(\mathbf{x}) + \epsilon \mathbf{u}^T \nabla f(\mathbf{x}) + O(\epsilon^2)$$

- What happens when we compose functions: $h(\mathbf{x}) = f(g(\mathbf{x}))$? Note that,

$$\mathbf{x} \in \mathbb{R}^n \quad \mathbf{g}: \mathbb{R}^n \rightarrow \mathbb{R}^m \quad \mathbf{f}: \mathbb{R}^m \rightarrow \mathbb{R} \quad h: \mathbb{R}^n \rightarrow \mathbb{R}$$

- The vector valued function $\mathbf{g}(\mathbf{x})$ can be seen as m independent

$$\text{function } \mathbf{g}(\mathbf{x}) = \begin{pmatrix} g_1(\mathbf{x}) \\ g_2(\mathbf{x}) \\ \vdots \\ g_m(\mathbf{x}) \end{pmatrix}$$

Taylor

- If we expand $f(\mathbf{x} + \epsilon \mathbf{u})$ to first order in ϵ

$$f(\mathbf{x} + \epsilon \mathbf{u}) = f(\mathbf{x}) + \epsilon \mathbf{u}^T \mathbf{g}(\mathbf{x}) + O(\epsilon^2)$$

$$\text{then } g_i(\mathbf{x}) = \frac{\partial f(\mathbf{x})}{\partial x_i}$$

- Recall we defined the vector of first order derivatives of $f(\mathbf{x})$ to be the gradient

$$\nabla f(\mathbf{x}) = \begin{pmatrix} \frac{\partial f(\mathbf{x})}{\partial x_1} \\ \frac{\partial f(\mathbf{x})}{\partial x_2} \\ \vdots \\ \frac{\partial f(\mathbf{x})}{\partial x_n} \end{pmatrix}$$

- Thus

$$f(\mathbf{x} + \epsilon \mathbf{u}) = f(\mathbf{x}) + \epsilon \mathbf{u}^T \nabla f(\mathbf{x}) + O(\epsilon^2)$$

This is the start of the high-dimensional Taylor expansion

Computing Gradients 2

- A slicker way is just to expand $f(\mathbf{x} + \epsilon \mathbf{u})$

- Consider $f(\mathbf{x}) = \mathbf{x}^T \mathbf{M} \mathbf{x} + \mathbf{a}^T \mathbf{x}$

$$\begin{aligned} f(\mathbf{x} + \epsilon \mathbf{u}) &= (\mathbf{x} + \epsilon \mathbf{u})^T \mathbf{M} (\mathbf{x} + \epsilon \mathbf{u}) + \mathbf{a}^T (\mathbf{x} + \epsilon \mathbf{u}) \\ &= f(\mathbf{x}) + \epsilon (\mathbf{u}^T \mathbf{M} \mathbf{x} + \mathbf{x}^T \mathbf{M} \mathbf{u} + \mathbf{a}^T \mathbf{u}) + O(\epsilon^2) \\ &= f(\mathbf{x}) + \epsilon \mathbf{u}^T (\mathbf{M} \mathbf{x} + \mathbf{M}^T \mathbf{x} + \mathbf{a}) + O(\epsilon^2) \end{aligned}$$

$$\text{using } \mathbf{x}^T \mathbf{M} \mathbf{u} = \mathbf{u}^T \mathbf{M}^T \mathbf{x} \text{ and } \mathbf{a}^T \mathbf{u} = \mathbf{u}^T \mathbf{a}$$

- But $f(\mathbf{x} + \epsilon \mathbf{u}) = f(\mathbf{x}) + \epsilon \mathbf{u}^T \nabla f(\mathbf{x}) + O(\epsilon^2)$ so

$$\nabla f(\mathbf{x}) = \mathbf{M} \mathbf{x} + \mathbf{M}^T \mathbf{x} + \mathbf{a}$$

Jacobians

- Perturbing $\mathbf{x}$ by $\epsilon \mathbf{u}$

$$g_i(\mathbf{x} + \epsilon \mathbf{u}) = g_i(\mathbf{x}) + \epsilon \mathbf{u}^T \nabla g_i(\mathbf{x}) + O(\epsilon^2)$$

so that

$$\mathbf{g}(\mathbf{x} + \epsilon \mathbf{u}) = \mathbf{g}(\mathbf{x}) + \epsilon \mathbf{J} \mathbf{u} + O(\epsilon^2)$$

where $\mathbf{J}$ is an $m \times n$ matrix with components $J_{ij} = \frac{\partial g_i(\mathbf{x})}{\partial x_j}$

- Note that $\mathbf{J} \mathbf{u}$ is a column vector of size m
- Ignoring terms of order ϵ^2

$$\begin{aligned} h(\mathbf{x} + \epsilon \mathbf{u}) &= f(\mathbf{g}(\mathbf{x} + \epsilon \mathbf{u})) = f(\mathbf{g}(\mathbf{x}) + \epsilon \mathbf{J} \mathbf{u}) \\ &= f(\mathbf{g}(\mathbf{x})) + \epsilon \mathbf{u}^T \mathbf{J}^T \nabla f(\mathbf{g}(\mathbf{x})) \end{aligned}$$

- Using $h(x) = f(g(x))$

$$h(x + \epsilon u) = h(x) + \epsilon u^T J^T \nabla_g f(g(x))$$

- We get that

$$\nabla_x h(x) = \nabla_x f(g(x)) = J^T \nabla_g f(g(x))$$

- Componentwise

$$\frac{\partial h(x)}{\partial x_i} = \frac{\partial f(g(x))}{\partial x_i} = \sum_{j=1}^m \frac{\partial g_j(x)}{\partial x_i} \frac{\partial f(g(x))}{\partial g_j}$$

Differentiating Matrices

- Often we have loss functions with respect to a matrix W , e.g.

$$L(W) = (a^T W b - c)^2$$

- We might want to find the minimum with respect to W
- This occurs at a point W^* where $L(W)$ does not increase as we change W in any way
- That is, we seek a W^* such that, for any matrices U

$$L(W^* + \epsilon U) - L(W^*) = O(\epsilon^2)$$

Example

- Suppose

$$L(W) = (a^T W b - c)^2$$

then

$$\begin{aligned} L(W + \epsilon U) &= (a^T (W + \epsilon U) b - c)^2 = (a^T W b + \epsilon a^T U b - c)^2 \\ &= L(W) + 2\epsilon (a^T W b - c) (a^T U b) + O(\epsilon^2) \end{aligned}$$

- Now

$$a^T U b = \sum_{ij} a_i U_{ij} b_j = \sum_{ij} U_{ji} a_j b_i = \text{tr } U^T a b^T$$

$$\text{Thus } \frac{\partial L(W)}{\partial W} = 2(a^T W b - c) a b^T$$

Quick Matrix Differentiation

- Let

$$\partial_U f(X) = \lim_{\epsilon \rightarrow 0} \frac{f(X + \epsilon U) - f(X)}{\epsilon} = \text{tr } U^T \frac{\partial f(X)}{\partial X}$$


- E.g.

$$\begin{aligned} \partial_U \text{tr } A X B &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} (\text{tr } A (X + \epsilon U) B - \text{tr } A X B) \\ &= \text{tr } A U B = \text{tr } B^T U^T A^T = \text{tr } U^T A^T B^T \end{aligned}$$

thus

$$\frac{\partial \text{tr } A X B}{\partial X} = A^T B^T$$

- Why Calculus?
- Differentiation
- Vector Calculus
- Matrix Calculus



Generalised Gradient

- We can generalise the idea of gradient to matrices

$$\frac{\partial L(W)}{\partial W} = \begin{pmatrix} \frac{\partial L(W)}{\partial W_{11}} & \frac{\partial L(W)}{\partial W_{12}} & \dots & \frac{\partial L(W)}{\partial W_{1m}} \\ \frac{\partial L(W)}{\partial W_{21}} & \frac{\partial L(W)}{\partial W_{22}} & \dots & \frac{\partial L(W)}{\partial W_{2m}} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial L(W)}{\partial W_{n1}} & \frac{\partial L(W)}{\partial W_{n2}} & \dots & \frac{\partial L(W)}{\partial W_{nm}} \end{pmatrix}$$

- From an identical argument we used for vectors

$$L(W + \epsilon U) = L(W) + \epsilon \text{tr } U^T \frac{\partial L(W)}{\partial W} + O(\epsilon^2)$$

where

$$\text{tr } U^T G = \sum_i [U^T G]_{ii} = \sum_{ij} U_{ji} G_{ji} = \sum_{ij} U_{ij} G_{ij} = \langle U, G \rangle$$

Traces

- The trace of a matrix is the sum of its diagonal elements

$$\text{tr } A = \text{tr } A^T = \sum_i A_{ii}$$

- Clearly $\text{tr } cA = c \text{tr } A$

- Also $\text{tr } (A + B) = \text{tr } A + \text{tr } B$

- We note that


$$\text{tr } AB = \sum_{i,j} A_{ij} B_{ji} = \sum_{i,j} B_{ij} A_{ji} = \text{tr } BA$$

- It follows that

$$\text{tr } ABCD = \text{tr } DABC = \text{tr } CDAB = \text{tr } BCDA$$

Summary

- With care you can differentiate most expressions
- The chain and product rule are incredibly powerful tools
- We can generalise differentiation to vectors and matrices
- There are a number of surprisingly useful results see [The Matrix Cookbook](#)
- When we look at [integration](#) it gets harder

Taylor's Expansion

Convergence, remainder, high-dimensions, Hessian

Approximating Functions



- In much of machine learning we are trying to approximate functions
- There are many different ways to do this
- Depending on your application some are much better than others
- All function approximation makes some assumptions about the regularity of the function they are approximating
- Taylor expansions are polynomial approximations of a smooth function around a point

Taylor Expansion


- This gives the Taylor expansion

$$f(x) \approx f(x^*) + f'(x^*)(x - x^*) + \frac{1}{2}f''(x^*)(x - x^*)^2 + \dots + \frac{1}{n!}f^{(n)}(x^*)(x - x^*)^n$$

- The left and right-hand side both have the same n derivatives at the point x^*

**Taylor Expansion of $\tan(x)$** 

1. **1-D Taylor Expansion**
2. **High Dimensions**
3. **Hessians**

**Polynomial Approximation**

- Suppose we wanted to approximate a function $f(x)$ around a point x^* using a polynomial

$$f(x) \approx a_0 + a_1(x - x^*) + a_2(x - x^*)^2 + \dots + a_n(x - x^*)^n$$

- Taking derivatives

$$f'(x) \approx a_1 + 2a_2(x - x^*) + 3a_3(x - x^*)^2 + \dots + na_n(x - x^*)^{n-1}$$

$$f''(x) \approx 2a_2 + 6a_3(x - x^*) + \dots + n(n-1)a_n(x - x^*)^{n-2}$$

- One way to decide on the coefficients is to set $x = x^*$ and assume equality

$$a_0 = f(x^*) \quad a_1 = f'(x^*) \quad a_2 = \frac{1}{2}f''(x^*) \quad a_i = \frac{1}{i!}f^{(i)}(x^*)$$

Convergence

- As we add more and more terms our local approximation becomes more accurate, provided the higher derivatives exist
- Polynomials have the nice property that they only diverge at ∞
- However, polynomials cannot capture functions that diverge for finite numbers, e.g.

$$\frac{1}{x} \quad \frac{1}{(x-5)^2} \quad \tan(x)$$

- This leads to a loss of global convergence


Complex Functions

- To understand convergence of Taylor series more generally we need to go into the complex plane
- We can generalise many functions by considering them as a function with a complex argument $z = x + iy$
- The value of the complex function will be complex valued
- Complex functions are a huge and beautiful field (almost all of which go beyond this course)

Analytic Functions

- An important class of functions are those where for almost all z the derivative

$$\frac{df(z)}{dz} = \lim_{\epsilon \rightarrow 0} \frac{f(z + \epsilon) - f(z)}{\epsilon}$$



does not depend on the direction ϵ

- This includes polynomials, but also includes exponentials, trigonometric functions, logarithms, ratios of polynomials, "special functions", etc.

The Hyperbolic Tangent

- The hyperbolic tangent, $\tanh$, is related to the normal tangent function through $\tanh(x) = -i \tanh(ix)$
- The function (along the real line) doesn't have any singularities
- But we know that $\tan(x)$ has singularities at $\frac{\pi}{2} + k\pi$ where k is any integer, so $\tanh(x)$ has singularities at $x = \pm i\frac{\pi}{2}, \pm i\frac{3\pi}{2}, \dots$ on the imaginary axis
- Thus a Taylor expansion of $\tanh(x)$ around $x = 0$ is guaranteed to converge in a circle of radius $\pi/2$, but not necessarily for larger values of x

Remainder

- Taylor expansions actually come with an error bound


$$f(x) = \sum_{i=0}^{n-1} \frac{(x-x^*)^i}{i!} f^{(i)}(x^*) + \frac{(x-x^*)^n}{n!} f^{(n)}(\theta_x)$$

θ_x is a number in the interval between x^* and x

- The last term, known as the **remainder**, can often be very useful


$$e^x = 1 + x + \frac{x^2}{2} e^{\theta_x}$$

$$e^x = 1 + x + \frac{x^2}{2} + \frac{x^3}{3!} e^{\theta_x}$$




Outline

- 1-D Taylor Expansion
- High Dimensions
- Hessians




Guaranteed Convergence of Taylor Series

- For analytic functions, Taylor series are guaranteed to converge inside a circle that excludes singular points



- E.g. the function $f(z) = \frac{1}{z^2+1}$ has a singularity at $z = \pm i$. The Taylor series of $f(z)$ around 0 is guaranteed to converge inside a circle of radius 1

Taylor Expansion of $\tanh(x)$




Useful Bounds

- A much used bound in mathematical analysis is $e^x \geq 1 + x$
- Many other bounds can be obtained similarly, e.g. expanding $\log(x)$ around $x = 1$

$$\log(x) = (x-1) - \frac{(x-1)^2}{2} \frac{1}{\theta^2}$$

giving

$$\log(x) \leq x-1 \quad \text{or} \quad \log(1+x) \leq x$$



- The proof of the remainder term comes from applying multiple integrals to the term $f^{(n)}(x)$. For the sake of time I will skip this

High Dimensional Spaces

- Just as we defined a Taylor expansion by approximating the function by a polynomial we can do the same for a function $f: \mathbb{R}^n \rightarrow \mathbb{R}$
- We can determine the coefficients by differentiating both sides and setting $x = x^*$

$$f(x^*) = c^{(0)}$$

$$\frac{\partial f(x^*)}{\partial x_i} = c_i^{(1)}$$

$$\frac{\partial^2 f(x^*)}{\partial x_i \partial x_j} = 2c_{ij}^{(2)} \quad (\text{choosing } c_{ij}^{(2)} = c_{ji}^{(2)})$$

- We find

$$f(\mathbf{x}) \approx f(\mathbf{x}^*) + (\mathbf{x} - \mathbf{x}^*)^T \nabla f(\mathbf{x}^*) + \frac{1}{2}(\mathbf{x} - \mathbf{x}^*)^T \mathbf{H}(\mathbf{x} - \mathbf{x}^*) + \dots$$

- Where the matrix $\mathbf{H}$ is known as the Hessian and has components

$$H_{ij} = \frac{\partial^2 f(\mathbf{x}^*)}{\partial x_i \partial x_j}$$

- Because

$$\frac{\partial^2 f(\mathbf{x})}{\partial x_i \partial x_j} = \frac{\partial^2 f(\mathbf{x})}{\partial x_j \partial x_i}$$

the Hessian is symmetric ($\mathbf{H} = \mathbf{H}^T$)

- For smooth functions in n dimensions the equation $f(\mathbf{x}) = c$ defines an $(n - 1)$ -dimensional hyper-surface (because we have imposed one constraint)
- This can be thought of as the high dimensional analogue of a contour line
- The plane that goes through a point $\mathbf{x}^*$ that lines up with the contour surface is known as the tangent plane
- Away from $\mathbf{x}^*$ the value of the function $f(\mathbf{x})$ differs from the tangent plane by a quantity of order $\|\mathbf{x} - \mathbf{x}^*\|^2$


Defining the Tangent Plane

- Thus if $\mathbf{x}$ is in the tangent plane $T_{\mathbf{x}^*}$ then

$$f(\mathbf{x}) = f(\mathbf{x}^*) + O(\|\mathbf{x} - \mathbf{x}^*\|^2)$$

- Or $(\mathbf{x} - \mathbf{x}^*)^T \nabla f(\mathbf{x}^*) = 0$
- That is the tangent plane at a point $\mathbf{x}^*$ is the set of points orthogonal to the gradient
- Notice that the gradient then points in the direction orthogonal to the tangent plane of the contour surface

Picture of Tangent Plane for $f(\mathbf{x}) = \|\mathbf{x}\|^2 = c$



Direction of Steepest Ascent


- We note that as

$$f(\mathbf{x}^* + \epsilon \mathbf{u}) = f(\mathbf{x}^*) + \epsilon \mathbf{u}^T \nabla f(\mathbf{x}^*) + O(\epsilon^2)$$


- For a step $\mathbf{u}$ of fixed length, the function locally grows fastest when $\mathbf{u}$ is in the direction $\nabla f(\mathbf{x}^*)$
- Likewise when $\mathbf{u}$ is in the direction $-\nabla f(\mathbf{x}^*)$ the function decreases as fast as possible
- This is only true locally (when $\epsilon \gg \epsilon^2$)

Outline

1. 1-D Taylor Expansion
2. High Dimensions
3. Hessians



2-D Taylor Expansion



The Hessian

- The second order term in the Taylor expansion tells us about the curvature of the function
- This is very relevant when we are at a stationary point, $\mathbf{x}^*$, where $\nabla f(\mathbf{x}^*) = 0$ so that

$$f(\mathbf{x}) = f(\mathbf{x}^*) + \frac{1}{2}(\mathbf{x} - \mathbf{x}^*)^T \mathbf{H}(\mathbf{x} - \mathbf{x}^*) + O(\|\mathbf{x} - \mathbf{x}^*\|^3)$$

- If we make a perturbation $\mathbf{x} = \mathbf{x}^* + \epsilon \mathbf{u}$ then

$$f(\mathbf{x}^* + \epsilon \mathbf{u}) = f(\mathbf{x}^*) + \frac{1}{2} \epsilon^2 \mathbf{u}^T \mathbf{H} \mathbf{u} + O(\epsilon^3)$$

- For sufficiently small perturbation the sign of change in the function depends on the sign of $\mathbf{u}^T \mathbf{H} \mathbf{u}$

Eigenvalues

- We will revisit eigen-systems, but recall that for a square matrix an eigenvector, $\mathbf{v}_i$, of a matrix $\mathbf{H}$ is defined such that

$$\mathbf{H}\mathbf{v}_i = \lambda_i \mathbf{v}_i$$

- For an $n \times n$ symmetric matrix (such as the Hessian), the eigenvalues are all real, and we can choose n orthonormal eigenvectors such that $\mathbf{v}_i^T \mathbf{v}_j = \delta_{ij} \iff [i = j]$
- We note that


$$\mathbf{v}_i^T \mathbf{H} \mathbf{v}_i = \mathbf{v}_i^T \lambda_i \mathbf{v}_i = \lambda_i \mathbf{v}_i^T \mathbf{v}_i = \lambda_i$$

- If we perturb in a direction $\mathbf{v}_i$ then

$$f(\mathbf{x}^* + \epsilon \mathbf{v}_i) = f(\mathbf{x}^*) + \frac{1}{2} \epsilon^2 \lambda_i + O(\epsilon^3)$$

Minima, Maxima and Saddle-Points

- If all the eigenvalues are positive, $\lambda_i > 0$, then for any $\mathbf{u}$ we have $\mathbf{u}^T \mathbf{H} \mathbf{u} > 0$ and $\mathbf{x}^*$ would be a local minimum. We say $\mathbf{H}$ is **positive definite** or $\mathbf{H} \succ 0$
- Similarly if all the eigenvalues are less than zero ($\mathbf{H}$ is **negative definite**, $\mathbf{H} \prec 0$), then $\mathbf{x}^*$ will be a local maximum
- If there are both positive and negative eigenvalues then $\mathbf{x}^*$ is a saddle-point (if some eigenvalues are zero we need higher order terms to decide)



Gradient Descent

- We usually proceed iteratively, starting from some $\mathbf{w}^{(0)}$

$$\mathbf{w}^{(t+1)} = \mathbf{w}^{(t)} - r \nabla L(\mathbf{w}^{(t)})$$

- This moves us in the direction that (locally) minimises the loss function $L(\mathbf{w})$
- As we get closer to the minimum $\nabla L(\mathbf{w}^{(t)})$ typically becomes smaller and we take smaller steps
- However, many nasty things can happen, and we have to be careful in how we perform gradient descent
- We will revisit this when we look at optimisation

General Directions

- We can write any vector $\mathbf{u}$ as a weighted sum of the eigenvectors (the eigenvectors form a basis) $\mathbf{u} = \sum_{i=1}^n c_i \mathbf{v}_i$
- Note that

$$\begin{aligned} \mathbf{u}^T \mathbf{H} \mathbf{u} &= \left(\sum_{i=1}^n c_i \mathbf{v}_i^T \right) \mathbf{H} \left(\sum_{j=1}^n c_j \mathbf{v}_j \right) \\ &= \left(\sum_{i=1}^n c_i \mathbf{v}_i^T \right) \left(\sum_{j=1}^n c_j \lambda_j \mathbf{v}_j \right) \\ &= \sum_{i,j=1}^n \lambda_j c_i c_j \mathbf{v}_i^T \mathbf{v}_j = \sum_{i=1}^n \lambda_i c_i^2 \end{aligned}$$

using the orthogonality of eigenvectors $\mathbf{v}_i^T \mathbf{v}_j = \delta_{ij}$


Finding Optima

- We are often given a loss function, $L(\mathbf{w})$, which we wish to minimise
- A sufficient condition for $\mathbf{w}^*$ to be a local minimum is that $\nabla L(\mathbf{w}^*) = \mathbf{0}$ and $\mathbf{H} \succ 0$ where $\mathbf{H}$ is the Hessian of $L(\mathbf{w})$ at the point $\mathbf{w}^*$
- Except for a few very simple loss functions, there are likely to be many (and often a huge number) of local minima
- Unless $L(\mathbf{w})$ is very simple we are unlikely to be able to solve $\nabla L(\mathbf{w}) = \mathbf{0}$ analytically

Lessons

- For functions that are smooth (at least almost everywhere) Taylor series are really useful in understanding the local behaviour of the function
- The gradient defines the direction in which a function increases fastest
- This will be orthogonal to the tangent plane, where locally the function is a constant
- The Hessian (not to be confused with the Jacobian) is the matrix of second derivatives defining the curvature away from the tangent plane

Vector Spaces



Vectors, metric spaces, norms

Matrices, Vectors and All That

- The language of machine learning is mathematics
- Sometimes we draw pretty pictures to explain the mathematics
- Much of the mathematics we will use involves vectors, matrices and functions
- You need to master the language of mathematics, otherwise you won't understand the algorithms
- I'm going to spend this lecture and the next revising the mathematics you need to know (but I'm going to use a slightly posher language than you are probably used to)

Vectors

- We often work with objects with many components (features)
- To help handle this we will use vector notation

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix}$$

- We represent vectors by bold symbols
- All our vectors are column vectors by default
- We treat them as $n \times 1$ matrices

- We write row vectors as transposes of column vectors


$$\mathbf{y}^T = (y_1, y_2, \dots, y_n)$$

Vector Space

- A vector space, $\mathcal{V}$, is a set of vectors which satisfies
 1. if $\mathbf{v}, \mathbf{w} \in \mathcal{V}$ then $a\mathbf{v} \in \mathcal{V}$ and $\mathbf{v} + \mathbf{w} \in \mathcal{V}$ (closure)
 2. $\mathbf{v} + \mathbf{w} = \mathbf{w} + \mathbf{v}$ (commutativity of addition)
 3. $(\mathbf{u} + \mathbf{v}) + \mathbf{w} = \mathbf{u} + (\mathbf{v} + \mathbf{w})$ (associativity of addition)
 4. $\mathbf{v} + \mathbf{0} = \mathbf{v}$ (existence of additive identity $\mathbf{0}$)
 5. $\mathbf{v} + (-\mathbf{v}) = \mathbf{0}$ (existence of additive inverse $-\mathbf{v}$)
 6. $1\mathbf{v} = \mathbf{v}$ (existence of multiplicative identity 1)
 7. $a(b\mathbf{v}) = (ab)\mathbf{v}$ (associativity)
 8. $a(\mathbf{v} + \mathbf{w}) = a\mathbf{v} + a\mathbf{w}$ (distributive properties)
 9. $(a + b)\mathbf{v} = a\mathbf{v} + b\mathbf{v}$

(You don't need to remember these)
- Just from these properties we can deduce other properties

1. Vector Spaces
2. Metrics (distances)
3. Norms



Scalars (Fields)


- Vector spaces involve fields (numbers)—aka scalars
- These are quantities we can add together ($a + b$) and multiply together ($a \times b$)
- Formally they form an Abelian group under addition with an identity 0 and excluding 0 an Abelian group under multiplication and they are distributive

$$a \times (b + c) = a \times b + a \times c$$


- Although this sounds rather daunting don't panic. They behave like numbers. The field might be rational numbers, reals, complex numbers or something a bit more exotic (not the integers: 2 has no integer inverse)—but we will almost always consider reals

Basic Vector Operations

- The basic vector operations are adding




- multiplying by a scalar (a number)



$$\mathbb{R}^n$$

- When we first learn about vectors we think of them as arrows in 3-D space
- If we centre them all at the origin then there is a one-to-one correspondence between vectors and points in space
- We call this vector space $\mathbb{R}^3$
- Any set of quantities $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$ which satisfy the axioms above form a vector space $\mathbb{R}^n$
- Of course, we can't so easily draw pictures of high-dimensional vectors



- Any objects that satisfy the axioms of a vector space are **vectors** — not just $v \in \mathbb{R}^n$
- Matrices satisfy all the conditions of a vector space
- Infinite sequences form a vector space
- Functions form a vector space
 - Let $C(a,b)$ be the set of functions defined on the interval $[a,b]$
 - Note that if $f(x), g(x) \in C(a,b)$ then $\alpha f(x) \in C(a,b)$ and $f(x) + g(x) \in C(a,b)$
- Bounded vectors in $\mathbb{R}^n$ **don't** form a vector space

Metrics

- Vector spaces become more interesting if we have a notion of distance
- We say $d(x,y)$ is a **distance** or **metric** if
 - $d(x,y) \geq 0$ (non-negativity)
 - $d(x,y) = 0$ iff $x = y$ (identity of indiscernibles)
 - $d(x,y) = d(y,x)$ (symmetry)
 - $d(x,y) \leq d(x,z) + d(z,y)$ (triangular inequality)
- There are typically many possible distances (e.g. Euclidean distance, Manhattan distance, etc.)
- If one or more of the conditions aren't satisfied we have a **pseudo-metric**

Lipschitz Function


- One way to characterise a well-behaved function, $f(x)$, is if there exists a number $K < \infty$ such that for all x and y

$$d(f(x), f(y)) \leq K d(x, y)$$
- This is known as a **Lipschitz condition** and the function is said to be K -Lipschitz or Lipschitz continuous
- Note that such functions cannot have any jumps (i.e. they are continuous)
- The size of K measures the limit on the amplifying effect of the function

Contractive Mappings

- An interesting class of functions are those for which $K < 1$
- These are said to be **contractive mappings**
- A famous theorem that applies to contractive mappings is the Banach fixed-point theorem which says there exists a unique fixed point such that $f(x) = x$, and repeatedly applying f converges to it
- This is used for example in showing that various algorithms will converge

- Vector Spaces
- Metrics (distances)**
- Norms



Mappings and Functions


- A function defines a mapping from one vector space to another (although the spaces might be the same) e.g.

$$f: \mathbb{R} \rightarrow \mathbb{R}$$

(f maps the reals to the reals, i.e. $f(x)$ takes a real x and gives you a new real number $y = f(x)$)

- We are often interested in functions that behave nicely
- E.g. They are continuous


Lipschitz Continuous Functions



- The definition applies to any function defined in a metric space

Outline

- Vector Spaces
- Metrics (distances)**
- Norms



Norms

- Vector spaces are even more interesting with a notion of length
- **Norms** provide some measure of the size of a vector
- To formalise this we define the **norm** of an object v as $\|v\|$ satisfying
 1. $\|v\| > 0$ if $v \neq 0$ and $\|0\| = 0$ (non-negativity)
 2. $\|av\| = |a|\|v\|$ (“linearity”)
 3. $\|u + v\| \leq \|u\| + \|v\|$ (triangular inequality)
- When some criteria aren’t satisfied we have a **pseudo-norm**
- Norms provide a metric $d(x, y) = \|x - y\|$ (they are metric spaces)

Norms for Regularisation

- Why so many norms? Here is one place where the choice matters
- When we fit a model many different weight vectors w explain the training data equally well
- We usually prefer a “small” w , so we minimise

$$\text{loss}(w) + \eta \|w\|$$

- Equivalently we minimise the loss subject to $\|w\| \leq c$, confining w to a **ball** of whichever norm we chose
- Each norm means something different by “small”
- $\|w\|_2$ small: no one component is large, the weight is spread out
- $\|w\|_1$ small: a few large components are tolerated provided the rest are tiny
- $\|w\|_\infty$ small: only the largest component is controlled
- **Choosing a norm is choosing what you mean by “small”**

Compatible Norms

- A vector and matrix norm are said to be compatible if

$$\|Mv\|_b \leq \|M\|_a \times \|v\|_b$$

(Spectral and Euclidean norms are compatible)

- Norms provide quick ways to bound the maximum growth of a vector under a mapping induced by the matrix
- We will see that a measure of the sensitivity of a mapping is in terms of the ratio of its maximum effect to its minimum effect on a vector
- This is known as the **conditioning**, given by $\|M\| \times \|M^{-1}\|$

The Lipschitz Constant of a Network

- We can now put the Lipschitz condition and the matrix norm together
- Recall f is Lipschitz with constant K if

$$\|f(x) - f(y)\| \leq K \|x - y\|$$

- If we compose two functions the constants **multiply**

$$K_{f \circ g} \leq K_f K_g$$

- A layer $x \mapsto L_n x$ has $K = \|L_n\|_2$, the **spectral norm**—this is just compatibility, $\|Lv\| \leq \|L\|_2 \|v\|$
- ReLUs and max-pooling never stretch distances (though they can shrink them): $K = 1$
- A network is just a composition of layers, so

$$K_{\text{net}} \leq \prod_n \|L_n\|_2$$

Vector Norms

- The familiar vector norm is the (Euclidean) two norm

$$\|v\|_2 = \sqrt{v_1^2 + v_2^2 + \dots + v_n^2}$$

- Other norms exist, such as the p -norm ($p \geq 1$)

$$\|v\|_p = \left(\sum_{i=1}^n |v_i|^p \right)^{1/p}$$

- Special cases include the 1-norm and the infinite norm

$$\|v\|_1 = \sum_{i=1}^n |v_i| \quad \|v\|_\infty = \max_i |v_i|$$

- The 0-norm is a pseudo-norm as it does not satisfy condition 2

$$\|v\|_0 = \text{number of non-zero components}$$

Matrix Norms

- We can define norms for other objects
- The norm of a matrix encodes how large the mapping is
- The Frobenius norm is defined by

$$\|A\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |A_{ij}|^2}$$

- Many other norms exist including 1-norm, max-norm, etc.
- Some, but not all, norms satisfy the inequality

$$\|AB\| \leq \|A\| \times \|B\|$$

Why Should You Care?

- Deep learning involves multiplying the input (which we can think of as a vector x) by many layers
- In CNNs we have convolutional layers and dense layers
- The effect of applying these layers can be represented by a matrix multiplication $x_n = L_n x_{n-1}$
- We also do other things like applying ReLUs or pooling that change the magnitude, $\|x_n\|$, of our representation
- If you are developing new architectures you want $\|x_n\|$ neither to blow up nor vanish
- This can be controlled by carefully choosing $\|L_n\|$

Why Care About K ?

- K tells us how much the output can move when we nudge the input

$$\|f(x + \delta) - f(x)\| \leq K \|\delta\|$$

- **Adversarial examples** are tiny δ that change the classification: a network with small K cannot be fooled this easily
- It is the same quantity that controls the signal: $K \gg 1$ and signals (and gradients) explode, $K \ll 1$ and they vanish
- Wasserstein GANs require the critic to have $K \leq 1$. One way to enforce this is **spectral normalisation**: divide each L_n by $\|L_n\|_2$
- Bounds on how well a network generalises also depend on K
- But there is a trade-off: make K too small and the network is too smooth to fit the data

- Functions can also have norms, for example, if $f(x)$ is defined on some interval $\mathcal{I}$

$$\|f\|_{L^2} = \sqrt{\int_{x \in \mathcal{I}} f^2(x) dx}$$

- The L^2 vector space is the set of functions for which $\|f\|_{L^2} < \infty$
- The L^1 -norm is given by $\|f\|_{L^1} = \int_{x \in \mathcal{I}} |f(x)| dx$
- The infinite-norm is given by $\|f\|_{\infty} = \max_{x \in \mathcal{I}} |f(x)|$

Summary

- Vector spaces with a distance (metric spaces) and vector spaces with a norm (normed vector spaces) are interesting objects
- They allow you to define a topology (open/closed sets, etc.)
- You can build up ideas about connectedness, continuity, contractive maps, fixed-point theorems, . . .
- For the most part we are going to consider an even more powerful vector space that has an inner-product defined


- Given a norm, there exists an induced distance

$$d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|$$

- Any normed space has an associated metric space
- Thus we can define the Frobenius distance between matrices

$$d_F(\mathbf{A}, \mathbf{B}) = \|\mathbf{A} - \mathbf{B}\|_F$$

- It is easy to see that non-negativity, symmetry and the triangular inequality for distances are all inherited from the same properties of norms



Inner products, operators

Recap


- We have looked at vector spaces (closed sets where we can add elements and multiply them by a scalar)
- Recall that vector spaces don't just apply to normal vectors ($\mathbb{R}^n$), but to matrices, functions, sequences, random variables, . . .
- Proper distances or metrics, $d(\mathbf{x}, \mathbf{y})$, allow us to construct ideas about geometry of the vector space
- Norms, $\|\mathbf{x}\|$, allow us to reason about the size of a vector
- Norms induce a distance, $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|$

Axioms of Inner Products

- An inner product satisfies
 1. $\langle \mathbf{x}, \mathbf{x} \rangle \geq 0$ for all $\mathbf{x} \in \mathcal{V}$
 2. $\langle \mathbf{x}, \mathbf{x} \rangle = 0$ if and only if $\mathbf{x} = \mathbf{0}$
 3. $\langle \alpha \mathbf{x}, \mathbf{y} \rangle = \alpha \langle \mathbf{x}, \mathbf{y} \rangle$
 4. $\langle \mathbf{x}, \mathbf{y} + \mathbf{z} \rangle = \langle \mathbf{x}, \mathbf{y} \rangle + \langle \mathbf{x}, \mathbf{z} \rangle$
 5. $\langle \mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{y}, \mathbf{x} \rangle$
- We can show that $\|\mathbf{x}\| = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$ satisfies the axioms of a norm, so that an inner-product space is a normed space
- The norm associated with the inner-product for vectors in $\mathbb{R}^n$ (i.e. $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^\top \mathbf{y}$) is the Euclidean norm $\|\mathbf{x}\| = \sqrt{\mathbf{x}^\top \mathbf{x}}$

Angles Between Vectors


- A natural interpretation of the inner product is in providing a measure of the angle between vectors



- Vectors are orthogonal if $\langle \mathbf{x}, \mathbf{y} \rangle = 0$
- We can extend this idea to functions

$$\langle f(x), g(x) \rangle = \int_{x \in \mathcal{I}} f(x)g(x)dx = \|f(x)\| \|g(x)\| \cos(\theta)$$
- Note that $\sin(x)$ and $\cos(x)$ are orthogonal in the interval $[0, 2\pi]$

1. Inner Products
2. Operators



Inner Products

- We will often consider objects with an *inner product*, $\langle, \rangle$
- For vectors in $\mathbb{R}^n$

$$\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^\top \mathbf{y} = \sum_{i=1}^n x_i y_i$$

- For functions

$$\langle f, g \rangle = \int_{x \in \mathcal{I}} f(x)g(x)dx$$

- For $m \times n$ matrices

$$\langle \mathbf{A}, \mathbf{B} \rangle = \text{Tr} \mathbf{A}^\top \mathbf{B} = \sum_{i=1}^m \sum_{j=1}^n A_{ij} B_{ij}$$

Cauchy-Schwarz Inequality

- One of the most important results of inner-product spaces, known as the **Cauchy-Schwarz inequality** is that

$$\langle \mathbf{x}, \mathbf{y} \rangle^2 \leq \langle \mathbf{x}, \mathbf{x} \rangle \langle \mathbf{y}, \mathbf{y} \rangle = \|\mathbf{x}\|^2 \|\mathbf{y}\|^2$$

- Or

$$|\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\| \|\mathbf{y}\|$$

- This is a very general result so for example

$$\left| \int f(x)g(x)dx \right| \leq \sqrt{\left(\int f^2(x)dx \right) \left(\int g^2(x)dx \right)}$$

Basis Functions

- Any set of vectors $\{\mathbf{b}_i | i = 1, \dots\}$ that span the space can be used as a basis or coordinate system
- The simplest and most useful case is when the vectors are orthogonal and normalised (i.e. $\|\mathbf{b}_i\| = 1$)

- In $\mathbb{R}^3$ we could use $\mathbf{b}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$, $\mathbf{b}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$, $\mathbf{b}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$


- This is not unique as we can rotate our basis vectors

- For an orthonormal basis we can write any vector as

$$\hat{\mathbf{x}} = \begin{pmatrix} \mathbf{x}^\top \mathbf{b}_1 \\ \mathbf{x}^\top \mathbf{b}_2 \\ \mathbf{x}^\top \mathbf{b}_3 \end{pmatrix}$$


Orthogonal Functions

- For functions we can use any ortho-normal set of functions as a basis
- The most familiar are the Fourier functions $\sin(n\theta)$ and $\cos(n\theta)$
- Any function in $C(0, 2\pi)$ can be represented by a point $f = \begin{pmatrix} \langle f(x), b_0(x) \rangle \\ \langle f(x), b_1(x) \rangle \\ \vdots \end{pmatrix}$
- There might be an infinite number of components
- This is analogous to points in $\mathbb{R}^n$ (for large n)



Projection Operator

- Given a vector a we can project it onto a vector b




- Using $\langle a, b \rangle = \|a\| \|b\| \cos(\theta)$ and $\langle b, b \rangle = \|b\|^2$
- A defining feature of a projection operator is

$$\begin{aligned} \text{proj}_b(\text{proj}_b(a)) &= \text{proj}_b\left(\frac{\langle a, b \rangle}{\langle b, b \rangle} b\right) \\ &= \left\langle \frac{\langle a, b \rangle}{\langle b, b \rangle} b, b \right\rangle \frac{b}{\langle b, b \rangle} = \frac{\langle a, b \rangle}{\langle b, b \rangle} b = \text{proj}_b(a) \end{aligned}$$

Constructing an Orthonormal Basis

- Given a set of linearly independent vectors $(a_1, a_2, \dots, a_n)$ we can construct an orthonormal basis



- We can iterate this

$$\begin{aligned} u_1 &= a_1 & e_1 &= \frac{u_1}{\|u_1\|} \\ u_2 &= a_2 - \text{proj}_{u_1}(a_2) & e_2 &= \frac{u_2}{\|u_2\|} \\ u_3 &= a_3 - \text{proj}_{u_1}(a_3) - \text{proj}_{u_2}(a_3) & e_3 &= \frac{u_3}{\|u_3\|} \end{aligned}$$

Gram-Schmidt Orthogonalisation


- The algorithm in the last slide is known as **Gram-Schmidt**
- We can perform it on any inner-product space
- For vectors in $\mathbb{R}^n$ it is prone to accumulate rounding errors and is not the method of choice
- Numerical packages use "Householder transformations" to perform orthogonalisation
- Nevertheless, it is nice to know that we can form an orthonormal basis

Algebraic Structure

- We have gone to these lengths as we want to show that many properties of vectors are shared by other objects (matrices, functions, etc.)
- The notions of distance (geometry), norms (size of vectors) and inner products (angles between vectors) provides a very rich set of concepts
- Vectors form the backbone of objects we will use repeatedly in machine learning
- The next piece of the jigsaw is to understand how we can transform these objects

Outline

- Inner Products
- Operators



Operators

- In machine learning we are interested in transforming our input vectors into some output predictions
- To accomplish this we will apply some mapping or operators on the vector $\mathcal{T} : \mathcal{V} \rightarrow \mathcal{V}'$
- This says that $\mathcal{T}$ maps some object $x \in \mathcal{V}$ to an object $y = \mathcal{T}[x]$ in a new vector space $\mathcal{V}'$
- This new vector space may or may not be the same as the original vector space
- Our objects may be any object in a vector space such as a function

Linear Operators

- Operators are in general very complicated, but a particularly nice set of operators are linear operators
- $\mathcal{T}$ is a linear operator if
 - $\mathcal{T}[ax] = a\mathcal{T}[x]$
 - $\mathcal{T}[x + y] = \mathcal{T}[x] + \mathcal{T}[y]$
- For normal vectors ($x \in \mathbb{R}^n$) the most general linear operation is

$$\mathcal{T}[x] = Mx$$

where M is a matrix

Matrix multiplication

- For an $\ell \times m$ matrix $\mathbf{A}$ and an $m \times n$ matrix $\mathbf{B}$ we can compute the $(\ell \times n)$ product, $\mathbf{C} = \mathbf{AB}$, such that

$$C_{ij} = \sum_{k=1}^m A_{ik} B_{kj} \quad \left(\begin{pmatrix} \equiv \end{pmatrix} \right) \left(\begin{pmatrix} \parallel \end{pmatrix} \right) = \left(\begin{pmatrix} \boxtimes \end{pmatrix} \right)$$

- Treating the vector $\mathbf{x}$ as an $n \times 1$ matrix then

$$\mathbf{y} = \mathbf{Ax} \Rightarrow y_i = \sum_j A_{ij} x_j \quad \left(\begin{pmatrix} \equiv \end{pmatrix} \right) \left(\begin{pmatrix} \parallel \end{pmatrix} \right) = \left(\begin{pmatrix} \parallel \end{pmatrix} \right)$$


- Using the same matrix notation we can define the inner product as

$$\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^\top \mathbf{y} = \sum_{i=1}^n x_i y_i \quad \left(\begin{pmatrix} \longleftarrow \end{pmatrix} \right) \left(\begin{pmatrix} \parallel \end{pmatrix} \right) = \left(\begin{pmatrix} \blacksquare \end{pmatrix} \right)$$


Non-commutativity

- In general $\mathbf{AB} \neq \mathbf{BA}$

$$\begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 0 & -1 & 0 \\ 0 & 0 & -1 \\ 1 & 0 & 0 \end{pmatrix}$$



Associativity of Mappings



- For all $\mathbf{x}$ we have $\mathbf{A(BC)x} = (\mathbf{AB})\mathbf{Cx}$
- This implies $\mathbf{A(BC)} = (\mathbf{AB})\mathbf{C}$

Kernels

- The equivalent of a matrix for functions (i.e. a linear operator) is known as a kernel $K(x, y)$

$$g(x) = \mathcal{T}[f] = \int_{y \in \mathcal{I}} K(x, y) f(y) dy$$

- Our domain does not need to be one dimensional, e.g.

$$g(x) = \mathcal{T}[f] = \int_{y \in \mathcal{I}} K(x, y) f(y) dy$$

- We shall soon see examples of high-dimensional kernels

Kernels in Machine Learning

- Kernels are used heavily in machine learning
- In kernel methods such as SVM, SVR, Kernel-PCA
- They are also used in Gaussian Processes
- In all these cases we consider symmetric, positive semi-definite kernels
- Sometimes they can be interpreted as covariance between random functions


$$K(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{f \sim \mathcal{P}}[(f(\mathbf{x}) - \mu(\mathbf{x}))(f(\mathbf{y}) - \mu(\mathbf{y}))]$$

Square Matrices

- We will spend a lot of time on operators that map from a vector space onto itself $\mathcal{T}: \mathcal{V} \rightarrow \mathcal{V}$
- For vectors in $\mathbb{R}^n$ such linear operators are represented by square matrices
- When there is a one-to-one mapping then we have a unique inverse
- We will study such mappings in detail in the next lecture

General Linear Mappings


- In general a linear operator will map vectors between different vector spaces
- E.g. $\mathbb{R}^3 \rightarrow \mathbb{R}^2$



Summary


- Mathematics is the language of machine learning and you have to get used to it
- Mathematics is like programming, if you don't understand the syntax and you can't write it down then it's meaningless
- We've taken a high level view of inner product spaces and operators, this will pay us back later as we look at kernel methods

Understand Mappings




Least Squares, Under-constrained Systems, Ill-conditioning

1. Least Squares
2. Under-constrained Systems
3. Ill-conditioning



Transforming Data

- In the last lecture we spent time developing a sophisticated view of vector spaces and operators
- At a mathematical level machine learning can be viewed as performing an inverse mapping



- Although our mappings are not necessarily linear in either direction we learn a lot by understanding linear operators

Inverse Problems


- Given m observations $\{(x_k, y_k) | k = 1, \dots, m\}$, find the p unknowns $w = (w_1, w_2, \dots, w_p)^T$ such that $x_k^T w = y_k$
- Define the *design matrix* as the matrix of feature vectors

$$X = \begin{pmatrix} x_1^T \\ x_2^T \\ \vdots \\ x_m^T \end{pmatrix} = \begin{pmatrix} x_{11} & x_{12} & \cdots & x_{1p} \\ x_{21} & x_{22} & \cdots & x_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ x_{m1} & x_{m2} & \cdots & x_{mp} \end{pmatrix}$$

- and the target vector $y = (y_1, y_2, \dots, y_m)^T$
- Then $y = Xw$ and if $m = p$ (and X is invertible) $w = X^{-1}y$

Linear Regression

- $x_k^T w$ depends on distance from the plane orthogonal to w



- If $m > p$ then X isn't square so doesn't have an inverse
- Worse, unless the data is exact, $y \approx Xw \Rightarrow$ no exact "solution"
- Problem solved by Gauss to predict the orbit of the asteroid Ceres

Linear Least Squares

- The error of input pattern x_k is

$$\epsilon_k = x_k^T w - y_k$$

- We can define the error vector

$$\epsilon = Xw - y$$

(note that $\epsilon_k = x_k^T w - y_k$)


- The squared loss

$$L(w|\mathcal{D}) = \sum_{k=1}^m (x_k^T w - y_k)^2 = \sum_{k=1}^m \epsilon_k^2 = \|\epsilon\|^2$$

- Minimising this loss is known as the least squares problem

Finding a Minimum

- At a minimum of a one-dimensional function, $f(x)$, we have $f'(x) = 0$



- At a minimum of an n -dimensional function $f(x)$ we have the set of equations

$$\frac{\partial f(x)}{\partial x_i} = 0 \quad \forall i = 1, \dots, n$$

Least Squares Solution

- The least squares solution is given by

$$\begin{aligned} \nabla L(w|\mathcal{D}) &= \nabla \|\epsilon\|^2 = \nabla \|Xw - y\|^2 \\ &= \nabla (w^T X^T X w - 2w^T X^T y + y^T y) \\ &= 2(X^T X w - X^T y) = 0 \end{aligned}$$

- Or

$$w = (X^T X)^{-1} X^T y = X^+ y$$

- $X^+ = (X^T X)^{-1} X^T$ is known as the pseudo-inverse

- Note that $\|a\|^2 = a^T a = \sum_i a_i^2$

$$\begin{aligned}\|Xw - y\|^2 &= (Xw - y)^T (Xw - y) = (w^T X^T - y^T)(Xw - y) \\ &= w^T X^T X w - 2w^T X^T y + y^T y\end{aligned}$$


- Where we have used $w^T X^T y = y^T X w$, $\sum_{i,j} w_i X_{ji} y_j = \sum_{i,j} y_i X_{ij} w_j$
- Also $\nabla w^T M w = M w + M^T w = (M + M^T) w$
- If $M = M^T$ (i.e. M is symmetric) then $\nabla w^T M w = 2M w$
- As $(AB)^T = B^T A^T$ then $(X^T X)^T = X^T X^{TT} = X^T X$ and $X^T X$ is symmetric

Solving Inverse Problems

- Gauss showed us how to solve **over-constrained** problems (we have more observations than parameters)
- We seek a solution which isn't necessarily exact but minimises an error
- But, what if we have more parameters than observations?
- That is, we are **under-constrained**
- Note that in some directions you might be over-constrained and in other directions under-constrained
- This is very typical of most machine learning problems

What is the Inverse?

- Many points can map to the same point




Handling Under-constrained Systems

- The issue with under-constrained systems is not that there is no solution
- There is a whole family of solutions, so which one should we pick
- William of Ockham in the 14th century postulated that the "simplest solution" is likely to be the correct solution
- How we determine the simplest solution is a bit of a black art
- However, for linear regression, adding an L_2 regulariser makes the solution unique


$$L_{full} = L_{fit}(w) + \eta \|w\|_2^2$$

- Least Squares
- Under-constrained Systems**
- Ill-conditioning**



Under-constrained Systems

- If we have fewer data points than parameters then there will be multiple solutions, e.g. fitting $y = w_0 + w_1 x + \dots + w_5 x^5$ to 4 points




- Every curve fits the data exactly, so which one should we choose?

Under-constrained Systems

- The system is **under-constrained**
- We have more unknowns than equations
- The inverse is not unique
- Solving the inverse problem ($w = (X^T X)^{-1} X^T y$) is said to be **ill-posed**
- The inverse $(X^T X)^{-1}$ doesn't exist
- If we have a complicated learning machine and not sufficient data we often end up with an ill-posed inverse problem (there are lots of sets of parameters that explain the data)

Outline

- Least Squares
- Under-constrained Systems**
- Ill-conditioning**



III-Conditioning

- Singular matrices are rare (although they occur when we don't have enough data), but matrices that are close to being singular are common
- If a matrix is close to singular it is ill-conditioned
- Ill-conditioned matrices have some eigenvalues that are tiny compared with the largest
- All points get squashed towards a lower-dimensional subspace (a plane in 3-D)
- Large matrices are very often ill-conditioned

A Worked Example

- Two data points and two features that are almost the same

$$\mathbf{X} = \begin{pmatrix} 1 & 1 \\ 1 & 1.001 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} 2 \\ 2.001 \end{pmatrix} \Rightarrow \mathbf{w} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

- Change y_2 by 0.001 (i.e. 0.05%)

$$\mathbf{y} = \begin{pmatrix} 2 \\ 2.002 \end{pmatrix} \Rightarrow \mathbf{w} = \begin{pmatrix} 0 \\ 2 \end{pmatrix}$$

- The prediction for a new input $\mathbf{x} = (1,0)^T$ jumps from 1 to 0
- The eigenvalues of $\mathbf{X}$ are about 2 and 0.0005
- $\mathbf{w}$ moved along $(-1,1)^T$, the direction with the tiny eigenvalue, which the data hardly constrains

Where Does Ill-Conditioning Come From?


- Too little data: about as many data points as parameters
- Nearly collinear features, e.g. floor area and number of rooms, or height measured in both cm and inches
 - ★ The columns of $\mathbf{X}$ are nearly parallel so the data cannot tell us how to share the weight between them
- Polynomial features $x, x^2, \dots, x^9$ for $x \in [0,1]$ all look alike: the eigenvalues of $\mathbf{X}^T\mathbf{X}$ span about 13 orders of magnitude
- Features on very different scales (income in pounds and age in years)—easily fixed by standardising the features

Three Regimes

	$m > p$ over-constrained	$m = p$	$m < p$ under-constrained
Exact fit?	usually not	yes	yes, many
Unique?	yes (usually)	yes (usually)	no
Method	least squares	$\mathbf{w} = \mathbf{X}^{-1}\mathbf{y}$	pick "simplest"

- Real problems mix these: over-constrained in some directions and (nearly) under-constrained in others
- The nearly under-constrained directions make the problem ill-conditioned

III-Conditioned Matrices



III-Conditioning in ML


- Ill-conditioning in machine learning occurs when a very small change in the learning data causes a large change in the predictions of the learning machine
- In linear regression the matrix $\mathbf{X}^T\mathbf{X}$ is typically ill-conditioned when we have roughly as many data points as parameters
- Much of machine learning is concerned with making learning machines better conditioned
- Adding regularisers is one approach to achieve this

Least Squares in Practice

- $\mathbf{w} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{y}$ is for understanding, not for computing
- Never compute $(\mathbf{X}^T\mathbf{X})^{-1}$; use a least squares solver such as `np.linalg.lstsq(X, y)`, which works with $\mathbf{X}$ directly
- Forming $\mathbf{X}^T\mathbf{X}$ squares the ratio of largest to smallest eigenvalues: 10^8 becomes 10^{16} , beyond the ~ 16 significant figures of a double
- Standardise features (zero mean, unit variance) before fitting
- When $m < p$, `lstsq` returns the solution with the smallest $\|\mathbf{w}\|$ —one version of the "simplest" solution (we will see why when we meet the SVD)

Summary

- Linear mappings are commonly used in machine learning algorithms such as regression
- Systems can be over-constrained, under-constrained or ill-conditioned—often all at once
- We will often meet the pseudo-inverse involving inverting $\mathbf{X}^T\mathbf{X}$
- Inverting an ill-conditioned $\mathbf{X}^T\mathbf{X}$ amplifies noise in the data


Eigensystems

Eigenvectors, Orthogonal Matrices, Eigenvector Decomposition, Rank

Eigenvector equation

- Eigen-systems help us to understand mappings
- A non-zero vector v is said to be an **eigenvector** if

$$Mv = \lambda v$$



- M is square (i.e. $n \times n$)
- Where the number λ is the **eigenvalue**
- Eigenvalues play a fundamental role in understanding operators

Proof of Orthogonality

- $(Mv_i = \lambda_i v_i)^T$ implies $v_i^T M^T = \lambda_i v_i^T$
- When M is symmetric then $Mv_i = \lambda_i v_i \Rightarrow v_i^T M = \lambda_i v_i^T$
- Consider two eigenvectors v_i and v_j of M

$$\begin{aligned} v_i^T M v_j &= (v_i^T M) v_j = \lambda_i v_i^T v_j \\ &= v_i^T (M v_j) = \lambda_j v_i^T v_j \end{aligned}$$

- So either $\lambda_i = \lambda_j$ or $v_i^T v_j = 0$
- If $\lambda_i = \lambda_j$ then any linear combination of v_i and v_j is an eigenvector ($M(av_i + bv_j) = \lambda_i(av_i + bv_j)$). So I can choose two eigenvectors that are orthogonal to each other.

Orthogonal Matrices


- We can construct an **orthogonal** matrix V from the normalised eigenvectors ($v_i^T v_i = 1$)

$$V = (v_1, v_2, \dots, v_n)$$

- Matrix V is an $n \times n$ matrix
- Because the vectors v_i are orthogonal and normalised

$$V^T V = \begin{pmatrix} v_1^T v_1 & v_1^T v_2 & \dots & v_1^T v_n \\ v_2^T v_1 & v_2^T v_2 & \dots & v_2^T v_n \\ \vdots & \vdots & \ddots & \vdots \\ v_n^T v_1 & v_n^T v_2 & \dots & v_n^T v_n \end{pmatrix} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix} = I$$

- Eigenvectors**
- Orthogonal Matrices
- Eigen Decomposition
- Low Rank Approximation
- Non-Normal Matrices



Symmetric Matrices

- If M is an $n \times n$ **symmetric** matrix then it has n real orthogonal eigenvectors with real eigenvalues
- We denote the i^{th} eigenvector by v_i and the corresponding eigenvalue by λ_i so that

$$Mv_i = \lambda_i v_i$$


- Orthogonal means that if $i \neq j$ then

$$v_i^T v_j = 0$$

- (We can always normalise eigenvectors if we want)

Outline

- Eigenvectors**
- Orthogonal Matrices**
- Eigen Decomposition
- Low Rank Approximation
- Non-Normal Matrices



The Other Way Around

- We have shown that $V^T V = I$
- Thus multiply both sides on the left by V

$$V V^T V = V$$

- V will have an inverse (we show this on the next slide), V^{-1} , such that $V V^{-1} = I$

- Multiplying the equation on the right by V^{-1}

$$\begin{aligned} (V V^T) V V^{-1} &= V V^{-1} \\ V V^T &= I \end{aligned}$$

- Note that, $V^{-1} = V^T$ (definition of orthogonal matrix)

Invertible Matrices

- A matrix, M , will be singular (not invertible) if there exists a vector $x \neq 0$ such that

$$Mx = 0$$

- Now if there exists such a vector such that $Vx = 0$ then multiplying by V^T we get

$$V^T Vx = V^T 0 \\ x = 0$$

since $V^T V = I$

- Thus V is invertible

Rotations

- Orthogonal matrices satisfy $V^T V = V V^T = I$
- As a consequence they define rotations (and possibly a reflection)
- Consider a vector x and $x' = Vx$, now

$$\|x'\|_2^2 = x'^T x' = (Vx)^T (Vx) = x^T V^T V x = x^T x = \|x\|_2^2$$


- Similarly if additionally $y' = Vy$ then

$$\langle x', y' \rangle = (Vx)^T (Vy) = x^T V^T V y = x^T y = \langle x, y \rangle = \|x\|_2 \|y\|_2 \cos(\theta)$$

- Rotations and reflections preserve lengths and angles

Outline

- Eigenvectors
- Orthogonal Matrices
- Eigen Decomposition**
- Low Rank Approximation
- Non-Normal Matrices



Matrix Decomposition

- Taking the matrix of eigenvectors, V , then

$$MV = M(v_1, v_2, \dots, v_n) = (\lambda_1 v_1, \lambda_2 v_2, \dots, \lambda_n v_n) = V\Lambda$$

- where $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n) = \begin{pmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{pmatrix}$


- Now

$$M = MVV^T = V\Lambda V^T$$

- Very important *similarity transform*

Mappings by Symmetric Matrices

$$M = \begin{pmatrix} 0.83 & -0.31 \\ -0.31 & 1.9 \end{pmatrix} = V\Lambda V^T = \begin{pmatrix} \cos(-75^\circ) & -\sin(-75^\circ) \\ \sin(-75^\circ) & \cos(-75^\circ) \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 0.75 \end{pmatrix} \begin{pmatrix} \cos(75^\circ) & -\sin(75^\circ) \\ \sin(75^\circ) & \cos(75^\circ) \end{pmatrix}$$



Inverses

- For any symmetric invertible matrix

$$M = V\Lambda V^T \quad M^{-1} = V\Lambda^{-1} V^T$$


- Where $\Lambda^{-1} = \text{diag}(\frac{1}{\lambda_1}, \frac{1}{\lambda_2}, \dots, \frac{1}{\lambda_n}) = \begin{pmatrix} \frac{1}{\lambda_1} & 0 & \dots & 0 \\ 0 & \frac{1}{\lambda_2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \frac{1}{\lambda_n} \end{pmatrix}$

- Since

$$MM^{-1} = (V\Lambda V^T)(V\Lambda^{-1} V^T) = V\Lambda(V^T V)\Lambda^{-1} V^T = V\Lambda\Lambda^{-1} V^T = VV^T = I$$

- I.e. small eigenvalues become large eigenvalues and vice versa

III-Conditioning Again




Condition Number

- Taking matrix inverses can be inherently unstable
- Any small error can be amplified by taking the inverse
- The stability of the inverse depends on the ratio of smallest eigenvalue to the largest eigenvalue (i.e. the biggest possible amplification compared to the smallest)
- For symmetric matrices the largest $|\lambda|$ is the spectral norm, $\|M\|_2$, from the vector spaces lecture
- The condition number is given by

$$\|M\|_2 \times \|M^{-1}\|_2 = \frac{|\lambda_{\max}|}{|\lambda_{\min}|}$$

- Large condition number implies very ill-conditioned


1. Eigenvectors
2. Orthogonal Matrices
3. Eigen Decomposition
4. Low Rank Approximation
5. Non-Normal Matrices



“Inverting” Rank Deficient Matrices

- Rank deficient matrices are non-invertible (i.e. we don't know the vector x such that $Mx = b$) as we don't know the component of the x in the null space
- Although we don't know x we can find a vector, x , that satisfies $Mx = b$ (provided one exists)
- Given a symmetric $n \times n$ matrix with k non-zero eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_k$ we can construct a “pseudo-inverse” M^+ as $V\Lambda^+V^T$ where $\Lambda^+ = \text{diag}(\frac{1}{\lambda_1}, \frac{1}{\lambda_2}, \dots, \frac{1}{\lambda_k}, 0, \dots, 0)$
- This finds the vector x with no component in the null space (it is the solution with the smallest norm)
- This and $(X^TX)^{-1}X^T$ are special cases of the pseudo-inverse we meet with SVD

1. Eigenvectors
2. Orthogonal Matrices
3. Eigen Decomposition
4. Low Rank Approximation
5. Non-Normal Matrices



Asymmetric Matrices

- For $n \times n$ real symmetric matrices we are guaranteed there are n orthogonal real eigenvectors
- For non-symmetric matrices the eigenvectors will generally not be orthogonal, they won't necessarily be real
- We are not even guaranteed there will be n linearly independent eigenvectors
- Matrices which don't have n linearly independent eigenvectors are said to be **defective**, e.g. $\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ has only one
- Randomly generated matrices will typically not be defective, but nevertheless they can arise in certain situations

- The rank of a symmetric matrix, M , is the number of non-zero eigenvalues
- The space spanned by the eigenvectors v_a, v_b , etc. with zero eigenvalue forms a **null space**
- Any vector in the null space will get mapped to the zero vector

$$M(av_a + bv_b + \dots) = 0$$

- A square matrix is said to be **rank deficient** if it has any eigenvectors with eigenvalue equal to 0
- This happens when the columns of the matrix are not linearly independent

Low Rank Approximation

- Recall that matrices with large and small eigenvalues are ill-conditioned so the inverse has the potential to greatly amplify any measurement error
- One work around is to set all small eigenvalues to zero and use the pseudo-inverse
- Setting small eigenvalues to zero reduces the rank of the matrix and is an example of a low rank approximation
- Low rank approximations are much used to obtain approximate models for arrays of data (we will revisit this when we look at SVD)

Non-Normal Matrices

- Symmetric matrices are part of a class of matrices known as **normal matrices**, which satisfy $MM^\dagger = M^\dagger M$ (for real matrices $M^\dagger = M^T$)
- Hermitian matrices, H , are complex matrices that equal their conjugate transpose: $H = (H^T)^* = H^\dagger$
- Hermitian matrices are also normal; they are important in quantum theory, less so in machine learning
- Normal matrices are exactly those with a full set of orthonormal (possibly complex) eigenvectors
- Most non-symmetric real matrices are not normal (rotations are an exception); their eigenvectors are more confusing

Characteristic Equation

- Since our eigenvectors are defined by $Mv = \lambda v$ then

$$(M - \lambda I)v = 0 \text{ has a solution } v \neq 0 \text{ only if } \det(M - \lambda I) = 0$$

- This is known as the characteristic equation and leads to an n^{th} order polynomial in λ
- An n^{th} order polynomial has n roots (counting repeats), which are in general complex
- Thus even for real (non-symmetric) matrices we can get complex eigenvalues

Complex Conjugate Pairs

- If $\mathbf{M}$ is real and λ is complex then $\mathbf{v}$ must be complex too, otherwise $\mathbf{M}\mathbf{v}$ would be real but $\lambda\mathbf{v}$ would not
- Furthermore

$$\mathbf{M}\mathbf{v} = \lambda\mathbf{v} \implies \mathbf{M}\mathbf{v}^* = \lambda^*\mathbf{v}^*$$

- So that eigenvectors and their corresponding eigenvalues exist in complex conjugate pairs

Eigen-Decomposition

- For a non-symmetric matrix, $\mathbf{M}$, provided it is not defective, we can decompose the matrix as

$$\mathbf{M} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$$

- $\mathbf{\Lambda}$ is a diagonal matrix of eigenvalues
- $\mathbf{V}$ is a matrix whose columns are the right-hand eigenvectors—in general it is not an orthogonal matrix
- $\mathbf{V}^{-1}$ is a matrix whose rows are the left-hand eigenvectors (scaled so that $\mathbf{u}_i^T \mathbf{v}_i = 1$)

Left and Right Eigenvectors

- For symmetric matrices if $\mathbf{v}$ is a right-hand eigenvector in the sense that $\mathbf{M}\mathbf{v} = \lambda\mathbf{v}$ then $\mathbf{v}^T$ is a left-hand eigenvector

$$\mathbf{v}^T \mathbf{M} = \lambda \mathbf{v}^T$$

- Earlier we used this property to show that the eigenvectors for symmetric matrices are orthogonal

- For asymmetric real matrices the left-hand eigenvectors $\mathbf{u}^T$ are different from the right-hand eigenvectors, although they share the same eigenvalues (they have the same characteristic equation)

$$\mathbf{M}\mathbf{v} = \lambda\mathbf{v} \quad \mathbf{u}^T \mathbf{M} = \lambda \mathbf{u}^T$$


- As before, $\mathbf{u}_i^T \mathbf{v}_j = 0$ if $\lambda_i \neq \lambda_j$, but in general $\mathbf{v}_i^T \mathbf{v}_j \neq 0$

Summary

- Linear mappings are commonly used in machine learning algorithms such as regression
- We can understand symmetric operators by looking at their eigenvectors
- Any symmetric matrix can be decomposed as $\mathbf{M} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$
 - ★ where $\mathbf{V}$ is an orthogonal matrix whose columns are the eigenvectors
 - ★ and $\mathbf{\Lambda}$ is a diagonal matrix of the eigenvalues
- This decomposition allows us to understand inverse mappings
- Non-symmetric matrices are more complex—we use the eigen-decomposition less because of this

Principal Component Analysis (PCA)


1.6 -1.1 -1.6 2.1 -0.52 2.8 0.72 0.7 -0.68 -0.41 -1.4 -1.5 -0.54 -0.62 1.3 -1.4 -0.27 0.74 0.77 -1



Covariance matrices, dimensionality reduction, PCA, Duality

Spread of Data


- Often data varies significantly in only some directions



- Reduce dimensions by projecting onto low dimensional subspace with maximum variation

Dimensionality Reduction

- Often helpful to consider only directions where data varies significantly
- Want to find directions along which data has its greatest variation

**Direction of Maximum Variation**


- Expanding the Lagrangian

$$\begin{aligned}\mathcal{L} &= \frac{1}{m-1} \sum_{k=1}^m (v^T(x_k - \mu))^2 - \lambda(\|v\|^2 - 1) \\ &= \frac{1}{m-1} \sum_{k=1}^m (v^T(x_k - \mu)(x_k - \mu)^T v) - \lambda(\|v\|^2 - 1) \\ &= v^T \left(\frac{1}{m-1} \sum_{k=1}^m (x_k - \mu)(x_k - \mu)^T \right) v - \lambda(\|v\|^2 - 1) \\ &= v^T C v - \lambda(v^T v - 1)\end{aligned}$$


- Extrema of the Lagrangian

$$\nabla \mathcal{L} = 2(Cv - \lambda v) = 0 \Rightarrow Cv = \lambda v$$

- Covariance Matrices
- Principal Component Analysis
- Duality


**Looking is not Enough**

Can't spot low dimensional data by looking at numbers

**Direction of Maximum Variation**

- Look for the vector v with $\|v\|^2 = 1$ to maximise

$$\sigma^2 = \frac{1}{m-1} \sum_{i=1}^m (v^T(x_i - \mu))^2$$




- This is a constrained optimisation problem
- Solve by maximising Lagrangian

$$\mathcal{L} = \frac{1}{m-1} \sum_{k=1}^m (v^T(x_k - \mu))^2 - \lambda(\|v\|^2 - 1)$$

- λ is a Lagrange multiplier

Direction of Maximum Variation

- The eigenvectors are directions that are extrema of the variance



- The variance in direction v is equal to

$$\begin{aligned}\sigma^2 &= \frac{1}{m-1} \sum_{i=1}^m (v^T(x_i - \mu))^2 \\ &= v^T C v = \lambda v^T v = \lambda\end{aligned}$$

- The variance is maximised by the eigenvector with the maximum eigenvalue

Covariance Matrix

- The **covariance matrix** is defined as

$$\mathbf{C} = \frac{1}{m-1} \sum_{k=1}^m (\mathbf{x}_k - \boldsymbol{\mu})(\mathbf{x}_k - \boldsymbol{\mu})^\top$$

- The components C_{ij} measure how the i^{th} and j^{th} components co-vary

$$C_{ij} = \frac{1}{m-1} \sum_{k=1}^m (x_{ik} - \mu_i)(x_{jk} - \mu_j)$$

- C.f. covariance of random variables

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

Matrix Form

- The covariance matrix is

$$\mathbf{C} = \frac{1}{m-1} \sum_{k=1}^m (\mathbf{x}_k - \boldsymbol{\mu})(\mathbf{x}_k - \boldsymbol{\mu})^\top$$

- Define the matrix

$$\mathbf{X} = \frac{1}{\sqrt{m-1}} (\mathbf{x}_1 - \boldsymbol{\mu}, \mathbf{x}_2 - \boldsymbol{\mu}, \dots, \mathbf{x}_m - \boldsymbol{\mu})$$

- We can write the covariance matrix as

$$\mathbf{C} = \mathbf{X}\mathbf{X}^\top$$

Eigenvalue Decomposition

- The eigenvectors of $\mathbf{C}$ with the largest eigenvalues are known as the **principal components**
- The eigenvalues are all greater than or equal to zero
- Recall an eigenvector $\mathbf{v}$ satisfies the equation

$$\mathbf{C}\mathbf{v} = \lambda\mathbf{v}$$


- Multiplying both sides by $\mathbf{v}^\top$

$$\mathbf{v}^\top \mathbf{C} \mathbf{v} = \lambda \mathbf{v}^\top \mathbf{v} = \lambda \|\mathbf{v}\|^2$$

but $\mathbf{v}^\top \mathbf{C} \mathbf{v} \geq 0$ and $\|\mathbf{v}\|^2 > 0$ so

$$\lambda = \frac{\mathbf{v}^\top \mathbf{C} \mathbf{v}}{\|\mathbf{v}\|^2} \geq 0$$

Ellipsoid and Eigen Space



Outer Product

- Remember that the outer-product of two vectors is defined as

$$\mathbf{x}\mathbf{y}^\top = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix} \begin{pmatrix} y_1 & y_2 & \dots & y_n \end{pmatrix} = \begin{pmatrix} x_1 y_1 & x_1 y_2 & \dots & x_1 y_n \\ x_2 y_1 & x_2 y_2 & \dots & x_2 y_n \\ \vdots & \vdots & \ddots & \vdots \\ x_n y_1 & x_n y_2 & \dots & x_n y_n \end{pmatrix}$$

- C.f. Inner product

$$\mathbf{x}^\top \mathbf{y} = \begin{pmatrix} x_1 & x_2 & \dots & x_n \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} = x_1 y_1 + x_2 y_2 + \dots + x_n y_n$$

Properties of Covariance Matrix

- The **quadratic form** of a vector and matrix is defined as

$$\mathbf{v}^\top \mathbf{M} \mathbf{v}$$

- The quadratic form of a covariance matrix is non-negative for any vector

$$\mathbf{v}^\top \mathbf{C} \mathbf{v} = \mathbf{v}^\top \mathbf{X} \mathbf{X}^\top \mathbf{v} = \mathbf{u}^\top \mathbf{u} = \|\mathbf{u}\|^2 \geq 0$$

where $\mathbf{u} = \mathbf{X}^\top \mathbf{v}$

- Matrices with non-negative quadratic forms are known as **positive semi-definite**

Surface Defined by Matrix

- The set of vectors $\mathbf{x}$ such that


$$\mathbf{x}^\top \mathbf{C}^{-1} \mathbf{x} = 1$$

defines a surface

- The surface is an ellipsoid, $\mathcal{E}$
- The eigenvectors point in the direction of the principal axes of the ellipsoid
- The radii of the principal axes are equal to the square root of the eigenvalues

Spanning Input Space

- A covariance matrix will have a zero eigenvalue only if there is no variation in the direction of the corresponding eigenvector
- A covariance matrix will have zero eigenvalues if the number of patterns are less than or equal to the number of dimensions
- A covariance matrix formed from $p+1$ patterns that do not all lie in a hyperplane (i.e. the p differences $\mathbf{x}_k - \mathbf{x}_1$ are linearly independent) will have no zero eigenvalues




- Matrices with no zero eigenvalues are called **full rank** matrices (as opposed to rank deficient)
- Full rank matrices are invertible, rank deficient matrices are singular and non-invertible
- Full rank covariance matrices have only positive eigenvalues and are said to be **positive definite**
- We would expect that when $m > p$ the covariance matrix will be positive definite unless there are some symmetries that linearly constrain the patterns


Principal Component Analysis

- PCA occurs as follows
 - Construct the covariance matrix
 - Find the eigenvalues and eigenvectors
 - Keep the eigenvectors with the largest eigenvalues (principal components)
 - Project the inputs into the space spanned by the principal components
- We then use the projected inputs as inputs to our learning machine


Subspace Projection



Eigenvectors



- Covariance Matrices
- Principal Component Analysis**
- Duality



Projection Matrix

- To project the inputs construct the projection matrix


$$P = \begin{pmatrix} v_1^T \\ v_2^T \\ \vdots \\ v_k^T \end{pmatrix}$$

- $k < p$ is the number of principal components we keep
- Given a p -dimensional input pattern x we can construct a k -dimensional representation z

$$z = P(x - \mu)$$

- Use z as our new inputs

Hand Written Digits




Reconstruction

- Projecting into a subspace of eigenvectors can be seen as approximating the inputs by


$$\hat{x}_i = \mu + \sum_{j=1}^k z_j^i v_j, \quad z_j^i = v_j^T (x_i - \mu), \quad \|v_j\| = 1$$

- Principal component analysis projects the data into the subspace of dimension k with the minimal approximation error $\mathbb{E}[\| \hat{x}_i - x_i \|^2]$
- The loss of “energy” (or squared error) is equal to the sum of the eigenvalues in the directions that are ignored

Eigenvalues for Digits




Reconstruction from Eigenvectors



Outline

1. Covariance Matrices
2. Principal Component Analysis
3. Duality



PCA for Images

- An image often contains around $p = 256 \times 256 = 64k$ pixels
- In standard PCA we would create an $p \times p$ matrix with over 4×10^9 elements
- This is intractable
- m images span at most a $m - 1$ dimensional subspace
- Usually this subspace will be much smaller than the space of all images $m \ll p$

Dual Matrix

- The covariance $\mathbf{C} = \mathbf{X}\mathbf{X}^T$ is a $p \times p$ matrix
- Consider the $m \times m$ matrix $\mathbf{D} = \mathbf{X}^T\mathbf{X}$
- Suppose $\mathbf{v}$ is an eigenvector of $\mathbf{D}$

$$\begin{aligned}\mathbf{D}\mathbf{v} &= \lambda\mathbf{v} \\ \mathbf{X}^T\mathbf{X}\mathbf{v} &= \lambda\mathbf{v} \\ \mathbf{X}\mathbf{X}^T\mathbf{X}\mathbf{v} &= \lambda\mathbf{X}\mathbf{v} \\ \mathbf{C}\mathbf{X}\mathbf{v} &= \lambda\mathbf{X}\mathbf{v} \Rightarrow \mathbf{C}\mathbf{u} = \lambda\mathbf{u}\end{aligned}$$

- $\mathbf{u} = \mathbf{X}\mathbf{v}$ (and $\mathbf{v} \propto \mathbf{X}^T\mathbf{u}$)


Dual Matrix

- Matrices $\mathbf{C} = \mathbf{X}\mathbf{X}^T$ and $\mathbf{D} = \mathbf{X}^T\mathbf{X}$ have the same non-zero eigenvalues
- Can use the dual $m \times m$ matrix $\mathbf{D}$ to find eigenvalues and eigenvectors of $\mathbf{C}$
- Note that $\mathbf{D} = \mathbf{X}^T\mathbf{X}$ has components $D_{kl} \propto (\mathbf{x}_k - \boldsymbol{\mu})^T(\mathbf{x}_l - \boldsymbol{\mu})$
- Takes $O(p \times m \times m)$ time to construct $\mathbf{D}$
- We work in a “dual space” which is the space spanned by the examples

What Does a Subspace Look Like?

- Consider $\mathbf{y}^1 = \begin{pmatrix} 2 \\ 4 \\ 5 \end{pmatrix}$, $\mathbf{y}^2 = \begin{pmatrix} 8 \\ 6 \\ 1 \end{pmatrix}$ with mean $\boldsymbol{\mu} = \begin{pmatrix} 5 \\ 5 \\ 3 \end{pmatrix}$
- Subtracting the mean $\mathbf{x}^i = \mathbf{y}^i - \boldsymbol{\mu}$ we can construct matrix

$$\mathbf{X} = \begin{pmatrix} x_1^1 & x_1^2 \\ x_2^1 & x_2^2 \\ x_3^1 & x_3^2 \end{pmatrix} = \begin{pmatrix} -3 & 3 \\ -1 & 1 \\ 2 & -2 \end{pmatrix}$$



Summary

- PCA allows us to reduce the dimensionality of the inputs
- We project the inputs into a sub-space where the data varies the most
- We can work in either the original space ($\mathbf{X}\mathbf{X}^T$) or the dual space ($\mathbf{X}^T\mathbf{X}$)
- When we have many more features than examples (i.e. $p \gg m$) then it is more efficient working in the dual space
- We will see examples of dual spaces again when we look at SVMs

Singular Value Decomposition (SVD)

$$\begin{pmatrix} 0 & X \\ X^T & 0 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = s \begin{pmatrix} u \\ v \end{pmatrix}$$

Singular Valued Decomposition, SVD, general linear maps

Singular Value Decomposition

- One of the most important results in linear algebra is that any matrix X can be decomposed as

$$X = USV^T$$

- If X is an $m \times p$ matrix then U is an $m \times m$ orthogonal matrix and V is a $p \times p$ orthogonal matrix i.e. $UU^T = U^T U = I$, $VV^T = V^T V = I$
- And S is an $m \times p$ matrix with non-zero diagonal elements $S_{ii} = s_i$ known as the **singular values**
- Where did that come from?

Eigenvectors

- Note that as $Xv = su$ and $X^T u = sv$ then

$$X(-v) = (-s)u \quad X^T u = (-s)(-v)$$

if $\begin{pmatrix} u \\ v \end{pmatrix}$ is an eigenvector of B with eigenvalue s then so is $\begin{pmatrix} u \\ -v \end{pmatrix}$ with eigenvalue $-s$

- If $m < p$ then $X^T X$ is not full rank so some eigenvalues are zero
- As a consequence $p - m$ vectors exist such that $Xv = 0$
- The eigenvalues and eigenvectors are

$$m \times \left(s_i, \begin{pmatrix} u_i \\ v_i \end{pmatrix} \right) \quad m \times \left(-s_i, \begin{pmatrix} u_i \\ -v_i \end{pmatrix} \right) \quad (p - m) \times \left(0, \begin{pmatrix} 0 \\ v_k \end{pmatrix} \right)$$

Normalisation Subtlety

$$\begin{pmatrix} 0 & X \\ X^T & 0 \end{pmatrix} = \begin{pmatrix} 0 & 2USV^T \\ 2VSU^T & 0 \end{pmatrix} = \begin{pmatrix} US & -US & 0 \\ VS & VS & 0 \end{pmatrix} = \begin{pmatrix} 0 & 2USV^T \\ 2VSU^T & 0 \end{pmatrix}$$

- Multiplying out we have $X = 2USV^T$ $X^T = 2VS^T U^T$
- Now the vectors u_i and v_i form orthogonal sets as they satisfy $X^T X v = s^2 v$ $XX^T u = s^2 u$
- But they are not normalised (since $\begin{pmatrix} u_i \\ v_i \end{pmatrix}$ is normalised). If we define $\tilde{U} = \sqrt{2}U$ and $\tilde{V} = \sqrt{2}V$ we find $X = \tilde{U}S\tilde{V}^T$ $X^T = \tilde{V}S^T\tilde{U}^T$

- Singular Value Decomposition
- General Linear Mappings
- Linear Regression Revisited

$$\begin{pmatrix} 0 & X \\ X^T & 0 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = s \begin{pmatrix} u \\ v \end{pmatrix}$$

Eigen Decomposition

- Consider an arbitrary $m \times p$ matrix X , and construct the $(m + p) \times (m + p)$ symmetric matrix, B ,

$$\begin{pmatrix} 0 & X \\ X^T & 0 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = s \begin{pmatrix} u \\ v \end{pmatrix}$$

$\begin{pmatrix} u \\ v \end{pmatrix}$ is an eigenvector of B with eigenvalue s

- We observe that

$$Xv = su \quad X^T u = sv$$

$$X^T X v = sX^T u = s^2 v \quad XX^T u = sXv = s^2 u$$

Matrix Decomposition

- Stacking the eigenvectors into a matrix

$$\begin{pmatrix} 0 & X \\ X^T & 0 \end{pmatrix} \begin{pmatrix} u & u & 0 \\ v & -v & v_0 \end{pmatrix} = \begin{pmatrix} u & u & 0 \\ v & -v & v_0 \end{pmatrix} \begin{pmatrix} s & 0 & 0 \\ 0 & -s & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

- Since the vectors $\begin{pmatrix} u_i \\ v_i \end{pmatrix}$ are eigenvectors of a symmetric matrix they form an orthogonal matrix if they are normalised.
- Multiply on the right by the transpose of the orthogonal matrix

$$\begin{pmatrix} 0 & X \\ X^T & 0 \end{pmatrix} = \begin{pmatrix} u & u & 0 \\ v & -v & v_0 \end{pmatrix} \begin{pmatrix} s & 0 & 0 \\ 0 & -s & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} u^T & v^T \\ u^T & -v^T \\ 0 & v_0^T \end{pmatrix}$$

SVD

- Any matrix, X , can be written as $X = USV^T$
 - U, V are orthogonal matrices
 - $S = \text{diag}(s_1, s_2, \dots, s_r)$ with $r = \min(m, p)$
- s_i can always be chosen to be non-negative and are known as **singular values**
- If X is symmetric and positive semi-definite (e.g. a covariance matrix) then $U = V$ and $S = \Lambda$ (the matrix of eigenvalues)—SVD is the same as eigen-decomposition (for a general symmetric matrix $s_i = |\lambda_i|$)
- Singular value decomposition applies to both square and non-square matrices—they describe general linear mappings

Spectrum of a Matrix

- The set of singular values (when ordered) is often referred to as the (singular value) spectrum of a matrix

- The largest singular value, s_{max} , provides a measure of the maximum change in the two norm of a vector

$$\mathbf{X}\mathbf{v}_{max} = \mathbf{U}\mathbf{S}\mathbf{V}^T\mathbf{v}_{max} = s_{max}\mathbf{u}_{max}$$

where $\mathbf{u}_{max}$ and $\mathbf{v}_{max}$ are the columns of the matrix $\mathbf{U}$ and $\mathbf{V}$ corresponding to s_{max} (note $\|\mathbf{v}_i\| = \|\mathbf{u}_i\| = 1$)

- $s_{max} = \|\mathbf{X}\|_2$ is known as the **spectral norm** or **2 norm** of a matrix satisfying

$$\|\mathbf{X}\mathbf{x}\|_2 \leq s_{max} \|\mathbf{x}\|_2$$

with equality when $\mathbf{x} = \mathbf{v}_{max}$

Economical Forms of SVD

- Often the rows or columns of the orthogonal matrices $\mathbf{U}$ and $\mathbf{V}$ that are not associated with a singular value are ignored

$$\begin{aligned} \begin{pmatrix} \mathbf{X} \\ \mathbf{0} \end{pmatrix} &= \begin{pmatrix} \mathbf{U} \\ \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{S} \\ \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{V}^T \\ \mathbf{0} \end{pmatrix} \\ \begin{pmatrix} \mathbf{X} \\ \mathbf{0} \end{pmatrix} &= \begin{pmatrix} \mathbf{U} \\ \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{S} \\ \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{V}^T \\ \mathbf{0} \end{pmatrix} \end{aligned}$$

- In numpy (which returns a vector $\mathbf{s}$ and $\mathbf{V}^T$)

```
U, s, Vt = np.linalg.svd(X) # full: U is m x m
# economical: U is m x r
U, s, Vt = np.linalg.svd(X, full_matrices=False)
```

General Matrix

- Recall that we can compute the SVD for any matrix, $\mathbf{X}$

- As matrices describe the most general linear mapping

$$\mathbf{v} \rightarrow \mathcal{T}[\mathbf{v}] = \mathbf{X}\mathbf{v}$$

- We can use SVD to understand any linear mapping
- Thus any linear mapping can be seen as a rotation followed by a squashing or expansion independently in each coordinate followed by another rotation

Determinants

- The determinant, $|\mathbf{M}|$ of a matrix $\mathbf{M}$ is defined for square matrices

- It describes the change in volume under the mapping

- Now for any two matrices $|\mathbf{AB}| = |\mathbf{A}||\mathbf{B}|$

- Thus $|\mathbf{M}| = |\mathbf{U}||\mathbf{S}||\mathbf{V}^T|$

- For an orthogonal matrix $|\mathbf{U}| = \pm 1$ since $\mathbf{U}\mathbf{U}^T = \mathbf{I} \Rightarrow |\mathbf{U}\mathbf{U}^T| = |\mathbf{I}| \Rightarrow |\mathbf{U}||\mathbf{U}^T| = 1$ or $|\mathbf{U}|^2 = 1$

- Thus

$$|\mathbf{M}| = \pm |\mathbf{S}| = \pm \prod_i s_i$$

Finding SVD

- Most libraries will compute the SVD for you

- In principle we can choose the smaller of the two matrices $\mathbf{X}\mathbf{X}^T$ and $\mathbf{X}^T\mathbf{X}$ and compute its eigenvalues (libraries avoid this: it squares the condition number)

- The singular values are the square root of the eigenvalues (notice that $\mathbf{X}\mathbf{X}^T$ and $\mathbf{X}^T\mathbf{X}$ are both positive semi-definite so the eigenvalues will be non-negative)

- We can then compute the $\mathbf{U}$ matrix or $\mathbf{V}$ matrix by multiplying through by $\mathbf{X}$ or $\mathbf{X}^T$ ($\mathbf{U} = \mathbf{X}\mathbf{V}\mathbf{S}^{-1}$ and $\mathbf{V} = \mathbf{X}^T\mathbf{U}\mathbf{S}^{-1}$)

- In practice to perform PCA most people subtract the mean from their data and then perform SVD


Outline

- Singular Value Decomposition
- General Linear Mappings
- Linear Regression Revisited

$$\begin{pmatrix} \mathbf{0} & \mathbf{X} \\ \mathbf{X}^T & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} = \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} \mathbf{s}$$

Matrices

$$\mathbf{M} = \begin{pmatrix} -0.45 & 1.9 \\ -0.77 & -0.025 \end{pmatrix} = \mathbf{U}\mathbf{S}\mathbf{V}^T = \begin{pmatrix} \cos(-175^\circ) & -\sin(-175^\circ) \\ \sin(-175^\circ) & \cos(-175^\circ) \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 0.75 \end{pmatrix} \begin{pmatrix} \cos(75^\circ) & -\sin(75^\circ) \\ \sin(75^\circ) & \cos(75^\circ) \end{pmatrix}$$



Non-Square Matrices

- When the matrices are non-square then the singular value matrix will either

★ Squash some directions to zero

★ Introduce new dimensions orthogonal to the vector

$$\begin{pmatrix} \mathbf{X} \\ \mathbf{0} \end{pmatrix} = \begin{pmatrix} \mathbf{U} \\ \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{S} \\ \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{V}^T \\ \mathbf{0} \end{pmatrix} \quad \begin{pmatrix} \mathbf{X} \\ \mathbf{0} \end{pmatrix} = \begin{pmatrix} \mathbf{U} \\ \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{S} \\ \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{V}^T \\ \mathbf{0} \end{pmatrix}$$

- The rank of an arbitrary matrix is the number of non-zero singular values (also number of linearly independent rows or columns)

Duality Revisited

- If $\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^T$ then

$$\begin{aligned}\mathbf{C} &= \mathbf{X}\mathbf{X}^T & \mathbf{D} &= \mathbf{X}^T\mathbf{X} \\ &= \mathbf{U}\mathbf{S}\mathbf{V}^T\mathbf{V}\mathbf{S}^T\mathbf{U}^T & &= \mathbf{V}\mathbf{S}^T\mathbf{U}^T\mathbf{U}\mathbf{S}\mathbf{V}^T \\ &= \mathbf{U}(\mathbf{S}\mathbf{S}^T)\mathbf{U}^T & &= \mathbf{V}(\mathbf{S}^T\mathbf{S})\mathbf{V}^T\end{aligned}$$

- If $\mathbf{X}$ is an $m \times p$ matrix then $\mathbf{S}\mathbf{S}^T$ is an $m \times m$ diagonal matrix with elements $S_{ii}^2 = s_i^2$
- $\mathbf{S}^T\mathbf{S}$ is a $p \times p$ matrix with elements $S_{ii}^2 = s_i^2$
- $\mathbf{U}$ and $\mathbf{V}$ are matrices of eigenvectors for $\mathbf{C}$ and $\mathbf{D}$
- The eigenvalues of $\mathbf{X}\mathbf{X}^T$ and $\mathbf{X}^T\mathbf{X}$ are $\lambda_i = S_{ii}^2 = s_i^2$

SS^T and S^TS

$$\mathbf{S} = \begin{pmatrix} s_1 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & s_2 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & s_m & 0 & 0 \dots & 0 \end{pmatrix}$$

$$\mathbf{S}^T\mathbf{S} = \begin{pmatrix} s_1^2 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & s_2^2 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & s_m^2 & 0 & 0 \dots & 0 \\ 0 & 0 & \dots & 0 & 0 & 0 \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & 0 & 0 \dots & 0 \end{pmatrix}$$

$$\mathbf{S}\mathbf{S}^T = \begin{pmatrix} s_1^2 & 0 & \dots & 0 \\ 0 & s_2^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & s_m^2 \end{pmatrix}$$

Having A Go

- It's really easy to verify this in Python

```
import numpy as np
X = np.random.rand(3, 2)
U, s, Vt = np.linalg.svd(X) # s: vector, Vt: V transposed
S = np.zeros_like(X)        # rebuild the 3 x 2 matrix S
np.fill_diagonal(S, s)
U @ S @ Vt                  # should equal X
U[:, 0] @ U[:, 1]           # columns of U are orthogonal
U.T @ U, U @ U.T            # both the identity
L, Va = np.linalg.eigh(X @ X.T) # eigenvalues, ascending
S @ S.T                      # the same values: s squared
```

- Test yourself!

Outline

1. Singular Value Decomposition
2. General Linear Mappings
3. Linear Regression Revisited

$$\begin{pmatrix} 0 & \mathbf{X} \\ \mathbf{X}^T & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} = s \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}$$

Linear Regression

- Given a set of data $\mathcal{D} = \{(\mathbf{x}_k, y_k) | k = 1, 2, \dots, m\}$
- In linear regression we try to fit a linear model

$$f(\mathbf{x}|\mathbf{w}) = \mathbf{x}^T\mathbf{w}$$

- Which we fit by minimising the squared error loss

$$L(\mathbf{w}) = \sum_{k=1}^m (f(\mathbf{x}_k|\mathbf{w}) - y_k)^2$$

Using SVD

- Using $\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^T$ then

$$\begin{aligned}\mathbf{X}^+ &= (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T \\ &= (\mathbf{V}\mathbf{S}^T\mathbf{S}\mathbf{V}^T)^{-1}\mathbf{V}\mathbf{S}^T\mathbf{U}^T \\ &= \mathbf{V}(\mathbf{S}^T\mathbf{S})^{-1}\mathbf{V}^T\mathbf{V}\mathbf{S}^T\mathbf{U}^T \\ &= \mathbf{V}(\mathbf{S}^T\mathbf{S})^{-1}\mathbf{S}^T\mathbf{U}^T = \mathbf{V}\mathbf{S}^+\mathbf{U}^T\end{aligned}$$

- If $m > p$

$$\mathbf{X}^T = \begin{pmatrix} | & | & | & | & | & | & | & | \\ s_1 & 0 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & s_2 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & 0 & s_3 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & s_p & 0 & 0 \dots & 0 \end{pmatrix}, \mathbf{S}^T = \begin{pmatrix} s_1 & 0 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & s_2 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & 0 & s_3 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & s_p & 0 & 0 \dots & 0 \end{pmatrix}$$

Matrix Form

- In matrix form we write $L(\mathbf{w}) = \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2$

$$\mathbf{X} = \begin{pmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_m^T \end{pmatrix} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{pmatrix}$$

- Then $\nabla L(\mathbf{w}^*) = 0$ implies

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

- This is known as the pseudo-inverse

Pseudo-Inverse of S

$$\mathbf{S}^T\mathbf{S} = \begin{pmatrix} s_1^2 & 0 & \dots & 0 \\ 0 & s_2^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & s_p^2 \end{pmatrix}, (\mathbf{S}^T\mathbf{S})^{-1} = \begin{pmatrix} s_1^{-2} & 0 & \dots & 0 \\ 0 & s_2^{-2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & s_p^{-2} \end{pmatrix}$$

$$\mathbf{S}^+ = (\mathbf{S}^T\mathbf{S})^{-1}\mathbf{S}^T = \begin{pmatrix} s_1^{-1} & 0 & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & s_2^{-1} & 0 & \dots & 0 & 0 & \dots & 0 \\ 0 & 0 & s_3^{-1} & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & s_p^{-1} & 0 & 0 \dots & 0 \end{pmatrix}$$

Moore–Penrose Pseudo-Inverse

- $(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$ needs every $s_i > 0$: it fails if $m < p$ or the columns of $\mathbf{X}$ are dependent
- The **Moore–Penrose** pseudo-inverse, $\mathbf{X}^+ = \mathbf{V} \mathbf{S}^+ \mathbf{U}^\top$, just skips the zero singular values: $S_{ii}^+ = 1/s_i$ if $s_i \neq 0$, and 0 otherwise
- **Over-constrained** ($m > p$, all $s_i > 0$): nothing changes, $\mathbf{w} = \mathbf{X}^+ \mathbf{y}$ minimises $\|\mathbf{X} \mathbf{w} - \mathbf{y}\|^2$
- **Under-constrained** ($m < p$): infinitely many $\mathbf{w}$ fit the data, and $\mathbf{X}^+ \mathbf{y}$ picks the one with the smallest $\|\mathbf{w}\|$
 - ★ If $s_i = 0$ then $\mathbf{X} \mathbf{v}_i = \mathbf{0}$: the data say nothing about $\mathbf{w}$ along $\mathbf{v}_i$, so that component is set to zero
- In numpy: `np.linalg.pinv` or `np.linalg.lstsq`

Regularisation

- Consider linear regression with a regulariser

$$\begin{aligned} \mathcal{L}(\mathbf{w}) &= \|\mathbf{X} \mathbf{w} - \mathbf{y}\|^2 + \eta \|\mathbf{w}\|^2 \\ &= \mathbf{w}^\top (\mathbf{X}^\top \mathbf{X} + \eta \mathbf{I}) \mathbf{w} - 2 \mathbf{w}^\top \mathbf{X}^\top \mathbf{y} + \mathbf{y}^\top \mathbf{y} \end{aligned}$$

- Thus


$$\nabla \mathcal{L}(\mathbf{w}) = 2(\mathbf{X}^\top \mathbf{X} + \eta \mathbf{I}) \mathbf{w} - 2 \mathbf{X}^\top \mathbf{y}$$

- and $\nabla \mathcal{L}(\mathbf{w}^*) = 0$ gives

$$\mathbf{w}^* = (\mathbf{X}^\top \mathbf{X} + \eta \mathbf{I})^{-1} \mathbf{X}^\top \mathbf{y}$$

Effect of Regularisation

- Without regularisation if $s_i = 0$ the problem would be ill-posed (even $\mathbf{S}^+$ does not exist since s_i^{-1} would be ill defined) and if s_i is small then $\mathbf{S}^+$ is ill-conditioned
- Using $\hat{\mathbf{S}}^+ = (\mathbf{S}^\top \mathbf{S} + \eta \mathbf{I})^{-1} \mathbf{S}^\top$ instead of $\mathbf{S}^+$ then



- Regularisation makes the machine much more stable (reduces the variance)

III-Conditioned Data Matrix

- Recall that

$$\mathbf{w}^* = \mathbf{X}^+ \mathbf{y} = \mathbf{V} \mathbf{S}^+ \mathbf{U}^\top \mathbf{y}$$

- If any of the singular values of $\mathbf{X}$ are small then $\mathbf{S}^+$ will magnify components in that direction
- Any errors in the target $\mathbf{y}$ will be magnified
- This leads to poor weights

Regularisation Continued

- Using $\mathbf{X} = \mathbf{U} \mathbf{S} \mathbf{V}^\top$



$$\begin{aligned} \mathbf{w}^* &= (\mathbf{X}^\top \mathbf{X} + \eta \mathbf{I})^{-1} \mathbf{X}^\top \mathbf{y} \\ &= \mathbf{V} (\mathbf{S}^\top \mathbf{S} + \eta \mathbf{I})^{-1} \mathbf{S}^\top \mathbf{U}^\top \mathbf{y} \end{aligned}$$

- where

$$(\mathbf{S}^\top \mathbf{S} + \eta \mathbf{I})^{-1} \mathbf{S}^\top = \begin{pmatrix} \frac{s_1}{s_1^2 + \eta} & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \frac{s_2}{s_2^2 + \eta} & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 0 & \frac{s_3}{s_3^2 + \eta} & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \frac{s_p}{s_p^2 + \eta} & 0 & 0 & \cdots & 0 \end{pmatrix}$$

Summary

- Any matrix can be decomposed as $\mathbf{X} = \mathbf{U} \mathbf{S} \mathbf{V}^\top$ where
 - ★ $\mathbf{U}$ and $\mathbf{V}$ are orthogonal (rotation matrices)
 - ★ $\mathbf{S} = \text{diag}(s_1, \dots, s_r)$ is a diagonal matrix of non-negative singular values
- This describes the most general linear transform
- The transform exploits the duality between $\mathbf{X} \mathbf{X}^\top$ and $\mathbf{X}^\top \mathbf{X}$
- In linear regression the pseudo-inverse involves the reciprocal of the singular values, which can lead to poor generalisation
- Regularisation improves the conditioning of the “inverse” matrix



$$z = e^{-(x+1)^2 - (y-1)^2} + 0.6e^{-(x-1)^2 - 0.5(y+1)^2 + 0.1(x-3)(y-3)}$$

Gradient descent, quadratic minima, differing length scales

ML = Optimisation

- Many learning machines can be thought of as functions of the form

$$\hat{y} = f(\mathbf{x}|\mathbf{w})$$

(or more generally $\hat{y} = f(\mathbf{x}|\mathbf{w})$)

- Given an input pattern (set of features) $\mathbf{x}$ the learning machine makes a prediction $\hat{y}$
- We try to choose the parameters $\mathbf{w}$ so that the predictions are good
- In practice training a learning machine comes down to optimising some loss function

Training

- Given a (labelled) training dataset

$$\mathcal{D} = \{(\mathbf{x}_k, y_k) | k = 1, \dots, m\}$$


- We define an error or loss function that we want to minimise

$$L(\mathbf{w}|\mathcal{D}) = \frac{1}{m} \sum_{k=1}^m (f(\mathbf{x}_k|\mathbf{w}) - y_k)^2$$


- We then use the machine with the weights $\mathbf{w}^*$ which minimise $L(\mathbf{w}|\mathcal{D})$

Outline

- Motivation
- Gradient Descent
- Why Gradient Descent is Difficult




- Motivation
- Gradient Descent
- Why Gradient Descent is Difficult



MLP

- We can depict a neural network such as an MLP by a diagram




- Stands for the function $(\hat{y} = f(\mathbf{x}|\mathbf{w}))$

$$\hat{y} = g(w_5^3 g(w_3^1 x_1 + w_3^2 x_2 + b_3) + w_5^4 g(w_4^1 x_1 + w_4^2 x_2 + b_4) + b_5)$$

where, for example, $g(V) = \frac{1}{1+e^{-V}}$

Computing Gradients

- $L(\mathbf{w}|\mathcal{D})$ is a complex function of the weights $\mathbf{w}$





- To minimise $L(\mathbf{w}|\mathcal{D})$ we compute the gradient $\nabla L(\mathbf{w}|\mathcal{D})$
- For neural networks an efficient algorithm for computing the gradient is known as back-propagation of errors (*back-prop*)

Gradient Optimisation

- At a maximum or minimum (or a saddle point) $\nabla L(\mathbf{w}|\mathcal{D}) = \mathbf{0}$
- E.g.


$$L = e^{-(x+1)^2 - (y-1)^2} + 0.6e^{-(x-1)^2 - 0.5(y+1)^2 + 0.1(x-3)(y-3)}$$



Gradient Descent


- For a simple function $L(w|\mathcal{D})$ we can solve $\nabla L(w|\mathcal{D}) = 0$ explicitly E.g. the linear perceptron
- For non-linear functions we usually can't solve this set of simultaneous equations
- We can find a maximum or minimum **iteratively**
- If we know the gradient then we can follow the gradient
 - Maximisation: $w \rightarrow w' = w + r \nabla L(w|\mathcal{D})$
 - Minimisation: $w \rightarrow w' = w - r \nabla L(w|\mathcal{D})$

Hill-Climbing



What Goes Right

- Almost all minima are quadratic (Morse's theorem)



- Taylor expanding around a minimum x^*

$$f(x) = f(x^*) + (x - x^*)f'(x^*) + \frac{1}{2}(x - x^*)^2 f''(x^*) + \dots$$

$$= f(x^*) + \frac{1}{2}(x - x^*)^2 f''(x^*) + \frac{1}{3!}(x - x^*)^3 f'''(x^*) + \dots$$
- If $x - x^*$ is sufficiently small the higher order terms are negligible

Newton's Method

- If we were in a quadratic minimum

$$f(x) = a + \frac{b}{2}(x - x^*)^2$$

- then

$$f'(x) = b(x - x^*), \quad f''(x) = b$$

- so

$$x - x^* = \frac{f'(x)}{b} = \frac{f'(x)}{f''(x)}$$

- or

$$x^* = x - \frac{f'(x)}{f''(x)}$$

Newton's Method

- This is Newton's method
- For non-quadratic functions Newton's method converges **quadratically** provided we are sufficiently close to a minimum
- If we are at a distance $x - x^* = \epsilon$ from the minimum then after one iteration we will be a distance of order ϵ^2 after two iterations of order ϵ^4 , etc.
- If we are too far from a minimum we might go anywhere
- We should follow the gradient until we are near the minimum

Taylor's Expansion in High Dimensions

- We can generalise these results to many dimensions
- The Taylor expansion of a function $f(x)$ about x_0

$$f(x) = f(x_0) + (x - x_0)^T \nabla f(x_0) + \frac{1}{2}(x - x_0)^T \mathbf{H}(x - x_0) + \dots$$

where $\mathbf{H}$ is the **Hessian** matrix with elements


$$H_{ij} = \frac{\partial^2 f(x_0)}{\partial x_i \partial x_j}$$

- Newton's method in high dimension is

$$x^* = x - \mathbf{H}^{-1} \nabla f(x)$$


Using the Second Derivative

- If we are optimising n parameters the Hessian is an $n \times n$ matrix
- It is time-consuming to compute (and prone to errors when coding) for deep learning it is impossible even to store the Hessian
- Away from minima they can be misleading



Outline

- Motivation
- Gradient Descent
- Why Gradient Descent is Difficult**




Step Size

- Gradient descent

$$x' = x - r \nabla f(x)$$


- Need to choose the learning rate or step size, r
- Too small a step takes lots of time
- Too large a step takes you away from a minimum

Step Size




Higher Dimensions

- In higher dimensions the problem is that there are some directions you need to move a long way




Getting There Quicker

- Increasing the step size speeds up convergence, but the direction of steepest descent doesn't point to the minimum




More Haste Less Speed

- Increasing the step size, just a little further, increases the rate of convergence in one direction, but ...




Linear and Quadratic Convergence

- Near a minimum, along each eigenvector of the Hessian, a gradient step multiplies the error by $(1 - r\lambda_i)$
- Stable only if $r < 2/\lambda_{\max}$, so the flattest direction shrinks by only about $1 - \lambda_{\min}/\lambda_{\max} = 1 - 1/\kappa$ per step (κ is the condition number)
- A constant factor per step is **linear** convergence; Newton squares the error: **quadratic** convergence




Line Minimisation

- We can systematically seek the minimum along a line of the gradient



Zig-Zag

- Note that in narrow valleys line minimisation zigzags (each step is at right angles to the last)



- If we computed the Hessian and used Newton's method we would jump straight to the minimum if we were in a quadratic potential
- However computing the Hessian is time consuming and misleading if we are not in a quadratic potential (i.e. far from the optimum)

- Good optimisation algorithms use second-order (curvature) information without computing the full Hessian
- E.g. Conjugate gradient
 - ★ Performs Line Minimisation
 - ★ Uses gradient, but does not go along it
 - ★ For a quadratic minimum in d dimensions it reaches the minimum in d steps
- E.g. Levenberg-Marquardt
 - ★ Used on least squares problem only
 - ★ Uses linear approximation of function to approximate Hessian
 - ★ Adapts from hill-climbing to Newton method
 - ★ Avoids line-minimisation

Levenberg-Marquardt

- Want to minimise $\|\epsilon(\mathbf{w})\|^2$ where $\epsilon_i(\mathbf{w}) = f(\mathbf{x}_i|\mathbf{w}) - y_i$
- Use linear approximation


$$\epsilon_i(\mathbf{w}) \approx \epsilon_i(\mathbf{w}^{(k)}) + (\mathbf{w} - \mathbf{w}^{(k)})^\top \nabla \epsilon_i(\mathbf{w}^{(k)})$$

$$\text{with } \nabla \epsilon_i(\mathbf{w}^{(k)}) = \nabla f(\mathbf{x}_i|\mathbf{w}^{(k)})$$

- Solve quadratic minimisation of approximate error $\text{argmin}_{\mathbf{w}} L_{\text{approx}}(\mathbf{w})$ with $\mathbf{J} = \nabla \epsilon(\mathbf{w}^{(k)})$

$$\begin{aligned} L_{\text{approx}}(\mathbf{w}) &= \|\epsilon(\mathbf{w}^{(k)}) + \mathbf{J}(\mathbf{w} - \mathbf{w}^{(k)})\|^2 \\ &= \epsilon(\mathbf{w}^{(k)})^\top \epsilon(\mathbf{w}^{(k)}) + 2(\mathbf{w} - \mathbf{w}^{(k)})^\top \mathbf{J}^\top \epsilon(\mathbf{w}^{(k)}) \\ &\quad + (\mathbf{w} - \mathbf{w}^{(k)})^\top \mathbf{J}^\top \mathbf{J} (\mathbf{w} - \mathbf{w}^{(k)}) \end{aligned}$$

Rosenbrock: $\epsilon_1 = 10(x_2 - x_1^2)$, $\epsilon_2 = 1 - x_1$



Summary

- There are some **non-gradient methods** (Nelder Mead, evolutionary strategies, Powell's method), but in very high dimensions these are not very competitive
- There are **gradient methods** (first order methods) that suffer from the problem of having to choose a single step size with conflicting requirements in different directions; they converge **linearly**
- **Newton's method** (a second order method) requires computing the Hessian matrix, converges **quadratically**, but can take you in the wrong direction if you are not sufficiently close to a minimum
- There exist a number of **quasi-Newton methods** (BFGS, conjugate gradient, Levenberg-Marquardt, etc.) that approximate Newton's method often without explicitly computing the Hessian

- BFGS (Broyden–Fletcher–Goldfarb–Shanno) is the most widely used quasi-Newton method
- It builds an approximation $\mathbf{B} \approx \mathbf{H}^{-1}$ from how the gradient changes from one step to the next

$$\mathbf{B}(\nabla f(\mathbf{x}_{t+1}) - \nabla f(\mathbf{x}_t)) = \mathbf{x}_{t+1} - \mathbf{x}_t$$

- Steps $\mathbf{x}_{t+1} = \mathbf{x}_t - \alpha \mathbf{B} \nabla f(\mathbf{x}_t)$, with a line search for α
- Converges **super-linearly** without ever computing the Hessian
- $\mathbf{B}$ is $N \times N$; **L-BFGS** keeps only the last ten or so steps instead, so its memory is $O(N)$

Trust Region

- Solution given by $\nabla_{\mathbf{w}} L_{\text{approx}}(\mathbf{w}) = 0$ gives

$$\mathbf{w}^{(k+1)} = \mathbf{w}^{(k)} - (\mathbf{J}^\top \mathbf{J})^{-1} \mathbf{J}^\top \epsilon(\mathbf{w}^{(k)})$$


- Can lead us in the wrong direction
- Instead use $\mathbf{w}^{(k+1)} = \text{argmin}_{\mathbf{w}} L_{\text{approx}}(\mathbf{w}) + \nu \|\mathbf{w} - \mathbf{w}^{(k)}\|^2$

$$\mathbf{w}^{(k+1)} = \mathbf{w}^{(k)} - (\mathbf{J}^\top \mathbf{J} + \nu \mathbf{I})^{-1} \mathbf{J}^\top \epsilon(\mathbf{w}^{(k)})$$

- ν limits the step size
- If predicted reduction in error is accurate then reduce ν , else if predicted reduction in error is very poor increase ν

Where Are These Used?

- For up to thousands of parameters (L-BFGS: millions), with gradients computed from all the data
- **Logistic regression**: scikit-learn's default solver is L-BFGS; Newton's method for generalised linear models
- **Gaussian processes**: fitting the kernel hyper-parameters with L-BFGS
- **Curve fitting** (non-linear least squares): Levenberg–Marquardt, e.g. SciPy's `curve_fit`
- **Support vector machines**: a quadratic programme (later lectures)
- With millions to billions of parameters and mini-batches of data we need different methods: next lecture



SGD, momentum, step size, RMSProp, Adam, learning-rate schedules

The Next Step

- In most machine learning problems we design a network and a loss function
- The rest requires simple optimisation
- Although Newton and Quasi-Newton methods ensure rapid convergence for many tasks this apparent advantage is illusory
- Since the rise of deep learning, gradient descent and its variants have become dominant


Stochastic Gradient Descent

- Stochastic gradient descent (SGD) follows the gradient of mini-batches
- Computationally this is efficient because it is much quicker to compute $\nabla L_B(w)$ than $\nabla L_D(w)$
- The batch gradients add noise (are stochastic) as we are only sampling the dataset
- This noise might help escape local-minima (but I'm not sure there is compelling evidence for this)
- Taking many smaller steps of approximate gradients reduces the likelihood of divergence

Convergence

- Asymptotic convergence of gradient descent (and SGD) is slower than quasi-Newton methods
- But there are three reasons why we don't really care
 - ★ For complex models we spend a lot of time before we are close to a quadratic minimum where asymptotic convergence kicks in
 - ★ These days we often use ReLUs (rectified linear units) which are not differentiable at 0. The convergence we discussed depends on a Taylor expansion that assumes a smooth (twice differentiable) loss
 - ★ We want to minimise the generalisation error, but are only minimising a poor surrogate, namely the training error

1. SGD
2. Momentum
3. Loss Landscapes



Mini-Batch Learning

- Rather than computing the gradient of the loss function, e.g.

$$\nabla L_D(w) = \frac{1}{|D|} \nabla \sum_{(x,y) \in D} L(f(x|w), y)$$

- We often compute an approximation

$$\nabla L_B(w) = \frac{1}{|B|} \nabla \sum_{(x,y) \in B} L(f(x|w), y)$$


- where $B \subset D$ is a randomly sampled subset of the training set
- Usually $|B| \ll |D|$

Automatic Differentiation

- For gradient descent we need to compute the gradient
- For most of my life this was just a pain
- You guys have it easy, modern deep learning frameworks do this for you automatically
- This is a *game changer*! We are no longer afraid of large models
- We can differentiate through algorithms allowing us to train incredibly sophisticated networks


Outline

1. SGD
2. Momentum
3. Loss Landscapes



Recap: One Step Size, Many Curvatures

- The Hessian has a spectrum of eigenvalues: some directions are steep, others nearly flat
- One step size r must satisfy $r < 2/\lambda_{\max}$, or we diverge (and quickly get NaN errors)
- The flat directions then shrink by only about $1 - 1/\kappa$ per step
- This seems hopelessly slow, yet plain SGD (nearly always with momentum) is widely used in deep learning



Momentum

- In high dimensions we zig-zag:
 - stepping consistently in directions with low curvature
 - jumping backwards and forwards past the minimum in directions with high curvature
- By introducing “momentum” we can increase our steps in low curvature directions and decrease them in high curvature directions

$$\mathbf{v}(t+1) = \gamma \mathbf{v}(t) - r \nabla f(\mathbf{x}(t))$$

$$\mathbf{x}(t+1) = \mathbf{x}(t) + \mathbf{v}(t+1)$$

$0 < \gamma < 1$ sets how much velocity is kept ($1 - \gamma$ acts as friction)

Nesterov Momentum


- Nesterov's variant takes the gradient where momentum is about to carry us

$$\mathbf{v}(t+1) = \gamma \mathbf{v}(t) - r \nabla f(\mathbf{x}(t) + \gamma \mathbf{v}(t))$$

$$\mathbf{x}(t+1) = \mathbf{x}(t) + \mathbf{v}(t+1)$$

- With well-tuned momentum the number of steps needed scales as $\sqrt{\kappa}$ rather than κ
- The convergence is still **linear**: momentum fixes the condition number, not the order of convergence
- On convex problems Nesterov reaches error $O(1/t^2)$ rather than $O(1/t)$, the best any gradient method can do (much more in COMP6260, Optimisation for Machine Learning)


Gradient Descent with Momentum



Adaptive Methods

- A major difficulty of high dimensional optimisation is the existence of different scales
- That is, there are some directions where we have to move a lot (low curvature), while other directions we have to make small steps
- In adaptive methods we use a dynamic algorithm to rescale the step size of each parameter (weight)
- We can think of this as a “regularisation” that makes our basins of attraction more spherical

Rescaling Co-ordinates



RMSProp

- Rescale each coordinate by the typical size of its gradient
- Keep a running average of the squared gradients

$$S_i(t+1) = (1 - \gamma) S_i(t) + \gamma \left(\frac{\partial L_B(\mathbf{w}(t))}{\partial w_i(t)} \right)^2$$

- and divide the step by its square root

$$w_i(t+1) = w_i(t) - \frac{r}{\sqrt{S_i(t+1) + \epsilon}} \frac{\partial L_B(\mathbf{w}(t))}{\partial w_i(t)}$$

- Steep directions take small steps and flat ones large steps: each weight moves about r per step, whatever the scale of the loss

Adding in Momentum

- RMSProp has no momentum (one argument is that rescaling has already “regularised the landscape”)
- Nevertheless we can learn a momentum term

$$M_i(t+1) = (1 - \beta) M_i(t) + \beta \frac{\partial L_B(\mathbf{w}(t))}{\partial w_i(t)}$$

$$S_i(t+1) = (1 - \gamma) S_i(t) + \gamma \left(\frac{\partial L_B(\mathbf{w}(t))}{\partial w_i(t)} \right)^2$$

- Starting from $M_i(0) = S_i(0) = 0$ biases these averages towards zero at first, which we correct

$$\hat{M}_i(t+1) = \frac{M_i(t+1)}{1 - (1 - \beta)^{t+1}} \quad \hat{S}_i(t+1) = \frac{S_i(t+1)}{1 - (1 - \gamma)^{t+1}}$$

ADAM


- Adaptive Moment Estimation (Adam) adapts the scale of the parameters, but also uses momentum
- Update weights

$$w_i(t+1) = w_i(t) - \frac{r}{\sqrt{\hat{S}_i(t+1) + \epsilon}} \hat{M}_i(t+1)$$

- Adam and its variants are very successful in deep learning: a network often learns with Adam when it would not with SGD
- AdamW** applies weight decay directly to the weights, $w_i \leftarrow w_i - r\eta w_i$, rather than through the loss, where Adam's rescaling would weaken it; the default for transformers


Correlated Weights

- Correlated weights will lead to skewed valleys
- Adaptive methods would be more efficient if they were rotationally invariant
- However, this would slow down the algorithm
- ADAM is a compromise between speed of implementation and effectiveness



Loss Landscape

- A useful concept for understanding optimisation is the loss landscape
- For every value of the weights w there is some loss associated with it
- Note that this landscape is very high dimensional, so our intuition can be a bit misleading
- The landscape is also huge




Noisy Optimisation

- Because we are using mini-batches we aren't actually following the correct gradients
- If we reduce our step size then our average gradient is closer to the true gradient
- Thus our optimisation is inherently noisy
- At least for deep learning there is no evidence that the weights stop changing even after learning for a huge amount of time

Covariant Equations


- Vector arithmetic, (matrix multiplication, addition, multiplying by a scale, computing gradients) has a covariant property that it is invariant to the coordinate system
- That is we can translate and rotate our coordinates and get the same result
- That is not true of element-wise multiplication of vectors
- Many ML algorithms do this (including RMSProp and Adam), but they aren't invariant if we rotate our coordinates



Can't rescale coordinates

Outline

- SGD
- Momentum
- Loss Landscapes




Global and Local Minima

- Our objective is to minimise our loss
- This would mean finding the global minimum
- For non-convex losses no efficient algorithm is guaranteed to find the global minimum
- The best we can hope is to find a local minimum by doing gradient descent
- In practice, our landscapes are so large that we are probably unable to find even a local minimum

Valleys in Valleys in Valleys


- My view of the loss landscape is that we have valleys inside bigger valleys inside bigger valleys, and so on
- If our noise is high (we use large step size), we can navigate away from local valleys and move towards the big valley
- But, we can't explore the small valleys
- Often when the rate of improvement slows down, practitioners will reduce their learning rate by a factor of 10 and the rate of improvement jumps up
- Some practitioners cycle through raising and lowering their learning rates which may speed up convergence

Reducing the Step Size



Learning-Rate Schedules


- Large r early: noisy steps cross the small valleys and find the big ones
- Small r late: settle into a valley



- A short warm-up stops the first steps diverging before Adam's running averages have settled; common for transformers

Symmetries

- When training neural networks (deep or shallow) there are typically many symmetries
- Examples come from permuting the neurons (or filters in CNNs) within a layer, together with their weights



More Symmetries


- With n hidden nodes there are $n!$ symmetries
- For odd activation functions such as $\tanh$, multiplying all the input and output weights of a node by -1 gives the same network
- There are also continuous symmetries. For linear networks with two layers performing a mapping $W_2 W_1$ then for any invertible matrix U (where $U U^{-1} = I$)

$$W_2 W_1 = W_2 I W_1 = W_2 (U U^{-1}) W_1 = (W_2 U) (U^{-1} W_1)$$

so a network with weights $W'_2 = W_2 U$ and $W'_1 = U^{-1} W_1$ will perform the same mapping

- There will be a huge space of equivalent networks


Skip Connections



- Breaks permutation symmetry. Used in Resnets, transformer, etc.


Symmetries

- The landscape potentially has a huge manifold with the same losses (neutral network)
- Find solution close to start
- We should think of optimisation more as a process of travelling than a process of arriving
- Of course even if we do minimise our loss we are not guaranteed to minimise our generalisation performance



Lessons

- SGD together with automatic differentiation has revolutionised machine learning
- There are a number of techniques to get the step size right: momentum (and Nesterov), adaptive step sizes (RMSProp), Adam and AdamW, learning-rate schedules
- Still need to understand that we are exploring a huge and complex loss landscape
- What works is problem specific (as always)



Lagrangians, Inequalities, KKT, Linear Programming, Quadratic Programming, Duality

Optimisation with Constraints

- There are a number of important applications where we wish to minimise an objective function subject to inequality constraints
- A prominent example of this is support vector machines
- More generally there are a large number of kernel models that involve constraints
- However, constraints are ubiquitous in machine learning (e.g. in Wasserstein GANs)

Conditions on Optimum

- The optimisation problem is


$$\max_{\alpha} \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \alpha) \quad \text{where} \quad \mathcal{L}(\mathbf{x}, \alpha) = f(\mathbf{x}) - \alpha g(\mathbf{x})$$

- Assuming differentiability


$$\begin{aligned} \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \alpha) &= \nabla_{\mathbf{x}} f(\mathbf{x}) - \alpha \nabla_{\mathbf{x}} g(\mathbf{x}) = 0 \\ \frac{\partial \mathcal{L}}{\partial \alpha} &= -g(\mathbf{x}) = 0 \end{aligned}$$

- The second condition is just the constraint
- But what about first condition: $\nabla_{\mathbf{x}} f(\mathbf{x}) = \alpha \nabla_{\mathbf{x}} g(\mathbf{x})$?

Constrained Optima



1. Constrained Optimisation
2. Inequalities
3. Duality



Solving Constrained Optimisation Problems

- Suppose we have a problem

$$\min_{\mathbf{x}} f(\mathbf{x}) \quad \text{subject to} \quad g(\mathbf{x}) = 0$$

- A standard procedure is to define the Lagrangian

$$\mathcal{L}(\mathbf{x}, \alpha) = f(\mathbf{x}) - \alpha g(\mathbf{x})$$

where α is known as a Lagrange multiplier

- In the extended space $(\mathbf{x}, \alpha)$ we have to solve

$$\max_{\alpha} \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \alpha)$$

Note on Gradients

- Note that for any function $f(\mathbf{x})$ we can Taylor expand around $\mathbf{x}_0$

$$f(\mathbf{x}) = f(\mathbf{x}_0) + (\mathbf{x} - \mathbf{x}_0)^T \nabla_{\mathbf{x}} f(\mathbf{x}_0) + \frac{1}{2} (\mathbf{x} - \mathbf{x}_0)^T \mathbf{H} (\mathbf{x} - \mathbf{x}_0) + \dots$$

where $\mathbf{H}$ is the matrix of second derivatives known as the Hessian

- If we consider the set of points perpendicular to $\nabla_{\mathbf{x}} f(\mathbf{x}_0)$ which go through $\mathbf{x}_0$ (the tangent plane), these will have values

$$f(\mathbf{x}) = f(\mathbf{x}_0) + O(\|\mathbf{x} - \mathbf{x}_0\|^2)$$

thus $\nabla_{\mathbf{x}} f(\mathbf{x})$ is always orthogonal to the contour lines

Example


- Minimise $f(\mathbf{x}) = x^2 + 2y^2 - xy$
- Subject to $g(\mathbf{x}) = x - 2y - 3 = 0$
- Writing $\mathcal{L} = f(\mathbf{x}) - \alpha g(\mathbf{x})$
- Condition for minima is $\nabla_{\mathbf{x}} \mathcal{L} = 0$

$$\nabla_{\mathbf{x}} f(\mathbf{x}) = \begin{pmatrix} 2x - y \\ -x + 4y \end{pmatrix} = \alpha \nabla_{\mathbf{x}} g(\mathbf{x}) = \alpha \begin{pmatrix} 1 \\ -2 \end{pmatrix}$$

$$\text{and } \frac{\partial \mathcal{L}}{\partial \alpha} = -g(\mathbf{x}) = -x + 2y + 3 = 0$$

- Solving simultaneous equations gives the minimum at $(x, y) = (\frac{3}{4}, -\frac{9}{8})$ with $\alpha = \frac{21}{8}$

Surface



Multiple Constraints

- Given an optimisation problem with multiple constraints

$$\min_{\mathbf{x}} f(\mathbf{x}) \quad \text{subject to } g_k(\mathbf{x}) = 0 \text{ for } k = 1, 2, \dots, m$$

- We introduce multiple Lagrange multipliers

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\alpha}) = f(\mathbf{x}) - \sum_{k=1}^m \alpha_k g_k(\mathbf{x})$$


- The condition for an optimum is $\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\alpha}) = 0$ which implies

$$\nabla_{\mathbf{x}} f(\mathbf{x}) = \sum_{k=1}^m \alpha_k \nabla_{\mathbf{x}} g_k(\mathbf{x})$$


plus the original constraints $\frac{\partial \mathcal{L}(\mathbf{x}, \boldsymbol{\alpha})}{\partial \alpha_k} = -g_k(\mathbf{x}) = 0$

Outline


- Constrained Optimisation
- Inequalities**
- Duality



Inside Region



Saddle-Point $y = -9/8$



Example

- Minimise $f(\mathbf{x}) = x^2 + 2y^2 + 5z^2 - xy - xz$ subject to $g_1(\mathbf{x}) = x - 2y - z - 3 = 0$ and $g_2(\mathbf{x}) = 2x + 3y + z - 2 = 0$

- Writing $\mathcal{L}(\mathbf{x}, \boldsymbol{\alpha}) = f(\mathbf{x}) - \alpha_1 g_1(\mathbf{x}) - \alpha_2 g_2(\mathbf{x})$

- Condition for minima is $\nabla_{\mathbf{x}} \mathcal{L} = 0$ or $\nabla_{\mathbf{x}} f(\mathbf{x}) = \sum_{k=1}^2 \alpha_k \nabla_{\mathbf{x}} g_k(\mathbf{x})$

$$\begin{pmatrix} 2x - y - z \\ -x + 4y \\ 10z - x \end{pmatrix} = \alpha_1 \begin{pmatrix} 1 \\ -2 \\ -1 \end{pmatrix} + \alpha_2 \begin{pmatrix} 2 \\ 3 \\ 1 \end{pmatrix}$$

and $\frac{\partial \mathcal{L}}{\partial \alpha_i} = -g_i(\mathbf{x}) = 0$

- Solving simultaneous equations gives the minimum at $(\frac{37}{20}, -\frac{11}{20}, -\frac{1}{20})$ with $\alpha_1 = 3$ and $\alpha_2 = \frac{13}{20}$

Inequality Constraints

- Suppose we have the problem


$$\min_{\mathbf{x}} f(\mathbf{x}) \quad \text{subject to } g(\mathbf{x}) \geq 0$$

- Looks much more complicated, but

- Only two things can happen

- ★ Either a minimum, $\mathbf{x}^*$, of $f(\mathbf{x})$ satisfies $g(\mathbf{x}^*) > 0$
 - * We then have an unconstrained optimisation problem
- ★ Otherwise, it satisfies $g(\mathbf{x}^*) = 0$
 - * We have a constrained optimisation problem

On the Boundary



KKT Conditions

- To minimise $f(x)$ subject to $g(x) \geq 0$

$$\mathcal{L}(x, \alpha) = f(x) - \alpha g(x)$$


- Then $\nabla_x \mathcal{L} = 0$ or

$$\nabla_x \mathcal{L} = \nabla_x f(x) - \alpha \nabla_x g(x) = 0$$

- with $\alpha \geq 0$ and $\alpha g(x) = 0$: either $\alpha = 0$ (the constraint plays no part) or $\alpha > 0$ and $g(x) = 0$ (the solution is on the boundary)
- These conditions are known as the Karush-Kuhn-Tucker conditions

Outline

- Constrained Optimisation
- Inequalities
- Duality



Dual Problem

- If $f(x)$ and $g_i(x)$ are simple we can sometimes find a set of variables $x^*(\alpha)$ that minimises the Lagrangian

$$\nabla_x \mathcal{L}(x^*(\alpha), \alpha) = 0$$

- This leaves us with the **dual problem**

$$\max_{\alpha} \mathcal{L}(x^*(\alpha), \alpha)$$

- If we had an inequality constraint $g_i(x) \geq 0$ then we would have the additional constraint in the dual problem $\alpha_i \geq 0$

Linear Programming Example

- Suppose we eat potatoes and rice and we want to ensure that we get enough vitamin A and C

	Potatoes	Rice	Daily Requirement
Vitamin A	3	5	20
Vitamin C	5	2	24
Price	5	4	

- We want to buy P kg potatoes and R kg of rice as cheaply as possible subject to fulfilling our vitamin requirement

$$\min_{P, R} 5P + 4R$$

$$\text{subject to } P, R \geq 0, \quad 3P + 5R \geq 20 \quad \text{and} \quad 5P + 2R \geq 24$$

Many Inequalities

- Given the problem

$$\min_x f(x) \quad \text{subject to } g_k(x) \geq 0 \text{ for } k = 1, 2, \dots, m$$

- We introduce multiple Lagrange multipliers

$$\mathcal{L}(x, \alpha) = f(x) - \sum_{k=1}^m \alpha_k g_k(x)$$

- The condition for an optimum is

$$\nabla_x f(x) = \sum_{k=1}^m \alpha_k \nabla_x g_k(x)$$

- Plus, for each k , $\alpha_k \geq 0$ and $\alpha_k g_k(x) = 0$

Solving the Lagrangian for x

- Consider minimising a function $f(x)$ subject to a set of constraints $g_i(x) = 0$ or $g_i(x) \geq 0$

- We can consider this a double optimisation problem

$$\max_{\alpha} \min_x \mathcal{L}(x, \alpha) = \max_{\alpha} \min_x \left(f(x) - \sum_i \alpha_i g_i(x) \right)$$

where $\alpha_i \geq 0$ for each inequality constraint

Linear Programming

- In linear programming we minimise a linear objective function $c^T x$ subject to linear constraints $g(x) = Mx - b = 0$ (or $g(x) \geq 0$ with $x \geq 0$)

- The Lagrangian becomes

$$\mathcal{L}(x, \alpha) = c^T x - \alpha^T (Mx - b)$$

- An equivalent way of writing the Lagrangian is

$$\mathcal{L}(x, \alpha) = b^T \alpha - x^T (M^T \alpha - c)$$


- An entirely equivalent interpretation is that we maximise an objective function $b^T \alpha$ subject to constraints $M^T \alpha - c = 0$ (or $M^T \alpha - c \leq 0$ with $\alpha \geq 0$)

Linear Programming

- Minimise $5P + 4R$

- Subject to

- $3P + 5R \geq 20$
- $5P + 2R \geq 24$
- $P, R \geq 0$



- We can write the problem as a Lagrange problem


$$\min_{P,R} \max_{A,C} 5P + 4R - A(3P + 5R - 20) - C(5P + 2R - 24)$$

- subject to $P, R, A, C \geq 0$
- A and C are Lagrange multipliers for vitamin A and C
- We can rearrange the Lagrangian to obtain

$$\max_{A,C} \min_{P,R} 20A + 24C - P(3A + 5C - 5) - R(5A + 2C - 4)$$

Dual Linear Programme

- Maximise $20A + 24C$
- Subject to
 - ★ $3A + 5C \leq 5$
 - ★ $5A + 2C \leq 4$
 - ★ $A, C \geq 0$



Quadratic Programming

- A quadratic programme involves minimising a quadratic function $x^T Q x$ (with $Q \succ 0$) subject to linear constraints $Mx = b$ (or $Mx \geq b$)
- We can define the Lagrangian

$$\mathcal{L}(x, \alpha) = x^T Q x - \alpha^T (Mx - b)$$
- Where the solution is given by $\max_{\alpha} \min_x \mathcal{L}(x, \alpha)$
- If the constraints are inequality constraints then $\alpha_i \geq 0$

Dual Quadratic Programming Problem

- Substituting $x^* = \frac{1}{2} Q^{-1} M^T \alpha$ into

$$\mathcal{L}(x, \alpha) = x^T Q x - \alpha^T (Mx - b)$$
- We get the dual problem

$$\max_{\alpha} -\frac{1}{4} \alpha^T M Q^{-1} M^T \alpha + \alpha^T b$$
- If the constraints were inequality constraints then we have $\alpha_i \geq 0$
- We have exchanged one quadratic programme for another, but sometimes that is very useful (e.g. SVMs)

- The Lagrangian

$$\max_{A,C} \min_{P,R} 20A + 24C - P(3A + 5C - 5) - R(5A + 2C - 4)$$

leads to the dual problem

$$\max_{A,C} 20A + 24C$$

$$\text{subject to } 3A + 5C \leq 5 \quad 5A + 2C \leq 4 \quad A, C \geq 0$$

- Consider someone selling vitamins A and C. They want to maximise the price of vitamins A and C, but their prices cannot exceed the price of the vitamins in potatoes or rice

Why?

- Why are we bothered about translating one linear programme into another?
- Sometimes one form is massively easier to solve than the other
- This is because the first linear programme depends on the dimensionality of x while the second linear programme depends on the number of constraints (or dimensionality of α)
- This is important, for example, in Wasserstein GANs

Solution to Quadratic Programming Problem

- Using

$$\mathcal{L}(x, \alpha) = x^T Q x - \alpha^T (Mx - b)$$

- Then

$$\nabla_x \mathcal{L}(x, \alpha) = 2Qx - M^T \alpha$$


- So $\nabla_x \mathcal{L}(x, \alpha) = 0$ implies

$$x^* = \frac{1}{2} Q^{-1} M^T \alpha$$

Lessons

- A useful tool for performing constrained optimisation is the introduction of Lagrange multipliers
- This is particularly useful for problems with unique solutions (it will work when there are multiple solutions, but finding many saddle points is a pain)
- For inequality constraints we need to satisfy KKT conditions
- For simple situations (linear and quadratic programming) we can eliminate the original variables to obtain the dual problem


Convexity



Convex sets, convex functions, Jensen's inequality

Convex Regions

- Convex regions are familiar



- For any two points x and y in a region $\mathcal{R}$ then for any $a \in [0, 1]$ if

$$z = ax + (1-a)y \in \mathcal{R}$$

- then $\mathcal{R}$ is a convex region



Positive Semi-Definite Matrices

- Recall that a matrix $\mathbf{M}$ is positive semi-definite if for any vector v


$$v^T \mathbf{M} v \geq 0$$

(i.e. any quadratic form of the matrix is non-negative)


- (We showed this also implies that all the eigenvalues are non-negative)
- We denote the fact that $\mathbf{M}$ is positive semi-definite by $\mathbf{M} \succeq 0$, and $\mathbf{M} \succ 0$ if it is positive definite
- The set of positive semi-definite (PSD) matrices (or kernels) form a convex set

Outline

- Convex sets
- Convex functions
- Jensen's inequality



- Convex sets
- Convex functions
- Jensen's inequality



Convex Sets

- For any set, $\mathcal{S}$, where addition and scalar multiplication is defined (e.g. a vector space) then:

If for any two elements $x, y \in \mathcal{S}$ and any $a \in [0, 1]$

$$z = ax + (1-a)y \in \mathcal{S}$$

then $\mathcal{S}$ is said to be a convex set

Proof

- Consider any two arbitrarily chosen PSD matrices $\mathbf{M}_1$ and $\mathbf{M}_2$ and any $a \in [0, 1]$ then let

$$\mathbf{M}_3 = a\mathbf{M}_1 + (1-a)\mathbf{M}_2$$

- Then for any vector v

$$\begin{aligned} v^T \mathbf{M}_3 v &= v^T (a\mathbf{M}_1 + (1-a)\mathbf{M}_2) v \\ &= av^T \mathbf{M}_1 v + (1-a)v^T \mathbf{M}_2 v \\ &= am_1 + (1-a)m_2 \end{aligned}$$


where $m_1 = v^T \mathbf{M}_1 v$ and $m_2 = v^T \mathbf{M}_2 v$

- But $m_1, m_2 \geq 0$ since $\mathbf{M}_1, \mathbf{M}_2 \succeq 0$. Thus $am_1 + (1-a)m_2 \geq 0$ and so $\mathbf{M}_3 \succeq 0$ $\square$

Convex Functions


- Any function $f(x)$ is said to be a **convex function** if for any two points x and y and any $a \in [0, 1]$

$$f(ax + (1-a)y) \leq af(x) + (1-a)f(y)$$



Epigraph

- The **epigraph** of a function is the area that lies above the function
- The epigraph of a convex function is a convex region



Convex-Down or Concave Functions

- Any function, $f(x)$, that satisfies the inverse inequality


$$f(ax + (1-a)y) \geq af(x) + (1-a)f(y)$$

for any points x and y and any $a \in [0,1]$ is said to be a **convex-down** or **concave** function


- Everything true for a convex(-up) function carries over to a convex-down function with a small modification
- If $f(x)$ is a convex-up function then $g(x) = -f(x)$ is a convex-down function
- The area that lies below a convex-down function is a convex region

Examples

Convex-Up Functions



Convex-Down Functions



Linear Functions


- Linear functions are given by

$$f(x) = mx + c$$

- They satisfy the **equality**

$$f(ax + (1-a)y) = af(x) + (1-a)f(y)$$

- As such they are both convex(-up) and convex-down function



- $|x|$ is a convex-up function

Strictly Convex Function

- Functions that satisfy the strict inequality (for $0 < a < 1$ and $x \neq y$)

$$f(ax + (1-a)y) < af(x) + (1-a)f(y)$$

are said to be **strictly convex functions**

- A strictly convex-down function satisfies the reverse strict inequality
- Strictly convex-(up or down) functions don't contain any linear regions

Convexity in High Dimensions

- If $f: \mathbb{R}^n \rightarrow \mathbb{R}$ (i.e. $f(x)$ maps high dimensional point $x \in \mathbb{R}^n$ to a real value) satisfies

$$f(ax + (1-a)y) \leq af(x) + (1-a)f(y)$$


for any $x, y \in \mathbb{R}^n$ and any $a \in [0,1]$ then $f(x)$ is a convex function

- $\|x\|_2^2 = \sum_i x_i^2$ is a (strictly) convex function
- $\|x\|_1 = \sum_i |x_i|$ is a convex function

Properties of Convex Functions

- Convex functions lie on or above any tangent line

$$f(x) \geq f(x^*) + (x - x^*)f'(x^*)$$



Second Derivatives

- As $f(x)$ lies on or above its tangent line then for any $\epsilon > 0$

$$f'(x + \epsilon) \geq f'(x)$$

therefore $f''(x) = \lim_{\epsilon \rightarrow 0} (f'(x + \epsilon) - f'(x))/\epsilon \geq 0$ at all points x

- In high dimensions a convex function lies above its tangent plane

$$f(x) \geq f(x^*) + (x - x^*)^T \nabla f(x^*)$$

- The matrix of second derivatives (the Hessian) must be positive semi-definite at all points x

$$\mathbf{H} = \begin{pmatrix} \frac{\partial^2 f(x)}{\partial x_1^2} & \frac{\partial^2 f(x)}{\partial x_1 \partial x_2} & \cdots \\ \frac{\partial^2 f(x)}{\partial x_2 \partial x_1} & \frac{\partial^2 f(x)}{\partial x_2^2} & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} \succeq 0$$


Proving Convexity

- $f(x) = x^2$ is strictly convex as $f''(x) = 2 > 0$
- $f(x) = e^{cx}$ is strictly convex as $f''(x) = c^2 e^{cx} > 0$
- $f(x) = \log(x)$ is strictly convex-down as $f''(x) = -\frac{1}{x^2} < 0$
- $f(x, y) = \frac{x^2}{y}$ is convex for $y > 0$ as

$$\mathbf{H} = \begin{pmatrix} \frac{\partial^2 f(x, y)}{\partial x^2} & \frac{\partial^2 f(x, y)}{\partial x \partial y} \\ \frac{\partial^2 f(x, y)}{\partial y \partial x} & \frac{\partial^2 f(x, y)}{\partial y^2} \end{pmatrix} = \begin{pmatrix} \frac{2}{y} & -\frac{2x}{y^2} \\ -\frac{2x}{y^2} & \frac{2x^2}{y^3} \end{pmatrix} = \frac{2}{y^3} \begin{pmatrix} y^2 & -xy \\ -xy & x^2 \end{pmatrix}$$
- Now $T = \text{tr} \mathbf{H} = \frac{2}{y^3}(x^2 + y^2)$, $D = \det(\mathbf{H}) = 0$
- $\lambda_{1,2} = T/2 \pm \sqrt{T^2/4 - D} = \{0, T\} = \{0, \frac{2(x^2+y^2)}{y^3}\} \geq 0 \Rightarrow \mathbf{H} \succeq 0$

Convex Functions Defined on Convex Sets


- All the properties we have discussed hold for functions defined on a convex set
- $\sin(x)$ is not generally neither convex up or down
- $\sin(x)$ for $x \in [0, \pi]$ is convex-down (its second derivative $-\sin(x)$ is less than or equal to 0 in this range)



- For a convex function defined on a non-convex set, $\mathcal{S}$, there exists points $\mathbf{x}, \mathbf{y} \in \mathcal{S}$ such that for some $a \in [0, 1]$ there will be points $\mathbf{z} = a\mathbf{x} + (1-a)\mathbf{y} \notin \mathcal{S}$ (the function isn't defined on such points)

Unique Minimum

- Strictly convex function have a unique minimum
- The existence of a local minimum would break convexity
 - ★ The line connecting a local minimum to a global minimum would be strictly decreasing
 - ★ Thus there are points next to the local minimum with lower values
 - ★ This is a contradiction
- This remains true if we consider convex functions that are constrained to live in a convex set



Linear Regression

- For linear regression the loss function

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

is convex

- Since the Hessian $\mathbf{H} = 2\mathbf{X}^T \mathbf{X} \succeq 0$ (positive semi-definite) (For any vector $\mathbf{v}$ then $\mathbf{v}^T \mathbf{H} \mathbf{v} = 2\mathbf{v}^T \mathbf{X}^T \mathbf{X} \mathbf{v} = 2\|\mathbf{X}\mathbf{v}\|^2 \geq 0$)
- If $\mathbf{H} \succ 0$ there will be a unique minima while if $\mathbf{H}$ has some zero eigenvalues there will be a family of solutions

Sums of Convex Functions

- For any set of convex functions $f_1(x), f_2(x), \dots$ and any set of non-negative scalars $a_1, a_2, \dots$ then

$$g(x) = \sum_i a_i f_i(x)$$

is convex

- Proof

$$g''(x) = \sum_i a_i f_i''(x)$$

but $f_i''(x) \geq 0$ so $g''(x)$ is a sum on non-negative terms

- This generalises to higher dimensions as the sum of PSD matrices times positive factors is a PSD matrix

Constraints

- Often we impose constraints on the set of points, e.g.

$$x_i > 0 \quad \mathbf{a}^T \mathbf{x} = b \quad \mathbf{x}^T \mathbf{M} \mathbf{x} \leq 1$$

- Linear constraints (e.g. $x_i > 0$ or $\mathbf{a}^T \mathbf{x} = b$ or $\mathbf{a}^T \mathbf{x} \leq b$) always define a convex region
- Multiple linear constraints always define a convex region
- Non-linear constraints may or may not define a convex region
 $\{\mathbf{x} \in \mathbb{R}^n | \mathbf{x}^T \mathbf{M} \mathbf{x} \leq 1, \mathbf{M} \succeq 0\}$ does while
 $\{\mathbf{x} \in \mathbb{R}^n | \mathbf{x}^T \mathbf{M} \mathbf{x} \geq 1, \mathbf{M} \succeq 0\}$ doesn't

Convex Set of Minima

- If $f(x)$ is **convex** but not **strictly convex** then there might exist a convex set $\mathcal{M} \subset \mathcal{X}$ of minima such that for all $\mathbf{x}, \mathbf{y} \in \mathcal{M}$ and any $\mathbf{z} \in \mathcal{X}$ we have $f(\mathbf{x}) = f(\mathbf{y}) \leq f(\mathbf{z})$
- This set of minima is convex, that is, if $\mathbf{x}, \mathbf{y} \in \mathcal{M}$ then for any $a \in [0, 1]$ the point $\mathbf{z} = a\mathbf{x} + (1-a)\mathbf{y} \in \mathcal{M}$
- The sum of a convex function, $f(x)$, and a strictly convex function $g(x)$ will always be strictly convex since

$$f''(x) + g''(x) > 0$$

as $f''(x) \geq 0$ and $g''(x) > 0$

Regularised Linear Regression

- In ridge regression we minimise a loss

$$L(\mathbf{w}) = \|\mathbf{X}\mathbf{w} - \mathbf{y}\|^2 + \eta \|\mathbf{w}\|^2 = \mathbf{w}^T (\mathbf{X}^T \mathbf{X} + \eta \mathbf{I}) \mathbf{w} - 2\mathbf{w}^T \mathbf{X}^T \mathbf{y} + \mathbf{y}^T \mathbf{y}$$


- Because $\|\mathbf{w}\|^2$ is strictly convex the loss function is strictly convex and so will have a unique solution
- Using an L_1 regulariser (Lasso)

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

again $\|\mathbf{w}\|_1$ is convex and so $L(\mathbf{w})$ will be convex

- Using an L_1 and an L_2 regulariser (elastic net) also gives a unique solution

1. Convex sets
2. Convex functions
3. Jensen's inequality



Proof

- We said before that a convex function must lie on or above its tangent plane at any point x^*

$$f(x) \geq f(x^*) + (x - x^*)^T \nabla f(x^*)$$



- Taking $x^* = \mathbb{E}[X]$

$$f(X) \geq f(\mathbb{E}[X]) + (X - \mathbb{E}[X])^T \nabla f(\mathbb{E}[X])$$

- Taking expectations of both sides

$$\begin{aligned} \mathbb{E}[f(X)] &\geq f(\mathbb{E}[X]) + (\mathbb{E}[X] - \mathbb{E}[X])^T \nabla f(\mathbb{E}[X]) \\ &= f(\mathbb{E}[X]) \quad \square \end{aligned}$$

Kullback-Leibler Divergence



$$\text{KL}(f\|g) = -\sum_{i=1}^n f_i \log\left(\frac{g_i}{f_i}\right) = 0.235$$

Lessons

- Although we haven't talked much about machine learning, convexity is heavily used in many machine learning applications
- A lot of ML algorithms involve convex functions e.g. SVMs
- As such they will have a unique minimum (or a convex set of minima)
- Convexity is an elegant idea which is relatively easy to prove theorems about
- One of the most useful tools is Jensen's inequality

- In proving many properties of learning machines inequalities are really useful
- One of the most useful inequalities involve expectations of convex functions, this is known as **Jensen's Inequality**

- If $f(x)$ is a convex(-up) function then

$$\mathbb{E}[f(X)] \geq f(\mathbb{E}[X])$$

- If $f(x)$ is a convex(-down) function then

$$\mathbb{E}[f(X)] \leq f(\mathbb{E}[X])$$

Simple Proofs with Jensen's Inequality

- Since $f(x) = x^2$ is convex by Jensen's inequality

$$\mathbb{E}[X^2] \geq (\mathbb{E}[X])^2 \quad \text{or} \quad \mathbb{E}[X^2] - \mathbb{E}[X]^2 \geq 0$$

(i.e. variance are non-negative)

- The KL-divergence $\text{KL}(f\|g)$ between two categorical probability distributions $(f_1, f_2, \dots)$ and $(g_1, g_2, \dots)$ is defined as

$$\text{KL}(f\|g) = -\sum_i f_i \log\left(\frac{g_i}{f_i}\right)$$

(note $f_i, g_i \geq 0$ and $\sum_i f_i = \sum_i g_i = 1$)


Proof of Gibbs' Inequality

- To show that $\text{KL}(f\|g) \geq 0$ (Gibbs' inequality) we note that since the logarithm is a convex-down function

$$\begin{aligned} \text{KL}(f\|g) &= -\sum_i f_i \log\left(\frac{g_i}{f_i}\right) = \mathbb{E}_f\left[-\log\left(\frac{g_i}{f_i}\right)\right] \\ &\geq -\log\left(\mathbb{E}_f\left[\frac{g_i}{f_i}\right]\right) \\ &= -\log\left(\sum_i f_i \frac{g_i}{f_i}\right) = -\log\left(\sum_i g_i\right) = -\log(1) = 0 \end{aligned}$$

- We will meet KL-divergences later on

Support Vector Machines



Support Vector Machines, maximum margins

Support Vector Machines

- Support vector machines, when used right, often have the best generalisation results
- They are typically used on numerical data, but can and have been adapted to text, sequences, etc.
- Although not as trendy as deep learning, they will often be the method of choice on small data sets
- They subtly regularise themselves, choosing a solution that generalises well from a host of different solutions

Extended Feature Space

- To increase the likelihood of linear-separability we often use a high-dimensional mapping

$$\mathbf{x} = (x_1, x_2, \dots, x_p)^T \rightarrow \phi(\mathbf{x}) = (\phi_1(\mathbf{x}), \phi_2(\mathbf{x}), \dots, \phi_r(\mathbf{x}))^T$$

$$r \gg p$$

- Finding the maximum margin hyper-plane is time consuming in “primal” form if r is large
- We can work in the “dual” space of patterns, then we only need to compute inner-products

$$\langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle = \phi(\mathbf{x}_i)^T \phi(\mathbf{x}_j) = \sum_{k=1}^r \phi_k(\mathbf{x}_i) \phi_k(\mathbf{x}_j)$$

Kernel Functions

- Kernel functions are symmetric functions of two variable
- Strong restriction: *positive semi-definite*
- Examples


Quadratic kernel: $K(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{x}_1^T \mathbf{x}_2)^2$

Gaussian (RBF) kernel: $K(\mathbf{x}_1, \mathbf{x}_2) = e^{-\gamma \|\mathbf{x}_1 - \mathbf{x}_2\|^2}$

- Consider the mapping


$$\mathbf{x}_i = \begin{pmatrix} x_i \\ y_i \end{pmatrix} \rightarrow \phi(\mathbf{x}_i) = \begin{pmatrix} x_i^2 \\ y_i^2 \\ \sqrt{2}x_i y_i \end{pmatrix}$$

- The Big Picture
- Maximum Margins
- Duality
- Practice



Linear Separation of Data

- SVMs classify linearly separable data



- Finds maximum-margin separating plane

Kernel Trick

- If we choose a **positive semi-definite** kernel function $K(\mathbf{x}, \mathbf{y})$ then there exists functions $\phi(\mathbf{x}) = (\phi_k(\mathbf{x}) | k = 1, 2, \dots, r)$, such that

$$K(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle$$


(like an eigenvector decomposition of a matrix)

- Never need to compute $\phi_k(\mathbf{x}_i)$ explicitly as we only need the inner-product $\langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle = K(\mathbf{x}_i, \mathbf{x}_j)$ to compute maximum margin separating hyper-plane
- Sometimes $\phi(\mathbf{x}_i)$ is an infinite dimensional vector so it is good we don't have to compute all the elements!

Non-linear Separation of Data


$$K(\mathbf{x}_1, \mathbf{x}_2) = \begin{pmatrix} x_1^2 & y_1^2 & \sqrt{2}x_1 y_1 \end{pmatrix} \begin{pmatrix} x_2^2 \\ y_2^2 \\ \sqrt{2}x_2 y_2 \end{pmatrix} = x_1^2 x_2^2 + y_1^2 y_2^2 + 2x_1 y_1 x_2 y_2$$

$$= (x_1 x_2 + y_1 y_2)^2 = (\mathbf{x}_1^T \mathbf{x}_2)^2$$



Outline

1. The Big Picture
2. Maximum Margins
3. Duality
4. Practice




Inner Product

- Recall the inner or dot product in $\mathbb{R}^n$

$$\langle x, y \rangle = x \cdot y = x^T y = \sum_{i=1}^n x_i y_i = \|x\| \|y\| \cos(\theta)$$

- If $\|w\| = 1$ then $\langle x, w \rangle = \|x\| \cos(\theta)$




Maximise Margin

- Consider a linearly separable set of data
 - $\mathcal{D} = \{(x_k, y_k)\}_{k=1}^m$
 - $y_k \in \{-1, 1\}$
- Our task is to find a separating plane defined by the orthogonal vector w and a threshold b such that

$$y_k \left(\frac{\langle w, x_k \rangle}{\|w\|} - b \right) \geq \Delta$$

where Δ is the margin

Distance to hyperplanes



Constrained Optimisation

- Wish to find w and b to maximise Δ subject to constraints

$$y_k \left(\frac{\langle w, x_k \rangle}{\|w\|} - b \right) \geq \Delta \quad \text{for all } k = 1, 2, \dots, m$$

- If we divide through by Δ

$$y_k \left(\frac{\langle w, x_k \rangle}{\Delta \|w\|} - \frac{b}{\Delta} \right) \geq 1 \quad \text{for all } k = 1, 2, \dots, m$$

- Define $\hat{w} = w / (\Delta \|w\|)$ and $\hat{b} = b / \Delta$

$$y_k \left(\langle \hat{w}, x_k \rangle - \hat{b} \right) \geq 1$$

Quadratic Programming Problem


- Note that as $\hat{w} = w / (\Delta \|w\|)$

$$\|\hat{w}\| = \left\| \frac{w}{\Delta \|w\|} \right\| = \frac{1}{\Delta \|w\|} \|w\| = \frac{1}{\Delta}$$

- Minimising $\|\hat{w}\|^2$ is equivalent to maximising the margin Δ
- Can write the optimisation problem as a *quadratic programming problem*

$$\min_{\hat{w}, \hat{b}} \frac{\|\hat{w}\|^2}{2} \quad \text{subject to } y_k \left(\langle \hat{w}, x_k \rangle - \hat{b} \right) \geq 1 \quad \text{for all } k = 1, 2, \dots, m$$


Quadratic Programming in SVMs



Quadratic Programming

- We have a quadratic programming problem for the weights $\hat{w} = (\hat{w}_1, \hat{w}_2, \dots, \hat{w}_p)$ and bias $\hat{b}$ and m constraints
- This is a classic but fiddly optimisation problems
- It is a **strongly convex** optimisation problem so will have a unique minimum
- It can be solved in $O(p^3)$ time (it involves inverting matrices) (pew it is not NP-complete!)
- We will see that there is an equivalent dual problem which allows us to use the kernel trick with time complexity $O(m^3)$

1. The Big Picture
2. Maximum Margins
3. Duality
4. Practice



Lagrangian

- In the extended feature space we can find a separating plane (given by w and b) with maximum margin by solving the problem

$$\min_{w,b} \frac{\|w\|^2}{2} \quad \text{subject to } y_k(\langle w, \phi(x_k) \rangle - b) \geq 1 \text{ for all } k = 1, 2, \dots, m$$

- We can write this as a Lagrange problem

$$\min_{w,b} \max_{\alpha \geq 0} \mathcal{L}(w, b, \alpha)$$

where

$$\mathcal{L}(w, b, \alpha) = \frac{\|w\|^2}{2} - \sum_{k=1}^m \alpha_k (y_k (\langle w, \phi(x_k) \rangle - b) - 1)$$

subject to $\alpha_k \geq 0$

The Dual Problem

- The dual problem is now to find α_k 's that maximise

$$\mathcal{L}(\alpha) = \sum_{k=1}^m \alpha_k - \frac{1}{2} \sum_{k,l=1}^m \alpha_k \alpha_l y_k y_l \langle \phi(x_k), \phi(x_l) \rangle$$

- subject to constraints

$$\sum_{k=1}^m \alpha_k y_k = 0 \quad \forall k = 1, 2, \dots, m \quad \alpha_k \geq 0$$

- The Hessian of $\mathcal{L}(\alpha)$ has elements $H_{kl} = -y_k y_l \langle \phi(x_k), \phi(x_l) \rangle$ so $v^T H v = -\|\sum_k v_k y_k \phi(x_k)\|^2 \leq 0$ (note this is negative semi-definite so there is a unique maximum)

Sequential Minimal Optimisation

- One of the most efficient techniques for training SVMs is *Sequential Minimal Optimisation* or SMO
- This takes two Lagrange multipliers α_i and α_j and adjusts them to maximise the dual objective function
- This is very quick as it can be done in closed form
- Note that because $\sum_{k=1}^m y_k \alpha_k = 0$ we have to change at least two variables at the same time
- A heuristic is used to choose good pairs of α 's to optimise
- Run until close to the optimum

- We can generalise the SVM if we map all our features vectors to an extended feature space

$$x \rightarrow \phi(x)$$

- The components of $\phi(x)$ will typically be (non-linear) functions of x (e.g. $\phi_1(x) = x_1^2, \phi_2(x) = x_2^2, \phi_3(x) = \sqrt{2}x_1x_2$)
- We are free to choose whatever mappings we like
- There may be many more components of $\phi(x)$ than of x making it easier to find a linear separation of the two classes
- But in the extended feature space (involving $\phi(x) = (\phi_1(x), \phi_2(x), \dots, \phi_r(x))$) the time complexity is $O(r^3)$

Obtaining the Dual Form of the Problem

- Differentiating the Lagrangian with respect to w

$$\mathcal{L}(w, b, \alpha) = \frac{1}{2} \|w\|^2 - \sum_{k=1}^m \alpha_k (y_k (\langle w, \phi(x_k) \rangle - b) - 1)$$

$$\nabla_w \mathcal{L} = w - \sum_{k=1}^m \alpha_k y_k \phi(x_k) = 0 \implies w^* = \sum_{k=1}^m \alpha_k y_k \phi(x_k)$$

$$\frac{\partial \mathcal{L}}{\partial b} = \sum_{k=1}^m \alpha_k y_k = 0$$

- Substituting back into the Lagrangian

$$\max_{\alpha \geq 0} \mathcal{L}(w^*, b^*, \alpha) = \max_{\alpha \geq 0} \sum_{k=1}^m \alpha_k - \frac{1}{2} \sum_{k,l=1}^m \alpha_k \alpha_l y_k y_l \langle \phi(x_k), \phi(x_l) \rangle$$

Kernel Trick

- We will show in the next lecture that if $K(x, y)$ is a positive semi-definite function then it can always be written as

$$K(x, y) = \langle \phi(x), \phi(y) \rangle$$

- As $\langle \phi(x_k), \phi(x_l) \rangle$ appears in the dual problem we can express the dual problem as finding α_k 's that maximise


$$\mathcal{L}(\alpha) = \sum_{k=1}^m \alpha_k - \frac{1}{2} \sum_{k,l=1}^m \alpha_k \alpha_l y_k y_l K(x_k, x_l)$$

- We therefore never have to compute $\phi(x)$

Soft Margins

- We can relax the margin constraints by introducing *slack variables*, $s_k \geq 0$

$$y_k(\langle x_k, w \rangle - b) \geq 1 - s_k$$



- Minimise $\frac{\|w\|^2}{2} + C \sum_{k=1}^n s_k$
- Larger C punishes slack variables more

Dual Problem with Slack Variables

- The Lagrangian with slack variables is

$$\mathcal{L} = \frac{1}{2} \|w\|^2 + C \sum_{k=1}^m s_k - \sum_{k=1}^m \alpha_k \left(y_k (\langle w, \phi(x_k) \rangle - b) - 1 + s_k \right) - \sum_{k=1}^m \beta_k s_k$$

where β_k are Lagrange multipliers that ensure $s_k \geq 0$ (note that $\beta_k \geq 0$ —this is the KKT condition)


- Now minimising with respect to s_i

$$\frac{\partial \mathcal{L}}{\partial s_i} = C - \alpha_i - \beta_i = 0$$

- Or $\alpha_i = C - \beta_i$. Since $\beta_i \geq 0$ the constraint is $\alpha_i \leq C$ (recall also $\alpha_i \geq 0$)

Getting SVMs to Work Well

- SVMs rely on distances between data points
- These will change relative to each other if we rescale some features but not other—giving different maximum-margin hyper-planes



- If we don't know what features are important (most often the case), then it is worth scaling each feature (for example, so their range is between 0 and 1 or their variance is 1)

Choosing the Right Kernel Function

- There are kernels design for particular data types (e.g. string kernels for text or biological sequences)
- For numerical data, people tend to look at using no kernel (linear SVM), a radial basis function (Gaussian) kernel or polynomial kernels
- Kernels often come with parameters, e.g. the popular radial basis function kernel

$$K(x, y) = e^{-\gamma \|x - y\|^2}$$


- Optimal γ values range over 2^{-15} – 2^3

SVM Libraries

- Although SVMs have unique solutions, they require very well written optimisers
- If you have a large data set they can be very slow
- There are good libraries out there: svm-lib, svm-lite, (now old), scikit-learn, etc.
- These will often automate normalisation of data and grid search for parameters

Outline

- The Big Picture
- Maximum Margins
- Duality
- Practice



Optimising C

- Recall that we can introduce soft-margins using slack variables where we minimise $\frac{\|w\|^2}{2} + C \sum_{k=1}^m s_k$ subject to constraints
- In practice it can make a huge difference to the performance if we change C
- Optimal C values changes by many orders of magnitude e.g. 2^{-5} – 2^{15}
- Typically optimised by a grid search (start from 2^{-5} say and double until you reach 2^{15})
- Measure performance on a validation set


Multi-Class Problems

- By construction SVMs separate only two classes
- If we have a multi-class problem we have to use multiple SVMs
- There are two major ways practitioners do this
 - One-versus-all:** for each class, train a separate classifier to determine that class versus all others
 - All-pairs:** train a classifier for all pairs of classes

- In both cases choose the class which the classifier is most certain about
- Beware SVMs don't like imbalanced datasets

Conclusions

- We've seen how SVMs work
- We've learnt how to use them
- We've seen that we can find the maximum margin hyper-plane by solving a quadratic programming problem (with a unique solution)
- This is a convex optimisation problem with a unique optimum
- The **dual problem** of an SVM is particularly simple, especially if we use a positive semi-definite kernel (we explore these in the next lecture)



The Kernel Trick, SVMs, Regression

SVM Kernels

- SVM Kernels are functions of two variables that can be factorised

$$K(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle = \sum_i \phi_i(\mathbf{x}) \phi_i(\mathbf{y})$$

- where $\phi(\mathbf{x}) = (\phi_1(\mathbf{x}), \phi_2(\mathbf{x}), \dots)^\top$ and $\phi_i(\mathbf{x})$ are real valued functions of $\mathbf{x}$
- $K(\mathbf{x}, \mathbf{y})$ will be positive semi-definite (because it is an inner-product)
- Furthermore, any positive semi-definite function will factorise
- This factorisation is not always obvious (we return to this later)

Classifying New Data

- Having trained the SVM we now have to use it
- Given a new input $\mathbf{x}$ we decide on the class

$$y = \text{sgn}(\langle \mathbf{w}, \phi(\mathbf{x}) \rangle - b) \quad \text{but} \quad \mathbf{w} = \sum_{k=1}^m \alpha_k y_k \phi(\mathbf{x}_k)$$

- In the dual representation this becomes

$$\text{sgn} \left(\sum_{k=1}^m \alpha_k y_k K(\mathbf{x}_k, \mathbf{x}) - b \right)$$

where we only need to sum over the non-zero α_k (i.e. the support vectors SVs)

Recap on Eigen Systems

- Recall for a symmetric ($n \times n$) matrix $\mathbf{M}$ an eigenvector, $\mathbf{v}$

$$\mathbf{M}\mathbf{v} = \lambda\mathbf{v}$$


- There are n independent eigenvectors $\mathbf{v}^{(i)}$ with real eigenvalues $\lambda^{(i)}$
- The eigenvectors are orthogonal so that $\mathbf{v}^{(i)\top} \mathbf{v}^{(j)} = 0$ if $i \neq j$
- Forming a matrix of eigenvectors $\mathbf{V} = (\mathbf{v}^{(1)}, \mathbf{v}^{(2)}, \dots, \mathbf{v}^{(n)})$ the matrix satisfies

$$\mathbf{V}^\top \mathbf{V} = \mathbf{V} \mathbf{V}^\top = \mathbf{I}$$

- Such matrices are said to be orthogonal

Outline

- The Kernel Trick
- Positive Semi-Definite Kernels
- Kernel Properties
- Beyond Classification



Dual Form

- Recall that the dual problem for an SVM is

$$\max_{\alpha} \sum_{k=1}^m \alpha_k - \frac{1}{2} \sum_{k,l=1}^m \alpha_k \alpha_l y_k y_l \langle \phi(\mathbf{x}_k), \phi(\mathbf{x}_l) \rangle$$


- subject to $\sum_{k=1}^m y_k \alpha_k = 0$ and $0 \leq \alpha_k \leq C$
- But since $K(\mathbf{x}_k, \mathbf{x}_l) = \langle \phi(\mathbf{x}_k), \phi(\mathbf{x}_l) \rangle$ the dual problem becomes

$$\max_{\alpha} \sum_{k=1}^m \alpha_k - \frac{1}{2} \sum_{k,l=1}^m \alpha_k \alpha_l y_k y_l K(\mathbf{x}_k, \mathbf{x}_l)$$

- This is the **kernel trick**—we never have to compute $\phi(\mathbf{x})$

Outline

- The Kernel Trick
- Positive Semi-Definite Kernels
- Kernel Properties
- Beyond Classification



Eigen Decomposition

- From the eigenvalue equation $\mathbf{M}\mathbf{v}^{(k)} = \lambda^{(k)}\mathbf{v}^{(k)}$

$$\mathbf{M}\mathbf{V} = \mathbf{V}\mathbf{\Lambda} \quad \text{where} \quad \mathbf{\Lambda} = \begin{pmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{pmatrix}$$

- Multiplying on the right by $\mathbf{V}^\top$ we get

$$\mathbf{M} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top = \sum_{k=1}^n \lambda^{(k)} \mathbf{v}^{(k)} \mathbf{v}^{(k)\top}$$

Or

$$M_{ij} = \sum_{k=1}^n \lambda^{(k)} v_i^{(k)} v_j^{(k)} = \sum_{k=1}^n u_i^{(k)} u_j^{(k)} = \langle \mathbf{u}_i, \mathbf{u}_j \rangle$$

$$u_i^{(k)} = \sqrt{\lambda^{(k)}} v_i^{(k)}$$

Eigenfunctions

- By analogy for a symmetric function of two variables we can define an *eigenfunction*

$$\int K(x, y) \psi(y) dy = \lambda \psi(x)$$

- In general there will be a denumerable set of eigenfunctions $\psi^{(k)}(x)$ where


$$K(x, y) = \sum_k \lambda^{(k)} \psi^{(k)}(x) \psi^{(k)}(y)$$

- This is known as Mercer's theorem

Linear Operators

- Recall a linear function $\mathcal{T}[f(x)]$ can be represented by a kernel


$$\mathcal{T}[f(x)] = \int_{y \in \mathcal{I}} K(x, y) f(y) dy \approx \Delta \sum_{j=1}^n K(x, y_j) f(y_j)$$



This is just a matrix equation with $M_{ij} = \Delta K(x_i, y_j)$

Outline

- The Kernel Trick
- Positive Semi-Definite Kernels
- Kernel Properties
- Beyond Classification



Properties of Positive Semi-Definiteness

- Since


$$K(x, y) = \langle \phi(x), \phi(y) \rangle$$

- An immediate consequence is that for any function $f(x)$

$$\begin{aligned} \int f(x) K(x, y) f(y) dx dy &= \int f(x) \langle \phi(x), \phi(y) \rangle f(y) dx dy \\ &= \left\langle \int f(x) \phi(x) dx, \int f(y) \phi(y) dy \right\rangle \\ &= \left\| \int f(x) \phi(x) dx \right\|^2 \geq 0 \end{aligned}$$

Limit Process

- Consider sampling a function at a set of points



- In the limit where the number of sample points goes to infinity the vector more closely approximates a function
- The k^{th} component of the vector corresponds to $f(x_k)$

SVM Kernels

- If we define $\phi^{(k)}(x) = \sqrt{\lambda^{(k)}} \psi^{(k)}(x)$ then

$$K(x, y) = \sum_k \lambda^{(k)} \psi^{(k)}(x) \psi^{(k)}(y) = \sum_k \phi^{(k)}(x) \phi^{(k)}(y)$$

- This is the definition of a SVM kernel we started with
- Note that for $\phi^{(k)}(x)$ to be real $\lambda^{(k)} \geq 0$ for all k
- If $\lambda^{(k)} < 0$ then $\langle \phi(x), \phi(x) \rangle = \|\phi(x)\|^2$ might be negative and "distance" between points in the extended feature space can be negative
- If we use a kernel that isn't positive semi-definite then the Hessian of the dual objective function will not be negative semi-definite and there will be a maximum where α diverges

Positive Semi-Definite Kernels

- Kernels (or matrices) that have eigenvalues $\lambda^{(k)} \geq 0$ are called positive semi-definite
- (If the eigenvalues are strictly positive $\lambda^{(k)} > 0$ the kernels or matrices are called positive definite)
- Positive semi-definite kernels can always be decomposed into a sum of real functions

$$K(x, y) = \langle \phi(x), \phi(y) \rangle$$

Positive Semi-Definiteness

- The following statements are equivalent

- $K(x, y)$ is positive semi-definite (written $K(x, y) \succeq 0$)
- The eigenvalues of $K(x, y)$ are non-negative
- The kernel can be written

$$K(x, y) = \langle \phi(x), \phi(y) \rangle$$

where the $\phi^{(k)}(x)$'s are real functions

- For any real function $f(x)$

$$\int f(x) K(x, y) f(y) dx dy \geq 0$$

Adding Kernels

- We can construct SVM kernels from other kernels

- If $K_1(x, y)$ and $K_2(x, y)$ are valid kernels then so is $K_3(x, y) = K_1(x, y) + K_2(x, y)$

$$\begin{aligned} Q &= \int f(x) K_3(x, y) f(y) dx dy \\ &= \int f(x) (K_1(x, y) + K_2(x, y)) f(y) dx dy \\ &= \int f(x) K_1(x, y) f(y) dx dy + \int f(x) K_2(x, y) f(y) dx dy \geq 0 \end{aligned}$$

- If $K(x, y)$ is a valid kernel so is $cK(x, y)$ for $c > 0$

Exponentiating Kernels

- If $K(x, y)$ is a valid kernel so is $K^n(x, y)$ (by induction)
 - ★ Assume $K(x, y) \succeq 0$ this satisfies base case
 - ★ If $K^{n-1}(x, y) \succeq 0$ then

$$K^n(x, y) = K^{n-1}(x, y) K(x, y) \succeq 0$$

- and $\exp(K(x, y))$ is also a valid kernel since

$$e^{K(x, y)} = \sum_{i=1}^{\infty} \frac{1}{i!} K^i(x, y) = 1 + K(x, y) + \frac{1}{2} K^2(x, y) + \dots$$

but each term in the sum is a kernel

Other Kernels

- The success of SVMs has meant that researchers try to increase the area of application
- The condition that a SVM kernel must be positive semi-definite is quite restrictive
- There has been a cottage industry of researchers finding smart kernels for solving complicated problems
- The key to finding new kernels is to use the properties of kernels to build more complicated kernels from simpler ones

Spectrum Kernel

- A simple way to compare documents is to collect a histogram of all occurrences of substrings of length p
- This is known as a p -spectrum
- A p -spectrum kernel counts the number of common substrings

$$\begin{aligned} s &= \text{statistics} & \mathcal{S}_3(s) &= \{\text{sta}, \text{tat}, \text{ati}, \text{tis}, \text{ist}, \text{sti}, \text{tic}, \text{ics}\} \\ t &= \text{computation} & \mathcal{S}_3(t) &= \{\text{com}, \text{omp}, \text{mpu}, \text{put}, \text{uta}, \text{tat}, \text{ati}, \text{tio}, \text{ion}\} \end{aligned}$$

- $K(s, t) = 2$ ("tat" and "ati")

Product of Kernels

- If $K_1(x, y)$ and $K_2(x, y)$ are valid kernels then so is $K_3(x, y) = K_1(x, y) K_2(x, y)$

- Writing

$$K_1(x, y) = \sum_i \phi_i^{(1)}(x) \phi_i^{(1)}(y), \quad K_2(x, y) = \sum_j \phi_j^{(2)}(x) \phi_j^{(2)}(y)$$

then

$$\begin{aligned} K_3(x, y) &= \sum_{i,j} \phi_i^{(1)}(x) \phi_i^{(1)}(y) \phi_j^{(2)}(x) \phi_j^{(2)}(y) \\ &= \sum_{i,j} \phi_{ij}^{(3)}(x) \phi_{ij}^{(3)}(y) = \left\langle \phi^{(3)}(x), \phi^{(3)}(y) \right\rangle \end{aligned}$$

$$\text{where } \phi_{ij}^{(3)}(x) = \phi_i^{(1)}(x) \phi_j^{(2)}(x)$$

RBF Kernel

- Now $x^T y = \langle x, y \rangle$ is a valid kernel because it is an inner product of functions $\phi(x) = x$
- For $\gamma > 0$ we have $2\gamma x^T y \succeq 0$
- Thus $\exp(2\gamma x^T y) \succeq 0$
- If $K(x, y) \succeq 0$ then $g(x) K(x, y) g(y) \succeq 0$

$$\int f(x) g(x) K(x, y) g(y) f(y) dx dy = \int h(x) K(x, y) h(y) dx dy \succeq 0$$

where $f(x) g(x) = h(x)$

$$e^{-\gamma x^T x} e^{2\gamma x^T y} e^{-\gamma y^T y} = e^{-\gamma \|x-y\|^2} \succeq 0$$

String Kernels

- One area where SVMs were very important is in document classification
- This requires comparing strings
- There are a large number of kernels developed to do this

All Subsequences Kernel


- A more sophisticated kernel is to count all of the common subsequences that occur in two documents
- Naively this would take an exponential amount of time to compute
- Using clever dynamic-programming techniques this can be done relatively efficiently
- This can even be extended to include sub-sequence matches with possible gaps between words

- String kernels for comparing subsequences are used in bioinformatics
- Kernels have been developed for comparing trees (e.g. for computer program evaluation, XML, etc.)
- Kernels have also been developed for comparing graphs (e.g. for comparing chemicals based on their molecular graph)

- In an attempt to build kernels that capture more domain knowledge, kernels are constructed from other learning machines
- An example of this are “Fisher kernels” whose features come from an Hidden Markov Model (HMM) trained on the data
- These tend to have better discriminative power than the underlying model (HMM), and has a better feature set than a SVM using a generic kernel


Outline

1. The Kernel Trick
2. Positive Semi-Definite Kernels
3. Kernel Properties
4. Beyond Classification



Error Functions

- Can introduce slack variables with different errors




- This can be transformed to a quadratic programming problem

Kernel Methods

- Kernel methods where we project into an extended feature space are used with other linear algorithms
 - ★ Kernel Fisher discriminant analysis (KFDA)
 - ★ Kernel principle component analysis (KPCA)
 - ★ Kernel canonical correlation analysis (KCCA)
 - ★ Gaussian Processes
- These are also extremely powerful machine learning algorithms

Regression with Margins

- SVMs can be modified to perform regression



Ridge Regression Using Kernels

- We can also solve regression problems without using margins
- To solve a regression problem once again the problem is set up as a quadratic programming problem


$$\min_w \lambda \|w\|^2 + \sum_{i=1}^m (y_i - w^T \phi(x_i))^2$$

- the $\|w\|^2$ is a regularisation term
- By assuming $w = \sum_i \alpha_i \phi(x_i)$ we obtain a quadratic equation for the α_i 's which we can solve

Summary


- SVMs require a positive definite kernel function
- These can be built from simpler function
- There was a cottage industry of people creating new kernels for different application
- SVMs are just one example of a host of machine that
 - ★ use the kernel trick
 - ★ often use linear constraints
 - ★ tend to be convex optimisation problems

Wasserstein GANs



GANs, Wasserstein distance, Duality, WGANs


1. GANs
2. Wasserstein Distance
3. Wasserstein GANs



Generative Adversarial Networks

- One of the applications of Deep Learning that has most excited the public are **Generative Adversarial Networks** or **GANs**
- Their aim is to generate new random samples from the same distribution as some training set, $\mathcal{D}$
- Their number of real world applications are questionable
- But nobody cares because they are cool
- *Out of date warning: someone invented diffusion models*

How GANs Work



Training GANs


- The loss of the generator depends on its ability to trick the discriminator
- The loss of the discriminator depends on its ability not to be tricked
- We try to train the two networks simultaneously
- We hope that over time the generator produces better and better fakes

Problems of GANs

- GANs are notoriously difficult to train
- The generator and discriminator training can decouple
- Often the discriminator becomes too good at correctly identifying the generated images
- Then there can be little gradient information to help the generator as every small change in parameters doesn't significantly change the discriminator decision
- To try to solve this problem we first make a seemingly unconnected diversion

Outline

1. GANs
2. Wasserstein Distance
3. Wasserstein GANs



Measuring Distances Between Distributions

- In many machine learning tasks we want to minimise the distance between two probability distributions
- This requires that we can measure distances between probability distributions
- One prominent measure is the Kullback-Leibler or KL divergence

$$\text{KL}(p||q) = \int p(x) \log \left(\frac{p(x)}{q(y)} \right) dx$$


- This is very commonly used in ML (e.g. VAEs, Variational Approximation)

Trouble with KL

- KL-divergences are non-negative quantities that are minimised when the two probability distributions are the same
- They are not distances (they aren't symmetric and they don't satisfy the triangular inequality)

We don't really care about this, but what

- we do care about is that if $q(x) = 0$ when $p(x) \neq 0$ then $\log\left(\frac{p(x)}{q(y)}\right)$ diverges




- We can therefore have distributions that seem very similar but their KL-divergence is huge (or infinite)

Wasserstein Distance


- A more benign measure of the differences between two probability functions is the **Wasserstein** or **Earth Moving** distance

This is a true distance, but more importantly, it measures distance in a very natural way so that distributions that are close has a small Wasserstein distance



- Although this seems contrived if our probability distribution represents the probability of a 128×128 matrix of real valued triples represents an image of dog, then it is easy to imagine that the Wasserstein distance may be more benign than the KL-divergence

High Probability Manifold




Transportation Policy

- But how do we formalise the Wasserstein distance?
- A good place to start is to define a transportation policy $\gamma(x, y)$ with

$$\int \gamma(x, y) dy = p(x) \quad \int \gamma(x, y) dx = q(y)$$

- This looks like a joint probability distribution, but we interpret $\gamma(x, y)$ as the amount of probability mass/density that we transfer from $p(x)$ to $q(y)$

Transportation Policy



The Cost of Transport


- We want to choose the transportation policy that minimises the amount of probability mass we need to move
- Let $d(x, y) = \|x - y\|$ be a distance measure then the cost of a transportation policy is

$$C(\gamma) = \int \int d(x, y) \gamma(x, y) dx dy = \mathbb{E}_{\gamma}[d(x, y)]$$

where we interpret $\gamma(x, y)$ as a probability distribution

- Usually we take $d(x, y)$ to be the Euclidean distance, but we can choose any distance

Transportation Cost



The Wasserstein Distance

- The Wasserstein distance $W(p, q)$ between two probability distributions is defined as

$$W(p, q) = \min_{\gamma \in \Lambda(p, q)} \mathbb{E}_{\gamma}[d(x, y)]$$

- Where $\Lambda(p, q)$ is the set of joint distributions $\gamma(x, y)$ such that

$$\int \gamma(x, y) dy = p(x) \quad \int \gamma(x, y) dx = q(y)$$

- To compute the Wasserstein distance we have to solve a minimisation task!
- This looks nasty, but it is a (continuous) linear programming problem!
- Suppose p and q were discrete distribution (i.e. x and y only take discrete points)
- Then we could treat each value of $\gamma(x, y)$ as an element of a vector γ and each value of $d(x, y)$ as an element of a vector D
- Our objective is to choose γ to minimise $D^T \gamma$

$$\sum_j \gamma(x_i, y_j) = p(x_i) \quad \sum_i \gamma(x_i, y_j) = q(y_j)$$

$$A\gamma = P$$

$$\begin{pmatrix} \gamma(x_1, y_1) \\ \gamma(x_2, y_1) \\ \vdots \\ \gamma(x_n, y_1) \\ \gamma(x_1, y_2) \\ \gamma(x_2, y_2) \\ \vdots \\ \gamma(x_n, y_2) \\ \vdots \\ \gamma(x_1, y_n) \\ \gamma(x_2, y_n) \\ \vdots \\ \gamma(x_n, y_n) \end{pmatrix} = \begin{pmatrix} q(y_1) \\ q(y_2) \\ \vdots \\ q(y_n) \\ p(x_1) \\ p(x_2) \\ \vdots \\ p(x_n) \end{pmatrix}$$

Lagrange Formulation

- For discrete distributions

$$\min_{\gamma} D^T \gamma$$

subject to $A\gamma = P, \quad \gamma \geq 0$

- Writing the Lagrangian

$$\mathcal{L}(\gamma, \alpha) = D^T \gamma - \alpha^T (A^T \gamma - P)$$

where α is a vector of Lagrange multipliers

- The solution to the discrete optimisation problem is given by

$$\min_{\gamma} \max_{\alpha} \mathcal{L}(\gamma, \alpha)$$

Dual Form

- We can rearrange

$$\begin{aligned} \mathcal{L}(\gamma, \alpha) &= D^T \gamma - \alpha^T (A^T \gamma - P) \\ &= P^T \alpha - \gamma^T (A^T \alpha - D) \end{aligned}$$

- We note that $\gamma \geq 0$ so the dual problem is to find a vector α that maximises $P^T \alpha$ subject to the constraints $A^T \alpha \leq D$
- Although the vector form allows us to make connections with our earlier discussion of linear programming, it is a little difficult to interpret

Explicit Form

- We can write a Lagrangian for the original problem

$$\begin{aligned} \mathcal{L} &= \sum_{i,j} d(x_i, y_j) \gamma(x_i, y_j) - \sum_i \alpha(x_i) \left(\sum_j \gamma(x_i, y_j) - p(x_i) \right) \\ &\quad - \sum_j \beta(y_j) \left(\sum_i \gamma(x_i, y_j) - q(y_j) \right) \end{aligned}$$

subject to $\gamma(x_i, y_j) \geq 0$ where $\alpha(x_i)$ and $\beta(y_j)$ are Lagrange multipliers (they are components of α)

- Rearranging

$$\mathcal{L} = \sum_i \alpha(x_i) p(x_i) + \sum_j \beta(y_j) q(y_j) - \sum_{i,j} \gamma(x_i, y_j) (\alpha(x_i) + \beta(y_j) - d(x_i, y_j))$$

- This is equivalent to maximising $\sum_i \alpha(x_i) p(x_i) + \sum_j \beta(y_j) q(y_j)$, subject to

$$\forall i, j \quad \alpha(x_i) + \beta(y_j) \leq d(x_i, y_j)$$

Continuous Form

- We can write a Lagrangian for the continuous problem

$$\begin{aligned} \mathcal{L} &= \iint d(x, y) \gamma(x, y) dx dy - \int \alpha(x) \left(\int \gamma(x, y) dy - p(x) \right) dx \\ &\quad - \int \beta(y) \left(\int \gamma(x, y) dx - q(y) \right) dy \end{aligned}$$

subject to $\gamma(x, y) \geq 0$ where $\alpha(x)$ and $\beta(y)$ are Lagrange multiplier functions

- Rearranging

$$\mathcal{L} = \int \alpha(x) p(x) dx + \int \beta(y) q(y) dy - \iint \gamma(x, y) (\alpha(x) + \beta(y) - d(x, y)) dx dy$$

- This is equivalent to maximising $\int \alpha(x) p(x) dx + \int \beta(y) q(y) dy$, subject to

$$\alpha(x) + \beta(y) \leq d(x, y)$$

Dual Form Constraint

- We note that $\alpha(x) + \beta(y) \leq d(x, y)$ for all x and y
- This has to be true when $x = y$ so that

$$\alpha(x) + \beta(x) \leq d(x, x) = 0$$

- So $\beta(x) = -\alpha(x) - \epsilon(x)$ where $\epsilon(x) \geq 0$
- But want to maximise

$$\int \alpha(x) p(x) dx + \int \beta(y) q(y) dy = \int \alpha(x) (p(x) - q(x)) dx - \int q(x) \epsilon(x) dx$$

- This is maximised when $\epsilon(x) = 0$ i.e. $\beta(x) = -\alpha(x)$

Dual Form

- Thus the dual problem is to find a function $\alpha(x)$ —or a vector of functions $(\alpha(x_i))_i$ —that maximises

$$\int \alpha(x) (p(x) - q(x)) dx$$

- Subject to the constraint

$$\alpha(x) - \alpha(y) \leq d(x, y) = \|x - y\|$$


- This is a continuity constraint on the Lagrange multiplier function $\alpha(x)$ known as Lipschitz-1

Lipschitz-1 Functions

- We note for a Lipschitz-1 function and any unit vector u

$$u^T \nabla \alpha(x) = \lim_{\epsilon \rightarrow 0} \frac{\alpha(x + \epsilon u) - \alpha(x)}{\epsilon} \leq \frac{\epsilon \|u\|}{\epsilon} = \frac{\epsilon}{\epsilon} = 1$$

- That is, at every point the gradient in all directions must be less than 1 (since the gradient defines the direction of greatest increase it is both necessary and sufficient for $\|\nabla \alpha(x)\| \leq 1$ everywhere)



Calculating the Wasserstein Distance

- To recall the big picture we want to compute the Wasserstein distance

$$W(p, q) = \min_{\gamma \in \Lambda(p, q)} \mathbb{E}_{\gamma} [d(x, y)]$$


- For high dimensional objects $\gamma(x, y)$ would be a huge object to approximate
- Instead we can compute the Wasserstein distance in the dual formulation

$$W(p, q) = \max_{\alpha(x)} \int \alpha(x) (p(x) - q(x)) dx = \max_{\alpha} \mathbb{E}_p[\alpha(X)] - \mathbb{E}_q[\alpha(X)]$$

subject to the constraint that $\alpha(x)$ is a Lipschitz-1 function

Outline

- GANs
- Wasserstein Distance
- Wasserstein GANs



Back to GANs

- What has this got to do with GANs?
- Suppose we want to minimise the distance between the distribution $p(x)$ of real images (of which $\mathcal{D}$ are samples) and the distribution $q(x)$ of images drawn from a generator
- We can use a normal GAN generator, $G(z, w_G)$, that generates an image when given a random variable $z \sim \mathcal{N}(0, I)$
- To do this we choose the weights, w_G of the generator to minimise

$$W(p, q) = \max_{\alpha(x)} (\mathbb{E}_{x \sim p}[\alpha(x)] - \mathbb{E}_{x \sim q}[\alpha(x)])$$

Estimating Expectations

- Although we can't compute $\mathbb{E}_p[\alpha(x)]$ and $\mathbb{E}_q[\alpha(x)]$ exactly, we can estimate them from samples

$$\mathbb{E}_p[\alpha(x)] \approx \frac{1}{|\mathcal{B}|} \sum_{x \in \mathcal{B}} \alpha(x), \quad \mathbb{E}_q[\alpha(x)] \approx \frac{1}{n} \sum_{i=1}^n \alpha(G(z_i, w_G))$$

- where $\mathcal{B} \subset \mathcal{D}$ is a minibatch of true images and $z_i \sim \mathcal{N}(0, I)$
- From this we can choose w_G to minimise

$$C = \frac{1}{|\mathcal{B}|} \sum_{x \in \mathcal{B}} \alpha(x) - \frac{1}{n} \sum_{i=1}^n \alpha(G(z_i, w_G))$$


The Critic

- For this quantity to approximate the Wasserstein distance we need to find a function $\alpha(x, w_\alpha)$ that maximises C
- To do this we learn a second network, the critic or discriminator whose job it is to maximise

$$C = \frac{1}{|\mathcal{B}|} \sum_{x \in \mathcal{B}} \alpha(x, w_\alpha) - \frac{1}{n} \sum_{i=1}^n \alpha(G(z_i, w_G), w_\alpha)$$

- The network $\alpha(x, w_\alpha)$ should be Lipschitz-1 (which we usually botched by, for example, by setting the spectral norm of the convolutional weight matrix to 1)

Wasserstein GANs



Lesson

- Wasserstein GANs are, at least for me, one of the most elegant pieces of theory in recent years
- By trying to minimise the Wasserstein distance between the distribution of a generator and a true distribution we arrive at optimising two adversarial networks just like a GAN
- This uses a rather beautiful dual formulation
- It is claimed that W-GANs solve many of the problems of traditional GANs



Stuff

First

- Start

Second

-

Third

-

1. **First**
2. **Second**
3. **Third**



Outline

1. **First**
2. **Second**
3. **Third**



Outline

1. **First**
2. **Second**
3. **Third**





Stuff

First

- Start

Second

-

Third

-

1. **First**
2. **Second**
3. **Third**



Outline

1. **First**
2. **Second**
3. **Third**



Outline

1. **First**
2. **Second**
3. **Third**




Probability

$Y = g(X)$				Ω
$y_{13} = g(x_{13})$	$y_{14} = g(x_{14})$	$y_{15} = g(x_{15})$	$y_{16} = g(x_{16})$	
$y_9 = g(x_9)$	$y_{10} = g(x_{10})$	$y_{11} = g(x_{11})$	$y_{12} = g(x_{12})$	
$y_5 = g(x_5)$	$y_6 = g(x_6)$	$y_7 = g(x_7)$	$y_8 = g(x_8)$	
$y_1 = g(x_1)$	$y_2 = g(x_2)$	$y_3 = g(x_3)$	$y_4 = g(x_4)$	

Probability, Random Variables, Expectations


Modelling Uncertainty

- To model a world with uncertainty we consider some set of **elementary events** or **outcomes** Ω
- For the outcome of rolling a dice $\Omega = \{1, 2, 3, 4, 5, 6\}$
- The elementary events ω_i are **mutually exclusive** $\omega_i \cap \omega_j = \emptyset$ and **exhaustive** $\bigcup_i \omega_i = \Omega$
- We consider **events** $\mathcal{E} = \bigcup_{i \in \mathcal{I}} \omega_i$
- E.g. For a dice throw $\mathcal{E} = \{2, 4, 6\}$



Random Variables

- We can define a **random variable**, X , by partition the set of outcomes Ω and assign a numbers to each partition
 - E.g. for a dice
- $$X = \begin{cases} 0 & \text{if } \omega \in \{1, 3, 5\} \\ 1 & \text{if } \omega \in \{2, 4, 6\} \end{cases}$$
- $\mathbb{P}(X = x_i) = \mathbb{P}(\mathcal{E}_i)$ where $\mathcal{E}_i$ is the event that corresponding to the partition with value x_i



Function of Random Variables

- Any function, $Y = g(X)$, of a random variable, X , is a random variable


$Y = g(X)$				Ω
$y_{13} = g(x_{13})$	$y_{14} = g(x_{14})$	$y_{15} = g(x_{15})$	$y_{16} = g(x_{16})$	
$y_9 = g(x_9)$	$y_{10} = g(x_{10})$	$y_{11} = g(x_{11})$	$y_{12} = g(x_{12})$	
$y_5 = g(x_5)$	$y_6 = g(x_6)$	$y_7 = g(x_7)$	$y_8 = g(x_8)$	
$y_1 = g(x_1)$	$y_2 = g(x_2)$	$y_3 = g(x_3)$	$y_4 = g(x_4)$	

- Random Variables
- Expectations
- Calculus of Probabilities

Ω				
x_{31}	x_{32}	x_{33}	x_{34}	x_{35}
x_{26}	x_{27}	x_{28}	x_{29}	x_{30}
x_{21}	x_{22}	x_{23}	x_{24}	x_{25}
x_{16}	x_{17}	x_{18}	x_{19}	x_{20}
x_{11}	x_{12}	x_{13}	x_{14}	x_{15}
x_6	x_7	x_8	x_9	x_{10}
x_1	x_2	x_3	x_4	x_5

Probabilities

- We attribute a **probability**, $\mathbb{P}(\mathcal{E})$, to an event, $\mathcal{E}$, with the requirements
 - $0 \leq \mathbb{P}(\mathcal{E}) \leq 1$
 - $\mathbb{P}(\mathcal{E}) + \mathbb{P}(\neg \mathcal{E}) = 1$ where $\neg \mathcal{E} = \Omega \setminus \mathcal{E}$
- In some cases we can interpret $\mathbb{P}(\mathcal{E})$ as the expected frequency of occurrence of a repetitive trial
- But $\mathbb{P}(\text{Pass COMP6208 exam})$ is something you do once
- Can think of probability as an informed belief that something might happen
- When our knowledge changes the probability changes



What's In A Name

- We denote random variables with capital letters, X, Y, Z , etc.
- The symbol denote an object that can take one of a number of different values, but which one is still to be decided by chance
- When we write $\mathbb{P}(X)$ we can view this as short-hand for

$$(\mathbb{P}(X = x) \mid x \in \mathcal{X}) = (\mathbb{P}(X = x_1), \mathbb{P}(X = x_2), \dots, \mathbb{P}(X = x_n))$$
 where $\mathcal{X}$ is the set of possible values that X can take
- We treat random variables very differently to normal numbers (scalars) when we consider taking expectations

Continuous Spaces

- If the space of elementary events is continuous (e.g. for darts $x = (x, y)$) then $\mathbb{P}(X = x) = 0$
- But if we consider a region, $\mathcal{R}$, then we can assign a probability to landing in the region $\mathbb{P}(X \in \mathcal{R})$
- It is useful to work with **probability densities function** (PDF)

$$f_X(x) = \lim_{\epsilon \rightarrow 0} \frac{\mathbb{P}(X \in \mathcal{B}(x, \epsilon))}{\text{Vol}(\mathcal{B}(x, \epsilon))}$$



where $\mathcal{B}(x, \epsilon)$ is a ball of radius ϵ around the point x and $\text{Vol}(\mathcal{B}(x, \epsilon))$ is the volume of the ball

- If we make a change of variable the volume $\text{Vol}(\mathcal{B}(x, \epsilon))$ might change so $f_X(x)$ will change

Change of Variables

- Consider a region $\mathcal{R}$ —we can describe this using different coordinate systems x or $y = g(x)$

- But

$$\mathbb{P}(X \in \mathcal{R}) = \int_{\mathcal{R}} f_X(x) dx = \mathbb{P}(Y \in \mathcal{R}) = \int_{\mathcal{R}} f_Y(y) dy$$

- As this is true for any region $\mathcal{R}$: $f_X(x)|dx| = f_Y(y)|dy|$

- Or

$$f_X(x) = f_Y(y) \left| \frac{dy}{dx} \right| = f_Y(g(x)) |g'(x)|$$

- A probability density measured in units of probability per cm is different to that measured in units of probability per inch

Meaning of Probability Densities

- Probability densities are not probabilities
- They are positive, but don't need to be less than 1
- Note that

$$f_X(x) = \lim_{\delta x \rightarrow 0} \frac{\mathbb{P}(x \leq X < x + \delta x)}{\delta x}$$

- We can think of $f_X(x)\delta x$ as $\mathbb{P}(x \leq X < x + \delta x)$
- Note that $f_X(x)\delta x \leq 1$

Outline

- Random Variables
- Expectations
- Calculus of Probabilities

Ω

x_{31}	x_{32}	x_{33}	x_{34}	x_{35}
x_{26}	x_{27}	x_{28}	x_{29}	x_{30}
x_{21}	x_{22}	x_{23}	x_{24}	x_{25}
x_{16}	x_{17}	x_{18}	x_{19}	x_{20}
x_{11}	x_{12}	x_{13}	x_{14}	x_{15}
x_6	x_7	x_8	x_9	x_{10}
x_1	x_2	x_3	x_4	x_5

Linearity of Expectation

- Because sums and integrals are linear operators

$$\sum_i (ax_i + by_i) = a \left(\sum_i x_i \right) + b \left(\sum_i y_i \right)$$

$$\int (af(x) + bg(x)) dx = a \left(\int f(x) dx \right) + b \left(\int g(x) dx \right)$$

then expectations are linear

$$\mathbb{E}[aX + bY] = a\mathbb{E}[X] + b\mathbb{E}[Y]$$

- Beware usually $\mathbb{E}[XY] \neq \mathbb{E}[X]\mathbb{E}[Y]$ (unless X and Y are independent)

Jacobian

- In high dimension if we make a change of variables $x \rightarrow y(x)$ (which can be seen as a change of random variables $X \rightarrow Y(X)$)
- Then

$$f_X(x) = f_Y(y) |\det(\mathbf{J})|$$

where $\mathbf{J}$ is the Jacobian matrix

$$\mathbf{J} = \begin{pmatrix} \frac{\partial y_1}{\partial x_1} & \frac{\partial y_1}{\partial x_2} & \cdots & \frac{\partial y_1}{\partial x_n} \\ \frac{\partial y_2}{\partial x_1} & \frac{\partial y_2}{\partial x_2} & \cdots & \frac{\partial y_2}{\partial x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial y_n}{\partial x_1} & \frac{\partial y_n}{\partial x_2} & \cdots & \frac{\partial y_n}{\partial x_n} \end{pmatrix}$$

- Ensures integrals over volumes are the same

Cumulative Distribution Functions

- We can define the **cumulative distribution function** (CDF)

$$F_X(x) = \mathbb{P}(X \leq x) = \begin{cases} \sum_{i: x_i \leq x} \mathbb{P}(X = x_i) \\ \int_{-\infty}^x f_X(y) dy \end{cases}$$

- This is a function that goes from 0 to 1 as x goes from $-\infty$ to ∞
- We note that for continuous random variables

$$f_X(x) = \frac{dF_X(x)}{dx}$$

Expectation

- We can define the expectation of $Y = g(X)$ as

$$\mathbb{E}_X[g(X)] = \begin{cases} \sum_{x \in \mathcal{X}} g(x) \mathbb{P}(X = x) \\ \int g(x) f_X(x) dx \end{cases}$$

- The expectation of a constant c is

$$\mathbb{E}_X[c] = \begin{cases} \sum_{x \in \mathcal{X}} c \mathbb{P}(X = x) = c \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) = c \\ \int c f_X(x) dx = c \int f_X(x) dx = c \end{cases}$$

- Note $\mathbb{E}_X[\mathbb{E}_X[g(X)]] = \mathbb{E}_X[g(X)]$

Indicator Functions

- An indicator function has the property

$$\mathbb{I}_{\text{predicate}} = \begin{cases} 1 & \text{if predicate is True} \\ 0 & \text{if predicate is False} \end{cases}$$

(sometimes written $\mathbf{I}_A(x)$ where $A(x)$ is the predicate)

- We can obtain probabilities from expectations

$$\mathbb{P}(\text{predicate}) = \mathbb{E}[\mathbb{I}_{\text{predicate}}]$$

- E.g. The CDF is given by

$$F_X(x) = \mathbb{P}(X \leq x) = \mathbb{E}[\mathbb{I}_{X \leq x}]$$

1. Random Variables
2. Expectations
3. Calculus of Probabilities

Ω

x_{31}	x_{32}	x_{33}	x_{34}	x_{35}
x_{26}	x_{27}	x_{28}	x_{29}	x_{30}
x_{21}	x_{22}	x_{23}	x_{24}	x_{25}
x_{16}	x_{17}	x_{18}	x_{19}	x_{20}
x_{11}	x_{12}	x_{13}	x_{14}	x_{15}
x_6	x_7	x_8	x_9	x_{10}
x_1	x_2	x_3	x_4	x_5

Marginalisation

- Probabilities are extremely easy to manipulate (although lots of people struggle)
- One of the most useful properties is known as **marginalisation**

$$\mathbb{P}(X) = \sum_{y \in \mathcal{Y}} \mathbb{P}(X, Y = y)$$

where $\mathcal{Y}$ is the set of values that the random variable Y takes

- Note that when we write $\mathbb{P}(X)$ we are saying this is true for all values that X can take
- Although obvious and easy this is extremely useful

Basic Calculus

- To obtain the joint probability we can use

$$\mathbb{P}(X, Y) = \mathbb{P}(X|Y)\mathbb{P}(Y) = \mathbb{P}(Y|X)\mathbb{P}(X)$$

- This generalises to more random variables

$$\mathbb{P}(X, Y, Z) = \mathbb{P}(X, Y|Z)\mathbb{P}(Z) = \mathbb{P}(X|Y, Z)\mathbb{P}(Y|Z)\mathbb{P}(Z)$$

- We can do this in a number of different ways

$$\mathbb{P}(X, Y, Z) = \mathbb{P}(Y, Z|X)\mathbb{P}(X) = \mathbb{P}(Z|Y, X)\mathbb{P}(Y|X)\mathbb{P}(X)$$

- Note that $\mathbb{P}(A, B | X, Y)$ means the probability of random variables A and B given that X and Y take particular values

Causality

- Conditional probabilities does not imply causality
- We might have causal relationships

$$\mathbb{P}(\text{pass} | \text{study}) = 0.9 \quad \mathbb{P}(\text{pass} | \neg \text{study}) = 0.2$$

- But if we know $\mathbb{P}(\text{study}) = 0.8$ then we can compute

$$\mathbb{P}(\text{pass}, \text{study}) = \mathbb{P}(\text{pass} | \text{study})\mathbb{P}(\text{study}) = 0.9 \times 0.8 = 0.72$$

$$\mathbb{P}(\text{pass}, \neg \text{study}) = \mathbb{P}(\text{pass} | \neg \text{study})\mathbb{P}(\neg \text{study}) = 0.2 \times 0.2 = 0.04$$

and

$$\begin{aligned} \mathbb{P}(\text{study} | \text{pass}) &= \frac{\mathbb{P}(\text{pass}, \text{study})}{\mathbb{P}(\text{pass})} \\ &= \frac{\mathbb{P}(\text{pass}, \text{study})}{\mathbb{P}(\text{pass}, \text{study}) + \mathbb{P}(\text{pass}, \neg \text{study})} = \frac{0.72}{0.72 + 0.04} \approx 0.947 \end{aligned}$$

- Often we want to model complex processes where we have multiple random variables

- We can define the joint probability

$$p_{X,Y}(x, y) = \mathbb{P}(X = x, Y = y)$$

i.e. the probability of the event where both $X = x$ and $Y = y$

- Clearly $\mathbb{P}(X, Y) = \mathbb{P}(Y, X)$

Conditional Probability

- We can also define the probability of an event X given that $Y = y$ has occurred

$$\mathbb{P}(X | Y = y) = \frac{\mathbb{P}(X, Y = y)}{\mathbb{P}(Y = y)}$$

- In constructing a model it is often much easier to specify conditional probabilities (because you know something) rather than joint probabilities
- When manipulating probabilities it is often easier to work with joint probabilities because we can simplify them by marginalising out random variables we are not interested in

Beware

- Conditional probabilities, $\mathbb{P}(X | Y)$ are probabilities for X , but not Y

$$\sum_{x \in \mathcal{X}} \mathbb{P}(X = x | Y) = 1$$

$$\sum_{y \in \mathcal{Y}} \mathbb{P}(X | Y = y) \neq 1$$

(in general)

- Note that

$$\begin{aligned} \mathbb{E}_Y[\mathbb{P}(X | Y)] &= \sum_{y \in \mathcal{Y}} \mathbb{P}(Y = y) \mathbb{P}(X | Y = y) \\ &= \sum_{y \in \mathcal{Y}} \mathbb{P}(X, Y = y) = \mathbb{P}(X) \end{aligned}$$

Independence

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

$$\mathbb{P}(X, Y) = \mathbb{P}(X)\mathbb{P}(Y)$$

- Because $\mathbb{P}(X, Y) = \mathbb{P}(X|Y)\mathbb{P}(Y)$ and $\mathbb{P}(X, Y) = \mathbb{P}(Y|X)\mathbb{P}(X)$ independence implies

$$\mathbb{P}(X|Y) = \mathbb{P}(X) \quad \mathbb{P}(Y|X) = \mathbb{P}(Y)$$

- Probabilistic independence implies a mathematical co-incident not necessarily causal independence
- However causal independence implies probabilistic independence
- If $X \in \{0, 1\}$ represents the outcome of tossing a coin and $Y \in \{1, 2, 3, 4, 5, 6\}$ the outcome of rolling a dice then X and Y are independent

Well Conducted Experiments

- In well conducted experiments we expect the results we obtain are independent
- Let $\mathcal{D} = (X_1, X_2, \dots, X_m)$ represents possible outcomes from a set of m well conducted experiments then

$$\mathbb{P}(\mathcal{D}) = \prod_{i=1}^m \mathbb{P}(X_i)$$

- Denoting a possible sentence I might say by $S = (W_1, W_2, \dots, W_m)$ then

$$\mathbb{P}(S) \neq \prod_{i=1}^m \mathbb{P}(W_i)$$

otherwise it's time I retired

Conclusion

- To work with probabilities you need to know
 - ★ How to go back and forward between joint probabilities and conditional probabilities
 - ★ How to marginalise out variables
- You need to understand that for continuous outcomes, it makes sense to talk about the probability density
- You need to know that expectations are linear operators and the expectation of a constant is the constant
- You need to understand independence

Conditional Independence

- Let $K(d)$ be a random variable measuring the amount you know about ML on day d of your revision
- From you revision schedule you can write down your belief


$$\mathbb{P}(K(d) \mid K(d-1), K(d-2), \dots, K(1))$$

- But a very reasonable model is

$$\mathbb{P}(K(d) \mid K(d-1), K(d-2), \dots, K(1)) = \mathbb{P}(K(d) \mid K(d-1))$$

what you are going to know today will just depend on what you knew yesterday

- We say that $K(d)$ is **conditionally independent** on $K(d-2)$, $K(d-3)$, etc. given $K(d-1)$



interpreting data, standard errors, confidence tests

Experiments and Reproducibility


- One of the fundamental requirements of science is that experiments can be reproduced
- This applies to machine learning and deep-learning as it does to particle physics, molecular biology, etc.
- We need to describe completely what procedures were carried out
- However, experiments have natural variability so what do we mean by repeatability?

Describing Your Experiments


- You need to tell the reader
 - ★ What your deep-learning architecture looks like
 - ★ Your training set
 - ★ How you partitioned your training set into train, validation and test sets
 - ★ Any preprocessing you performed
 - ★ The optimiser you used, mini-batch size, learning rate
 - ★ The number of epochs you trained for
 - ★ Any hyper-parameter tuning you carried out
 - ★ How long it took in clock time on the platform you are using
- You can put some of this into an appendix if necessary

Outline

1. Outline
2. Experimental Error
3. Confidence Estimates



1. Outline
2. Experimental Error
3. Confidence Estimates



Describing Research

- When describing research you need to find a story (narrative)—why is the work you are doing interesting
- You need to give a general motivation as well as a more detailed explanation of what you are specifically going to test
- You then need to give enough details that someone can take your report/paper and reproduce what you are doing
- These days it is also common practice to put the model you are using on-line—but you should still describe everything in your report

Results

- A paper builds a story (thesis) based on experimental evidences
- It is not enough to present results, it is also your job to interpret the results
- Your results provide evidence in support of your thesis
- But results all come with uncertainty. If you tested on a different test set, or used a different training set then the answer would be different
- Results are meaningless unless the uncertainty is quantified

Experimental Error

- Any empirical result is useless without some estimate of error
- If you tell me algorithm A has a performance of 95% and algorithm B has a performance of 97%, how do I know that this didn't happen by chance?
- Maybe algorithm A is actually better, but you got lucky (or unlucky)
- We need to get a handle on the estimated experimental error

- We assume that our test set consists of typical examples that we might encounter and that they are independent (e.g. we are not seeing the same example over and over again)■
- These are assumptions that are not always true■
- We need to think about it and if they are not true you may need to collect new data or be very cautious of the claims that you are making■
- However, if the person curating the training set did a good job then usually these are reasonable assumptions■

Variance in Mean Estimator

- Consider the mean squared error (MSE) of the estimator $\hat{\mu}_n$

$$\begin{aligned}\mathbb{E}[(\hat{\mu}_n - \mu)^2] &= \mathbb{E}\left[\left(\frac{1}{n}\sum_{i=1}^n X_i - \mu\right)^2\right] = \mathbb{E}\left[\left(\frac{1}{n}\sum_{i=1}^n (X_i - \mu)\right)^2\right] \\ &= \frac{1}{n^2} \mathbb{E}\left[\left(\sum_{i=1}^n (X_i - \mu)\right)\left(\sum_{j=1}^n (X_j - \mu)\right)\right] \\ &= \frac{1}{n^2} \mathbb{E}\left[\sum_{i=1}^n (X_i - \mu)^2 + \sum_{i=1}^n \sum_{j=1, j \neq i}^n (X_i - \mu)(X_j - \mu)\right] \\ &= \frac{1}{n^2} \sum_{i=1}^n \left(\mathbb{E}[(X_i - \mu)^2] + \sum_{j=1, j \neq i}^n \mathbb{E}[X_i - \mu] \mathbb{E}[X_j - \mu]\right) \\ &= \frac{1}{n^2} \sum_{i=1}^n \sigma^2 = \frac{1}{n} \sigma^2\end{aligned}$$

Repeated Bernoulli Trials

- Suppose we are measuring a binary outcome, success or failure■
- Assuming the trials are independent (so called Bernoulli trials) then the only unknown is the probability, p , of a success■
- If we make n trials then probability of k successes is given by the binomial distribution

$$\mathbb{P}(K = k) = \binom{n}{k} p^k (1-p)^{n-k}$$

- The mean and variance of K are

$$\mathbb{E}[K] = np \quad \mathbb{E}[(K - np)^2] = np(1-p)$$

Standard Error for Bernoulli Trials

- The estimated error in the mean is given by the root mean squared error (RMSE) $\sqrt{\mathbb{E}[(\hat{p} - p)^2]} = \sqrt{\frac{p(1-p)}{n}}$ ■
- We don't know p , but we can use an estimate, $\hat{p} = \frac{K}{n}$

$$\Delta \approx \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

- We can also use a worst-case bound (when $p = \frac{1}{2}$)

$$\Delta \approx \frac{1}{2\sqrt{n}}$$

This is pessimistic if p is small, although the previous estimate is optimistic if $\hat{p} = 0$ ■

- Given some experimental measurements $\{X_i | i = 1, 2, \dots, n\}$ we can estimate the mean and variance

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n X_i \quad \hat{\sigma}_n^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - \hat{\mu}_n)^2$$

- These are not the true mean, $\mu = \mathbb{E}[X_i]$, and true variance $\sigma^2 = \mathbb{E}[(X_i - \mu)^2]$, but they are unbiased estimates in that

$$\mathbb{E}[\hat{\mu}_n] = \mu \quad \mathbb{E}[\hat{\sigma}_n^2] = \sigma^2$$

Estimated Error in the Mean

- The root mean squared error (RMSE) is equal to

$$\sqrt{\mathbb{E}[(\hat{\mu}_n - \mu)^2]} = \frac{\sigma}{\sqrt{n}}$$

- The **estimated error in the mean** or **standard error** is defined as $\Delta = \frac{\hat{\sigma}_n}{\sqrt{n}}$ where $\hat{\sigma}_n$ is the empirical estimate of the standard deviation■
- Traditionally we would write an experimental result as $\hat{\mu}_n \pm \Delta$ ■
- Also we only quote $\hat{\mu}_n$ up to the number of significant figures consistent with Δ ■
- Not 1.6582356 ± 0.122543 but 1.7 ± 0.1 ■

Mean Squared Error

- For the case of n Bernoulli trials we can estimate the mean success rate


$$\hat{p} = \frac{K}{n}$$

- The mean squared error (MSE) in this estimate is


$$\begin{aligned}\mathbb{E}[(\hat{p} - p)^2] &= \mathbb{E}\left[\left(\frac{K}{n} - p\right)^2\right] = \frac{1}{n^2} \mathbb{E}[(K - np)^2] \\ &= \frac{1}{n^2} np(1-p) = \frac{p(1-p)}{n}\end{aligned}$$

Errorbars

- There are different conventions on errorbars■
- Traditionally errorbars show the estimated error in the mean (standard error, $\Delta = \hat{\sigma}_n / \sqrt{n}$)■
- This means that about 30% of the time the true mean would be outside the errorbar■
- Some people use 95% confident intervals as errorbars, saying 95% of the time the true mean lies within the errorbars—this corresponds to errorbars showing twice the standard error■
- Some people show the standard deviation as errorbars■ this tells you nothing about the confidence of the point you plot—in my opinion it is just wrong■




1. Outline
2. Experimental Error
3. Confidence Estimates



Know What's Best

- How do we know if model B is better than model A?



The more data we have the less likely we will make a mistake

Quantifying Uncertainty

- We want a way of quantifying the chance that model B is better than model A
- To do this we can consider a **null hypothesis**
 - We suppose that model A and B have the same performance
- We then ask, under the null hypothesis, what is the chance of obtaining a gap in performance at least as big as the one we observed
- This will depend on the number of observations and the empirical variance

T-tests

- The original t-test was invented by William Gosset who worked for the Guinness brewery in Dublin, but published under the pseudonym *Student*
- He worked with the null hypothesis that the distributions for A and B are normal with the same means and standard deviations
- The slightly more common test is Welch's t-test where the null hypothesis is that the two datasets are normally distributed with the same means but different standard deviations
- The gap is equal to the t value

$$t = \frac{\hat{\mu}_A - \hat{\mu}_B}{\hat{\sigma}_{A,B}}$$

Statistical Significance

- It is common to say that A is *statistically significantly* better than B if the probability of the observed t value occurring given the null hypothesis is less than 0.05
- Statistically significant difference in performance does not mean that the difference is anything to care about (that is a judgement call)
- We are assuming the data is normally distributed. We can check but no one does. Usually its not a bad approximation if the data is not normal but there are exceptions. . .

Welch's T-Test

- For Welch's t-test


$$t = \frac{\hat{\mu}_A - \hat{\mu}_B}{\hat{\sigma}_{A,B}} \quad \hat{\sigma}_{A,B} = \sqrt{\frac{\hat{\sigma}_A^2}{n} + \frac{\hat{\sigma}_B^2}{m}}$$


- Where p depends on t and the *number of degrees of freedom*

$$\nu = \frac{\hat{\sigma}_{A,B}^2}{(\hat{\sigma}_A^2/n)^2/(n-1) + (\hat{\sigma}_B^2/m)^2/(m-1)}$$

- The beauty of the t -test is that from t and ν we can compute p
- Small p is good. It shows the null hypothesis is unlikely

Example





- If you have enough data points then when the errorbars (1 standard error) are separate then the two distributions are significantly different

- Sometimes the fluctuations in performance will come both from the test example and from the difference in performance
- For example model A might always be better than model B, but the performance fluctuates depending on the example
- A normal t-test might not show any significant difference
- Using a paired t-test we instead use


$$\hat{\sigma}_{X+Y}^2 = \frac{\hat{\sigma}_X^2 + \hat{\sigma}_Y^2 - \hat{c}_{X,Y}}{n}$$

$$\hat{c}_{X,Y} = \frac{1}{n-1} \sum_{i=1}^n (X_i - \hat{\mu}_X)(Y_i - \hat{\mu}_Y)$$

Paired T-Test Example

$$X \sim Z + X_1 \quad Y \sim Z + Y_1$$

$$Z \sim \mathcal{N}(0, 0.8) \quad X_1 \sim \mathcal{N}(0.4, 0.2) \quad Y_1 \sim \mathcal{N}(0.6, 0.2)$$



Cheating Significant Tests

- You should only run your t-test once
- If you run multiple t-tests until you get one with $p < 0.05$ then you have cheated
- With a p value of 0.05 you expect to find the observed t value given the null hypothesis once every 20 trials—if you do the t-test twice you should expect this to happen every 10 trials (i.e. $p = 0.1$)
- If you run t -tests until you find a p value of 0.05 you are cheating
- If you want to compare many different models then you should use ANOVA instead

Computing Significance Ab-Initio


- The t -test results follow from a beautiful, but difficult calculation
- However, you can simulate the null-hypothesis multiple times and measure empirically the number of times the t -value exceeds the observed t -values
- This isn't necessary under the assumptions that the data is normally distributed (done by Student and Welch)
- However, there are other times where you want to show that what you have found has not happened by chance. Often it is easy to simulate a null hypothesis

Lessons

- You need to understand the errors in your data
- Computing the standard error is always necessary
- In some cases you can quantify your certainty that A is better than B using statistical confidence tests— t -tests (other tests are available)
- You need to apply common sense in interpreting data
- The person you cheat when you don't do this is yourself

Good, Bad and the Ugly

- What I've told you is true of data that follows well behaved distributions
- Just occasionally you come across data that has long (thick) tails



- Here taking the average of many samples might not give you reliable estimates of the mean
- This isn't that common, but when nothing seems to make sense plot a histogram of the data and check for thick tails




Bayes, Conjugate Priors, Uninformative Priors

Dealing with Uncertainty

- In machine learning we are attempting to make inference under uncertainty
- The natural language for discussing uncertainty is probability
- The natural framework for making inferences is Bayesian statistics
- However, this requires that we encode our prior knowledge of the problem and specify a likelihood
- In consequence, probabilistic methods tend to be bespoke, rather than general purpose black boxes

Solving Inverse Problems



- We want the posterior $\mathbb{P}(\mathcal{H}_i|\mathcal{D})$ (i.e. the probability of what happened given some evidence)
- The Bayesian formalism converts this into the forward problem

$$\mathbb{P}(\mathcal{H}_i|\mathcal{D}) = \frac{\mathbb{P}(\mathcal{D}|\mathcal{H}_i)\mathbb{P}(\mathcal{H}_i)}{\mathbb{P}(\mathcal{D})}$$

Evidence

- The normalisation term

$$\mathbb{P}(\mathcal{D}) = \sum_{i=1}^n \mathbb{P}(\mathcal{H}_i, \mathcal{D}) = \sum_{i=1}^n \mathbb{P}(\mathcal{D}|\mathcal{H}_i)\mathbb{P}(\mathcal{H}_i)$$

tells you how likely the data is (given the prior and likelihood function)

- It is called the **marginal likelihood** or **evidence**
- At one level this is just a normalisation condition which is often a pain to compute
- But it also plays an important role in **model selection**



1. Bayes' Rule
2. Conjugate Priors
3. Uninformative Priors

Revision on Bayes

- Bayes' rule

$$\mathbb{P}(\mathcal{H}_i|\mathcal{D}) = \frac{\mathbb{P}(\mathcal{D}|\mathcal{H}_i)\mathbb{P}(\mathcal{H}_i)}{\mathbb{P}(\mathcal{D})}$$

- ★ $\mathbb{P}(\mathcal{H}_i|\mathcal{D})$ is the **posterior** probability of a hypothesis $\mathcal{H}_i$ (i.e. the probability of $\mathcal{H}_i$ after we see the data)
- ★ $\mathbb{P}(\mathcal{D}|\mathcal{H}_i)$ is the **likelihood** of the data given the hypothesis. Note, that we calculated this from the forward problem
- ★ $\mathbb{P}(\mathcal{H}_i)$ is the **prior** probability (i.e. the probability of $\mathcal{H}_i$ before we see the data)
- ★ $\mathbb{P}(\mathcal{D})$ is the **evidence** or **marginal likelihood**

$$\mathbb{P}(\mathcal{D}) = \sum_{i=1}^n \mathbb{P}(\mathcal{H}_i, \mathcal{D}) = \sum_{i=1}^n \mathbb{P}(\mathcal{D}|\mathcal{H}_i)\mathbb{P}(\mathcal{H}_i)$$

Bayesian Inference

- Bayes' rule says $\mathbb{P}(\mathcal{H}_i|\mathcal{D}) = \frac{\mathbb{P}(\mathcal{D}|\mathcal{H}_i)\mathbb{P}(\mathcal{H}_i)}{\mathbb{P}(\mathcal{D})}$
- We calculate the likelihood $\mathbb{P}(\mathcal{D}|\mathcal{H}_i)$ (i.e. assuming the hypothesis, what is the chance of obtaining the data?)
- We consider the process of how the data is generated
- This uses the data we have (doesn't care about missing data)
- But we also need to know the prior $\mathbb{P}(\mathcal{H}_i)$
- Also, this can get difficult when we have many hypotheses

Model Selection

- Suppose we have two models of the likelihood or prior; M_1 and M_2 . Then the posterior will depend on the model

$$\mathbb{P}(\mathcal{H} | \mathcal{D}, M_i) = \frac{\mathbb{P}(\mathcal{D} | \mathcal{H}, M_i)\mathbb{P}(\mathcal{H} | M_i)}{\mathbb{P}(\mathcal{D} | M_i)}$$

- We can use $\mathbb{P}(\mathcal{D} | M_1)$ or $\mathbb{P}(\mathcal{D} | M_2)$ to select which model is most likely
- E.g. is my likelihood best captured by a normal or gamma distribution
- This also allows us to select hyperparameters for a model

Probability Density

- When we are working with continuous variables it is more natural to work with probability densities

$$f_X(x) = \lim_{\delta x \rightarrow 0} \frac{\mathbb{P}(x \leq X < x + \delta x)}{\delta x}$$

- Note that densities are non-negative, but can be greater than 1 (they are not probabilities)
- However

$$\mathbb{P}(a \leq X \leq b) = \int_a^b f_X(x) dx$$

is a probability and is less than or equal to 1

Practical Bayesian Inference

- Often consider learning parameters θ

$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta)p(\theta)}{p(\mathcal{D})}$$

- This can be hard for large data sets as the posterior, $p(\theta|\mathcal{D})$, is often a mess
- If we are lucky and have a simple likelihood then if we choose the right prior we end up with a posterior of the same form as the prior
- This occurs in some classic probabilistic inference problems, but as we will see soon it is also true for Gaussian Processes

Learning a Probability

- Suppose we have a coin and we want to establish the probability of a head
- We want to learn this from a series of independent trials
- (Independent trials with two possible outcomes are known in probability theory as Bernoulli trials)
- Let X_i equal 1 if the i^{th} trial is a head and 0 otherwise
- If the probability of a head is p then the **likelihood** of a X_i is

$$\mathbb{P}(X_i|p) = p^{X_i}(1-p)^{1-X_i} = \begin{cases} p & \text{if } X_i = 1 \\ (1-p) & \text{if } X_i = 0 \end{cases}$$

Uninformative Prior

- Suppose we have no idea about p what should we do?
- Laplace (one of the first Bayesian's) suggested giving equal weighting to all values of p
- This corresponds to a beta distribution with $a = b = 1$
- (Surprisingly other arguments suggest using $a = b = 0$ which provides a strong bias towards $p = 0$ and $p = 1$)
- Given enough data the prior is not so important and we will stick with Laplace for now

Densities and Bayes

- Bayes' rule also applies to densities

$$\mathbb{P}(x \leq X < x + \delta x|Y) = \frac{\mathbb{P}(Y|x)\mathbb{P}(x \leq X < x + \delta x)}{\mathbb{P}(Y)}$$

- Dividing by δx and taking the limit $\delta x \rightarrow 0$

$$f_{X|Y}(x|Y) = \frac{\mathbb{P}(Y|x)f_X(x)}{\mathbb{P}(Y)}$$

- Similarly if X is discrete and Y continuous

$$\mathbb{P}(X|y) = \frac{f_{Y|X}(y|X)\mathbb{P}(X)}{f_Y(y)}$$

- If both X and Y are continuous

$$f_{X|Y}(x|y) = \frac{f_{Y|X}(y|x)f_X(x)}{f_Y(y)}$$

Outline

- Bayes' Rule
- Conjugate Priors
- Uninformative Priors



Prior

- We may have a prior belief (e.g. we have made a few trials or we see the coin looks like a normal penny)
- We will suppose we can model our prior belief in terms of a **Beta distribution**

$$f(p) = \text{Beta}(p|a,b) = \frac{p^{a-1}(1-p)^{b-1}}{B(a,b)}$$

- $B(a,b)$ is just a normalisation constant

$$B(a,b) = \int_0^1 p^{a-1}(1-p)^{b-1} dp = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}$$

- This is a useful function for modelling the distribution of a random variable in the range 0 to 1

Independent Trials

- Using Bayes' rule

$$f(p|\mathcal{D}) = \frac{\mathbb{P}(\mathcal{D}|p)f(p)}{\mathbb{P}(\mathcal{D})}$$

- Assuming the trials are independent (a reasonably fair assumption for tossing coins) then the likelihood factorises

$$\begin{aligned} \mathbb{P}(\mathcal{D}|p) &= \prod_{i=1}^n p^{X_i}(1-p)^{1-X_i} \\ &= p^{X_1}(1-p)^{1-X_1} p^{X_2}(1-p)^{1-X_2} \dots p^{X_n}(1-p)^{1-X_n} \\ &= p^{\sum_i X_i} (1-p)^{\sum_i (1-X_i)} = p^s (1-p)^{n-s} \end{aligned}$$

$$s = \sum_i X_i \text{ (number of successes/heads)}$$

Posterior

- Plugging in a prior $f(p) = \text{Beta}(p|a_0, b_0)$

$$f(p|\mathcal{D}) = \frac{\mathbb{P}(\mathcal{D}|p)f(p)}{\mathbb{P}(\mathcal{D})} = \frac{p^s(1-p)^{n-s} \times p^{a_0-1}(1-p)^{b_0-1}}{\mathbb{P}(\mathcal{D})B(a_0, b_0)}$$

- The denominator is a normalising factor

$$\begin{aligned}\mathbb{P}(\mathcal{D}) &= \int_0^1 \mathbb{P}(\mathcal{D}|p)f(p)dp = \int_0^1 \frac{p^{s+a_0-1}(1-p)^{n-s+b_0-1}}{B(a_0, b_0)}dp \\ &= \frac{B(s+a_0, n-s+b_0)}{B(a_0, b_0)}\end{aligned}$$

Incremental Updating

- For independent data we can update incrementally
 $\mathcal{D} = (X_1, X_2, \dots, X_n)$

$$\begin{aligned}f(p|X_1) &= \frac{\mathbb{P}(X_1|p)f(p)}{\mathbb{P}(X_1)} \\ f(p|X_1, X_2) &= \frac{\mathbb{P}(X_2|p)f(p|X_1)}{\mathbb{P}(X_2)} \\ &\vdots \\ f(p|X_1, X_2, \dots, X_n) &= \frac{\mathbb{P}(X_n|p)f(p|X_1, \dots, X_{n-1})}{\mathbb{P}(X_n)}\end{aligned}$$

- The posterior becomes the prior for the next piece of data
- For our problem the posterior is always Beta distributed

Estimating Prediction Errors

- A full Bayesian treatment gives a prediction of its own error
- Assuming $f(p|\mathcal{D}) = \text{Beta}(p|a, b)$
- The expected value of p is given by $a/(a+b) = 23/32 = 0.72$
- The standard deviation is

$$\sqrt{\frac{ab}{(a+b)^2(a+b+1)}} = 0.08$$

Using Bayes

- Let us assume a Gamma distributed prior

$$p(\mu) = \Gamma(\mu|a_0, b_0) = \frac{b_0^{a_0} \mu^{a_0-1} e^{-b_0\mu}}{\Gamma(a_0)}$$

- We will assume that we know nothing. The uninformative prior is $a_0 = b_0 = 0$
- The data is $\mathcal{D} = \{N_1, N_2, \dots, N_n\}$
- The likelihood is $\text{Pois}(N_i|\mu)$

Conjugate Priors


- The posterior distribution is Beta distribution

$$f(p|\mathcal{D}) = \frac{p^{s+a_0-1}(1-p)^{n-s+b_0-1}}{B(s+a_0, n-s+b_0)} = \text{Beta}(p|s+a_0, n-s+b_0)$$

- Something rather nice happened
- Starting with a beta distributed prior $f(p) = \text{Beta}(p|a_0, b_0)$ for a set of Bernoulli trials we obtain a beta distributed posterior $f(p|\mathcal{D}) = \text{Beta}(p|a_0+s, b_0+n-s)$
- This is not always the case (often the posterior will be very complicated) but it happens for a few likelihoods and priors
- When the posterior is the same as the prior then the likelihood and prior distributions are said to be **conjugate**

Example (p=0.7)

$$\mathcal{D} = \{H, H, H, H, T, T, T, H, H, H, H, T, H, T, H, H, H, H, H, H, T, H, H, H, H, T, H, H, H, T\}$$



Poisson Likelihoods

- Let's look at a second example of conjugate priors
- Suppose we want to find the rate of traffic along a road between 1:00pm and 2:00pm
- We assume the number of cars is given by a Poisson distribution

$$\mathbb{P}(N) = \text{Pois}(N|\mu) = \frac{\mu^N}{N!} e^{-\mu}$$


- μ is the rate of traffic per hour which we want to infer from observation taken on different days

Posterior

- The posterior after seeing the first piece of data is

$$\begin{aligned}p(\mu|N_1) &\propto \mathbb{P}(N_1|\mu)p(\mu) \\ &\propto \frac{\mu^{N_1}}{N_1!} e^{-\mu} \mu^{a_0-1} e^{-b_0\mu} \\ &\propto \mu^{N_1+a_0-1} e^{-(b_0+1)\mu}\end{aligned}$$

- The posterior is also a Gamma distribution $\Gamma(\mu|a_1, b_1)$ with $a_1 = a_0 + N_1$, $b_1 = b_0 + 1$



Uninformative Priors

- What if we have no prior knowledge, what should we do?
- OK usually we know whether we should make a measurement using a micrometer, ruler or car mileage, but we might still know almost nothing
- This led to Bayesian statistics being labelled as *subjective*
- However Ed. Jaynes (the greatest proponent of Bayesian methods) argued that we could answer this using symmetry arguments

Benford's Law

- Numbers occurring in life (physical constants, amounts of money) should not depend on the units (scale) measuring them
- They should then be distributed as $p(x) \propto 1/x$
- A curious consequence of this is that the significant figure has a distribution

$$\begin{aligned} \mathbb{P}(\text{most s.f. of } x = n) &= \frac{\int_n^{n+1} \frac{1}{x} dx}{\int_1^{10} \frac{1}{x} dx} = \frac{\int_{10^n}^{10^{n+1}} \frac{1}{x} dx}{\int_{10^0}^{10^{100}} \frac{1}{x} dx} \\ &= \frac{\log(n+1) - \log(n)}{\log(10)} = \log_{10}\left(\frac{n+1}{n}\right) \end{aligned}$$

Conclusion


- Bayesian inference provides a coherent framework which we can use for machine learning
- However, it requires a model of what is happening
- In practice Bayesian methods are easy if the data is generated from a likelihood with a conjugate prior distribution—we have to be clever to choose the right prior
- We will see in the next lecture that much more frequently we will have likelihoods with no conjugate prior and we have to work much harder
- When we have no knowledge there are consistent ways to express our ignorance

- Bayes' Rule
- Conjugate Priors
- Uninformative Priors



Uninformative Priors for Scale Parameter

- Why did we choose $a_0 = b_0 = 0$ implying a prior $p(\mu) = 1/\mu$?




- That is, we have no idea on what scale to measure μ


$$\int_A^B p(\mu) d\mu = \int_{A/c}^{B/c} p(\mu) d\mu = \int_A^B \frac{1}{c} p\left(\frac{\nu}{c}\right) d\nu = \int_A^B \frac{1}{c} p\left(\frac{\mu}{c}\right) d\mu$$

making a change of variables $\mu = \nu/c$

- Or $p(\mu) = \frac{1}{c} p\left(\frac{\mu}{c}\right)$ implying $p(\mu) \propto \frac{1}{\mu}$

Population Size of 238 Countries






Riemann Integration, integration by parts, gaussian integrals

Riemann Integral

- Integrals represent area beneath a curve



Fundamental Law of Calculus

- Let

$$I(a, x) = \int_a^x f(z) dz = \lim_{\delta z_i \rightarrow 0} \sum_{i=0}^{n-1} (z_{i+1} - z_i) f(z_i)$$

- Now for small δx

$$I(a, x + \delta x) = \int_a^{x+\delta x} f(z) dz = \lim_{\delta z_i \rightarrow 0} \sum_{i=0}^{n-1} (z_{i+1} - z_i) f(z_i) + \delta x f(x)$$

- Thus


$$\frac{dI(a, x)}{dx} = \lim_{\delta x \rightarrow 0} \frac{I(x + \delta x) - I(x)}{\delta x} = \lim_{\delta x \rightarrow 0} \frac{\delta x f(x)}{\delta x} = f(x)$$

Indefinite Integrals

- So far we have considered **definite integrals** where we integrate between two points (a and b)
- However, when think about integration as an anti-derivative, it is useful to think of a function $F(x) = \int f(x) dx$
- So that $F'(x) = f(x)$
- However the function $F(x)$, $F(x) + 1$, $F(x) + \pi$, etc. all have the same derivative so $F(x)$ is only defined up to an additive constant
- Note that the definite integral is given by

$$\int_a^b f(x) dx = F(b) - F(a)$$

- Defining Integrals
- Doing Integrals
- Gaussian Integrals



Linearity of Integration

- Integration is a linear operator

$$\begin{aligned} \int_a^b (r f(x) + s g(x)) dx &= \lim_{\delta x_i \rightarrow 0} \sum_{i=0}^{n-1} (x_{i+1} - x_i) (r f(x_i) + s g(x_i)) \\ &= \lim_{\delta x_i \rightarrow 0} \left(\sum_{i=0}^{n-1} (x_{i+1} - x_i) r f(x_i) + \sum_{i=0}^{n-1} (x_{i+1} - x_i) s g(x_i) \right) \\ &= \lim_{\delta x_i \rightarrow 0} \left(r \sum_{i=0}^{n-1} (x_{i+1} - x_i) f(x_i) + s \sum_{i=0}^{n-1} (x_{i+1} - x_i) g(x_i) \right) \\ &= r \lim_{\delta x_i \rightarrow 0} \sum_{i=0}^{n-1} (x_{i+1} - x_i) f(x_i) + s \lim_{\delta x_i \rightarrow 0} \sum_{i=0}^{n-1} (x_{i+1} - x_i) g(x_i) \\ &= r \int_a^b f(x) dx + s \int_a^b g(x) dx \end{aligned}$$

The Other Way Around

- Consider


$$\begin{aligned} \int_a^b \frac{df(x)}{dx} dx &= \int_a^b \lim_{\delta x \rightarrow 0} \frac{f(x + \delta x) - f(x)}{\delta x} dx \\ &= \lim_{x_{i+1} - x_i \rightarrow 0} \sum_{i=0}^{n-1} (x_{i+1} - x_i) \frac{f(x_{i+1}) - f(x_i)}{x_{i+1} - x_i} \\ &= \lim_{x_{i+1} - x_i \rightarrow 0} \sum_{i=0}^{n-1} (f(x_{i+1}) - f(x_i)) \\ &= (f(x_1) - f(x_0)) + (f(x_2) - f(x_1)) + (f(x_3) - f(x_2)) + \dots \\ &\quad + (f(x_{n-1}) - f(x_{n-2})) + (f(x_n) - f(x_{n-1})) \\ &= f(x_n) - f(x_0) = f(b) - f(a) \end{aligned}$$

- We can think of integration as an **anti-derivative**—it undoes differentiation

Multiple Integrals

- For functions involving many independent variables (e.g. $f(x, y)$, $f(x, y, z)$, $f(x)$) we can integrate over multiple dimensions
- For example


$$\iint_A f(x, y) dx dy$$



- It gets tedious writing multiple integral signs and I tend to write just one

$$\int \dots \int f(x_1, x_2, \dots, x_n) dx_1 dx_2 \dots dx_n = \int f(\mathbf{x}) d\mathbf{x}$$

1. Defining Integrals
2. Doing Integrals
3. Gaussian Integrals



Is Integration Straightforward?

- We saw, due to the product and chain rules, that we can differentiate almost anything. Given integration is the anti-derivative can we integrate anything?

- Products and compositions

$$\int f(x)g(x)dx = ? \quad \int f(g(x))dx = ?$$

- Unfortunately, unlike differentiation we don't have a small parameter we can expand in
- In general integration is hard

Example of Integration by Parts

- Consider

$$\begin{aligned} \Pi(z) &= \int_0^\infty x^z e^{-x} dx = \int_0^\infty x^z \frac{d(-e^{-x})}{dx} dx \\ &= [x^z (-e^{-x})]_0^\infty - \int_0^\infty \frac{dx^z}{dx} (-e^{-x}) dx \\ &= \int_0^\infty (zx^{z-1}) e^{-x} dx = z \int_0^\infty x^{z-1} e^{-x} dx = z \Pi(z-1) \end{aligned}$$

- Thus $\Pi(z) = z \Pi(z-1)$, but

$$\Pi(0) = \int_0^\infty e^{-z} dz = [-e^{-z}]_0^\infty = -e^{-\infty} - (-e^0) = 1$$

- Now

$$\Pi(n) = n \Pi(n-1) = n(n-1) \Pi(n-2) = n(n-1)(n-2) \dots 1 = n!$$

Example of Integration by Substitution

- We consider $I(n) = \int_0^\infty x^n e^{-x^2/2} dx$
- Let $u(x) = x^2/2$ or $x(u) = \sqrt{2u}$ so that

$$\frac{dx(u)}{du} = \frac{1}{\sqrt{2u}} \quad u(0) = 0 \quad u(\infty) = \infty$$

- Thus

$$\begin{aligned} I(n) &= \int_0^\infty (\sqrt{2u})^n e^{-u} \frac{1}{\sqrt{2u}} du \\ &= 2^{\frac{n-1}{2}} \int_0^\infty u^{\frac{n-1}{2}} e^{-u} du = 2^{\frac{n-1}{2}} \Pi\left(\frac{n-1}{2}\right) \end{aligned}$$

- $I(1) = 1$, $I(3) = 2 \times 1! = 2$, $I(5) = 2^2 \times 2! = 8$ but $I(0) = \Pi(-1/2)/\sqrt{2}$, $I(2) = \sqrt{2}\Pi(1/2) = \Pi(-1/2)/\sqrt{2}$

- A key method for performing integrals is through knowledge of the anti-derivative
- If we know $F'(x) = f(x)$ then $F(x) + c = \int f(x) dx$
- E.g. we know that $dx^n/dx = nx^{n-1}$ therefore

$$\int x^{n-1} dx = \frac{1}{n} \int \frac{dx^n}{dx} dx = \frac{x^n}{n} + c$$

and

$$\int_a^b x^{n-1} dx = \frac{b^n}{n} - \frac{a^n}{n}$$

Integration by Parts

- Recall the product rule $\frac{d(f(x)g(x))}{dx} = \frac{df(x)}{dx}g(x) + f(x)\frac{dg(x)}{dx}$
- Integrating we get

$$\begin{aligned} \int_a^b \frac{df(x)g(x)}{dx} dx &= \int_a^b \frac{df(x)}{dx} g(x) dx + \int_a^b f(x) \frac{dg(x)}{dx} dx \\ &= [f(x)g(x)]_a^b = f(b)g(b) - f(a)g(a) \end{aligned}$$

- Unfortunately we get two integrals, but we can turn this around

$$\int_a^b f(x) \frac{dg(x)}{dx} dx = [f(x)g(x)]_a^b - \int_a^b \frac{df(x)}{dx} g(x) dx$$

whether this is helpful depends on $f(x)$ and $g(x)$

Substitution

- We can make a transformation from x to $u = u(x)$

$$\begin{aligned} \int_a^b f(x) dx &= \lim_{\delta x_i \rightarrow 0} \sum_{i=0}^{n-1} f(x_i)(x_{i+1} - x_i) \\ &= \lim_{\delta u_i \rightarrow 0} \sum_{i=0}^{n-1} f(x(u_i)) \frac{x(u_{i+1}) - x(u_i)}{u_{i+1} - u_i} (u_{i+1} - u_i) \\ &= \int_{u(a)}^{u(b)} f(x(u)) \frac{dx(u)}{du} du \end{aligned}$$

★ where u_i is such that $x(u_i) = x_i$ or $u_i = u(x_i)$ where $u(x)$ is the inverse of $x(u)$

★ using $\lim_{\delta u_i \rightarrow 0} \frac{x(u_{i+1}) - x(u_i)}{u_{i+1} - u_i} = \frac{dx(u_i)}{du}$

Changing Variables in Multidimensional Space


- When changing variables in many dimensions $\mathbf{x} \rightarrow \mathbf{u}$ the change of variables involves the Jacobian

$$\int f(\mathbf{x}) d\mathbf{x} = \int f(\mathbf{x}(\mathbf{u})) |\det(\mathbf{J})| d\mathbf{u}, \quad \mathbf{J} = \begin{pmatrix} \frac{\partial x_1}{\partial u_1} & \frac{\partial x_1}{\partial u_2} & \dots & \frac{\partial x_1}{\partial u_n} \\ \frac{\partial x_2}{\partial u_1} & \frac{\partial x_2}{\partial u_2} & \dots & \frac{\partial x_2}{\partial u_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial x_n}{\partial u_1} & \frac{\partial x_n}{\partial u_2} & \dots & \frac{\partial x_n}{\partial u_n} \end{pmatrix}$$

- E.g. transforming from Cartesian coordinates (x, y) to polar coordinates (r, θ) then $x = r \cos(\theta)$ and $y = r \sin(\theta)$

$$\begin{aligned} |\det(\mathbf{J})| &= \left| \det \begin{pmatrix} \frac{\partial r \cos(\theta)}{\partial r} & \frac{\partial r \cos(\theta)}{\partial \theta} \\ \frac{\partial r \sin(\theta)}{\partial r} & \frac{\partial r \sin(\theta)}{\partial \theta} \end{pmatrix} \right| = \left| \det \begin{pmatrix} \cos(\theta) & -r \sin(\theta) \\ \sin(\theta) & r \cos(\theta) \end{pmatrix} \right| \\ &= r(\cos^2(\theta) + \sin^2(\theta)) = r \end{aligned}$$

- That is, $dx dy = r dr d\theta$



Cumulant Generating Function

- Note that $e^{\ell x} = 1 + \ell x + \frac{1}{2}\ell^2 x^2 + \frac{1}{3!}\ell^3 x^3 + \dots$
- So

$$Z(\ell) = \int_{-\infty}^{\infty} e^{\ell x} f_X(x) dx = 1 + \ell M_1 + \frac{1}{2}\ell^2 M_2 + \frac{1}{3!}\ell^3 M_3 + \dots$$
- Now using $\log(1 + \epsilon) = \epsilon - \frac{1}{2}\epsilon^2 + \frac{1}{3}\epsilon^3 + \dots$

$$G(\ell) = \log(Z(\ell)) = \ell M_1 + \frac{1}{2}\ell^2 (M_2 - M_1^2) + \frac{1}{3!}\ell^3 (M_3 - 3M_2 M_1 + 2M_1^3) + \dots$$
- So that $\kappa_n = G^{(n)}(0)$, with $\kappa_1 = M_1$ (the mean), $\kappa_2 = M_2 - M_1^2$ (the variance), $\kappa_3 = M_3 - 3M_2 M_1 + 2M_1^3$ (the third cumulant related to the skewness)

Special Functions

- There are integrals with no known closed form solution
- We saw that $\Pi(z) = \int_0^{\infty} x^z e^{-x} dx$ satisfies $\Pi(z) = z\Pi(z-1)$
- For integer n then $\Pi(n) = n!$ but for general z , the integral $\Pi(z)$ can't be written in terms of elementary functions
- We consider $\Pi(z)$ as a special function in its own right
- Although, history has left us with the gamma function instead

$$\Gamma(z) = \int_0^{\infty} x^{z-1} e^{-x} dx = \Pi(z-1)$$

- Other special function defined by integrals exist (e.g. the Bessel, Aire, hypergeometric, elliptic, error functions, . . .)

Gaussian Integrals

- Gaussian integrals are integrals involving e^{-x^2} , e.g.

$$\int_{-\infty}^{\infty} e^{-x^2} dx \quad \int_{-\infty}^{\infty} x^4 e^{-ax^2-bx} dx$$

- They are important in computing integrals with respect to the normal distribution

$$\mathcal{N}(x|\mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/(2\sigma^2)}$$

- The great news is that these integrals are all doable
- The bad news is that they are quite tricky to do

- A trick that sometimes works is differentiating through an integral, e.g. consider finding moments

$$M_n = \mathbb{E}[X^n] = \int_{-\infty}^{\infty} x^n f_X(x) dx$$

- We can define a momentum generating function

$$Z(\ell) = \int_{-\infty}^{\infty} e^{\ell x} f_X(x) dx$$

- Then $M_n = Z^{(n)}(0)$ (the n^{th} derivative of $Z(\ell)$ at $\ell = 0$)


$$\left. \frac{d^n Z(\ell)}{d\ell^n} \right|_{\ell=0} = \int_{-\infty}^{\infty} \left. \frac{d^n e^{\ell x}}{d\ell^n} \right|_{\ell=0} f_X(x) dx = \int_{-\infty}^{\infty} x^n f_X(x) dx = M_n$$

More Integration

- Although we have a few tricks, integration is hard
- Surprisingly integration sometimes is easier when carried out in the complex plane
- This is a beautiful part of mathematics (due largely to Cauchy) but beyond the scope of this course
- Interestingly, also there is an algorithm that allows us to integrate a lot of function. It is sufficiently complicated that you need to write a computer algorithm of considerable complexity to implement it. Most symbolic manipulation packages (e.g. Mathematica) have implemented some part of this algorithm

Outline

- Defining Integrals
- Doing Integrals
- Gaussian Integrals



The Gaussian Integral

- The integral over a Gaussian is surprisingly difficult

$$I_1 = \int_{-\infty}^{\infty} e^{-x^2/2} dx$$

- There is a nice trick which is to consider

$$I_1^2 = \int_{-\infty}^{\infty} e^{-x^2/2} dx \int_{-\infty}^{\infty} e^{-y^2/2} dy = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-(x^2+y^2)/2} dx dy$$

- Making the change of variables $r = \sqrt{x^2 + y^2}$ and $\theta = \arctan(y/x)$ (so that $x = r \cos(\theta)$, $y = r \sin(\theta)$ and $x^2 + y^2 = r^2$)

$$I_1^2 = \int_0^{2\pi} d\theta \int_0^{\infty} r e^{-r^2/2} dr = 2\pi \int_0^{\infty} r e^{-r^2/2} dr$$

The Gaussian Integral Continued

- From before

$$I_1^2 = 2\pi \int_0^\infty r e^{-r^2/2} dr$$

- Finally let $u = r^2/2$ so that $du/dr = r$ or $du = r dr$ we get

$$I_1^2 = 2\pi \int_0^\infty e^{-u} du = 2\pi$$

- So that $I_1 = \sqrt{2\pi}$
- Incidentally, $I_1 = \sqrt{2}\Pi(-1/2)$ so $\Pi(-1/2) = \Gamma(1/2) = \sqrt{\pi}$

Multi-dimensional Gaussians

- Consider

$$I_3 = \int_{-\infty}^\infty \dots \int_{-\infty}^\infty e^{-\frac{1}{2}\|\mathbf{x}\|_2^2} dx_1 \dots dx_n$$

where $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$

- Note that $\|\mathbf{x}\|_2^2 = x_1^2 + x_2^2 + \dots + x_n^2$ and using $e^{\sum_i a_i} = \prod_i e^{a_i}$

$$\begin{aligned} I_3 &= \int_{-\infty}^\infty \dots \int_{-\infty}^\infty e^{-\frac{1}{2}\sum_{i=1}^n x_i^2} dx_1 \dots dx_n \\ &= \prod_{i=1}^n \int_{-\infty}^\infty e^{-x_i^2/2} dx_i = \prod_{i=1}^n \sqrt{2\pi} = (2\pi)^{n/2} \end{aligned}$$

Determinants

- Using the facts, that $\Xi = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$ then

$$\det(\Xi) = \det(\mathbf{V}\mathbf{\Lambda}\mathbf{V}^T) = \det(\mathbf{V})\det(\mathbf{\Lambda})\det(\mathbf{V}^T) = \det(\mathbf{\Lambda}) = \prod_{i=1}^n \lambda_i$$

using $\det(\mathbf{AB}) = \det(\mathbf{A})\det(\mathbf{B})$ and $\det(\mathbf{V}) = 1$

- Recall $I_4 = \prod_i \sqrt{2\pi\lambda_i} = (2\pi)^{n/2} \sqrt{\det(\Xi)}$
- We note for an $n \times n$ matrix $\mathbf{M}$ then $\det(c\mathbf{M}) = c^n \det(\mathbf{M})$ so that

$$I_4 = (2\pi)^{n/2} \sqrt{\det(\Xi)} = \sqrt{\det(2\pi\Xi)}$$

- Finally, we get that for the PDF of a normal to integrate to 1

$$\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \Xi) = \frac{1}{\sqrt{\det(2\pi\Xi)}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \Xi^{-1}(\mathbf{x}-\boldsymbol{\mu})}$$

1

Normal Distribution

- We consider

$$I_2 = \int_{-\infty}^\infty e^{-(x-\mu)^2/(2\sigma^2)} dx$$

- Making the change of variables $z = (x-\mu)/\sigma$ so that $dz = dx/\sigma$ or $dx = \sigma dz$. Then

$$I_2 = \sigma \int_{-\infty}^\infty e^{-z^2/2} dz = \sigma I_1 = \sqrt{2\pi}\sigma$$

- Note that the *probability density function* (PDF) for a normally distributed random variable is given by

$$\mathcal{N}(x \mid \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/(2\sigma^2)}$$

Full Multi-variate Normal

- Consider

$$I_4 = \int_{-\infty}^\infty \dots \int_{-\infty}^\infty e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \Xi^{-1}(\mathbf{x}-\boldsymbol{\mu})} dx_1 \dots dx_n$$

- Let $\Xi^{-1} = \mathbf{V}\mathbf{\Lambda}^{-1}\mathbf{V}^T$ and make the change of variables $\mathbf{y} = \mathbf{V}^T(\mathbf{x}-\boldsymbol{\mu})$
- The Jacobian $\mathbf{J}$ has elements (note that $\mathbf{x} = \mathbf{V}\mathbf{y} + \boldsymbol{\mu}$)


$$J_{ij} = \frac{\partial x_i}{\partial y_j} = \frac{\partial}{\partial y_j} \left(\sum_{k=1}^n V_{ik} y_k + \mu_i \right) = V_{ij}$$

- So that $\mathbf{J} = \mathbf{V}$ and consequently $|\det(\mathbf{J})| = |\det(\mathbf{V})| = 1$ then

$$I_4 = \int_{-\infty}^\infty \dots \int_{-\infty}^\infty e^{-\frac{1}{2}\mathbf{y}^T \mathbf{\Lambda}^{-1} \mathbf{y}} dy_1 \dots dy_n = \prod_{i=1}^n \int_{-\infty}^\infty e^{-y_i^2/(2\lambda_i)} dy_i = \prod_i \sqrt{2\pi\lambda_i}$$

Summary

- Integration is extra-ordinarily useful as a tool of analysis
- It occurs when you work with probabilities densities for continuous random variables
- Integration is beautiful, but hard—often impossible
- Normal distributions lucky almost always give raise to integrals that can be computed in closed form although often it requires quite a bit of work
- Making friends with integration will give you a super-power that not too many people share



Gaussian Processes, regression

Machine Learning with Gaussians

- Due to the fact that we can analytically integrate over multivariate Gaussians they are heavily used of machine learning
- Examples include multivariate modelling, Kalman filters and Gaussian Processes
- Despite being analytically tractable, they are very fiddly
- This is mainly because it is easy to specify means, μ , and covariance, C , but Gaussians involve the inverse covariance matrix

$$\mathcal{N}(x | \mu, C) = \frac{e^{-\frac{1}{2}(x-\mu)^T C^{-1}(x-\mu)}}{\sqrt{\det(2\pi C)}}$$

Moment Generating Function

- Again through completing the square we can show

$$M(a) = \int_{-\infty}^{\infty} e^{a^T x} \mathcal{N}(x | \mu, C) dx = e^{a^T \mu + \frac{1}{2} a^T C a}$$


- The first moment is given by

$$\mathbb{E}[x] = \nabla_a M(a)|_{a=0} = (\mu + C a) e^{a^T \mu + \frac{1}{2} a^T C a} \Big|_{a=0} = \mu$$

- The second moment is


$$\mathbb{E}[x_i x_j] = \frac{\partial^2 M(a)}{\partial x_i \partial x_j} \Big|_{a=0} = \mu_i \mu_j + C_{ij}$$

Multivariate Normals



$$\mathcal{N}\left(\begin{pmatrix} x \\ y \end{pmatrix} \middle| \begin{pmatrix} \mu^x \\ \mu^y \end{pmatrix}, \begin{pmatrix} A & B \\ B^T & C \end{pmatrix}\right) = \frac{e^{-\frac{1}{2} \begin{pmatrix} x-\mu^x \\ y-\mu^y \end{pmatrix}^T \begin{pmatrix} A & B \\ B^T & C \end{pmatrix}^{-1} \begin{pmatrix} x-\mu^x \\ y-\mu^y \end{pmatrix}}}{\sqrt{\det(2\pi \begin{pmatrix} A & B \\ B^T & C \end{pmatrix})}}$$

- ML with Gaussians
- Gaussian Processes
- Bayesian Inference
- Hyper-parameters



Completing the Square

- One of the key tools for performing Gaussian integrals is “completing the square”

$$\begin{aligned} -\frac{1}{2} x^T C^{-1} x + a^T x &= -\frac{1}{2} x^T C^{-1} x + \frac{1}{2} a^T x + \frac{1}{2} x^T a \\ &= -\frac{1}{2} x^T C^{-1} x + \frac{1}{2} a^T C C^{-1} x + \frac{1}{2} x^T C^{-1} C a \\ &= -\frac{1}{2} (x^T - a^T C) C^{-1} (x - C a) + \frac{1}{2} a^T C a \end{aligned}$$

- This gives us the Hubbard–Stratonovich transformation

$$\begin{aligned} \int_{-\infty}^{\infty} \frac{e^{-\frac{1}{2} x^T C^{-1} x + a^T x}}{\sqrt{|\det(2\pi C)|}} dx &= \int_{-\infty}^{\infty} \frac{e^{-\frac{1}{2} (x^T - a^T C) C^{-1} (x - C a) + \frac{1}{2} a^T C a}}{\sqrt{|\det(2\pi C)|}} dx \\ &= \int_{-\infty}^{\infty} \mathcal{N}(x | C a, C) e^{\frac{1}{2} a^T C a} dx = e^{\frac{1}{2} a^T C a} \end{aligned}$$

Covariance

- The covariance is given by

$$\mathbb{E}[(x_i - \mu_i)(x_j - \mu_j)] = C_{ij}$$

$$\text{or } \mathbb{E}[(x - \mu)(x - \mu)^T] = C$$

- The cumulant generating function is given by

$$C(a) = \log(M(a)) = a^T \mu + \frac{1}{2} a^T C a$$

- The mean is equal to $\nabla C(a)|_{a=0} = \mu$
- The covariance is given by $\frac{\partial^2 C(a)}{\partial a_i \partial a_j} \Big|_{a=0} = C_{ij}$
- Normal distributions have no high cumulants (skewness, kurtosis)

Marinalising Multivariate Normals

- I can define a moment generating function,

$$\begin{aligned} M(a, b) &= \int e^{(a^T, b^T) \begin{pmatrix} x-\mu^x \\ y-\mu^y \end{pmatrix}} \mathcal{N}\left(\begin{pmatrix} x \\ y \end{pmatrix} \middle| \begin{pmatrix} \mu^x \\ \mu^y \end{pmatrix}, \begin{pmatrix} A & B \\ B^T & C \end{pmatrix}\right) dx dy \\ &= e^{\frac{1}{2} (a^T, b^T) \begin{pmatrix} A & B \\ B^T & C \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix}} = e^{\frac{1}{2} (a^T A a + b^T C b + 2 a^T B b)} \end{aligned}$$

- If I set $b = 0$ this is equivalent to marginalising out y
- But then we just get $M(a, 0) = e^{\frac{1}{2} a^T A a}$
- This is the same as the moment generating function for $\mathcal{N}(x | \mu^x, A)$

- An important and non-trivial result is

$$\int \mathcal{N}\left(\begin{pmatrix} x \\ y \end{pmatrix} \middle| \begin{pmatrix} \mu^x \\ \mu^y \end{pmatrix}, \begin{pmatrix} A & B \\ B^T & C \end{pmatrix}\right) dy = \mathcal{N}(x \mid \mu^x, A)$$

- You would think this is easy to prove by straightforward integration, but it ain't!
- It can be done, but it is a page of work!

Gaussian Processes

- Gaussian processes (GPs) are a mathematically defined ensemble of functions
- They can be combined with Bayesian inference to give one of the most powerful regression techniques
- Although Bayesian they can be used in a black-box fashion due to the ubiquity of the prior
- Mathematically they are a bit complicated (because Gaussians involve the inverse of matrices which are a real pain to work with)
- In practice they aren't that difficult to use

Priors on Functions

- We can think of a solution as a function $f(x)$
- We can put a prior probability distribution, $p(f)$, on a function, f , that prefers smooth functions
- We can then compute a posterior probability distribution on functions given the data, $p(f|D)$
- As a likelihood, $p(y_i|f(x_i))$, we use the probability of observing y_i given the true function value is $f(x_i)$
- In general, this would be next to impossible to compute except in the special case where everything is Gaussian (normally distributed)

Meaning of GP


- To understand GP's we can discretise space, x , into a lattice of points $\{x_i\}$
- Then (assuming $m(x) = 0$)

$$p(f|m, k) \propto \prod_i e^{-\frac{f_i^2 k^{-1}(x_i, x_i)}{2}} + f_i \sum_j k^{-1}(x_i, x_j) f_j$$

where $f_i = f(x_i)$

- We see that the value of the function at each point is normally distributed with a mean that depends on functions at neighbouring points

1. ML with Gaussians
2. Gaussian Processes
3. Bayesian Inference
4. Hyper-parameters



Regression

- In regression we try to fit a multi-dimensional function to our data
- (You can use Gaussian Processes for classification, e.g. by inferring the probabilities of being in a class, but we ignore this as regression is where GP excel)
- In regression we have some p dimensional feature vectors x_i and some target $y_i \in \mathbb{R}$
- Our task is to fit a function through all the data points

Gaussian Processes

- Gaussian Processes are probability distributions over functions
 - (Functions can be viewed as vectors in an infinite dimensional vector space)
 - In the Gaussian Process, $\mathcal{GP}(m, k)$, the probability of a function, f , is proportional
- $$p(f|m, k) \propto e^{-\frac{1}{2} \int (f(x) - m(x)) k^{-1}(x, y) (f(y) - m(y)) dx dy}$$
- The function $m(x)$ is the mean $\mathbb{E}[f(x)]$ (usually taken to be zero in most inference problems)

Covariance function


- $k(x, y)$ is a covariance function

$$\mathbb{E}[(f(x) - m(x))(f(y) - m(y))] = k(x, y)$$

- This is sometimes known as a **kernel**—it must be positive semi-definite (just like in SVMs)
- It is a free “parameter” that the user gets to choose (although we can learn its parameters too)
- If $k(x, y)$ is a function of $x - y$ it is “**stationary**”
- If $k(x, y)$ is a function of $\|x - y\|$ it is also “**isometric**”

Popular Choices of GP Kernel Function

- Constant: $k_C(\mathbf{x}, \mathbf{y}) = C$
- Gaussian noise: $k_{GN}(\mathbf{x}, \mathbf{y}) = \sigma^2 \delta_{\mathbf{x}, \mathbf{y}}$
- Squared exponential: $k_{SE}(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|^2}{2\ell^2}\right)$
- Ornstein–Uhlenbeck: $k_{OU}(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|}{\ell}\right)$




Adam Prügel-Bennett

Mathematics of Machine Learning

//ecs-vlc.github.io/aice3001/

17

Gaussian Process Worlds




Adam Prügel-Bennett

Mathematics of Machine Learning

//ecs-vlc.github.io/aice3001/

18

2-D Gaussian Processes



Adam Prügel-Bennett


Mathematics of Machine Learning

//ecs-vlc.github.io/aice3001/

19

Outline

1. ML with Gaussians
2. Gaussian Processes
3. Bayesian Inference
4. Hyper-parameters



Adam Prügel-Bennett

Mathematics of Machine Learning

//ecs-vlc.github.io/aice3001/

20

Observed Gaussian Processes

- Given some data points $\mathcal{D} = ((\mathbf{x}_i, y_i) | i = 1, \dots, m)$ the likelihood (assuming Gaussian error are independence of the data point) is given by

$$p(\mathcal{D}|f) = \prod_{i=1}^m \mathcal{N}(y_i | f(\mathbf{x}_i), \sigma^2)$$

- Using a Gaussian Process prior we can compute a posterior using Bayes's rule
- The posterior is a Gaussian Process with a shifted mean and variance depending on the data-points
- This direct Bayesian derivation gives the answer involving the inverse matrix of the correlation function, $k^{-1}(\mathbf{x}, \mathbf{y})$ —this is a pain to work with

Adam Prügel-Bennett

Mathematics of Machine Learning

//ecs-vlc.github.io/aice3001/

21

Alternative Derivation

- Denoting the target values as a vector $\mathbf{y}$ with elements y_i
- Denoting the matrices of covariances between data points as $\mathbf{K}$ with elements $k(\mathbf{x}_i, \mathbf{x}_j)$
- Denoting the covariance between the data points and a particular position, $\mathbf{x}_*$ as $\mathbf{k}_*$ with elements $k(\mathbf{x}_i, \mathbf{x}_*)$
- Denoting the variance a point $\mathbf{x}_*$ as $k_* = k(\mathbf{x}_*, \mathbf{x}_*)$
- Then the distribution of function values at points at $\mathbf{x}_i$ and $\mathbf{x}_*$ is

$$p(\mathbf{y}, f_*) = \mathcal{N}\left(\begin{pmatrix} \mathbf{y} \\ f_* \end{pmatrix} \middle| \mathbf{0}, \begin{pmatrix} \mathbf{K} + \sigma^2 \mathbf{I} & \mathbf{k}_* \\ \mathbf{k}_*^T & k_* \end{pmatrix}\right)$$

Adam Prügel-Bennett

Mathematics of Machine Learning

//ecs-vlc.github.io/aice3001/

22

Conditional Probability

- To compute the posterior $p(f_* | \mathbf{y})$ we use

$$p(f_* | \mathbf{y}) = \frac{p(f_*, \mathbf{y})}{p(\mathbf{y})}$$

- where $p(\mathbf{y}) = \int p(f_*, \mathbf{y}) df_*$
- Because all integrals are Gaussian we can compute the integral to obtain

$$p(f_* | \mathbf{y}) = \mathcal{N}\left(f_* \middle| \mathbf{k}_*^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}, k_* - \mathbf{k}_*^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_*\right)$$

- Looks complicated, but numerically easy to evaluate


Adam Prügel-Bennett

Mathematics of Machine Learning

//ecs-vlc.github.io/aice3001/

23

$$K(x, x') = \exp(-(x - x')^2 / (2\ell^2))$$



Adam Prügel-Bennett

Mathematics of Machine Learning

//ecs-vlc.github.io/aice3001/

24

- I've shown a 1-D regression example because it is easy to visualise
- This might be used with a time series
- The much more typical situation in machine learning is for x to have many features so we are doing multi-dimensional regression
- Gaussian process inference were first used in spatial problems where it was known as **kriging**
- It was re-invented by the machine learning community who call it Gaussian Processes (GP)

Choosing the Correct Covariance Function

- Choosing the correct covariance function is critical
- Most covariance functions include a continuous **hyper-parameter** (e.g. the correlation length ℓ) that we have to choose correctly
- This is typical of many Bayesian problems where we have some set of hyper-parameters, ϕ , describing the model
- These are different to the normal parameters we learn (e.g. weights w or in GP the functions $f(x)$)
- In Bayesian inference we learn the posterior for these normal parameters

$$p(f|\mathcal{D}, \phi) = \frac{p(\mathcal{D}|f, \phi)p(f|\phi)}{p(\mathcal{D}|\phi)}$$

Maximum-Likelihood-II


- The integral

$$p(f|\mathcal{D}) = \int p(f|\mathcal{D}, \phi)p(\phi|\mathcal{D})d\phi$$


usually can't be computed analytically and we have to use Monte Carlo methods (see later lecture)

- An alternative is to use the most likely hyper-parameter
- We can find this by using gradient search of $p(\mathcal{D}|\phi)$
- This is sometimes referred to as ML-II
- Normally even this can be difficult, but for GP its not too difficult

Example (slightly pathological)



1. ML with Gaussians
2. Gaussian Processes
3. Bayesian Inference
4. **Hyper-parameters**



Evidence Framework

- The normalisation factor, $p(\mathcal{D}|\phi)$ is known as the **marginal likelihood** or **evidence**

$$p(\mathcal{D}|\phi) = \int p(\mathcal{D}|f, \phi)p(f|\phi)df$$

- We can perform a Bayesian calculation at a second level by putting a prior on ϕ

$$p(\phi|\mathcal{D}) = \frac{p(\mathcal{D}|\phi)p(\phi)}{p(\mathcal{D})}$$

- From this we can now marginalise out the hyper-parameters

$$p(f|\mathcal{D}) = \int p(f|\mathcal{D}, \phi)p(\phi|\mathcal{D})d\phi$$

Evidence for GP

- For GP the (log)-evidence can be computed in closed form

$$\log(p(\mathcal{D}|\phi)) = -\frac{1}{2}\mathbf{y}^T(\mathbf{K} + \sigma^2\mathbf{I})\mathbf{y} - \frac{1}{2}\log(|\mathbf{K} + \sigma^2\mathbf{I}|) - \frac{m}{2}\log(2\pi)$$


- ★ First term measures goodness of fit
- ★ Second term measure complexity of model
- ★ Last term is a common normalisation constant

- Can efficiently compute derivatives and find best parameters
- Could overfit!

Conclusions

- Gaussian processes are very powerful for regression (and classification?)
- Because all calculations involve Gaussian integrals we can compute everything in closed form
- (Actually its a pain to do the mathematics because you end up working with inverse of matrices)
- Fairly generic (black-box) technique because the prior captures many continuity constraints
- We can use the evidence framework (probability of data) to do model selection and hyper-parameter optimisations

Probabilistic Inference



Hierarchical Models, Mixture of Gaussians, Expectation Maximisation

Building Probabilistic Models

- To describe a system with uncertainty we use random variables, X, Y, Z , etc.■
- We use the convention of writing random variables in capitals (this is sometimes confusing as when you observe a random variable it is no longer random)■
- The variables are described by probability mass function $\mathbb{P}(X, Y, Z)$ or if our variables are continuous, but probability densities $f_{X, Y, Z}(x, y, z)$ ■
- A major rule of probability is

$$\sum_X \mathbb{P}(X, Y, Z) = \mathbb{P}(Y, Z) \blacksquare$$

Discriminative Models


- We often think of our observations as given and the predictions as random variables■
- For example we might be given some features $\mathbf{x}$ and we wish to predict a class $C \in \mathcal{C}$ ■
- Our objective is then to find the probability $\mathbb{P}(C|\mathbf{x})$ ■
- This is known as a **discriminative model**■
- E.g. in *foundations of machine learning* you learnt how to find the Bayes' optimal discrimination surface■

Latent Variables

- Sometimes we have models that involve random variables that we don't observe and we don't care about■
- These are called **latent variables**■
- If we have a latent variable Z and observed variable $\mathbf{X}$ and we are predicting a variable Y then we would **marginalise** over the latent variable

$$\mathbb{P}(\mathbf{X}, Y) = \sum_Z \mathbb{P}(\mathbf{X}, Y, Z) \blacksquare$$

1. **Building Probabilistic Models**
2. Mixture of Gaussians
3. Expectation Maximisation



Conditional Probabilities

- When developing models it is often useful to consider conditional probabilities e.g. $\mathbb{P}(X, Y|Z)$ or $f_{X|Y, Z}(x|y, z)$ ■
- A second major rule in probabilistic modelling is

$$\mathbb{P}(X, Y) = \mathbb{P}(X|Y)\mathbb{P}(Y) = \mathbb{P}(Y|X)\mathbb{P}(X) \blacksquare$$

- This is a mathematical identity that does not imply causality (it defines conditional probability)■
- It is the origins of Bayes' rule: $\mathbb{P}(X|Y) = \frac{\mathbb{P}(Y|X)\mathbb{P}(X)}{\mathbb{P}(Y)}$ ■

Generative Models

- Sometimes it is easy to think about the joint process of generating the features and outputs together■
- This leads to a joint distribution $\mathbb{P}(\mathbf{X}, Y)$ where $\mathbf{X}$ are your features and Y is your output you are trying to predict■
- This is known as a **generative model**■
- Generative models are often more natural to think about■
- We can use them to do discrimination using

$$\mathbb{P}(Y|\mathbf{X}) = \frac{\mathbb{P}(\mathbf{X}, Y)}{\mathbb{P}(\mathbf{X})} = \frac{\mathbb{P}(\mathbf{X}, Y)}{\sum_Y \mathbb{P}(\mathbf{X}, Y)} \blacksquare$$

Modelling Virus

- Suppose we want to estimate the number of hospitalisation from Corona virus in the next month■
- Our observable is the number of reported cases■
- In our model we might want to estimate the number of actual cases■
- This would be a latent variable (it is not an observable or our final target, but it is very useful intermediate in our model)■
- This will be a random variable (it is uncertain, but we can build a probabilistic model giving a distribution of number of actual cases)■

- Of course, if I was really modelling the spread of a disease I would care about the probability, $f(C|A,V)$, of catching the disease, C , given the persons age A and the variant of the disease V
- I would want to know the distribution of ages $f(A)$ and try to infer the probability of different variants $\mathbb{P}(V)$
- I would care about the probability, $f(R|A,V)$, of cases being reported given age and variant
- And the probability, $f(H|A,V)$, of hospitalisation given A and V
- This would involve an elaborate (hierarchical) model with a large number of latent variables

Problem with Bayes

- Bayes is problematic because it is often hard
- The posterior is often not expressible as a nice probability function
- We need to compute the *evidence* or *margin likelihood* we use

$$\mathbb{P}(\mathcal{D}) = \sum_{\Theta} \mathbb{P}(\mathcal{D}|\Theta)\mathbb{P}(\Theta)$$

- But sometimes the number of values that Θ can take are so large that we cannot easily compute this
- Nevertheless we can usually do this using Monte Carlo techniques

Maximum Likelihood

- When we assume a uniform prior then the MAP solution is just maximising the likelihood
- Weirdly this hack was accepted as part of mainstream statistics even when Bayesian statistics was considered unscientific
- Maximum likelihood is often sufficient for *government work*, but it isn't the best you can do
- In high-dimensional problems using a non-uniform prior can make a big difference
- And, of course, doing a full probabilistic calculation has real advantages (no overfitting, estimate of error, etc.)

Mixture of Gaussians

- Suppose we were observing the decays from two types of short-lived particle, A or B
- We observe the half life, X_i , but not the particle type
- We assume X_i is normally distributed with unknown means and variances: $\Theta = \{\mu_A, \sigma_A^2, \mu_B, \sigma_B^2\}$
- Let $Z_i \in \{0,1\}$ be an indicator that particle i is of type A
- The probability of X_i is given by

$$f(X_i|Z_i, \Theta) = Z_i \mathcal{N}(X_i|\mu_A, \sigma_A^2) + (1 - Z_i) \mathcal{N}(X_i|\mu_B, \sigma_B^2)$$

- We can use Bayes' rules to learn a set of parameter Θ that occur in our likelihood function

$$\mathbb{P}(\Theta|\mathcal{D}) = \frac{\mathbb{P}(\mathcal{D}|\Theta)\mathbb{P}(\Theta)}{\mathbb{P}(\mathcal{D})}$$

- This provides us a full probabilistic description of the parameters
- It doesn't overfit (we are not choosing the best)
- Bayesian inference provides a description of its own uncertainty
- We need to specify a likelihood and prior, but this is usually not difficult

Maximum A Posteriori (MAP) Solution


- One work around is to compute the mode of the posterior

$$\Theta_{\text{MAP}} = \underset{\Theta}{\operatorname{argmax}} f(\mathcal{D}|\Theta)f(\Theta) = \underset{\Theta}{\operatorname{argmax}} \log(f(\mathcal{D}|\Theta)) + \log(f(\Theta))$$

- We don't need to calculate $f(\mathcal{D})$ or explicitly calculate the posterior distribution
- But it is not Bayesian (despite what you are sometime told) — its not properly probabilistic
- You can overfit and you don't get an estimate of the error in your inference

Outline


1. Building Probabilistic Models
2. Mixture of Gaussians
3. Expectation Maximisation



Data

- Note that

$$\begin{aligned} f(X_i|\Theta) &= \sum_{Z_i \in \{0,1\}} f(X_i, Z_i|\Theta) = \sum_{Z_i \in \{0,1\}} f(X_i|Z_i, \Theta)\mathbb{P}(Z_i) \\ &= \mathbb{E}_{Z_i}[f(X_i|Z_i, \Theta)] = p\mathcal{N}(X_i|\mu_A, \sigma_A^2) + (1-p)\mathcal{N}(X_i|\mu_B, \sigma_B^2) \end{aligned}$$



- To solve the model as a Bayesian we would have to assign priors to our parameters $\Theta = (\mu_A, \sigma_A, \mu_B, \sigma_B, p)$
- This is doable, but complicated (we would also end up with a distribution for our parameters)
- Often we only want a reasonable estimate for some of our parameters (e.g. the half-lives μ_A and μ_B)
- A reasonable approach is to seek those parameters that maximise the likelihood of our observed data

$$f(\mathcal{D}|\Theta) = \prod_{X \in \mathcal{D}} f(X|\Theta)$$

Maximum Likelihood with Latent Variables

- The maximum likelihood is a non-linear function of the parameters so cannot be immediately maximised
- If we knew which type of particle a data-point belongs to (Z_i) then it would be straightforward to maximise the likelihood
- As we don't we need to estimate $\mathbb{P}(Z_i = 1)$, but this depends on $\mu_A, \sigma_A^2, \mu_B, \sigma_B^2$ and p
- We could use a standard optimiser, but this is slightly inelegant

Why EM Algorithm Works

- The argument around why this works is quite involved
- Note that at each step we maximise

$$Q(\Theta|\Theta^{(t)}) = \sum_{\mathbf{Z} \in \{0,1\}^m} \mathbb{P}(\mathbf{Z}|\mathcal{D}, \Theta^{(t)}) \log(f(\mathcal{D}|\mathbf{Z}, \Theta))$$

- We can show that the maximum, $\Theta^{(t+1)}$, is such that

$$\log(f(\mathcal{D}|\Theta^{(t+1)})) - \log(f(\mathcal{D}|\Theta^{(t)})) \geq Q(\Theta^{(t+1)}|\Theta^{(t)}) - Q(\Theta^{(t)}|\Theta^{(t)}) \geq 0$$
- The details are given in the supplemental notes

EM for Mixture of Gaussians


- Maximise with respect to parameters θ

$$\begin{aligned} Q(\theta|\theta^{(t)}) &= \sum_{\mathbf{Z}} \mathbb{P}(\mathbf{Z}|\mathcal{D}, \theta^{(t)}) \log(f(\mathcal{D}|\mathbf{Z}, \theta)) = \sum_{i=1}^m \sum_{Z_i} \mathbb{P}(Z_i|\mathcal{D}, \theta^{(t)}) \log(f(X_i|Z_i, \theta)) \\ &= \sum_{i=1}^m \sum_{Z_i \in \{0,1\}} \mathbb{P}(Z_i|X_i, \theta^{(t)}) (Z_i \log(p) + (1 - Z_i) \log(1 - p) \\ &\quad - \frac{(X_i - \mu_{Z_i})^2}{2\sigma_{Z_i}^2} - \log(\sqrt{2\pi}\sigma_{Z_i})) \end{aligned}$$

- Compute update equations

$$\frac{\partial Q(\theta|\theta^{(t)})}{\partial \mu_k} = 0, \quad \frac{\partial Q(\theta|\theta^{(t)})}{\partial \sigma_k} = 0, \quad \frac{\partial Q(\theta|\theta^{(t)})}{\partial p} = 0$$

1. Building Probabilistic Models
2. Mixture of Gaussians
3. Expectation Maximisation



EM Algorithm

- Instead we can use an **expectation-maximisation algorithm** usually known as an **EM algorithm**
- We proceed iteratively by maximising the expected log-likelihood with respect to the current set of parameters

$$\Theta^{(t+1)} = \underset{\Theta}{\operatorname{argmax}} \sum_{\mathbf{Z}} \mathbb{P}(\mathbf{Z}|\mathcal{D}, \Theta^{(t)}) \log(f(\mathcal{D}|\mathbf{Z}, \Theta))$$

- It isn't obvious why this works

Conditional Latent Variables

- We need to compute the distribution of latent variables conditioned on the data and current estimated parameters
- For our problem

$$\mathbb{P}(\mathbf{Z}|\mathcal{D}, \Theta^{(t)}) = \prod_{i=1}^m \mathbb{P}(Z_i|X_i, \Theta^{(t)})$$

where

$$\begin{aligned} \mathbb{P}(Z_i = 1|X_i, \Theta^{(t)}) &= \frac{p^{(t)} \mathcal{N}(X_i|\mu_A^{(t)}, \sigma_A^{2(t)})}{p^{(t)} \mathcal{N}(X_i|\mu_A^{(t)}, \sigma_A^{2(t)}) + (1 - p^{(t)}) \mathcal{N}(X_i|\mu_B^{(t)}, \sigma_B^{2(t)})} \\ \mathbb{P}(Z_i = 0|X_i, \Theta^{(t)}) &= 1 - \mathbb{P}(Z_i = 1|X_i, \Theta^{(t)}) \end{aligned}$$

Update Equations

- Means


$$\mu_{Z_i}^{(t+1)} = \frac{\sum_{i=1}^n \mathbb{P}(Z_i|X_i, \theta^{(t)}) X_i}{\sum_{i=1}^n \mathbb{P}(Z_i|X_i, \theta^{(t)})}$$

- Variances


$$(\sigma_{Z_i}^{(t+1)})^2 = \frac{\sum_{i=1}^n \mathbb{P}(Z_i|X_i, \theta^{(t)}) (X_i - \mu_{Z_i}^{(t+1)})^2}{\sum_{i=1}^n \mathbb{P}(Z_i|X_i, \theta^{(t)})}$$

- Probability of being type 1

$$p^{(t+1)} = \frac{1}{n} \sum_{i=1}^n \mathbb{P}(Z_i = 1 | X_i, \theta_i^{(t)})$$



- Building probabilistic models is an intricate process
- Identifying random variables that describe the system is the first step
- Often we need to introduce variables that we don't observe and need to be marginalised out
- The EM algorithm provide one approach to maximising likelihoods or MAP solutions when we have latent variables
- It often gives nice update equations, but convergence can be slow



Monte Carlo methods, MCMC, Variational Methods

Bayesian Inference Gets Hard

- We saw that in some cases if we had a simple likelihood (normal, binomial, Poisson, multinomial) you can choose a conjugate prior (gamma-normal/Wishart, beta, gamma, Dirchlet) so that the posterior has the same form as the prior
- Very often we are working with more complex models where no conjugate prior exists
- The posterior is not described by a known distribution
- We have to work a lot harder—particularly with multivariate distributions

Complex Likelihoods


- In the classical problems with parameters, θ , that are well known statistics such as the mean and variance
- But suppose we are trying to distinguish between different models of cosmological inflation (physics models of the early Universe)
- Our data is the Cosmic Microwave Background measure by a satellite telescope
- Our parameter, θ , might represent unknown parameters of the early universe (width of surface of scattering, number of e-foldings, primordial perturbations, amount of cool dark matter)
- Here the likelihood $\mathbb{P}(\text{CMB} | \theta)$ is far from a well known probability distribution

Sample Estimation


- If we can draw independent samples (aka **deviates** or **variates**), Θ_i , from our posterior distribution then we can obtain an estimate of our expectation

$$\mathbb{E}[g(\Theta)] \approx \frac{1}{n} \sum_{i=1}^n g(\Theta_i)$$

- Provided our posterior distribution is well behaved the relative error in our estimate will drop off as $1/\sqrt{n}$



1. **Sampling**
2. **Random Number Generation**
3. **MCMC**
4. **Quiz**



Bayesian Inference

- Recall our problem is that we are given some data $\mathcal{D}$
- Our posterior is given by

$$\mathbb{P}(\theta|\mathcal{D}) = \frac{\mathbb{P}(\mathcal{D}|\theta)\mathbb{P}(\theta)}{\mathbb{P}(\mathcal{D})} \quad \text{or} \quad f(\theta|\mathcal{D}) = \frac{f(\mathcal{D}|\theta)f(\theta)}{f(\mathcal{D})}$$
- Where θ are the parameters we are trying to infer
- But our likelihood (and/or prior) might be quite complicated
- Typically we don't have a closed form representation for our posterior distribution


Histograms, Samples and Means

- We could represent our posterior as a histogram, although for multivariate distributions (i.e. when we are modelling more than one variable) a histogram can be unwieldy
- A sample from the posterior distribution is often sufficient e.g. in our topic models (LDA) a typical set of topics is what we are after
- However, when samples vary a lot, often the most useful quantities are expectation, e.g.

$$\begin{aligned} \mathbb{E}[\Theta] & \quad \mathbb{E}[\Theta_i^2] - \mathbb{E}[\Theta_i]^2 \\ \mathbb{E}[\Theta_i \Theta_j] - \mathbb{E}[\Theta_i]\mathbb{E}[\Theta_j] & \quad \mathbb{E}[\Theta \Theta^T] - \mathbb{E}[\Theta]\mathbb{E}[\Theta]^T \end{aligned}$$


Outline

1. **Sampling**
2. **Random Number Generation**
3. **MCMC**
4. **Quiz**



Drawing Random Samples


- Drawing (pseudo) random variables from a distribution is known as **Monte Carlo**
- For some very simple distributions we can use the **transformation methods** to transform a uniform distribution



Rejection Method

- The transformation method only works when we can easily compute the inverse *cumulative distribution function* (CDF)
- A more general technique is the **rejection method** where we generate deviates from $g_Y(y)$ such that $c g_Y(x) \geq f_X(x)$
- To draw deviates from $f_X(x)$ we draw a deviate $Y \sim g_Y$ and then accept the deviate with probability $f_X(Y)/(c g_Y(Y))$
- The expected rejection rate is $c - 1$
- Need to choose a good distribution $g_Y(y)$

Drawing Normal Deviates




Problems with Rejection

- The rejection method is very general and often the method of choice (although for normal deviates there is a clever transformation method which is faster)
- However, for complicated probability distributions it can be difficult to find a good proposal distribution $g_Y(y)$
- This is particular true for multivariate distributions
- If the proposal distribution is poor c might be very high and the number of rejections is stupidly high

Outline

- Sampling
- Random Number Generation
- MCMC**
- Quiz



Detailed Balance

- Suppose we have a set of states $\mathcal{S}$ and want to draw sample from a probability distribution $\pi = (\pi_i | i \in \mathcal{S})$
- We invent a dynamical system with a transition probability M_{ij} from state j to state i such that

$$M_{ij}\pi_j = M_{ji}\pi_i$$

- This is known as **detailed balance**
- Summing both sides over j

$$\sum_j M_{ij}\pi_j = \sum_j M_{ji}\pi_i = \pi_i \quad \mathbf{M}\pi = \pi$$

Convergence of MCMC

- Suppose we start from a state $x(0) = \sum_i c_i v^{(i)}$ where the $v^{(i)}$'s are an eigenvectors of the transition matrix $\mathbf{M}$ with eigenvalues λ_i
- It l apply $\mathbf{M}$ many times then

$$x(t) = \mathbf{M}^t x(0) = \mathbf{M}^t \sum_i c_i v^{(i)} = \sum_i \lambda_i^t c_i v^{(i)}$$

- And $\lim_{t \rightarrow \infty} x(t) = v^*$ where v^* is the eigenvector with the maximum eigenvalue
- Now $\|\mathbf{M}v\|_1 \leq \|\mathbf{M}\|_1 \|v\|_1 = \|v\|_1$ so the maximum eigenvalue is 1 with eigenvector π ($\mathbf{M}$ is known as a **stochastic matrix**)

Metropolis Algorithm

- A very easy way to achieve detailed balance is starting from state j choose a "neighbouring" state, i with equal probability
- We accept the move if either
 - $\pi_i > \pi_j$ or
 - we make the move with a probability π_i/π_j
- If $\pi_i > \pi_j$ then $M_{ij} = 1$ and $M_{ji} = \pi_j/\pi_i$. Thus

$$M_{ij}\pi_j = \pi_j \quad M_{ji}\pi_i = \frac{\pi_j}{\pi_i}\pi_i = \pi_j$$

- Note that we require the state i to have the same number of neighbours as state j so that detailed balance is satisfied

- If we are working with continuous variables θ then the equation for detailed balance for the transition probability $W(\theta \rightarrow \theta')$ is

$$W(\theta \rightarrow \theta')\pi(\theta) = W(\theta' \rightarrow \theta)\pi(\theta')$$

- where $\pi(\theta)$ is the probability distribution we wish to sample from
- The update rule is to choose a nearby value θ' , compute $r = \pi(\theta')/\pi(\theta)$ and accept the update with probability $\min(1, r)$
- We require that the probability of choosing θ from θ' is the same as the reverse

What Makes MCMC Nasty

- It can take a long time until our states occur with the probability π (i.e. we have forgotten our initial state)
- We don't even know how long we have to wait
- Even when we have reached this *equilibration time* each sample is correlated with the previous sample
- To get a good approximation to the posterior expectation requires running for many times the equilibration time
- Note, if we are just finding sample averages then we can use all samples after equilibrating even if they are not independent

Proposals and Metropolis-Hastings

- We have some freedom in choosing a new proposal θ' from our current position θ —a good choice can increase the acceptance rate making the MCMC more efficient
- We define the proposal distribution $p(\theta'|\theta)$
- For the standard Metropolis algorithm to work we require $p(\theta'|\theta) = p(\theta|\theta')$
- In some cases (e.g when $\theta_i \geq 0$) this can be hard to achieve
- We can modify our update rule to accept a move with probability

$$\min\left(1, \frac{p(\theta|\theta')f(\mathcal{D}|\theta')f(\theta')}{p(\theta'\theta)f(\mathcal{D}|\theta)f(\theta)}\right)$$


Proposal Distribution

- If we can choose our proposal distribution $p(\mu'|\mu)$ to be close to the posterior distribution then our acceptance rate would be close to 1
- We choose $p(\mu'|\mu) = \text{Gam}(\mu'|\mu, \mu^2)$ which has $\mathbb{E}[\mu'] = \mu$ and variance 1
- We update with probability $\min(1, r)$ where

$$\begin{aligned} r &= \frac{\text{Gam}(\mu|\mu'^2, \mu') \frac{1}{\mu'} \prod_{i=1}^n \text{Poi}(N_i|\mu')}{\text{Gam}(\mu'|\mu^2, \mu) \frac{1}{\mu} \prod_{i=1}^n \text{Poi}(N_i|\mu)} \\ &= \frac{\mu \text{Gam}(\mu|\mu'^2, \mu')}{\mu' \text{Gam}(\mu'|\mu^2, \mu)} e^{-n(\mu' - \mu) + \sum_{i=1}^n N_i \log\left(\frac{\mu'}{\mu}\right)} \end{aligned}$$

- Because we are free to choose where we move (and choose close by neighbours) $\pi(\theta') \approx \pi(\theta)$ so that moves are not too infrequent
- Also very importantly the updates depend only on the ratio $\pi(\theta')/\pi(\theta)$
- We only need to know our probabilities up to a multiplicative scaling factor
- For sampling from the posterior we only need to know the likelihood and prior $\mathbb{P}(\mathcal{D}|\theta)\mathbb{P}(\theta)$ (or $f(\mathcal{D}|\theta)f(\theta)$)
- We don't need to know $\mathbb{P}(\mathcal{D})$ which we generally don't know

Burn-In



Traffic Rate


- Consider monitoring the flow of traffic where we have data

$$\mathcal{D} = (N_1, N_2, \dots, N_n)$$

where N_i is the number of car that past on day i

- We assume $N_i \sim \text{Poi}(\mu)$ and want to infer μ
- The Poisson distribution has a beta conjugate prior
- We don't have any prior knowledge on μ so we use a non-informative prior $\text{Gam}(\mu|0,0) = 1/\mu$
- Note that we can solve this problem exactly—however, lets compare with MCMC

MCMC in Practice



- To compute correct histograms you need to count samples where no move is made multiple times
- On modern computers its quite quick to compute millions of samples
- The code is not very difficult to write (although care is need to get everything correct)
- This can be used on complicated problems such as topic models (LDA) with thousands of parameters
- The accuracy of MCMC is slow if it takes a long time to sample the posterior distribution

Conclusions

- As soon as we use complex models we are no longer able to compute the posterior in closed form
- Monte Carlo techniques and particularly MCMC are a very general method for computing samples from the posterior
- These techniques have been highly developed, but very frequently even simple implementations are sufficient to do good inference
- MCMC is very easy to use, although it can be slow
- There exist other methods such as variational approximations that can be much quicker, but are approximations

Question 1

- How does the transformation method work to generate random deviates?
 1. You transform from one distribution to another
 2. You generate a uniformly distributed random number and apply the inverse probability distribution function to it
 3. You generate a uniformly distributed random number and apply the inverse cumulative distribution function to it
 4. You generate a uniformly distributed random number and apply the probability distribution function to it
- Answer: 3. You generate a uniformly distributed random number and apply the inverse cumulative distribution function to it


Question 3

- What is detailed balance?
 1. It means you visit all possible values of the random variable equally often
 2. It is a Markov Chain that converges to a particular distribution
 3. It is the condition that the transition matrix is stochastic
 4. It is the equation $M_{ij}\pi_j = M_{ji}\pi_i$
- Answer: 4. It is the equation $M_{ij}\pi_j = M_{ji}\pi_i$

- MCMC provides a means to accurately sample from very complex models
- There have been many advanced techniques developed to improve MCMC performance
- E.g. hybrid MCMC simulates a dynamics to find good proposals with similar probability far from the starting point
- Often it seems that MCMC is complicated because there are so many optimisations, but often simple implementations are sufficient

Outline

1. Sampling
2. Random Number Generation
3. MCMC
4. Quiz



Question 2

- How does the rejection method work to generate random deviates?
 1. You generate from a nearby distribution and apply a transformation to the actual distribution
 2. You generate a random deviate, Y with a distribution g_Y such that $\forall x, cg_Y(x) \geq f_X(x)$ and then choose a uniform deviate U and accept Y if $Ucg_Y(Y) \leq f_X(Y)$
 3. You generate a random deviate, Y with a distribution g_Y such that $\forall x, cg_Y(x) \leq f_X(x)$ and then choose a uniform deviate U and accept Y if $Ucg_Y(Y) \geq f_X(Y)$
- Answer: 2. You generate a random deviate, Y with a distribution g_Y such that $cg_Y(x) \geq f_X(x)$ and then choose a uniform deviate U and accept Y if $Ucg_Y(Y) \leq f_X(Y)$

Question 4


- What statement is **false** about the transition matrix M_{ij} ?
 1. Each column sums to one, i.e. it is a stochastic matrix
 2. The absolute size of the eigenvalues are all less than or equal to 1
 3. Both its rows and columns sum to one, i.e. it is doubly stochastic
 4. The equilibrium state correspond to the eigenvectors with eigenvalue 1
- Answer: 3. Both its rows and columns sum to one, i.e. it is doubly stochastic—**this is not true**

Question 5

- What is the difference between the Metropolis algorithm and Metropolis-Hastings
 1. They are the same algorithm
 2. In Metropolis-Hastings you chose the proposal move from a non-symmetric distribution which you compensate for in the acceptance criteria
 3. You choose a transition probability that does not satisfy detailed balance
 4. You perform a molecular dynamics simulation to obtain the proposal move
- Answer: 2. In Metropolis-Hastings you chose the proposal move from a non-symmetric distribution which you compensate for in the acceptance criteria

Question 6

- What statement is *false* about MCMC
 1. Once you have reached an equilibrium state every state after that is an independent deviate
 2. It is useful for Bayesian inference because you don't have to compute the marginal likelihood
 3. Generally it is difficult to determine if you have reached an equilibrium state
 4. To compute expectation you can use every state once you have reached equilibrium, but you need to count states if they are accepted or rejected, as well as run for a long time
- Answer: 1. Once you have reached an equilibrium state every state after that is an independent deviate—this is **false**

Entropy*Entropy, Coding, Maximum Entropy***Measuring Uncertainty**

- What is more uncertain tossing a coin three times or throwing a dice?
- The answer depends on whether you care about the order of the coin tosses?
- But, how do we answer such a question?
- Let X be a random variable denoting the possible outcomes
- Interestingly, Shannon entropy give a precise answer

$$H_X = - \sum_{x \in \mathcal{X}} \mathbb{P}(X = x) \log_2(\mathbb{P}(X = x))$$

Unordered Coin Toss


- What if we don't care about the order of the out-come then $\mathbb{P}(HHH) = \mathbb{P}(TTT) = 1/8$, $\mathbb{P}(HHT) = \mathbb{P}(HTT) = 3/8$ so

$$H_U = -\frac{1}{4} \log_2\left(\frac{1}{8}\right) - \frac{3}{4} \log_2\left(\frac{3}{8}\right) \approx 1.811 \text{ bits}$$


- This seems reasonable, although it is not obvious how you would determine this without using entropy
- But why Shannon entropy?

Outline

1. Measuring Uncertainty
2. Code Length
3. Maximum Entropy



1. Measuring Uncertainty
2. Code Length
3. Maximum Entropy

**Let's Calculate**

- For an honest dice $D \in \{1, 2, 3, 4, 5, 6\}$ and $\mathbb{P}(D = i) = 1/6$ so

$$H_D = - \sum_{i=1}^6 \frac{1}{6} \log_2\left(\frac{1}{6}\right) = -\log_2\left(\frac{1}{6}\right) = \log_2(6) \approx 2.584 \text{ bits}$$

- For an honest coin where we care about the order so $C \in \{000, 001, \dots, 111\}$ the $\mathbb{P}(C = i) = \frac{1}{8}$ and

$$H_C = - \sum_{i=0}^7 \frac{1}{8} \log_2\left(\frac{1}{8}\right) = -\log_2\left(\frac{1}{8}\right) = \log_2(8) = 3 \text{ bits}$$

- This clearly makes sense: there are more possible outcomes; all equally likely

Additive Entropy

- If H_X and H_Y is the uncertainty of two independent random variable X and Y , what is the uncertainty of the combined event (X, Y) ?

$$\begin{aligned} H_{(X,Y)} &= - \sum_{X,Y} \mathbb{P}(X,Y) \log_2(\mathbb{P}(X,Y)) \\ &= - \sum_{X,Y} \mathbb{P}(X) \mathbb{P}(Y) \log_2(\mathbb{P}(X) \mathbb{P}(Y)) \\ &= - \sum_{X,Y} \mathbb{P}(X) \mathbb{P}(Y) (\log_2(\mathbb{P}(X)) + \log_2(\mathbb{P}(Y))) \\ &= - \sum_X \mathbb{P}(X) \log_2(\mathbb{P}(X)) - \sum_Y \mathbb{P}(Y) \log_2(Y) = H_X + H_Y \end{aligned}$$

- Shannon's entropy is one of the few functions that satisfy this condition

Why Measure Entropy in Bits

- Suppose we had to communicate a message with 2^n equally likely outputs (e.g. the result of n -coin tosses)
- We can do this with a binary string with n bits (011..0)
- If there were 5 possible outcomes I could do this with 3 bits, but waste $3/8$ of the message
- However if we have a batch of 3 independent messages each with 5 outcomes then there are 125 possible outcomes. We could communicate this with 8 bits. This would waste $3/128$ of the message
- By batching together enough messages with N outcomes then we asymptotically need just $\log_2(N)$ bits

Different Probabilities


- We “showed” that if we had N events, X_i , each with probability, $\mathbb{P}(X_i) = 1/N$, we can code the outcomes with a message of length $-\log_2(\mathbb{P}(X_i)) = \log_2(N)$
- With a shorter message we would not be able to distinguish all possible outcomes from the message
- What happens if some of outcomes occur with a different probability

X_i :	1	2	3	4	5	6
$p(X_i)$:	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{4}$	$\frac{1}{4}$
Code:	000	001	010	011	10	11
$L = -\log_2(p(X_i))$:	3	3	3	3	2	2

Real Codes


- Those of you with some computer science background will realise that we can't actually use different length strings in a code without paying some price
- We won't know where a code word ends so we can't decode the message
- An optimal solution is to use Huffman encoding where we associate the leaf of a tree with each code word
- Using the tree we can decode any message constructed using the tree
- There is a greedy algorithm for constructing the optimal tree

Coding Normals to Accuracy Δx



Outline

- Measuring Uncertainty
- Code Length
- Maximum Entropy




Shannon's Entropy

- If the probabilities are not equal to 2^{-n} we can still find a code with a length very close to $-\log_2(\mathbb{P}(X))$ per message by transmitting a large number of messages
- The length of the message measures the amount of **surprise** on receiving the message
- Shannon's entropy is the expected length of the message to communicate a random variable X

$$H_X = \mathbb{E}_X[-\log_2(\mathbb{P}(X))] = -\sum_{x \in \mathcal{X}} \mathbb{P}(X = x) \log_2(\mathbb{P}(X = x))$$

- The expected length is a measure of the uncertainty (how much information on average we need to convey the outcome)

Coding Normals



bits and nats

- We have measured entropy in **bits** using

$$H_X = -\sum_{x \in \mathcal{X}} \mathbb{P}(X = x) \log_2(\mathbb{P}(X = x))$$


- Sometimes it is easier to use natural logarithms

$$H_X = -\sum_{x \in \mathcal{X}} \mathbb{P}(X = x) \ln(\mathbb{P}(X = x))$$

- In this case the entropy is measured in **nats** with 1 nat equal to $\log_2(e)$ bits
- This is often easier when we want to do calculus on entropy

Number of States

- Suppose I have N balls I then put in K boxes so that n_i balls finish up in box i



- Counting the number of ways to do this

$$\mathbb{P}(\mathbf{n}) = \frac{N!}{n_1! n_2! \dots n_K!}$$

Stirling's Approximation

- We can approximate the factorial $n!$ using **Stirling's approximation**

$$n! \approx \sqrt{2\pi n} n^n e^{-n}$$


$$\log(n!) = n \log(n) - n + \frac{1}{2} \log(2\pi n)$$

- Using this in our formula for $\mathbb{P}(\mathbf{n})$ we have

$$\mathbb{P}(\mathbf{n}) = \frac{N!}{n_1! n_2! \dots n_K!} = e^{\log(N!) - \sum_i \log(n_i!)}$$

$$\approx C e^{N \log(N) - \sum_i n_i \log(n_i)} \approx C e^{-N \sum_i \frac{n_i}{N} \log\left(\frac{n_i}{N}\right)}$$

Order n terms cancel as $N = \sum_i n_i$



Maximum Entropy Method

- When we are trying to infer a distribution given some observations then we can maximise the entropy subject to constraints—the entropy acts as a prior
- This is known as the **maximum entropy method**
- We can rationalise this as this is by far the most likely set of configurations consistent with the observations
- Alternatively we can see this as maximising our uncertainty given what we know—being as unbiased as possible
- It only gives a good approximation if all possibilities are equally likely

Normal Distribution

- We have three constraints

$$\int e^{-1+\lambda_0+\lambda_1 x+\lambda_2 x^2} dx = 1$$

$$\int e^{-1+\lambda_0+\lambda_1 x+\lambda_2 x^2} x dx = \mu$$

$$\int e^{-1+\lambda_0+\lambda_1 x+\lambda_2 x^2} x^2 dx = \mu^2 + \sigma^2$$

- Solving for λ_0 , λ_1 and λ_2 then

$$f_X(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/(2\sigma^2)}$$

- That is, the normal distribution is the maximum entropy distribution given we know the mean and variance

Historic Entropy

- Historically entropy was first introduced in statistical physics by Rudolf Clausius in 1865 (although Macquorn Rankine discussed it in 1850)
- Its interpretation as the number of states was introduced by Ludwig Boltzmann
- The person who got it all right was Josiah Willard Gibbs (and James Clerk Maxwell)
- Claude Shannon invented information theory based on entropy around 1948 (more on that in the next lecture)
- Ed Jaynes was the first to understand that statistical physics can be seen as an inference problem

Number of States and Entropy

- Let $p(x_i) = n_i/N$ be the proportion of balls in bin i then

$$\mathbb{P}(\mathbf{n}) \approx C e^{N H_X}$$

where

$$H_X = -\sum_i p(x_i) \log(p(x_i))$$

- That is, the “entropy” can be seen as a measure of the logarithm of the number of configurations
- When the number of balls, $N \rightarrow \infty$ the overwhelmingly likely configurations is the one that maximises the entropy subject to the observed mean and variance

Knowing the Mean and Variance

- Consider a continuous random variable, X , with a known mean and second moment

$$\mathbb{E}[X] = \mu, \quad \mathbb{E}[X^2] = \mu^2 + \sigma^2$$

- To maximise the entropy subject to constraints consider

$$\mathcal{L}(f) = -\int f_X(x) \log(f_X(x)) dx + \lambda_0 \left(\int f_X(x) dx - 1 \right)$$

$$+ \lambda_1 \left(\int f_X(x) x dx - \mu \right) + \lambda_2 \left(\int f_X(x) x^2 dx - \mu^2 - \sigma^2 \right)$$

- Thus

$$\frac{\delta \mathcal{L}(f)}{\delta f_X(x)} = -\log(f_X(x)) - 1 + \lambda_0 + \lambda_1 x + \lambda_2 x^2 = 0$$

- Or


$$f_X(x) = e^{-1+\lambda_0+\lambda_1 x+\lambda_2 x^2}$$

Using Maximum Entropy

- Maximum entropy is often used to infer distributions
- It can be very effective, but it might not work well if there are other constraints that we have not included
- The place that they work superbly well is in statistical physics
- The whole of statistical physics is about inferring distributions making observations of volume, pressure, etc.
- Temperature appears rather strangely as a Lagrange multiplier

Conclusion

- Entropy provides a measure of the disorder or uncertainty in a system
- It forms the basis of information theory which we will look at in the next lecture
- $-\log(\mathbb{P}(X = x))$ can be seen as the minimum length of a message to communicate x
- This will be used as the basis of the minimum description length formalism also discussed in the next lecture
- Entropy can be used as a prior, which we often maximise subject to constraints to obtain an unbiased estimate



Information, KL-divergence, Minimum Description Length

Communicating Via a Noisy Channel

- Information theory considers communicating down a (noisy) channel

$$X \sim \mathbb{P}(X) \xrightarrow{\text{noisy channel}} Y \sim \mathbb{P}(Y | X)$$

- We send a message X (with probability $\mathbb{P}(X)$) and receive a message Y with probability $\mathbb{P}(Y | X)$
- The uncertainty of the message sent, given we received a message y is

$$H_{X|Y=y} = - \sum_{x \in \mathcal{X}} \mathbb{P}(X = x | Y = y) \log(\mathbb{P}(X = x | Y = y))$$

- The expected uncertainty in the message sent is

$$H_{X|Y} = \sum_{y \in \mathcal{Y}} \mathbb{P}(Y = y) H_{X|Y=y} = - \sum_{x,y} \mathbb{P}(X = x, Y = y) \log(\mathbb{P}(X = x | Y = y))$$

Mutual Information

- The amount of uncertainty about the message being sent, X , before receiving the message is $H_X = -\mathbb{E}_X[\log \mathbb{P}(X)]$
- Shannon define the *mutual information* to be the expected loss in uncertainty when we receive a message

$$I_{X;Y} = H_X - H_{X|Y}$$

- Since $H_X - H_{X|Y} = H_Y - H_{Y|X}$ it follows

$$I_{X;Y} = I_{Y;X}$$

Independent Noise


- The simplest model of a noisy channel is a binary channel where each symbol is corrupted independently with a probability f

$$\mathbb{P}(X = 1 | Y = 0) = \mathbb{P}(X = 0 | Y = 1) = f$$


- An elementary calculations shows that

$$H_{X_i|Y_i} = -(1-f) \log(1-f) - f \log(f) = H(f)$$

- For a message of length n , $H_{X|Y} = nH(f)$



1. Information Theory
2. KL-Divergence
3. Minimum Description Length
4. Variational Auto-Encoders



Joint Entropy

- We can define the **joint entropy**

$$H_{X,Y} = - \sum_{x,y} P_{X,Y}(x,y) \log(P_{X,Y}(x,y))$$

- If the message we receive is independent of the message that is sent then $H_{X,Y} = H_X + H_Y$ (we saw this in the last lecture)

- $H_{X,Y} \neq H_X + H_Y$ if X and Y are correlated


- Since $\mathbb{P}(X,Y) = \mathbb{P}(Y|X)\mathbb{P}(X) = \mathbb{P}(X|Y)\mathbb{P}(Y)$ it follows

$$H_{X,Y} = H_X + H_{Y|X} = H_Y + H_{X|Y}$$

- Or $H_X - H_{X|Y} = H_Y - H_{Y|X}$

Channel Capacity

- We can summarise these relationships diagrammatically




- Shannon defined the *capacity* of a noisy channel as

$$C = \max_{\mathbb{P}(X)} I_{X;Y}$$

- That is, you choose the probability distribution of the message to maximise the information gain

Error Correcting Codes

- To reduce the chance of misinterpreting a message we need to build an error correcting code
- We can do this dividing the space of binary messages into a set of Hamming balls



- A Hamming ball $B(x, r)$ is the set of strings that differ from n -dimensional binary string, x , by at most r digits

- The expected number of errors in a string of length n given an error rate of f is nf
- For sufficiently large n we would expect all errors are smaller than $(f + \epsilon)n$ (for $\epsilon > 0$)
- If we make the radius of the Hamming ball $r = (f + \epsilon)n$ ($\epsilon > 0$) then we would expect no error for sufficiently large n
- An upper bound on the number of code words we can send in a string of length n is

$$\frac{2^n}{|B(x_i, r)|} = \frac{2^n}{\binom{n}{r}} = 2^n e^{-\log(n!) + \log(r!) + \log((n-r)!)} = c\sqrt{n}2^{I_{X;Y}}$$

Using Mutual Information

- Mutual information is used quite often in machine learning
 - Wikipedia mentions 14 applications
- Suppose we want to align two sets of images through some non-linear transformations
- One way of doing this is to choose the non-linear transformations that maximise the mutual information (or normalised mutual information) between the two sets of images

KL-Divergence

- We have met the Kullback-Leibler divergence

$$\text{KL}(p||q) = \mathbb{E}_{X \sim p(X)} \left[\log \left(\frac{p(X)}{q(X)} \right) \right]$$

$$= \mathbb{E}_{X \sim p(X)} [-\log(q(X))] - H_X$$
- Recall $-\log(q(X=x))$ is the length of code need to send a message x with a probability $q(X=x)$
- Thus $-\mathbb{E}_{X \sim p(X)} [\log(q(X))]$ is the expected length of message needed to code $X \sim p(X)$ using the optimal code for the distribution $q(X)$ rather than $p(X)$
- $\text{KL}(p||q)$ is also known as the **relative entropy** and measures the expected extra length in coding $X \sim p(X)$ if we use the wrong distribution $q(X)$

Evidence Lower Bound (ELBO)

- We can re-write the variational free energy as

$$\Phi(\phi) = - \int g(\theta | \phi) \log \left(\frac{g(\theta | \phi)}{(f(\theta, \mathcal{D})/f(\mathcal{D})) f(\mathcal{D})} \right) d\theta$$


$$= - \int g(\theta | \phi) \left(\log \left(\frac{g(\theta | \phi)}{f(\theta | \mathcal{D})} \right) - \log(f(\mathcal{D})) \right) d\theta$$

$$= -\text{KL}(g(\theta | \phi) || f(\theta | \mathcal{D})) + \log(f(\mathcal{D}))$$
- If we maximise $\Phi(\phi)$, we end up minimising the KL divergence between g and f so that $g \approx f$ and $\Phi(\phi) \approx \log(f(\mathcal{D}))$
- That is, we choose the parameters of our simple factorised distribution so that it is close to the true posterior

- Shannon also showed that choosing $2^{I_{X;Y}}$ random strings of length n the Hamming distance between balls would be at least f with high probability
- This means that we can send information at rate of $I_{X;Y}$
- The maximum rate is given by the channel capacity $\max I_{X;Y}$
- If $f = 0.1$ then $C = I_{X;Y} = 0.469$ bits so we need codes of just over twice as long to communicate accurately over a noisy channel with a 10% corruption rate
- Unfortunately, we can't efficiently decode random code positions, so although we know Shannon's bound is achievable we don't have practical codes that do this

Outline

- Information Theory
- KL-Divergence
- Minimum Description Length
- Variational Auto-Encoders



Variational Approximation

- Recall we use MCMC in Bayesian inference because the posterior distribution is too complicated to write down in closed form
- In the variational approximation we approximate the posterior distribution by a simpler (typically factored distribution), e.g.

$$f(\theta | \mathcal{D}) \approx g(\theta | \phi) = \prod_i g(\theta_i | \phi_i)$$

- The standard method for solving this is to maximise the **variational free energy**

$$\Phi(\phi) = - \int g(\theta | \phi) \log \left(\frac{g(\theta | \phi)}{f(\theta, \mathcal{D})} \right) d\theta$$

Put Another Way

- We can rewrite the variational free energy as

$$\Phi(\phi) = - \int g(\theta | \phi) \log \left(\frac{g(\theta | \phi)}{f(\theta, \mathcal{D})} \right) d\theta = L_g(\phi) + H_g(\phi)$$
 where

$$L_g(\phi) = \int g(\theta | \phi) (\log(f(\mathcal{D}) f(\theta))) d\phi$$

acts like an expected posterior that is maximised when the data is well modelled (we put $g(\theta | \phi)$ where $f(\theta, \mathcal{D})$ is large, c.f. MAP)

- The second term is an entropy

$$H_g(\phi) = - \int g(\theta | \phi) \log(g(\theta | \phi)) d\phi$$

That is, we maximise the uncertainty of the distribution $g(\theta | \phi)$

- Variational methods can be much faster than MCMC (although they tend to involve some iterations to maximise the variation free energy)■
- They can produce very good approximations, although this is not guaranteed (depends on the problem)■
- They can be extended (e.g. by minimising $\text{KL}(g\|f)$ rather than $\text{KL}(f\|g)$ —this is known as *belief propagation*)■
- MCMC is less elegant, but is a controlled approximation (we get better results by increasing the number of iterations)■
- MCMC is slower, but on modern computers this isn't usually a problem■

Compression and Model Selection

- Outside of the Bayesian framework it is difficult to do model selection—most of ML isn't Bayesian■
- When is it better to accept a more complex model for a better fit and when are we just over-fitting?■
- Usually we answer this using a validation set, but this is not always possible■
- One principled approach is to use the model that allows us to maximally compress the data■
- If we are compressing the data then we are capturing features of the data■

Description Length

- The **description length** for $\{y_i \mid i = 1, 2, \dots, m\}$ is then the cost of transmitting θ plus the cost of transmitting the errors


$$L = \sum_{k=1}^n \ell(\theta_k) - \sum_{i=1}^m \left(\log(p_\delta(y_i - \hat{f}(x_i|\theta))) + \log(\Delta) \right)$$

where $\ell(\theta_k)$ is the number of bits need to communicate θ_k (we get to choose the accuracy it is worth encoding the parameters)■


- To select between models we choose the model with the **minimum description length**■
- Note that the accuracy Δ will lead to the same cost, $-m \log(\Delta)$, for all models so doesn't affect which model is selected■

Outline

1. Information Theory
2. KL-Divergence
3. Minimum Description Length
4. **Variational Auto-Encoders**



1. Information Theory
2. KL-Divergence
3. **Minimum Description Length**
4. **Variational Auto-Encoders**



Alice and Bob

- Suppose Alice has data $\mathcal{D} = \{(x_i, y_i) \mid i = 1, 2, \dots, m\}$ while Bob has only the feature vectors $\{x_i \mid i = 1, 2, \dots, m\}$ ■
- Alice wants to communicate y_i to Bob as efficiently as possible■
- We suppose Alice & Bob have available a model $\hat{f}(x|\theta)$ ■
- Rather than sending the complete list $\{y_i \mid i = 1, 2, \dots, m\}$ Alice can send Bob the parameter θ and the errors


$$\delta_i = y_i - \hat{f}(x_i|\theta)$$

- Assuming the δ_i 's have a distribution p_δ then the cost of communicating an error to accuracy Δ is $-\log(p_\delta(\delta_i) \times \Delta)$ ■

Minimum Description Length (MDL) Method


- The minimum description length method can be a powerful way of choosing between models■
- Often it is the only principled method available■
- It allows you to trade model accuracy against model complexity■
- It can be fiddly as we need to determine the accuracy to which we should store the parameters of our model■
- This isn't something we usually think about, but often we can get very good models even when we truncate the parameters to low precision■

Variational Auto-Encoders VAE



$$\mathcal{L} = \mathbb{E}_{x \sim \mathcal{D}} [\text{KL}(q_\phi(z|x) \| \mathcal{N}(\mathbf{0}, \mathbf{I})) - \log(p_\theta(x|z(x)))]$$

Latent Space



$$\text{KL}(q_\theta(z|x) \parallel \mathcal{N}(\mathbf{0}, \mathbf{I}))$$

Understanding the Loss Function

- The original paper derived the loss function as a variational approximation to maximising some posterior
- This is difficult to understand (at least, for me)
- It has a very natural explanation in terms of minimum description length
- Alice wants to communicate the images to Bob
- Alice uses the encoder to derive a (latent) code $q(z|x)$ which she communicates to Bob
- She also communicates the errors $\delta = x - \bar{x}$
- Bob uses the decoder to decode $q(z|x)$ and δ to repair the images

Description Length

- The loss


$$\mathcal{L} = \mathbb{E}_{x \sim \mathcal{D}} [\text{KL}(q_\theta(z|x) \parallel \mathcal{N}(\mathbf{0}, \mathbf{I})) - \log(p_\theta(x|z(x)))]$$

can be interpreted as

- ★ The cost of communicating the code $\text{KL}(q_\theta(z|x) \parallel \mathcal{N}(\mathbf{0}, \mathbf{I}))$
- ★ Plus the cost to send the repair $\log(p_\theta(x|z(x)))$
- We minimise the loss function equivalent to MDL
- What is really clever is that we can choose the accuracy of the code we send $q_\theta(z|x)$ to minimise the over-all cost

Conclusions

- Information theory has regularly been used in machine learning
- It requires some understanding and care to do it properly
- The KL-divergence (or relative entropy) is often used to make two probability distribution more alike
- The minimum description length is a powerful principle for model selection
- Variational Auto-Encoders have a very natural interpretation in terms of minimising a description length



When ML Works, Bias Variance

What Makes a Good Learning Machine?

- We want to understand why some machine learning techniques work well and other don't
- To understand why these work we need to understand what makes a good learning machine
- For this we have to get conceptual and think about **generalisation** performance

generalisation: how well do we do on unseen data as opposed to the training data

Least Squared Errors

- Suppose we want to learn some output y for a feature vector x
- We construct a learning machine that makes a prediction $\hat{f}(x|\theta)$
- We typically choose the machine to minimise a *training loss*


$$L_T(\mathcal{D}) = \sum_{(x,y) \in \mathcal{D}} (\hat{f}(x|\theta) - y)^2 = \sum_{i=1}^m (\hat{f}(x_i|\theta) - y_i)^2$$

where $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^m$ is a set of size m , sampled from a probability distribution $\mu(x, y)$


- We call this machine $\hat{f}(x|\theta_{\mathcal{D}})$

Too Simple or Too Complex?

- Fit $\hat{f}(x|\theta_{\mathcal{D}})$ to data



1. What Makes a Good Learning Machine?
2. Bias-Variance Dilemma




What Makes Machine Learning Hard?

- Typically we work in high dimensions (i.e. have many features)
- The problem can be over-constrained (i.e. we have conflicting data to deal with) — solve by minimising an error function
- The problem can be under-constrained (i.e. there are many possible solutions that are consistent with the data) — need to choose a plausible solution
- Typically in machine learning the data will be over-constrained in some dimensions and under-constrained in others
- We can't visualise the data to know what is going on

Generalisation Error

- We want to minimise the *generalisation loss* which in this case is

$$L_G(\theta_{\mathcal{D}}) = \sum_{(x,y) \in \mathcal{Z}} \mu(x,y) (\hat{f}(x|\theta_{\mathcal{D}}) - y)^2$$




(we can estimate this if we have some labelled examples (x_i, y_i) which we have not trained on)


- We want to minimise $L_G(\theta_{\mathcal{D}})$ but in practice we are minimising $L_T(\mathcal{D})$ what could possibly go wrong?

Measuring Generalisation Error for Regression


- Consider the regression example. The root mean squared error is




1. What Makes a Good Learning Machine?
2. **Bias-Variance Dilemma**



Approximation and Estimation Errors



Regression Example $n = 1$



Bias and Variance

Consider the expected generalisation for data sets of size $|\mathcal{D}| = m$

$$\begin{aligned}
 \bar{L}_G &= \mathbb{E}_{\mathcal{D}}[L_G(\theta_{\mathcal{D}})] = \mathbb{E}_{\mathcal{D}} \left[\sum_{x \in \mathcal{X}} \mu(x, y) \left(\hat{f}(x|\theta_{\mathcal{D}}) - y \right)^2 \right] \\
 &= \sum_{x \in \mathcal{X}} \mu(x, y) \mathbb{E}_{\mathcal{D}} \left[\left(\hat{f}(x|\theta_{\mathcal{D}}) - y \right)^2 \right] \\
 &= \sum_{x \in \mathcal{X}} \mu(x, y) \mathbb{E}_{\mathcal{D}} \left[\left(\left(\hat{f}(x|\theta_{\mathcal{D}}) - \hat{f}_m(x) \right) + \left(\hat{f}_m(x) - y \right) \right)^2 \right] \\
 &= \sum_{x \in \mathcal{X}} \mu(x, y) \left(\mathbb{E}_{\mathcal{D}} \left[\left(\hat{f}(x|\theta_{\mathcal{D}}) - \hat{f}_m(x) \right)^2 \right] + \left(\hat{f}_m(x) - y \right)^2 \right. \\
 &\quad \left. + 2 \mathbb{E}_{\mathcal{D}} \left[\left(\hat{f}(x|\theta_{\mathcal{D}}) - \hat{f}_m(x) \right) \left(\hat{f}_m(x) - y \right) \right] \right)
 \end{aligned}$$

- Our generalisation performance will depend on our training set, $\mathcal{D}$
- To reason about generalisation we can ask what is the *expected generalisation loss*, when we average over all different data sets of size m drawn independently from $\mu(x, y)$
- For each data set, $\mathcal{D}$, we would learn a different approximator $\hat{f}(x|\theta_{\mathcal{D}})$
- Note that in practice we only get one data set. We might be lucky and do better than the expected generalisation or we might be unlucky and do worse

Mean Machine


- To help understand generalisation we can consider the mean prediction with respect to machines trained with all data sets of size m

$$\hat{f}_m(x) = \mathbb{E}_{\mathcal{D}} \left[\hat{f}(x|\theta_{\mathcal{D}}) \right]$$

- We can define the **bias** to be generalisation performance of the mean machine

$$B = \sum_{x \in \mathcal{X}} \mu(x, y) \left(\hat{f}_m(x) - y \right)^2$$

Regression Example $n = 5$



Cross Term

- The cross term vanishes

$$\begin{aligned}
 C &= \mathbb{E}_{\mathcal{D}} \left[\left(\hat{f}(x|\theta_{\mathcal{D}}) - \hat{f}_m(x) \right) \left(\hat{f}_m(x) - y \right) \right] \\
 &= \mathbb{E}_{\mathcal{D}} \left[\left(\hat{f}(x|\theta_{\mathcal{D}}) - \hat{f}_m(x) \right) \right] \left(\hat{f}_m(x) - y \right) \\
 &= \left(\mathbb{E}_{\mathcal{D}} \left[\hat{f}(x|\theta_{\mathcal{D}}) \right] - \hat{f}_m(x) \right) \left(\hat{f}_m(x) - y \right) \\
 &= \left(\hat{f}_m(x) - \hat{f}_m(x) \right) \left(\hat{f}_m(x) - y \right) = 0
 \end{aligned}$$

- Thus

$$\bar{L}_G = \sum_{x \in \mathcal{X}} \mu(x, y) \mathbb{E}_{\mathcal{D}} \left[\left(\hat{f}(x|\theta_{\mathcal{D}}) - \hat{f}_m(x) \right)^2 + \left(\hat{f}_m(x) - y \right)^2 \right]$$

- We can write the expected generalisation loss as

$$\mathbb{E}_{\mathcal{D}}[L_G(\boldsymbol{\theta}_{\mathcal{D}})] = \sum_{\mathbf{x} \in \mathcal{X}} \mu(\mathbf{x}, y) \mathbb{E}_{\mathcal{D}} \left[\left(\hat{f}(\mathbf{x} | \boldsymbol{\theta}_{\mathcal{D}}) - \hat{f}_m(\mathbf{x}) \right)^2 \right] \\ + \sum_{\mathbf{x} \in \mathcal{X}} \mu(\mathbf{x}, y) \left(\hat{f}_m(\mathbf{x}) - y \right)^2 = V + B$$

- Where B is the bias and V is the variance defined by

$$V = \sum_{\mathbf{x} \in \mathcal{X}} \mu(\mathbf{x}, y) \mathbb{E}_{\mathcal{D}} \left[\left(\hat{f}(\mathbf{x} | \boldsymbol{\theta}_{\mathcal{D}}) - \hat{f}_m(\mathbf{x}) \right)^2 \right]$$

- Warning some of you will see a similar, but subtly different calculation elsewhere—here $\bar{L}_G = B + V$

Balancing Bias and Variance

- We want to choose a learning machine that is complex enough to capture the underlying function we are trying to learn, but otherwise as simple as possible
- There are a number of tricks to achieve this balance
- Some require us to preprocess the data to reduce the number of inputs
- Some machines cleverly adjust their own complexity
- This course looks at machines that achieve this balance

- The bias measure the generalisation performance of the *mean machine* and is large if the machine is too simple to capture the changes in the function we want to learn
- The variance measures the variation in the prediction of the machines as we change the data set we train on

$$V = \sum_{\mathbf{x} \in \mathcal{X}} \mu(\mathbf{x}, y) \mathbb{E}_{\mathcal{D}} \left[\left(\hat{f}(\mathbf{x} | \boldsymbol{\theta}_{\mathcal{D}}) - \hat{f}_m(\mathbf{x}) \right)^2 \right]$$

- The variance is usually large if we have a complex machine
- Striking the right balance is often the key to getting good results

Lessons

- This course is about understanding machine learning techniques that work well
- Which one to use will depend on the data set
- One of the most useful intuitions about what works is the bias-variance framework
- The bias is high for simple machines that can't capture the data
- The variance is high for complex machines that are sensitive to the training set
- Good machines are powerful enough to capture complex data sets, but they can control their own capacity (ability to (over-)fit the data)

Over-Fitting

Overfitting, regularisation, feature selection

Over-fitting

- Complex machine can **over-fit**
over-fitting: fitting the training data well at the cost of getting poorer generalisation performance
- Three red cars. . .
- If we used an infinitely flexible machine we can fit our training data perfectly, but would have no generalisation ability

1. **Over-fitting?**
2. Controlling Complexity
3. Hidden structure
4. Regularisation



Binary Classification Task for You



Class 1

Class 2

Which Category?

- Which category does the following image belong to?



Spurious Rules

- You ask a learning machine to solve a task based on data
- It will find a rule that does this, but not necessary the rule you had in mind—machine learning can't read your mind
- Infinitely flexible machines have an infinity of spurious rules they can learn—they are useless
- What should we do?

All Binary Functions

$$x_0 = 000 \quad y_0 = \begin{cases} 0 \\ 1 \end{cases}$$

$$x_1 = 100 \quad y_1 = \begin{cases} 0 \\ 1 \end{cases} \quad \text{unseen}$$

$$x_2 = 010 \quad y_2 = \begin{cases} 0 \\ 1 \end{cases}$$

$$x_3 = 110 \quad y_3 = \begin{cases} 0 \\ 1 \end{cases}$$

$$x_4 = 001 \quad y_4 = \begin{cases} 0 \\ 1 \end{cases} \quad \text{seen}$$

$$x_5 = 101 \quad y_5 = \begin{cases} 0 \\ 1 \end{cases}$$

$$x_6 = 011 \quad y_6 = \begin{cases} 0 \\ 1 \end{cases}$$

$$x_7 = 111 \quad y_7 = \begin{cases} 0 \\ 1 \end{cases} \quad \text{unseen}$$

$$\mathcal{D} = \{(000, 0), (010, 1), (110, 1), (001, 0), (101, 0)\}$$

Are MLPs Universal Approximators?

- Yes and No
- Yes: If you give me any function, I can find an MLP that approximates that function to any desired accuracy
- No: If you give me an MLP, I can find a function with an arbitrary high approximation error
- Would an MLP that could approximate any function be useful?
- **Absolutely not!**

1. Over-fitting?
2. **Controlling Complexity**
3. Hidden structure
4. Regularisation



Training Examples

- As we increase the number of training examples, we make it hard to find a spurious rule
- Bigger data sets allow us to use more complicated machines
- Part of the success of deep learning is because they use huge training sets—but this is only a part of their success
- (Labelled) data is often expensive to collect so we sometimes have no choice but to use a small training set
- One of the limitations of using deep learning comes because we often have limited data

Preprocessing

- Structure might often be obscure to the learning machine
- If we are trying to predict the spread of disease then a list of place names might be a lot less useful than their coordinates
- Imposing an ordering on an unordered set might **not** be useful

$\{\text{"blue"} : 0, \text{"brown"} : 1, \text{"green"} : 2, \text{"black"} : 3\}$

- Choosing an encoding that reflect meaningful structure is essential to successful machine learning

Outline

1. Over-fitting?
2. Controlling Complexity
3. **Hidden structure**
4. Regularisation



- Infinitely flexible machine don't generalise (because any unseen data could have any value)
- **Machine learning only works because there is some structure in the data**
- A successful machine should capture this structure
- Even deep learning machines with millions or billions of parameters only work because they successfully capture the structure of images or text
- Different learning machines have different performance on different problems because the data has different structure

Identifying Structure

- In some cases we know *a priori* some of the structure in the data
- In images we believe the identity of an object is invariant to translation and scaling
- The success of *convolutional neural networks* (CNNs) in deep learning is in large part because the convolutions respect translational invariance

Automatic Preprocessing

- One view of deep learning is that each layer (particularly in CNNs) acts as a preprocessor
- That is, it finds filters that captures features salient to the problem being tackled
- For both images and texts we expect salient features to be spatially localised (CNN finds localised filter)
- The deep structure allows ever more complicated features to be captured—that is, we can find spatially localised features on different scales
- Having very large datasets and simple filters (the number of weights in the CNN layers tends to be small) stops overfitting

Hidden Structure

- Often the structure of data is invisible to us
- A very successful strategy is to try many different machine learning techniques and choose the best (stupid but effective)
- Often learning machines have adjustable parameters (hyper-parameters) that we have to set (they are the same for all input data, but change with the problem)
- We need to choose the hyper-parameters to fit the data in our problem
- Fine tuning hyper-parameter is important and almost always required (true in SVMs, MLP, deep learning)

Measuring Generalisation Performance

- Recall, we want to predict **unseen** data
- You cannot use data that you have trained on!**—you will overfit
- Need to split your data up into training and validation set
- Use the validation set to choose the hyper-parameters
- You need a separate testing set if you want to measure your generalisation performance

Cross Validation

- If you want to use more data for training then you can use cross validation
- K -fold cross validation splits the data into K groups

$$\mathcal{D} = \{D_i\}_{i=1}^P, \quad D_i = (\mathbf{x}_i, y_i)$$

$D_1, D_2, D_3, D_4, D_5, D_6, D_7, D_8, D_9, D_{10}, D_{11}, D_{12}, D_{13}, D_{14}, D_{15}, D_{16}, D_{17}, D_{18}, D_{19}, D_{20}$
 Termination condition: $\text{Sum}(\text{Set}) \leq \text{Sum}(\text{Set})$
 Leaf node: $\text{Sum}(\text{Set})$

$$\frac{5.5 \pm 5.5 \pm 1.8 \pm 1.8 \pm 3.7 \pm 3.6 \pm 7.4 \pm 4.6 \pm 0.99 \pm 4.5 \pm 4.6 \pm 5.4 \pm 6.2 \pm 3.3 \pm 2.7}{9.9} = 4.3$$

- Leave-one-out cross-validation is extreme case

Dimensionality Reduction


- We can sometimes simplify our machines by using less features
- We can project our data onto a lower dimensional sub-space (e.g. one with the maximum variation in the data: PCA)
- We can use clustering to find exemplars and recode our data in terms of distances from the exemplars (radial basis functions)
- Whether this helps depends on whether the information we discard is pertinent to the task we are trying to perform

Normalising Features

- Measuring a feature in millimeters or kilometers is going to make a lot of difference to the size of that feature
- Many learning algorithms are sensitive to the size of a feature (larger features are more important)
- If we don't know how important different features are then it makes sense to normalise our features E.g.


$$x_i^\alpha \leftarrow \frac{x_i^\alpha - \hat{\mu}_i}{\hat{\sigma}_i}, \quad \hat{\mu} = \frac{1}{m} \sum_{\beta=1}^m x_i^\beta, \quad \hat{\sigma}_i^2 = \frac{1}{m-1} \sum_{\beta=1}^m (x_i^\beta - \hat{\mu}_i)^2$$

The Overfitting Game



Hidden Structure

Can't spot low dimensional data by looking at numbers



Feature Selection

- Spurious features will allow us to find spurious rules (**over-fitting**)
- We can try different combinations of features to find the best set, although it rapidly becomes intractable to do this in all ways
- We can use various heuristics to decide which features to keep, but no heuristic is fail-safe method to find the best set of features
- Feature selection however can be powerful, often we can get very good results by keeping only a few variables
- As well as possibly improving generalisation we also get a more **interpretable** rule

Outline

1. Over-fitting?
2. Controlling Complexity
3. Hidden structure
4. **Regularisation**



Explicit Regularisation

- As you've seen in the foundations of ML course, we can modify our error function to choose smoother functions

$$L = \sum_{k=1}^m (w^T x_k - y_k)^2 + \nu \|w\|^2$$

(Good to normalise features)

- Second term is minimised when $w_i = 0$
- If w_i is large then

$$f(x|w) = w^T x = \sum_{i=1}^p w_i x_i$$

varies rapidly as we change x_i

Implicit Regularisation

- In the last two examples we added an explicit regularisation term that made the function we learnt simpler
- Some learning machines do this less explicitly
- Some deep learning architectures do subtle averaging
- Sometimes the architecture biases the machine to find a simple solution

Success of SVMs


- SVMs regularise themselves by choosing the machine with the largest margin
- This ensures maximum stability to noise on the data
- It leads to very good generalisation on small datasets—usually beats everything else
- But you still need to normalise the features
- You also need to tune its hyper-parameters (C and sometimes γ)

Lasso

- We can use other regularisers


$$L = \sum_{k=1}^m (w^T x_k - y_k)^2 + \nu \sum_{i=1}^p |w_i|$$

- Spurious features (e.g. colour of flag on energy consumption) will give us a small improvement in training error



Maximum Margin Machines

- Perceptrons have many options to separate data



- SVMs choose the machine with the biggest margins

Lessons

- Machine learning isn't magic
- It works when the learning machine is well attuned to the problem
- Sometimes you can help by preprocessing your data
- Sometimes there is a regularisation term that helps select a simpler machine
- Most machines have hyper-parameter that you tune to match the machine to the data
- Really clever machines try to do this matching automatically

Symmetry



Inductive Bias, Symmetry, Invariance, Group theory

Inductive Bias

- Learning machines differ in how they split up the space of inputs
- Perceptrons split the space using a hyperplane, RBFs split the space depending on some set of centres, MLPs successive split the world at different layers
- These differences are known as the **inductive bias** of the machine
- The different inductive bias lead to different performance
- How can we exploit the inductive bias to get good performance?

Smoothness

- A natural assumption is smoothness (the classification surface or regression surface changes slowly as we change the input)
- Often there is a hyperparameter that controls this
 - ★ The kernel parameter γ in SVMs (or ℓ in Gaussian Processes)
 - ★ k in k -nearest neighbours
 - ★ The size of the L_2 regularisation term in neural networks
- We often have to choose these using trial and error (using a validation set)

Symmetry

- There are sometimes symmetries in the world we can exploit
- In image classification we consider an object should be classified the same wherever it is in the image
- The classification should be invariant to translations (scaling and perhaps rotations)
- The classification function possesses a translational symmetry
- Learning machines that respect translational invariance have far less propensity to overfit—they don't need to learn that the same object in a different part of the image is the same


Outline

1. **Inductive Bias**
2. **Invariance**
3. **Group Theory**



The Real World

- The problems that we care about come from a highly structured world where the solutions inherit some of the properties of that world
- Usually similar input have similar outputs (we don't expect to change one pixel in an image and that it should be classified differently)
- The classes are often quite well separated




Outline

1. **Inductive Bias**
2. **Invariance**
3. **Group Theory**




Convolutional Neural Networks




- CNNs are not translationally invariant
- The output will be different if we shift the input
- They are (up to edge effects) **equivariant** in that we would get the same result if we shifted both the input and the output in the same way
- They are not perfectly equivariant (they break at the pixel level and at the full image level—due to edges)
- But they are sufficient good that CNNs have dominated computer vision for years (ignoring visual transformers)

Scale Symmetry

- CNNs are not scale invariant
- There are various pyramid image processing techniques to try and build in approximate scale invariance
- However, in many image classification tasks we concentrate on the object in the foreground
- Many datasets rescale objects in the image so they are roughly the same size



Augmentations



- Data augmentation is standard practice in deep learning for images and makes a very considerable difference

Normalisation

- Sometimes clever normalisation can make a lot of difference to the performance of an algorithm
- We might do this by feature engineering (divide the length on the nose by the width between the persons eyes)
- We might do this by normalising the output—use cosine similarity rather than distance

$$\cos(\theta) = \frac{\mathbf{a}^T \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}$$

- These make the results invariant to the distance of the camera or the length of $\mathbf{a}$ and $\mathbf{b}$

- The identity of an object should not change if we move the camera further away or rotate the camera
- CNNs are not naturally scale or rotationally invariant (although there are attempts to remedy this)
- Rotation invariance doesn't always reflect "the world"—people usually stand on their feet not on their heads
- Although you want rotational invariance when interpreting aerial images

Learning Invariance

- Currently we are not smart enough to come up with networks that are scale invariant or rotationally invariant or invariant to small distortions
- We therefore hope that our networks learn this by training on objects of different sizes, orientations, lighting conditions, etc.
- Knowing this is an invariance we want, it is now common practice to do data augmentation where we perform transformations that replicate physical plausible changes in the data
- We are removing spurious rules by (artificially) increasing our dataset size

Other Symmetries Exist

- In NLP it is now common to translate sentences to a foreign language and back to the original language to learn invariance to minor changes in the words used
- There are cases where we care about sets of objects, but not their ordering
- An example might be sets of objects in an image or vertices in a graph
- Here constructing operators that are equivariant (the same if we permute the inputs and outputs) can again reduce overfitting

Outline

1. Inductive Bias
2. Invariance
3. Group Theory




- Symmetries are very well studied in mathematics
- The language of symmetries are **groups**
- These capture the structure of transformations that represent some symmetry
- Group theory is highly abstract and it is easy to lose sight of its connection to symmetry, but it has proved itself extraordinarily powerful in revealing structures

Tossing a coin

- We consider a coin with two actions
 - ★ flip the coin: f
 - ★ leave the coin alone: e
- These form a group (the cyclic group C_2)
 - ★ Closure is obvious $f \cdot e = f$, $f \cdot f = e$, $e \cdot e = e$ and $e \cdot f = f$
 - ★ Associativity: e.g. $(e \cdot f) \cdot f = e \cdot (f \cdot f)$
 - ★ e is the identity
 - ★ $e^{-1} = e$ and $f^{-1} = f$

Non-commutativity

- Not all groups commute
- For 3-d rotations : $R_x \cdot R_z \neq R_z \cdot R_x$




Permutations

- The set of permutations of n objects form a group (known as the symmetric group S_n)
- The composition of a permutation is a permutation
- There is an identity (do nothing)
- For every permutation there is an inverse permutation
- Obviously S_n has $n!$ elements (you can reorder elements in $n!$ ways)
- Permutations have a lot of structure (in terms of elements cycling, subgroups, etc.)

- Groups consists of a set (we can think of the elements as transformations), $\mathcal{G} = \{a, b, \dots\}$
- And an operator “ $\cdot$ ” (which we can often think of as composition)
- That is, $a \cdot b$ we can interpret as apply transformation b and then apply transformation a
- The set and transform form a group if they satisfy four axioms
 - ★ closure: for all $a, b \in \mathcal{G}$ then $a \cdot b \in \mathcal{G}$
 - ★ associativity: $a \cdot (b \cdot c) = (a \cdot b) \cdot c$
 - ★ existence of an identity e such that $a \cdot e = a$ for all a
 - ★ Every element, a , has an inverse a^{-1} such that $a^{-1} \cdot a = e$

The Four Group

- A slightly less trivial example represents the set of transformations that leaves a rectangle invariant



Associativity

- We can think of the elements of a group as transformations acting on some object (e.g. rotations of a book)
- We consider $\cdot$ to denote composition so $a \cdot b$ denotes apply action b then action a
- When we say $c = a \cdot b$ we mean that action c is equivalent to apply action b followed by action a
- Suppose $g = a \cdot b$ and $h = b \cdot c$ then

$$g \cdot c = a \cdot h$$

$$(a \cdot b) \cdot c = a \cdot (b \cdot c)$$

- These correspond to an equivalent set of actions

Infinite Groups

- The set of integers under addition form a group
- Its closed, has an identity, 0, and each element n has an inverse $-n$
- The group describes the set of discrete shifts
- The set of rational number under addition also forms a group
- Nothing new here, you've known about these since you were a toddler

- The set of reals under addition form a group and describes translations in one dimension
- The set of reals excluding 0 form a group under multiplication with an identity 1 and inverse of x being $1/x$
- These describe scale transformation
- These are not terrifically interesting groups (or at least groups you are so familiar with that they are now boring)
- Things get more interesting in high dimensional space

Rotations $SO(n)$

- The set of $n \times n$ orthogonal matrices (matrices, $\mathbf{M}$, such that $\mathbf{M}^T \mathbf{M} = \mathbf{M} \mathbf{M}^T = \mathbf{I}$) form a group $O(n)$
- Note that $\det(\mathbf{M}) = \pm 1$
- These correspond to rotations and reflections
- The group of matrices that don't correspond to reflections (where $\det(\mathbf{M}) = 1$) form a subgroup $SO(n)$
- $SO(2)$ consist of the 2-d rotation matrices

$$\begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}$$

- Continuous groups are known as Lie Groups
- They describe things like general translational invariance, rotational invariance, relativistic invariance
- The transformations can be represented by a set of matrices
- To study these one looks at infinitesimal transformations and the algebra between these infinitesimal transformations (Lie algebras)
- This is much studied in physics and occasionally studied in machine learning, although more on the sidelines

Summary

- Invariances are important as if a machine properly captures a true invariance it can massively limit the space of functions being approximated—reducing the chances of finding spurious rules
- Often it is hard to build-in invariances into the machine learning model and practitioners will often resort to data augmentation in the hope that these invariances can be learned
- There is a mathematical language around describing the structure of symmetries: **group theory**
- You should know *group theory* exists, it is occasionally used in machine learning, but not that much
- I'm not going to expect you to know any details about group theory