
INTRODUCTION TO PROBABILISTIC MACHINE LEARNING

a handbook for innumerate computer scientists

JACOB MOSS

University of Cambridge

Last updated: December 17, 2020

Contents

| | |
|---|-----------|
| Preface | 3 |
| 1 Introduction | 4 |
| 1.1 Moments | 5 |
| 1.2 Distributions | 6 |
| 1.2.1 Continuous | 6 |
| 1.2.2 Discrete | 6 |
| 1.2.3 Mapping Between Random Variables | 7 |
| 1.3 Softmax Transformation | 8 |
| 1.4 Types of Models | 8 |
| 1.5 Graphical Models | 8 |
| 1.5.1 Directed Graphical Models | 8 |
| 1.5.2 Markov Random Fields | 9 |
| 1.6 Keep in mind | 9 |
| 2 Maximum Likelihood | 10 |
| 2.1 Least Squares Regression | 10 |
| 2.2 Regression with Additive Independent Gaussian Noise | 11 |
| 3 Bayesian Inference | 12 |
| 3.1 Maximum a posteriori | 12 |
| 3.2 Conjugate Priors | 12 |
| 3.3 Gaussian Likelihood | 13 |
| 3.4 Marginal Likelihood | 14 |
| 4 Gaussian Processes | 15 |
| 4.1 Non-parametric Modelling | 15 |
| 4.2 Definition | 15 |
| 4.3 Hyperparameters & Covariance Functions | 16 |
| 4.3.1 Hilbert Spaces | 17 |
| 4.3.2 Positive Definiteness of Inner Products in Hilbert Spaces | 17 |
| 4.3.3 Reproducing Kernel Hilbert Space | 18 |
| 4.3.4 Squared Exponential (RBF) | 18 |
| 4.3.5 Periodic | 18 |
| 4.3.6 Reproducing Kernel for Vector-Valued Functions | 19 |
| 4.4 Coregionalization | 19 |
| 4.5 Relation with Linear in the Parameters Model | 20 |
| 4.6 Eigenvectors and Relation to PCA | 21 |
| 4.7 Cholesky Decomposition | 21 |
| 4.8 Gaussian Process Example | 21 |
| 4.9 Scalability | 22 |
| 5 Gaussian Process Classification | 23 |
| 6 Monte Carlo | 24 |
| 6.1 Monte Carlo | 24 |
| 6.2 Markov Chains & Gibbs Sampling | 24 |
| 6.3 Example: Step Model | 24 |
| 7 Ranking | 27 |
| 7.1 Towards Probabilistic Ranking | 27 |
| 7.2 Gibbs Sampling in TrueSkill | 27 |
| 7.3 Message Passing on Factor Graphs | 28 |

| | | |
|-----------|---|-----------|
| 8 | Expectation-Maximisation | 30 |
| 8.1 | Gaussian Mixture Models | 30 |
| 8.2 | Mathematical Machinery | 31 |
| 8.2.1 | Jensen’s Inequality | 31 |
| 8.2.2 | Kullback-Leibler Divergence | 31 |
| 8.2.3 | Mutual Information | 31 |
| 8.3 | Expectation Maximisation Algorithm | 31 |
| 8.4 | Example: K-Means | 33 |
| 8.5 | Example: Topic Modelling | 33 |
| 8.5.1 | Latent Dirichlet Allocation | 35 |
| 9 | Variational Inference | 36 |
| 9.1 | Amortizing Variational Inference | 36 |
| 9.2 | General Form | 36 |
| 9.2.1 | Mean Field Approximation | 36 |
| 9.3 | Examples | 37 |
| 9.3.1 | Inducing Point Approximation for Gaussian Processes | 37 |
| 9.3.2 | Deep Gaussian Processes | 39 |
| 9.4 | Implementation Details | 40 |
| 10 | Stochastic Calculus | 42 |
| 10.1 | Probability Generating Function | 42 |
| 10.2 | Stochastic Process | 43 |
| 10.3 | Brownian Motion | 43 |
| 10.3.1 | Brownian Motion as a GP | 43 |
| 10.4 | Stochastic Differential Equations | 44 |
| 10.4.1 | Itô vs. Stratonovich | 44 |
| 11 | Advanced Generative Models | 45 |
| 11.1 | Variational Autoencoders | 45 |
| 11.1.1 | Reparameterisation Trick | 46 |
| 11.2 | Generative Adversarial Networks | 46 |
| 11.3 | Normalising Flows | 47 |
| A | Derivations | 49 |
| A.1 | Double Integral of a Mean Function | 49 |
| A.2 | Expansion of Strange Gaussian Identity | 49 |
| A.3 | Derivation of Biased Variance Estimator | 50 |
| A.4 | Derivations of Means and Variances | 50 |
| B | Kullback Leibler divergence | 51 |
| C | Backup | 51 |

Todo List:

1. Beta distribution derivation, arising from Bayesian inference;
2. Monte Carlo
3. Things they don’t tell us...

Preface

The purpose of this booklet is to give the foundations and intuitions for probabilistic machine learning. The **targeted audience** are Computer Scientists who might have missed out on some critical components in their mathematical education. When I started my Master's in machine learning in 2019, I found I had considerable gaps which made reading the literature both intimidating and inefficient. This is why I set out to create this booklet. I am now a PhD student at the University of Cambridge—a Bayesian stronghold—which has only solidified my passion for the subject, and I can only hope this helps others find their way into this field.

Probabilistic machine learning is a fascinating subject, and also incredibly useful in practice. The field is growing rapidly, so I will regularly update this document with new material, clarifications, and corrections. Please do contact me at jm2311@cam.ac.uk if you spot any mistakes or have any requests.

Disclaimer: Some of the structure is inspired by courses including Carl Rasmussen's Probabilistic ML (4f13) course at Cambridge and Dariush Hosseini's data mining course at UCL.

1 Introduction

In machine learning, we typically try to fit a model to a dataset. This model may be parameterised by θ . In a Bayesian model, we assign some **prior** distribution over parameters. We also have a likelihood: the probability of the data given a particular parameter setting. For example, the likelihood of a coin toss is $p(x = 1|\theta) = \theta$ where $\theta = 0.5$ for an even-weighted coin. This is an example of a discrete distribution. In a Bayesian setting, however, this θ is not fixed—it is a distribution. In the next few sections, it will become clear how a Bayesian mindset enables us to apply our **prior** knowledge that θ is around 0.5!

Colour Coding

In order to make equations easier to read, we use the following colour coding throughout:

- **prior distribution**
- **posterior distribution**
- **marginal distribution**
- **variational distribution**

n.b. the terms above will be defined at their first point of use

The key thing to understand in Bayesian statistics is the *posterior update*; using the **prior** and **likelihood**, the posterior is updated according to the data. This is achieved using Bayes' theorem: the root of all Bayesian statistics,

$$p_{\Theta}(\theta|\mathbf{x}) = \frac{p_D(\mathbf{x}|\theta)p_{\Theta}(\theta)}{p_D(\mathbf{x})},$$

where $p_{\Theta}(\theta|\mathbf{x})$ is the **posterior**, $p_D(\mathbf{x}|\theta)$ is the **likelihood**, $p_{\Theta}(\theta)$ is the **prior**, and $p_D(\mathbf{x})$ is the **evidence** or the **marginal likelihood**. The **marginal likelihood** is very important, see Section 3.4.

Note that some literature may use $lik(\theta|\mathbf{x})$, which is the *likelihood function*: it is neither a conditional probability nor a probability density. Likelihood is the probability of observing the dataset given the parameters.

Types of Probability

- **Discrete variables** have probability **mass** assigned to each event in the **countable** sample space. The Poisson distribution is an example.
- **Continuous variables** have probability **density** assigned to an **uncountable** set. The Gaussian distribution is an example.
- **Marginal:** Marginalisation is simply summing out/integrating over the probability of a R.V as follows:

$$p(x) = \sum_y p(x, y) \text{ for discrete distributions, } \int p(x, y) dy \text{ for continuous distributions}$$

- **Joint:** $p_{X,Y}(x, y) = p(X = x \text{ and } Y = y)$
- **Conditional:** $p(y|x)$

Another useful result is the law of total probability:

$$p(x) = \int p(x|y)p(y) dy$$

1.1 Moments

Moments summarise key features of distributions, for example the mean, variance, kurtosis, and skewness. We will discuss a few of these below, in addition to deriving some useful identities.

| MOMENT | EQUATION | NAME |
|--------|--|----------|
| 0 | $\int p_X(x) dx = 1$ | |
| 1 | $\mathbb{E}[X] = \int xp_X(x) dx$ | Mean |
| 2 | $\text{Var}[X] = \int (x - \mu)^2 p_X(x) dx$ | Variance |

Table 1: Moments

Expectation is a linear operator, so for **linear** f , $\mathbb{E}[f(X)] = f(\mathbb{E}[X])$. Furthermore, for linear operations the variance exhibits the following results:

$$\begin{aligned}
 \text{Var}(X) &= \mathbb{E}((X - \mathbb{E}(X))^2) = \mathbb{E}(X^2 + \mathbb{E}(X)^2 - 2X\mathbb{E}(X)) \\
 &= \mathbb{E}(X^2) - \mathbb{E}(X)^2 \\
 \implies \text{Var}(aX + b) &= \mathbb{E}((aX + b)^2) - \mathbb{E}(aX + b)^2 \\
 &= \mathbb{E}(a^2X^2 + b^2 + 2abX) - (a\mathbb{E}(X) + b)^2 \\
 &= a^2\mathbb{E}(X^2) + 2ab\mathbb{E}(X) + b^2 - a^2\mathbb{E}(X)^2 - 2ab\mathbb{E}(X) - b^2 \\
 &= a^2(\mathbb{E}(X^2) - \mathbb{E}(X)^2) \\
 &= a^2\text{Var}(X)
 \end{aligned}$$

Note that the sample mean $\hat{\mu} = \frac{1}{N} \sum_{i=1}^N X_i$ is an **estimator** of the mean, **not** a moment! The same goes for other estimators such as sample variance. These estimators are, in fact, random variables. For example:

$$\mathbb{E}[\hat{\mu}] = \mathbb{E}\left[\frac{1}{N} \sum_{i=1}^N X_i\right] = \frac{1}{N} \sum_{i=1}^N \mathbb{E}(X_i) = \mathbb{E}(X_i)$$

Therefore, as the sample size increases, the error drops proportionally to the square root.

Let us now derive¹ the same for the variance:

$$\begin{aligned}
 \mathbb{E}[s^2] &= \mathbb{E}\left[\frac{1}{N} \sum_{i=1}^N \left(X_i - \frac{1}{N} \sum_{j=1}^N X_j\right)^2\right] \\
 &= \frac{N-1}{N} \text{Var}[X_i]
 \end{aligned}$$

Because of the multiplying factor, this is called a *biased estimator*. We can make it an unbiased estimator by multiplying by $\frac{1}{N-1}$. How bad are these estimators? Well, we can interpret the standard deviation of the estimator as the error in the estimator; e.g., for the mean:

$$\text{Var}[\hat{\mu}] = \text{Var}\left[\frac{1}{N} \sum_{i=1}^N X_i\right] = \frac{1}{N} \sum_{i=1}^N \text{Var}[X_i] = \frac{1}{N} \sigma_X \implies \sigma_{\hat{\mu}} = \frac{1}{\sqrt{N}} \sigma_X$$

¹The full derivation of the estimator for variance is found in Appendix A.3.

1.2 Distributions

Probability distributions are functions that outputs probabilities for given events/data. Distributions are nice to work, but note that you may encounter distributions of an unknown analytical form. For example, products of known distributions are not always tractable. Some well-known distributions are discussed below.

1.2.1 Continuous

Gaussian Distribution

Also called the Normal distribution. Very special distribution, used for measuring *magnitudes*, for example heights or temperatures. The Central Limit Theorem states that the sum of independent random variables (distributed however) tend to a Gaussian.

A joint Gaussian distribution is the same as a multivariate Gaussian.

1.2.2 Discrete

Counts

Suppose we have an experiment where we toss a coin n times, observing k heads. A suitable distribution would be the **binomial**. What is π , the probability of getting head? If we pursue a frequentist approach of maximum likelihood:

$$p(k|\pi, n) \propto \pi^k(1 - \pi)^{n-k} \quad (\text{binomial pdf})$$
$$\arg \max_{\pi} p(k|\pi, n) = \arg \max_{\pi} \log p(k|\pi, n) = \arg \max_{\pi} k \log \pi + (n - k) \log(1 - \pi)$$
$$\frac{\partial \log p(k|\pi, n)}{\partial \pi} = \frac{k}{\pi} - \frac{n - k}{1 - \pi} = 0 \implies \pi = \frac{k}{n}$$

If we approach instead in a Bayesian setting, we apply a prior distribution on our probability π ... a distribution over probabilities?

$$\text{Beta}(\pi|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \pi^{\alpha-1} (1 - \pi)^{\beta-1}$$

α and β are shape parameters and relate to the pseudo-count +1. The Beta distribution is defined over $[0, 1]$ and is actually a special case of the Dirichlet distribution, which will be discussed in a later section. The Beta is a conjugate prior to the Binomial distribution (the posterior is Beta, likelihood is Binomial).

Continuing with the Bayesian approach; let our data consist of a single coin toss: $\mathcal{D} = \{k = 1\}$ with $n = 1$. Our likelihood:

$$p(\mathcal{D}|\pi) = \pi$$

Our posterior:

$$p(\pi|\mathcal{D}) = \frac{p(\pi|\alpha, \beta)p(\mathcal{D}|\pi)}{p(\mathcal{D})} \propto \text{Beta}(\pi|\alpha, \beta)\pi$$
$$\propto \pi^{\alpha}(1 - \pi)^{\beta-1} \propto \text{Beta}(\alpha + 1, \beta)$$

Word Counts For example a word frequency bar chart. Zipf's law states that the frequency of any word is inversely proportional to its index in the ordered frequency table. Just counting words, not taking into account order, is very simple and not very useful. Words like 'the' are the most frequent but do not give much information about meaning.

Multinomial Distribution

Dirichlet Distribution The Dirichlet distribution is, as mentioned, a generalisation of the Beta distribution, defined on the $m - 1$ dimensional simplex. A simplex is a higher-dimensional triangle, defined by $m - 1$ points in an m -dimensional space. The Dirichlet is the *conjugate prior* for the multinomial.

- The **binomial** to the **Bernoulli** is the **multinomial** to the **categorical**. (multiple trials)
- The **Dirichlet** to the **multinomial** is the **Beta** to the **binomial**.

$$\text{Dir}(\pi|\alpha_1, \dots, \alpha_m) = \frac{\Gamma(\sum_{i=1}^m \alpha_i)}{\prod_{i=1}^m \Gamma(\alpha_i)} \prod_{i=1}^m \pi^{\alpha_i - 1}$$

The α_i s are the shape parameters, and $\mathbb{E}(\pi_j) = \frac{\alpha_j}{\sum_i \alpha_i}$ is the mean for the j -th element.




Figure 1: Symmetric Dirichlet plot with various settings of α

1.2.3 Mapping Between Random Variables

We will discuss a method called Inverse Transform Sampling here, which is used to sample from other distributions using the c.d.f. We discuss more complicated sampling methods in Section 6. This is crucial for some advanced models such as the Variational Autoencoder in Section 11.1. Suppose we have a continuous r.v. U with p.d.f. $p_U(u)$ and c.d.f. $F_U(u)$, and we seek a function $f : U \rightarrow X$ where X is our desired distribution. For maintaining the order, f must be a monotonically increasing function.

The conservation of probability, $p_U(u)du = p_X(x)dx$, means:

$$\begin{aligned} p_U(f^{-1}(x))du &= p_X(x)dx \\ \implies p_X(x) &= p_U(f^{-1}(x)) \frac{du}{dx} = p_U(f^{-1}(x)) \frac{df^{-1}(x)}{du} \end{aligned}$$

The same applies to the c.d.f.:

$$\begin{aligned} F_X(x) &= \int_{-\infty}^x p_X(t)dt = \int_{-\infty}^x p_U(f^{-1}(t)) \frac{df^{-1}(t)}{du} dt \\ \implies F_X(x) &= F_U(f^{-1}(x)) \end{aligned}$$

Let U be a uniform distribution s.t. $F_U(u) = u$.

$$\implies F_X(x) = F_U(f^{-1}(x)) = f^{-1}(x) = u \implies x = F_X^{-1}(u)$$

Therefore, if we have the inverse c.d.f. then we can sample from the distribution. The same applies to discrete distributions. This is the foundation for Normalising Flows [Rezende and Mohamed, 2015], a generative model that we will discuss in Section 11.3

1.3 Softmax Transformation

The softmax function is very useful in machine learning, it takes a length K vector and outputs a normalised probability distribution (adds to 1) consisting of K probabilities. It is used in multi-class classification, and can be used in multi-variate optimisation. For example, optimise

$$f(x_1, x_2, x_3) = 0.2x_1 + 0.3x_2 + 0.5x_3$$

where $x_1, x_2, x_3 \in [0, 1]$ and $x_1 + x_2 + x_3 = 1$

We can set $x_3 = 1 - x_1 - x_2$. The softmax is normalised, which will satisfy the constraint. Also due to this constraint, we have only two free variables.

$$x_1 = \frac{e^{\xi_1}}{e^{\xi_1} + e^{\xi_2} + 1}, x_2 = \frac{e^{\xi_2}}{e^{\xi_1} + e^{\xi_2} + 1}, x_3 = \frac{1}{e^{\xi_1} + e^{\xi_2} + 1}$$

todo finish

1.4 Types of Models

The **Generative** model estimates the joint distribution, and from that computes the posterior $p_y(y|x)$ to make predictions. Often this is done by inferring the likelihood $p_x(x|y)$ and prior $p_y(y)$.

Learning the posterior directly is termed **Discriminative** classification.

1.5 Graphical Models

Graphical models are sometimes treated as a separate topic, but I prefer to view them as a tool for visualising and constructing probabilistic models. A graphical model is essentially a dependency graph, where there is an arrow from node A to node B , where A, B are r.v.s, if B is conditioned on A .

1.5.1 Directed Graphical Models

They are directed acyclic graphs, where for each node we assign a random variable X_i and a probability density function $f_i(x_i, x_{\pi_i})$, where π_i is the set of parents of X_i , which are those on which X_i conditions. The graphical model is possible using the chain rule of probability:

$$p_{X_1, \dots, X_n}(x_1, \dots, x_n) = p_{X_1}(x_1)p_{X_2|X_1}(x_2|x_1) \cdots p_{X_n, \dots, X_{n-1}}(x_n|x_1, \dots, x_{n-1}) \quad (1)$$

Equation 1 creates a fully-connected DAG, which can represent any probability distribution. The joint probability distribution for N variables where each variable can take on a value in \mathcal{X} , requires a table of size $|\mathcal{X}|^N$. When the conditional distributions involved do not depend on all conditioning variables some edges can be removed. Sparse graphs can lead to more efficient inference.



Figure 2: Directed and undirected graphical models

1.5.2 Markov Random Fields

Markov Random Fields are **undirected graphical models** which satisfy the Markov property: the future and past are conditionally independent given the present, i.e., $X \perp\!\!\!\perp Y|Z$ (independent) iff all paths from a node in X to a node in Y pass through a node in Z . In (b), we see that nodes in A are indep. of nodes in B.

1.6 Keep in mind

When marginalising out a variable, the limits are often not explicit.

$$\mathcal{L}_i(\Theta) = -\log \left(\prod_{j=1}^K \int p_{ij}^{y_{ij} + \alpha_{ij} - 1} dp_{ij} \right)$$

$$\mathcal{L}_i(\Theta) = -\log \left(\prod_{j=1}^K \int_0^1 p_{ij}^{y_{ij} + \alpha_{ij} - 1} dp_{ij} \right)$$

2 Maximum Likelihood

Before we discuss Bayesian models of inference, we will discuss an incredibly important (and simple) concept: maximum likelihood estimation (MLE). Most neural networks carry out MLE: they minimise some cost function which is often the negative log likelihood—in other words, maximising the log likelihood! MLE is exactly this: maximise the likelihood. Equivalently, often we maximise the log-likelihood, since the logarithm is an increasing function and it often makes the analytical solution easier to derive.

$$\theta^* = \max_{\theta} \text{lik}(\theta|\mathbf{x}) = \max_{\theta} p_{\mathbf{x}}(\mathbf{x}|\theta)$$

where $p_{\mathbf{x}}$ is either the probability mass for discrete variables or the density for continuous variables.


For example, with a binomial likelihood with n trials:

$$\begin{aligned} \theta^* &= \max_{\theta} \log \binom{n}{x} \theta^x (1-\theta)^{n-x} \\ &= \max_{\theta} \underbrace{[x \log \theta + (n-x) \log(1-\theta)]}_f \\ \frac{\partial f}{\partial \theta} &= \frac{x}{\theta} - \frac{n-x}{1-\theta} = 0 \\ x &= \theta^*(n-x+x) \\ \theta^* &= \frac{x}{n} \end{aligned}$$

Therefore $\theta^* = \frac{x}{n}$ is our MLE, i.e. the probability of success is the number of successful trials over the total number of trials—very intuitive. We will look at this example again through a Bayesian lens in Section 3.1.

2.1 Least Squares Regression

Another example of MLE is least squares regression²: a method of fitting a line through a set of points. To start with, how can we fit a line through the points on the right? Perhaps a polynomial function would be good? The question then is: what order should the polynomial be? Suppose we take an M -order polynomial, $f_{\mathbf{w}}(x)$. We denote $\{\phi_j(x)\}_{j=0}^M$ as the *basis functions*. Such a model is *linear in parameters* (\mathbf{w}) but **not** linear in variables.



$$\begin{aligned} \mathbf{y}^* &= f_{\mathbf{w}}(x) = w_0 + w_1x + w_2x^2 + \dots + w_Mx^M \\ &= \sum_{j=0}^M w_j \phi_j(x) \quad \text{where } \phi_j(x) = x^j \\ \mathbf{y}^* &= \mathbf{X}\mathbf{w} \end{aligned}$$

where \mathbf{X} is the design matrix $[\phi_0(x), \phi_1(x), \dots]^T$ and \mathbf{w} is the row vector of weights.

We then attempt to minimise the error, which is the discrepancy between the actual points \mathbf{y} and the model estimates \mathbf{y}^* .

²In fact, least squares regression can be derived from several directions: projection matrices and maximum likelihood are two such starting points.

$$\begin{aligned}
\mathbf{e} &= \mathbf{y} - \mathbf{y}^* \\
E(\mathbf{w}) &= \|\mathbf{e}\|^2 = (\mathbf{y} - \mathbf{y}^*)^\top (\mathbf{y} - \mathbf{y}^*) \\
&= (\mathbf{y} - \mathbf{X}\mathbf{w})^\top (\mathbf{y} - \mathbf{X}\mathbf{w}) = \mathbf{y}^\top \mathbf{y} - \mathbf{w}^\top \mathbf{X}^\top \mathbf{y} + (\mathbf{X}\mathbf{w})^\top (\mathbf{X}\mathbf{w}) \\
\frac{\partial E(\mathbf{w})}{\partial (\mathbf{w})} &= 2\mathbf{X}^\top \mathbf{X}\mathbf{w} - 2\mathbf{y}^\top \mathbf{X} = 0 \\
\implies \mathbf{X}^\top \mathbf{X}\mathbf{w} &= \mathbf{y}^\top \mathbf{X} \\
\mathbf{w} &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \quad \text{which is the normal equation}
\end{aligned}$$

2.2 Regression with Additive Independent Gaussian Noise

An interesting result is that the MLE and OLS regression as shown above yields the same result if we model the function as $y^{(i)} = f_w(x^{(i)}) + \varepsilon^{(i)}$ where $\varepsilon^{(i)} \sim \mathcal{N}(0, \sigma^2)$, σ^2 is the noise variance. Thus $\boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, which is called an isotropic multivariate Gaussian. Isotropic because there is only variance along the diagonal, therefore it will be a hypersphere.

Since we are assuming noise is independent, and there are N terms ($\text{len}(\boldsymbol{\varepsilon}) = N$), the following result follows simply from the Gaussian PDF:

$$p(\boldsymbol{\varepsilon}) = \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}) = \prod_{n=1}^N p(\varepsilon^{(n)}) = \left(\frac{1}{\sigma\sqrt{2\pi}} \right)^N \exp\left(-\frac{\boldsymbol{\varepsilon}^\top \boldsymbol{\varepsilon}}{2\sigma^2}\right)$$

Since $\mathbf{y} = \mathbf{y}^* + \boldsymbol{\varepsilon}$, we can work out the probability of the data \mathbf{y} given the model estimate \mathbf{y}^* , by constructing the same normal distribution with the mean as the estimate and variance as the variance of the noise! In other words; **centered around the point, with the spread at that point.**

$$\boldsymbol{\varepsilon}^\top \boldsymbol{\varepsilon} = \|\boldsymbol{\varepsilon}\|^2 = \|\mathbf{y} - \mathbf{y}^*\|^2 = E(\mathbf{w}) \quad \text{from the sum of squared errors from before}$$

$$p(\mathbf{y}|\mathbf{y}^*) = \mathcal{N}(\mathbf{y}^*, \sigma^2 \mathbf{I}) = \left(\frac{1}{\sigma\sqrt{2\pi}} \right)^N \exp\left(-\frac{\|\mathbf{y} - \mathbf{y}^*\|^2}{2\sigma^2}\right)$$

Since $\mathbf{y}^* = \mathbf{X}\mathbf{w}$, for a given \mathbf{X} , $p(\mathbf{y}|\mathbf{y}^*) = p(\mathbf{y}|\mathbf{w})$. Ahah! The likelihood function of \mathbf{w} has fallen out. Thus we attempt to maximise the likelihood $\mathcal{L}(\mathbf{w}) \propto p(\mathbf{y}|\mathbf{w})$:

$$\mathbf{w}^* = \arg \max_{\mathbf{w}} \mathcal{L}(\mathbf{w}) = \arg \max_{\mathbf{w}} \exp\left(-\frac{\|\mathbf{y} - \mathbf{y}^*\|^2}{2\sigma^2}\right) = \arg \min_{\mathbf{w}} E(\mathbf{w})$$

Above is possible since \exp is an increasing function. Note how this is the same result as with the least squares method. Overfitting is still an issue, with MLE, more complex models overfit the data. Next we bring Bayes into it.

3 Bayesian Inference

3.1 Maximum a posteriori

The MAP estimate is the point estimate of a parameter θ , which is essentially the MLE but taking into account the *prior*. Suppose we have observations x from the distribution $p(x|\theta)$. Start with the MLE. We therefore need the likelihood: $\text{lik}(\theta|x) = p(x|\theta)$

We can calculate the posterior using Bayes' theorem. We have a prior belief in the form of a distribution over θ , $p_{\Theta}(\theta)$

$$p(\theta|x) = \frac{p(x|\theta)p_{\Theta}(\theta)}{p(x)}$$

We now maximise the posterior, noting that the marginal likelihood does not affect the maximisation of θ and is always positive so we can ignore it:

$$\theta^* = \arg \max_{\theta} p(\theta|x) = \arg \max_{\theta} \frac{p(x|\theta)p_{\Theta}(\theta)}{p(x)} = \arg \max_{\theta} p(x|\theta)p_{\Theta}(\theta)$$

3.2 Conjugate Priors

If the prior and posterior distribution are in the same family then they are said to be *conjugate* and the prior is the *conjugate prior* to the likelihood. We will continue our example from Section 2, where we have a binomial likelihood: $p(x|\theta) = \binom{n}{x}\theta^x(1-\theta)^{n-x}$. Since we are doing Bayesian inference, our parameters must be given prior distributions. In this case, we need a prior for θ : the probability of success. We therefore need a **distribution of probabilities**. The beta distribution is one such distribution, which also happens to be the conjugate prior to the binomial and Bernoulli likelihoods. It is defined as:

$$p(\theta) = \frac{1}{B(\alpha, \beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}$$

where α and β can be interpreted as one plus the *pseudo-counts*. For example, if you have seen 5 successes and 4 failures, set $\alpha = 5, \beta = 4$. $B(\alpha, \beta)$ is defined as:

$$B(\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)}$$

where $\Gamma(x) = (x+1)!$ Now we will show that the beta distribution is the conjugate prior. Starting with Bayes' rule:

$$p(\theta|x) = \frac{p(x|\theta)p_{\Theta}(\theta)}{p(x)} = \frac{\binom{n}{x}\theta^x(1-\theta)^{n-x} \frac{1}{B(\alpha, \beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}}{p(x)} \tag{2}$$

$$= \frac{\binom{n}{x}B(\alpha, \beta)^{-1}\theta^{x+\alpha-1}(1-\theta)^{n-x+\beta-1}}{p(x)} \tag{3}$$

$$= \frac{\frac{n!}{x!(n-x)!} \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}\theta^{x+\alpha-1}(1-\theta)^{n-x+\beta-1}}{p(x)} \tag{4}$$

$$= \frac{\overbrace{\frac{\Gamma(n+1)\Gamma(\alpha)\Gamma(\beta)}{\Gamma(x+1)\Gamma(n-x+1)\Gamma(\alpha+\beta)}}^{\text{g}}\theta^{x+\alpha-1}(1-\theta)^{n-x+\beta-1}}{p(x)} \tag{5}$$

We pause here and calculate the marginal:

$$p(x) = p(x|\theta)p_{\Theta}(\theta) \tag{6}$$

$$= \int \mathcal{G} \cdot \underbrace{\theta^{x+\alpha-1}(1-\theta)^{n-x+\beta-1}}_{\text{unnormalised Beta}(\alpha+x, n-x+\beta)} d\theta \tag{7}$$

$$= \mathcal{G} \cdot B(x+\alpha, n-x+\beta) \underbrace{\int \frac{1}{B(x+\alpha, n-x+\beta)} \theta^{x+\alpha-1}(1-\theta)^{n-x+\beta-1} d\theta}_{\text{normalises to 1}} \tag{8}$$

where we multiplied and divided by the $B(\cdot)$ in order to normalise the integrand in Eq. 8. Plugging this marginal into the intermediate result Eq. 5:

$$\begin{aligned} p(\theta|x) &= \frac{\mathcal{G} \cdot \theta^{x+\alpha-1}(1-\theta)^{n-x+\beta-1}}{\mathcal{G} \cdot B(x+\alpha, n-x+\beta)} \\ &= \text{Beta}(x+\alpha, n-x+\beta) \end{aligned}$$

Since our posterior is also beta, it is conjugate with the beta prior. We can finish this example by calculating the MAP and comparing it to the MLE! We maximise as follows:

$$\begin{aligned} \theta^* &= \arg \max_{\theta} \log p(\theta|x) = \arg \max_{\theta} \log \text{Beta}(x+\alpha, n-x+\beta) \\ &= \arg \max_{\theta} \log \frac{\theta^{x+\alpha-1}(1-\theta)^{n-x+\beta-1}}{B(x+\alpha, n-x+\beta)} \\ &= \arg \max_{\theta} (x+\alpha-1) \log \theta + (n-x+\beta-1) \log(1-\theta) \\ \frac{\partial}{\partial \theta} &= \frac{x+\alpha-1}{\theta} - \frac{n-x+\beta-1}{1-\theta} \\ \theta_{\text{MAP}}^* &= \frac{x+\alpha-1}{n+\alpha+\beta-2} \end{aligned}$$

Compare this to our MLE: $\theta_{\text{MLE}}^* = \frac{x}{n}$. They are equivalent if $\alpha = \beta = 1$, and otherwise they reflect whatever prior information we give with our setting of α and β . Very satisfying!

3.3 Gaussian Likelihood

Take a model \mathcal{M} , representing a choice of model structure and parameter values. Let the structure be $y = f_w(x) + \varepsilon$, **the probability of the data is conditional:** $p(\mathbf{y}|\mathbf{x})$. The Gaussian likelihood is:

$$p(\mathbf{y}|\mathbf{x}, \mathbf{w}, \mathcal{M}) \propto \prod_{i=1}^N \exp\left(-\frac{(y^{(i)} - f_w(x^{(i)}))^2}{2\sigma^2}\right)$$

Fit the model by optimising for the weights (MLE):

$$\mathbf{w}^* = \arg \max_w p(\mathbf{y}|\mathbf{x}, \mathbf{w}, \mathcal{M})$$

After maximisation, make predictions from $p(y|x, \mathbf{w}^*, \mathcal{M})$. Notice how it now uses the fitted weights.

With a certain likelihood distribution, the conjugate prior is the distribution such that the posterior will be tractable (there is a closed form). Here we will have a Gaussian likelihood $p(\mathbf{y}|\mathbf{x}, \mathbf{w}, \mathcal{M}) = \mathcal{N}(\mathbf{X}\mathbf{w}, \sigma^2\mathbf{I})$ and Gaussian prior $p(\mathbf{w}|\mathcal{M}) = \mathcal{N}(\mathbf{0}, \sigma_w^2)$.

Let's now apply Bayes' rule as before to include the prior. We first calculate the posterior:

$$\begin{aligned}
p(\mathbf{w}|\mathbf{x}, \mathbf{y}, \mathcal{M}) &= \frac{p(\mathbf{y}|\mathbf{x}, \mathbf{w}, \mathcal{M})p(\mathbf{w}|\mathcal{M})}{p(\mathbf{y}|\mathbf{x}, \mathcal{M})} = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \\
&\propto \exp\left(-\frac{1}{2\sigma^2}|\mathbf{Y} - \mathbf{X}\mathbf{w}|^2\right) \exp\left(-\frac{1}{2}\mathbf{w}^\top \boldsymbol{\Sigma}_w^{-1} \mathbf{w}\right) \\
&= \exp\left(-\frac{1}{2}((\mathbf{X}\mathbf{w})^\top \mathbf{X}\mathbf{w}\sigma^{-2} - 2\mathbf{y}^\top (\mathbf{X}\mathbf{w})\sigma^{-2} + \mathbf{w}^\top \boldsymbol{\Sigma}_w^{-1} \mathbf{w})\right) \\
&= \exp\left(-\frac{1}{2}(\mathbf{w}^\top (\mathbf{X}^\top \mathbf{X}\sigma^{-2} + \boldsymbol{\Sigma}_w^{-1})\mathbf{w} - 2\sigma^{-2}\mathbf{w}^\top \mathbf{X}^\top \mathbf{y})\right) \tag{9}
\end{aligned}$$

Notice equation (1), remember completing the square with a scalar equation e.g. $x^2 + xy + y^2$. With matrices, the result is similar:

$$x^\top Ax + x^\top b + c = (x - \mu)^\top A(x - \mu) + k$$

where $\mu = -\frac{1}{2}A^{-1}b$, $k = c - \frac{1}{4}b^\top A^{-1}b$. Apply this result to the exponent as follows:

$$\begin{aligned}
&= \exp\left(-\frac{1}{2}(\mathbf{w} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{w} - \boldsymbol{\mu})\right) \\
\text{where } \boldsymbol{\Sigma} &= (\sigma^{-2}\mathbf{X}^\top \mathbf{X} + \boldsymbol{\Sigma}_w^{-1})^{-1}, \\
\boldsymbol{\mu} &= -\frac{1}{2} \times -2\sigma^{-2}\boldsymbol{\Sigma}^{-1}\mathbf{X}^\top \mathbf{y} = \sigma^{-2}\boldsymbol{\Sigma}\mathbf{X}^\top \mathbf{y}
\end{aligned}$$

Reminder of the marginal: $p(x) = \int p(x, y)dy = \int p(x|y)p(y)dy$ we need to marginalise out the parameters in order to make predictions. The following is the **predictive distribution**:

$$p(y|x, \mathbf{x}, \mathbf{y}, \mathcal{M}) = \int p(y, \mathbf{w}|x, \mathbf{x}, \mathbf{y}, \mathcal{M}) d\mathbf{w} = \int p(y|x, \mathbf{w}, \mathbf{x}, \mathbf{y}, \mathcal{M})p(\mathbf{w}|\mathbf{x}, \mathbf{y}, \mathcal{M}) d\mathbf{w}$$

3.4 Marginal Likelihood

We marginalise out \mathbf{w} from the marginal likelihood $p(\mathbf{y}|\mathbf{x}, \mathcal{M})$. If there is confusion regarding \mathbf{w} not being conditioned on, I believe the intuition is that the distribution of \mathbf{w} is in some way parameterised by \mathbf{x} . This is an application of the law of total probability.

$$p(\mathbf{y}|\mathbf{x}, \mathcal{M}) = \int p(\mathbf{w}|\mathbf{x}, \mathcal{M})p(\mathbf{y}|\mathbf{x}, \mathbf{w}, \mathcal{M}) d\mathbf{w}$$

One then optimises the marginal likelihood to tune the hyperparameters.

To demonstrate how the marginal likelihood assists in model selection, we apply Bayes' rule again:

$$\begin{aligned}
p(\mathbf{y}|\mathbf{x}, \mathcal{M}) &= \frac{p(\mathcal{M}|\mathbf{y}, \mathbf{x})p(\mathbf{y}|\mathbf{x})}{p(\mathcal{M})} \\
p(\mathcal{M}|\mathbf{y}, \mathbf{x}) &= \frac{p(\mathbf{y}|\mathbf{x}, \mathcal{M})p(\mathcal{M})}{p(\mathbf{y}|\mathbf{x})} \propto p(\mathbf{y}|\mathbf{x}, \mathcal{M})p(\mathcal{M})
\end{aligned}$$

Since the probability of a model given the data is proportional to the marginal likelihood, this provides some hand-wavy intuition as to why the marginal likelihood gives some kind of score to the model.

4 Gaussian Processes

4.1 Non-parametric Modelling

Previously the models discussed have been parametric. These parameters are marginalised over to yield a predictive distribution.

$$p(y^*|x^*, \mathbf{x}, \mathbf{y}) = \int p(y^*|x^*, \mathbf{w}, \mathbf{x}, \mathbf{y})p(\mathbf{w}|\mathbf{x}, \mathbf{y}) d\mathbf{w}$$

Observe the last term, the $p(\mathbf{w}|\mathbf{x}, \mathbf{y})$ is a bottleneck for the model; there are a fixed number of parameters for the model. In fact, the distribution over the parameters implies a distribution over functions. A non-parametric model would work directly with such a distribution. Moreover, the predictive distribution is usually an intractable integral. Which distribution is easily integrable?

Gaussians

Suppose we have the joint probability:

$$p(\mathbf{x}_1, \mathbf{x}_2) = \mathcal{N}\left(\begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}, \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}\right)$$

- **What is the marginal distribution, $p(\mathbf{x}_1)$?**

We integrate out \mathbf{x}_2 :

$$p(\mathbf{x}_1) = \int p(\mathbf{x}_1, \mathbf{x}_2) d\mathbf{x}_2 = \mathcal{N}(\mu_1, \Sigma_{11})$$

- **What is the conditional distribution, $p(\mathbf{x}_1|\mathbf{x}_2)$?**

This has the solution $p(\mathbf{x}_1|\mathbf{x}_2) = \mathcal{N}(\mu_c, \Sigma_c)$ where:

$$\begin{aligned} \mu_c &= \mu_1 + \Sigma_{12}\Sigma_{22}^{-1}(\mathbf{x}_2 - \mu_2) \\ \Sigma_c &= \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21} \end{aligned}$$

Recall a model using basis functions: $f_w(x) = \sum_{m=0}^M w_m \phi_m(x)$. A **prior** over the weights $p(\mathbf{w})$ induces a **prior** distribution over functions, $p(\mathbf{f})$.

How do we make predictions from such a distribution over functions? Given a parametric family of functions, $f(\mathbf{x}|\mathbf{f})$ and a prior over \mathbf{f} , Bayesian modelling helps us predict given the data, $p(\mathbf{f}|\mathbf{y}, \mathbf{x})$:

$$p(\mathbf{f}|\mathbf{y}, \mathbf{x}) = \frac{p(\mathbf{y}|\mathbf{x}, \mathbf{f})p(\mathbf{f})}{p(\mathbf{y}|\mathbf{x})}$$

4.2 Definition

Definition: a Gaussian process is a collection of random variables, every finite subset of which are jointly Gaussian.

where $m(x)$ is the mean function and $k(x, x')$ is the covariance function. A GP is a Gaussian distribution with an infinitely long mean vector and infinite covariance matrix. We can't really reason with an infinitely long mean vector and covariance matrix, so we restrict ourselves to a finite subset, and rely on the marginalisation property to marginalise out the infinite. So now we've learnt how to draw **random functions**, but this is not useful, we want to somehow model the data.

$$f \sim \mathcal{GP}(m, k)$$




Figure 3: Polynomials are not a great prior over functions: after 1 they experience rapid growth.

Since we are in a non-parametric model, the parameters are the function itself. Let's go through the same Bayesian inference procedure as in Section 3.3 but plug in the function r.v. instead of the parameter r.v. We first find an expression for the likelihood:

$$p(\mathbf{y}|\mathbf{x}, \mathbf{f}) \sim \mathcal{N}(\mathbf{f}, \sigma^2 \mathbf{I})$$

where \mathbf{f} is the vector of function applied to the input data. We can do this instead of a Gaussian distribution over $\mathbf{y}|\mathbf{x}$ since we only know what the function is doing at the observations.

We have a Gaussian process prior for the functions, as we saw before: $p(\mathbf{f}) \sim \mathcal{GP}(m = 0, \mathbf{k})$

The posterior is calculated by multiplying the likelihood and the prior. Product of two Gaussians is a Gaussian. This leads to an infinite Gaussian, which is a Gaussian process.

$$p(\mathbf{f}|\mathbf{x}, \mathbf{y}) \sim \mathcal{GP}(\mathbf{m}_{\text{post}}, \mathbf{k}_{\text{post}})$$

And the following predictive distribution:

$$p(y_*|x_*, \mathbf{x}, \mathbf{y}) = \int p(y_*, \mathbf{f}|\mathbf{x}, \mathbf{y}, x_*) d\mathbf{f} = \int p(y_*|x_*, \mathbf{f}, \mathbf{x}, \mathbf{y}) p(\mathbf{f}|\mathbf{x}, \mathbf{y}) d\mathbf{f}$$

What does the marginal likelihood look like? The evidence is just the probability of observations where the function has been marginalised out. This yields a *closed form* solution:

$$\log p(\mathbf{y}|\mathbf{x}) = -\frac{1}{2} \mathbf{y}^\top [K + \sigma_n^2 I]^{-1} \mathbf{y} - \frac{1}{2} \log[K + \sigma_n^2 I] - \frac{n}{2} \log(2\pi)$$

where the **data fit** measures how well the model fits the training data, and the **complexity penalty** penalises how big the model class that we're using. The nice thing is that there are no hyperparameters here like with regularisers; Occam's razor is automatic.

4.3 Hyperparameters & Covariance Functions

Kernel methods are useful since they can express an infinite amount of features in closed form using the dot product. This section will outline how we can obtain such a kernel. First, **what is positive definiteness?**

A symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is positive definite if $\mathbf{x}^\top \mathbf{A} \mathbf{x} > 0 \quad \forall \mathbf{x} \neq 0$. Let's build an intuition of what this means.

$$f(x, y) = \begin{bmatrix} x & y \end{bmatrix} \begin{bmatrix} 1 & 3 \\ 3 & 8 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = x^2 + 6xy + 8y^2$$

What are the roots of the equation $f(x, y) = 0$? To calculate them we can set $x = \alpha$ and $y = 1$.³ Now our equation is $x^2 + 6x + 8$ with roots $x = -2$, $x = -4$. Plug in any value for x in between our two roots:

$$f(-3, 1) = 9 - 18 + 8 = -1$$

shows our matrix is **not** positive definite.

³We have essentially scaled y by α : let $x = \beta$ and $y = \gamma$. Now set $x = \beta/\gamma = \alpha$ and $y = 1$




Figure 4: The arrows are eigenvectors of Σ scaled by the sqrt of their eigenvalues.




Figure 5: $x^2 + 6x + 8$

Let us now apply this in the case of a covariance matrix. Suppose we have two variables jointly Gaussian:

$$(x, y) \sim \mathcal{N}\left(\begin{bmatrix} \mu & \nu \end{bmatrix}, \Sigma = \begin{bmatrix} \sigma^2 & \alpha \\ \alpha & \rho^2 \end{bmatrix}\right)$$

We will now see that for the covariance matrix to have any meaning, it must be positive semi-definite. Suppose $Z = aX + bY$

What is the distribution of Z ? Gaussian. The mean is trivial: $\mathbb{E}(aX + bY) = a\mathbb{E}(X) + b\mathbb{E}(Y) = a\mu + b\nu$. The variance is also easy:

$$\begin{aligned} \text{Var}(aX + bY) &= \text{Var}(aX) + \text{Var}(bY) + 2\text{Cov}(aX, bY) = a^2\sigma^2 + b^2\rho^2 + 2\mathbb{E}((aX - a\mu)(bY - b\nu)) \\ &= a^2\sigma^2 + b^2\rho^2 + 2(ab\mathbb{E}(XY) - ab\mathbb{E}(X)\mathbb{E}(Y)) = a^2\sigma^2 + b^2\rho^2 + 2ab\text{Cov}(X, Y) \\ &= a^2\sigma^2 + b^2\rho^2 + 2ab\alpha \end{aligned}$$

Common sense tells us that $\text{Var}(Z) > 0$ (variances can't be negative). Thus the only covariance matrices that are permissible are those that are positive semi-definite.

$$\begin{aligned} Z &\sim \mathcal{N}\left(a\mu + b\nu, \begin{bmatrix} a & b \end{bmatrix} \begin{bmatrix} \sigma^2 & \alpha \\ \alpha & \rho^2 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix}\right) \\ &\sim \mathcal{N}(a\mu + b\nu, \mathbf{w}^\top \Sigma \mathbf{w}) \quad \text{where } \mathbf{w}^\top \Sigma \mathbf{w} \geq 0 \end{aligned}$$

Notice the similarity to the result $\text{Var}(aX) = a^2\text{Var}(X)$

4.3.1 Hilbert Spaces

Let \mathcal{H} be a vector space over \mathbb{R} . The **inner product** $\langle \cdot, \cdot \rangle : \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}$ iff (i) $\langle \alpha f_1 + \beta f_2, g \rangle = \alpha \langle f_1, g \rangle + \beta \langle f_2, g \rangle$; (ii) $\langle f, g \rangle = \langle g, f \rangle$; (iii) $\langle f, f \rangle \geq 0$ and is 0 only if $f = 0$.

A **Hilbert space** is a space where the inner product is defined in addition to another technical condition relating to Cauchy sequences, which is not necessary to go into here. A **kernel** is defined as follows:

The function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a kernel if there exists a Hilbert space and a map $\phi : \mathcal{X} \rightarrow \mathcal{H}$ s.t. $\forall x, x' \in \mathcal{X}, k(x, x') = \langle \phi(x), \phi(x') \rangle$

4.3.2 Positive Definiteness of Inner Products in Hilbert Spaces

Let \mathcal{H} be any Hilbert space, \mathcal{X} a non-empty set and ϕ a feature mapping. This implies that $k(x, y) = \langle \phi(x), \phi(y) \rangle$ is a positive semidefinite function.

Proof. The norm is written as

$$\|f\| = \sqrt{\langle f, f \rangle}$$

$$\begin{aligned} \sum_i \sum_j a_i a_j k(x_i, x_j) &= \sum_j \langle a_i \phi(x_i), a_j \phi(x_j) \rangle \\ &= \left\langle \sum_i a_i \phi(x_i), \sum_j a_j \phi(x_j) \right\rangle \\ &= \left\| \sum_i a_i \phi(x_i) \right\|^2 \geq 0 \end{aligned}$$

□

The reverse direction also holds. A positive semidefinite function is guaranteed to be the inner product in a Hilbert space. Thus positive semidefiniteness is a way of proving a function is a kernel.

4.3.3 Reproducing Kernel Hilbert Space

We now have kernels on feature spaces. We want to define what our functions on \mathcal{X} look like. The space of these functions is known as a reproducing kernel Hilbert space.

Suppose our feature map $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ is, for example, $\phi \left(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \right) = \begin{bmatrix} x_1 \\ x_2 \\ x_1x_2 \end{bmatrix}$ with a dot-product kernel $k(\mathbf{x}, \mathbf{y}) = x^\top y$. This feature space is denoted by \mathcal{H}

Let $f(\mathbf{x}) = ax_1 + bx_2 + cx_1x_2$ or $f(\cdot) = \begin{bmatrix} a \\ b \\ c \end{bmatrix}$

We can also express f as $f(\mathbf{x}) = f(\cdot)^\top \phi(\mathbf{x}) = \langle f(\cdot), \phi(\mathbf{x}) \rangle$. $\phi(\mathbf{x})$ is a function mapping $\mathbb{R}^2 \rightarrow \mathbb{R}^3$ and defines the parameters of a function mapping $\mathbb{R}^2 \rightarrow \mathbb{R}$. To illustrate this further, take

$$k(\cdot, \mathbf{y}) = \begin{bmatrix} y_1 \\ y_2 \\ y_1y_2 \end{bmatrix} = \phi(\mathbf{y})$$

For every \mathbf{y} , there is a vector $k(\cdot, \mathbf{y})$ s.t. $\langle k(\cdot, \mathbf{y}), \phi(\mathbf{x}) \rangle = ax_1 + bx_2 + cx_1x_2$ where $a = y_1, b = y_2, c = y_1y_2$

This is equivalent to

$$\langle k(\cdot, \mathbf{x}), \phi(\mathbf{y}) \rangle = uy_1 + vy_2 + wy_1y_2 = k(\mathbf{x}, \mathbf{y})$$

So we can write $\phi(\mathbf{x}) = k(\cdot, \mathbf{x})$ and $\phi(\mathbf{y}) = k(\cdot, \mathbf{y})$ without ambiguity.

This shows that:

- every feature mapping is in the feature space: $\forall \mathbf{x} \in \mathcal{X}, k(\cdot, \mathbf{x}) \in \mathcal{H}$;
- $\forall \mathbf{x} \in \mathcal{X}, \forall f \in \mathcal{H}, \langle f, k(\cdot, \mathbf{x}) \rangle = f(\mathbf{x})$

This last property is the **reproducing property**. It yields another appealing property: the norm in an RKHS is a natural measure of how complex a function is.

4.3.4 Squared Exponential (RBF)

$$k(x, x') = v^2 \exp \left(-\frac{(x - x')^2}{2l^2} \right) + \sigma^2 \delta_{xx'}$$

Example: l is length scale for a squared exponential covariance function. It's roughly the length between inputs before the covariance is lower. You can have anisotropic RBF kernels where there are multiple length scales, one for each direction in the input space. That way you can accommodate input features that are on different scales.

4.3.5 Periodic

$$k(x, x') = \exp \left(-\frac{2 \sin^2(\pi(x - x'))}{l^2} \right)$$

The x 's are first mapped to $u = [\sin(x), \cos(x)]^\top$ and then distances are measured in the u -space. If the length scale is larger than the period, then the covariance is high, whereas in the inverse case, there can be a lot of action within the period.

4.3.6 Reproducing Kernel for Vector-Valued Functions

Our reproducing kernel is now defined as a symmetric function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^{D \times D}$ s.t. $k(x, x')$ yields a positive semidefinite **matrix**.

Let \mathcal{H} be the vector-valued RKHS over functions $f : \mathcal{X} \rightarrow \mathbb{R}^D$. This means that $\forall \mathbf{x} \in \mathcal{X}, \forall f \in \mathcal{H}, \forall \mathbf{c} \in \mathbb{R}^D, f(\mathbf{x}') = k(\mathbf{x}, \mathbf{x}')\mathbf{c}$ belongs to \mathcal{H}

and the reproducing property is written now as: $\langle f, k(\cdot, \mathbf{x}) \rangle = f(\mathbf{x})^\top \mathbf{c}$

4.4 Coregionalization

Coregionalization originated in geostatistics literature, where it is known as *cokriging*. In order for covariance functions to be valid kernels, as seen before, they must be positive semidefinite. Suppose we have a multi-output problem with D outputs. In the *linear model of coregionalization*, these outputs are written as a linear combination of Q independent latent functions which have zero mean and a covariance function.

For each $d \in \{1, \dots, D\}$, the output is determined by the function $f_d(\mathbf{x})$ with p -dimensional input vector \mathbf{x} .

$$f_d(\mathbf{x}) = \sum_{q=1}^Q a_{d,q} u_q(\mathbf{x})$$

where $u_q(\mathbf{x})$ are the latent functions with covariance $\text{Cov}(u_q(\mathbf{x}), u_{q'}(\mathbf{x}')) = \begin{cases} k_q(\mathbf{x}, \mathbf{x}') & q = q' \\ 0 & \text{otherwise} \end{cases}$, due to independence. Some of these latent functions can share the same covariance kernel and can be grouped:

$$f_d(\mathbf{x}) = \sum_{q=1}^Q \sum_{i=1}^{R_q} a_{d,q}^i u_q^i(\mathbf{x})$$

where $u_q^i(\mathbf{x})$ have covariance $\text{Cov}(u_q^i(\mathbf{x}), u_{q'}^{i'}(\mathbf{x}')) = k_q(\mathbf{x}, \mathbf{x}')$ for $i = i', q = q'$. There are now Q groups of functions, within each one sharing a covariance function.

We can now write the cross-covariance between functions as:

$$\begin{aligned} \text{Cov}(f_d(\mathbf{x}), f_{d'}(\mathbf{x}')) &= (\mathbf{K}(\mathbf{x}, \mathbf{x}'))_{d,d'} = \sum_{q=1}^Q \sum_{q'=1}^Q \sum_{i=1}^{R_q} \sum_{i'=1}^{R_{q'}} a_{d,q}^i a_{d',q'}^{i'} \cdot \text{Cov}(u_q^i(\mathbf{x}), u_{q'}^{i'}(\mathbf{x}')) \\ &= \sum_{q=1}^Q \sum_{i=1}^{R_q} a_{d,q}^i a_{d',q}^i \cdot k_q(\mathbf{x}, \mathbf{x}') \quad \text{by independence} \\ &= \sum_{q=1}^Q b_{d,d'}^q k_q(\mathbf{x}, \mathbf{x}') \end{aligned}$$

where $b_{d,d'}^q = \sum_{i=1}^{R_q} a_{d,q}^i a_{d',q}^i$ which forms a $D \times D$ matrix \mathbf{B}_q called the **coregionalisation matrix**. The rank of \mathbf{B}_q , the number of linearly independent row or column vectors, is intuitively determined by R_q .

Writing our kernel $\mathbf{K}(\mathbf{x}, \mathbf{x}') = \sum_{q=1}^Q \mathbf{B}_q k_q(\mathbf{x}, \mathbf{x}')$, one can intuit that this is a sum of the products of two kernels (called *separable kernels*), one that models the output dependencies (\mathbf{B}_q) and one that models the input dependencies (k_q).

4.5 Relation with Linear in the Parameters Model

Consider $f(x) = ax + b$ where $a \sim \mathcal{N}(0, \alpha)$, $b \sim \mathcal{N}(0, \beta)$. We can work out the mean function (see A.1 for derivation):

$$\mu(x) = \mathbb{E}(f(x)) = \int \int (ax + b)p(a)p(b) da db = \int axp(a) da + \int bp(b) db = 0$$

And now the covariance function:

$$\begin{aligned} k(x, x') &= \mathbb{E}[(f(x) - 0)(f(x') - 0)] = \int \int (ax + b)(ax' + b)p(a)p(b) da db \\ &= \int \int (a^2xx' + b^2 + ab(x + x'))p(a)p(b) da db \\ &= \int_b p(b)xx' \int_a a^2p(a) + (x' + x) \int_a \int_b ap(a)bp(b) + \int_a p(a) \int_b b^2p(b) \\ &= \alpha xx' + \beta \end{aligned}$$

So we have, in a very overly complicated way, constructed a linear model. We can now take this finite linear model and go to a Gaussian process (infinite):

$$f(x) = \sum_{m=1}^M w_m \phi_m(x) \quad p(\mathbf{w}) = \mathcal{N}(\mathbf{0}, \mathbf{A})$$

The joint distribution of any vector $\mathbf{f} = [f(x_1), \dots, f(x_N)]$ is a multivariate Gaussian and therefore a Gaussian process. The mean function is 0. The covariance:

$$\begin{aligned} k(x_i, x_j) &= \text{Cov}_{\mathbf{w}}(f(x_i), f(x_j)) = \mathbb{E}(f(x_i)f(x_j)) - \mathbb{E}(f(x_i))\mathbb{E}(f(x_j)) = \mathbb{E}(f(x_i)f(x_j)) \\ &= \int \dots \int \left(\sum_{k=1}^M \sum_{l=1}^M w_k w_l \phi_k(x_i) \phi_l(x_j) \right) p(\mathbf{w}) d\mathbf{w} = \dots \\ &= \sum_{k=1}^M \sum_{l=1}^M \phi_k(x_i) \phi_l(x_j) \int \int w_k w_l p(w_k, w_l) dw_k dw_l \end{aligned}$$

This shows that any linear in the parameters model with Gaussian prior over weights, is also a Gaussian process. Mercer's theorem states that every GP also corresponds to a linear in the parameters model but not necessarily a finite one.

We will now show a very cool result that a squared exponential covariance function corresponds to a linear in parameters model with infinitely many Gaussian bumps.

Consider the following *Gaussian bump basis function*:

$$f(x) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{-N/2}^{N/2} \gamma_n \exp \left(- \left(x - \frac{n}{\sqrt{N}} \right)^2 \right) \quad \text{where } \gamma_n \sim \mathcal{N}(0, 1)$$

We use the sum (limited to infinity) to place bumps everywhere along x .

But $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{-N/2}^{N/2} = \int_{-\infty}^{\infty}$, so

$$\int_{-\infty}^{\infty} \gamma(u) \exp(-(x - u)^2) du$$




Figure 6: Example Gaussian bump

$$\begin{aligned}\mu(x) &= \mathbb{E}(f(x)) \\ &= \int_{-\infty}^{\infty} \exp(-(x-u)^2) \int_{-\infty}^{\infty} \gamma(u) p(\gamma(u)) d\gamma(u) du = 0\end{aligned}$$

4.6 Eigenvectors and Relation to PCA

Let's visualise what a 2 variable joint Gaussian distribution (2-D multivariate Gaussian) would look like.




Figure 7: $(x, y) \sim \mathcal{N}\left(\mathbf{0}, \Sigma = \begin{bmatrix} 0.2 & 0.4 \\ 0.4 & 0.2 \end{bmatrix}\right)$




Figure 8: $(x, y) \sim \mathcal{N}\left(\mathbf{0}, \Sigma = \begin{bmatrix} 0.2 & 0.2 \\ 0.2 & 0.2 \end{bmatrix}\right)$

4.7 Cholesky Decomposition

Due to the marginalisation/consistency property (or even the definition of a GP), the marginal (and the conditional) distributions are also Gaussian.

Recall $p(\mathbf{f}) = \int p(\mathbf{f}, \mathbf{y}) d\mathbf{y}$, but now suppose \mathbf{y} is infinitely long.

$$p(\mathbf{x}, \mathbf{y}) = \mathcal{N}\left(\begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix}, \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{C} \end{bmatrix}\right) \implies p(\mathbf{x}) = \mathcal{N}(\mathbf{a}, \mathbf{A})$$

Cholesky factorisation decomposes a positive-definite matrix $\mathbf{A} = \mathbf{R}^\top \mathbf{R}$ where \mathbf{R} is upper triangular. Essentially the “square root”.

We can sample now from a D-dimensional joint Gaussian with mean vector \mathbf{m} and covariance matrix \mathbf{K} .

$$\begin{aligned}\mathbf{z} &= \text{randn}(D, 1) \\ \mathbf{y} &= \text{chol}(\mathbf{K})^\top \mathbf{z} + \mathbf{m}\end{aligned}$$

$$\mathbb{E}((\mathbf{y} - \mathbf{m})(\mathbf{y} - \mathbf{m})^\top) = \mathbb{E}(\mathbf{R}^\top \mathbf{z} \mathbf{z}^\top \mathbf{R}) = \mathbf{R}^\top \mathbb{E}(\mathbf{z} \mathbf{z}^\top) \mathbf{R} = \mathbf{R}^\top \mathbf{I} \mathbf{R} = \mathbf{K}$$

4.8 Gaussian Process Example

$$y = f + \epsilon$$

\mathbf{y} are our training outputs.

where the likelihood for our data \mathbf{y} is $p(\mathbf{y}|\mathbf{f}) = \mathcal{N}(\mathbf{f}, \sigma^2 \mathbf{I})$

Any set of function variables \mathbf{f} has $p(\mathbf{f}|\mathbf{X}) = \mathcal{N}(\mathbf{0}, \mathbf{K})$

The marginal likelihood is $p(\mathbf{y}) = \int p(\mathbf{y}|\mathbf{f}) p(\mathbf{f}) d\mathbf{f} = \mathcal{N}(\mathbf{0}, \mathbf{K} + \sigma^2 \mathbf{I})$

For prediction, consider joint training and test marginal likelihood:

$$p(\mathbf{y}, \mathbf{y}^*) = \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K_{xx} & K_{xt} \\ K_{tx} & K_{tt} \end{bmatrix}\right)$$

Conditioning on training outputs:

$$p(\mathbf{y}^* | \mathbf{y}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$$

where

$$\boldsymbol{\mu} = K_{tx}[K_{xx} + \sigma^2\mathbf{I}]^{-1}\mathbf{y}$$

$$\boldsymbol{\Sigma} = K_{tt} - K_{tx}K_{xx}^{-1}K_{xt}$$

4.9 Scalability

GPs scale poorly due to the $O(n^3)$ matrix inversion. In Section 8.1, we will discuss methods of scaling GPs using inducing point approximations.

5 Gaussian Process Classification

Now, the likelihood is categorical and we have a new likelihood:

$$p(y|x) = \sigma(f(x))^y (1 - \sigma(f(x))^{1-y})$$

where σ is the sigmoid function.

The integral is no longer tractable.


finish

6 Monte Carlo

How can we integrate an intractable function?

We want to find approximate expectations of a function $\phi(\mathbf{x})$ w.r.t. probability $p(\mathbf{x})$.

$$\mathbb{E}_{p(x)}[\phi(x)] = \bar{\phi} = \int \phi(\mathbf{x})p(\mathbf{x}) \, d\mathbf{x}$$



We could lay out a grid and compute

$$\int \phi(\mathbf{x})p(\mathbf{x}) \, d\mathbf{x} \approx \sum_{\tau=1}^T \phi(\mathbf{x}^{(\tau)})p(\mathbf{x}^{(\tau)})\Delta\mathbf{x}$$

. However, this requires too many points if the dimensionality increases (suppose we have a grid of 10 points along each dimension).

6.1 Monte Carlo

If we choose the points from the distribution $p(x)$ then

$$\mathbb{E}_{p(x)}[\phi(x)] \simeq \frac{1}{T} \sum_{\tau=1}^T \phi(\mathbf{x}^{(\tau)}) \quad \text{where } \mathbf{x}^{(\tau)} \sim p(\mathbf{x})$$

Furthermore, due to the central limit theorem, the sum of the independent samples yields an unbiased estimate: $\mathbb{V}[\hat{\phi}] = \frac{\mathbb{V}[\phi]}{T}$ so this variance is independent of the dimensionality of \mathbf{x} .

This leads to the question of how we sample from $p(\mathbf{x})$? What if it is intractable?

6.2 Markov Chains & Gibbs Sampling

What if we are trying to find the posterior $p(\theta|x) = \frac{p(x|\theta)p(\theta)}{p(x)}$. We don't need the denominator to get the shape of the posterior - we just need the relative value of the posterior at one point versus all others. This involves (for continuous variables) an infinite number of calculations. This is where dependent sampling comes in, and in particular Markov chains. A Markov chain is a sequence defined by a transition function $q(x'|x)$. Gibbs sampling states that if you generate a new x but keep all other components the same, by sampling it from this distribution:

$$x'_i \sim p(x_i|x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_D)$$

If you iterate this over all the indices, and repeat this process many times then the sample will have the correct distribution. Of course this conditional distributions must be known.

<https://www.youtube.com/watch?v=ER3DDBFzH2g>

6.3 Example: Step Model

Suppose

$$x_{i \leq \theta} \sim \text{Poisson}(\lambda), \quad x_{i > \theta} \sim \text{Poisson}(\mu)$$

Such that

$$\begin{aligned} p(x_i) &= \frac{e^{-\lambda}\lambda^x}{x!} \\ &= \exp(x \log \lambda - \lambda - \log(x!)) \iff \log \frac{e^{-\lambda}\lambda^x}{x!} = -\lambda + x \log \lambda - \log(x!) \end{aligned}$$

with priors:

$$\lambda \sim \text{Gamma}(a, b), \quad \mu \sim \text{Gamma}(a, b), \quad \theta \sim \mathcal{U}(1851, 1961)$$

$$p(\lambda) = \frac{1}{\Gamma(a)} b^a \lambda^{a-1} e^{-b\lambda}$$

But since $\log\left(\frac{1}{\Gamma(a)} b^a \lambda^{a-1} e^{-b\lambda}\right) = -\log \Gamma(a) + a \log b + (a-1) \log \lambda - b\lambda$:

$$\begin{aligned} p(\lambda) &= a \log b - \log \Gamma(a) + (a-1) \log \lambda - b\lambda \\ p(\mu) &= a \log b - \log \Gamma(a) + (a-1) \log \mu - b\mu \end{aligned}$$

We wish to find the posterior, so we can use Bayes' theorem:

$$\begin{aligned} p(\theta, \lambda, \mu | x) &\propto p(x | \theta, \lambda, \mu) p(\theta, \lambda, \mu) \\ &= p(x_{i \leq \theta} | \theta, \lambda) \cdot p(x_{i > \theta} | \theta, \mu) \cdot p(\theta, \lambda, \mu) \\ &= \prod_{i \leq \theta} \exp(x_i \log \lambda - \lambda - \log(x_i!)) \prod_{i > \theta} \exp(x_i \log \mu - \mu - \log(x_i!)) \cdot p(\theta, \lambda, \mu) \\ \log p(\theta, \lambda, \mu | x) &\propto \sum_{i \leq \theta} (x_i \log \lambda - \lambda - \log(x_i!)) + \sum_{i > \theta} (x_i \log \mu - \mu - \log(x_i!)) + \log p(\theta, \lambda, \mu) \end{aligned}$$

To find the conditional posteriors for the parameters, we ignore terms which do not include that parameter.

$$\begin{aligned} \log p(\lambda | x, \theta, \mu) &= \sum_{i \leq \theta} (x_i \log \lambda - \lambda - \log(x_i!)) + a \log b - \log \Gamma(a) + (a-1) \log \lambda - b\lambda \\ &\propto \sum_{i \leq \theta} (x_i \log \lambda - \lambda) + (a-1) \log \lambda - b\lambda \\ &\propto (a-1 + \sum_{i \leq \theta} (x_i)) \log \lambda - \sum_{i \leq \theta} \lambda - b\lambda \\ &\propto (a-1 + \sum_{i \leq \theta} (x_i)) \log \lambda - (\theta + b)\lambda \\ &\propto \log \text{Gamma}(a + \sum_{i \leq \theta} (x_i), \theta + b) \end{aligned}$$

Similarly;

$$\begin{aligned} \log p(\mu | x, \lambda, \theta) &\propto (a-1 + \sum_{i > \theta} (x_i)) \log \mu - (N - \theta + b)\mu \\ &\propto \log \text{Gamma}(a + \sum_{i > \theta} (x_i), N - \theta + b) \end{aligned}$$

It is worth noting at this point that the Gamma and Poisson are conjugates.

$$\begin{aligned} \log p(\theta | x, \lambda, \mu) &\propto \sum_{i \leq \theta} (x_i \log \lambda - \lambda - \log(x_i!)) + \sum_{i > \theta} (x_i \log \mu - \mu - \log(x_i!)) \\ &\propto \sum_{i \leq \theta} (x_i \log \lambda) - \theta\lambda + \sum_{i > \theta} (x_i \log \mu) - (N - \theta)\mu \end{aligned}$$

This is not a standard distribution, but simple enough to sample. Sample it from $i = 0 : N$ and use it to construct a multinomial distribution

finish

7 Ranking

7.1 Towards Probabilistic Ranking

Suppose we have player 1 and player 2, with skills w_1, w_2 respectively. The skill difference is therefore $s = w_1 - w_2$. The performance may not be perfectly consistent though, so we can add noise:

$$t = s + n \quad \text{where } n \sim \mathcal{N}(0, 1)$$

The game outcome is given as $y = \text{sign}(t) = \begin{cases} +1 & \text{player 1 wins} \\ -1 & \text{player 2 wins} \end{cases}$

$$p(t|w_1, w_2) = \mathcal{N}(w_1 - w_2, 1)$$

$$p(y = 1|w_1, w_2) = p(t > 0|w_1, w_2) = \Phi(w_1 - w_2) \quad \text{where } \Phi \text{ is the cum. dist.}$$

We can now construct the likelihood:

$$p(\mathbf{y}|w_1, w_2) = \Phi(y(w_1 - w_2))$$

We can also write the likelihood as this kind of chain:

$$p(\mathbf{y}|w_1, w_2) = \iint p(y|t)p(t|s)p(s|w_1, w_2) dt ds$$

Recall Bayes' rule:

$$\begin{aligned} p(\mathbf{w}|\mathbf{y}) &= \frac{p(\mathbf{y}|\mathbf{w})p(\mathbf{w})}{p(\mathbf{y})} \\ p(w_1, w_2|y) &= \frac{p(y|w_1, w_2)p(w_1)p(w_2)}{p(y)} && \text{our data is } y \\ &= \frac{p(y|w_1, w_2)p(w_1)p(w_2)}{\iint p(w_1)p(w_2)p(y|w_1, w_2) dw_1 dw_2} \\ &= \frac{\Phi(y(w_1 - w_2))\mathcal{N}(w_1|\mu_1, \sigma_1^2)\mathcal{N}(w_2|\mu_2, \sigma_2^2)}{\iint \Phi(y(w_1 - w_2))\mathcal{N}(w_1|\mu_1, \sigma_1^2)\mathcal{N}(w_2|\mu_2, \sigma_2^2) dw_1 dw_2} \end{aligned}$$


where we set the prior to be $p(w_i) = \mathcal{N}(w_i|\mu_i, \sigma^2)$

Everytime one has a game and an outcome, then skills are correlated, since one player will win, so their skill will be higher. This posterior does not have a closed form, since it is not Gaussian.

The prior over y , the normalising constant, the model evidence, the marginal likelihood, does have a closed form:

$$p(y) = \Phi\left(\frac{y(\mu_1 - \mu_2)}{\sqrt{1 + \sigma_1^2 + \sigma_2^2}}\right)$$

Consider if we are very uncertain about the skills, then the argument to the cumulative distribution gets closer to zero, so the probability of the outcome gets closer to 50%.



7.2 Gibbs Sampling in TrueSkill

Suppose for games $g = 1, \dots, G$ we have the variable I_g and J_g for the id of the first and second player. The outcome is $y = \begin{cases} +1 & I_g \text{ wins,} \\ -1 & \text{otherwise} \end{cases}$.

We first initialise \mathbf{w} from a prior $p(\mathbf{w})$. We then need the performance differences:

$$p(t_g|w_{I_g}, y_g) \propto \delta(y_g - \text{sign}(t_g))\mathcal{N}(w_{I_g} - w_{J_g}, 1)$$

The conditional distribution of the performance differences as above is a univariate truncated Gaussian, sampled either using rejection sampling or "inverse transformation" method.

Jointly sample the skills:

$$p(\mathbf{w}|\mathbf{t}, \mathbf{y}) = p(\mathbf{w}|\mathbf{t}) \propto p(\mathbf{w}) \prod_{g=1}^G p(t_g|w_{I_g}, w_{J_g})$$

But $t = s + \mathcal{N}(0, 1) \implies p(t_g|w_{I_g}, w_{J_g}) \propto \mathcal{N}(\mathbf{w}; w_{I_g} - w_{J_g}, 1)$

Let $\mu_g = w_{I_g} - w_{J_g}$ and $t_g = \mu_1 - \mu_2$. Refer to appendix for an expansion which doesn't show much.

For Gibbs sampling, we then iterate back to calculating the performance differences.

When one conditions on a random variable, the value of the random variable is fixed, which simplifies things.

The product of two Gaussians yields an unnormalised Gaussian:

$$\begin{aligned} \mathcal{N}(\mu_a, \Sigma_a)\mathcal{N}(\mu_b, \Sigma_b) &= z_c \mathcal{N}(\mu_c, \Sigma_c) \\ \Sigma_c^{-1} &= \Sigma_a^{-1} + \Sigma_b^{-1} & \mu_c &= \Sigma_c(\Sigma_a^{-1}\mu_a + \Sigma_b^{-1}\mu_b) \end{aligned}$$

Suppose $p(\mathbf{w}) \sim \mathcal{N}(\mu_0, \Sigma_0)$. Using the above results of multiplying Gaussians:

$$\begin{aligned} \implies \Sigma^{-1} &= \Sigma_0^{-1} + \sum_{g=1}^G \Sigma_g^{-1} \\ \mu &= \end{aligned}$$

$$\begin{aligned} p(t_g|w_{I_g}, w_{J_g}) &\propto \exp\left(-\frac{1}{2}(w_{I_g} - w_{J_g} - t_g)^2\right) \\ &\propto \mathcal{N}\left(-\frac{1}{2} \begin{bmatrix} w_{I_g} - \mu_1 & w_{J_g} - \mu_2 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} w_{I_g} - \mu_1 \\ w_{J_g} - \mu_2 \end{bmatrix}, 1\right) \end{aligned}$$


An alternative to Gibbs sampling for TrueSkill is message passing on graphs.

7.3 Message Passing on Factor Graphs

Factor graphs are a type of *probabilistic graphical model*.

We can lay out the probabilities in a factor graph. Suppose:

$$p(v, w, x, y, z) = f_1(v, w)f_2(w, x)f_3(x, y)f_4(x, z)$$



We can now ask questions like what are the marginal distributions, conditional distributions, $p(w)$?

$$p(w) = \sum_v \sum_x \sum_y \sum_z f_1(v, w)f_2(w, x)f_3(x, y)f_4(x, z)$$

Computing this is K^4 sums (where K is the number of values the variables can take) and K possible values of w therefore $\mathcal{O}(K^5)$. We can break up the factor graph above into two subgraphs split at w .

$$\begin{aligned} p(w) &= \sum_v f_1(v, w) \sum_x \sum_y \sum_z f_2(w, x) f_3(x, y) f_4(x, z) \quad \text{and while we're here...} \\ &= \sum_v f_1(v, w) \sum_x f_2(w, x) \sum_y f_3(x, y) \sum_z f_4(x, z) \end{aligned}$$

We call these components **messages**

$$\begin{aligned} m_{f_1 \rightarrow w}(w) &= \sum_v f_1(v, w) \\ m_{f_2 \rightarrow w}(w) &= \sum_x \sum_y \sum_z f_2(w, x) f_3(x, y) f_4(x, z) \\ &= \sum_x f_2(w, x) \sum_y \sum_z f_3(x, y) f_4(x, z) \\ &= \sum_x f_2(w, x) m_{x \rightarrow f_2}(x) \\ p(w) &= m_{f_1 \rightarrow w}(w) \cdot \sum_x m_{f_3 \rightarrow x}(x) \cdot m_{f_4 \rightarrow x}(x) \end{aligned}$$

So nodes take incoming messages and passes them on.

In summary message passing involves three update equations:

- Marginals are the product of all incoming messages from neighbour factors

$$p(t) = \prod_{f \in F_t} m_{f \rightarrow t}(t)$$

- Messages from factors sum out all variables except the receiving one
- Messages from variables are the product of all incoming messages except the message from the receiving factor

$$m_{t \rightarrow f}(t) = \frac{p(t)}{m_{f \rightarrow t}(t)}$$

The benefits of this are partial and localised computations.

8 Expectation-Maximisation

8.1 Gaussian Mixture Models

In a GMM, the parameters are $\theta = \{\mu_j, \sigma_j^2, \pi_j\}_{j=1..k}$. We have one latent variable for each datapoint z_i , an assignment to a class. GMMs can represent any distribution. We wish to find:

$$\arg \max_{\theta} p(X|\theta) = \prod_i^N p(x_i|\theta) = \prod_i^N \pi_1 \mathcal{N}(x_i|\mu_1, \Sigma_1) + \dots$$

subject to $\sum_c^K \pi_c = 1$ and $\Sigma_i \succ 0$. This psd constraint is quite tough. So we introduce the EM algorithm for efficiently training these models.

We represent the GMM as a latent variable problem, where we introduce a latent variable z for each datapoint such that $p(z = c|\theta) = \pi_c$. We assign a Gaussian prior on data. We now marginalise away the latent variable:

$$p(x|z = c, \theta) = \mathcal{N}(x|\mu_c, \Sigma_c)$$

$$p(x|\theta) = \sum_{c=1}^K p(x|z = c, \theta)p(z = c|\theta)$$

which gives us the same likelihood result as without the latent variable, z , which we refer to as the source.




Figure 9: Graphical model for the GMM.

The idea of the EM algorithm is that we keep iterating between keeping the parameters and the sources fixed.

Keeping the sources fixed, we can calculate the parameters:

$$p(x|z = 1, \theta) = \mathcal{N}(x|\mu_1, \sigma_1^2)$$

where $p(z_i = j|\theta) = \pi_j$. We will derive μ_1, σ_1 for “soft” assignments in the next sections, which would result in $\mu_1 = \frac{\sum_i p(z_i=1|x_i, \theta)x_i}{\sum_i p(z_i=1|x_i, \theta)}$. For now, suppose we choose a hard assignment: $\mu_1 = \frac{\sum_{i \in \mathcal{J}_1} x_i}{|\mathcal{J}_1|}$, $\sigma_1^2 = \frac{\sum_{i \in \mathcal{J}_1} (x_i - \mu_1)^2}{|\mathcal{J}_1|}$.

Keeping the parameters fixed, we can work out the sources. Suppose our parameters are $p(x|z = 1, \theta) = \mathcal{N}(-2, 1)$, and we want $p(z = 1|x, \theta)$,

$$p(z = 1|x, \theta) = \frac{p(z = 1|\theta)p(x|z = 1, \theta)}{p(x|\theta)}$$

The normalising constant can be determined explicitly: it is marginalising over just K terms.

8.2 Mathematical Machinery

8.2.1 Jensen's Inequality

A function f is concave is such that all function values between any two points are greater or equal to the value of the line at that point which joints the two points:

$$\forall a, b, \alpha : f(\alpha a + (1 - \alpha)b) \geq \alpha f(a) + (1 - \alpha)f(b),$$

where $\alpha \in [0, 1]$ represents the distance along the line between a and b . This generalises to multiple α s:

$$f(\mathbb{E}_{p(z)}z) \geq \mathbb{E}_{p(z)}f(z) \quad (\text{Jensen's inequality})$$

8.2.2 Kullback-Leibler Divergence

$$\mathcal{KL}(p||q) = \int p(x) \log \frac{p(x)}{q(x)} dx$$

This is an asymmetric and non-negative metric. Proof:

$$-\mathcal{KL}(p||q) = \mathbb{E}_p \left(-\log \frac{p}{q} \right) = \mathbb{E}_p \left(\log \frac{q}{p} \right)$$

$$\mathbb{E}_p \left(\log \frac{p}{q} \right) \leq \log \left(\mathbb{E} \frac{p}{q} \right) = \log \int p(x) \frac{p(x)}{q(x)} dx = 0 \quad \text{using Jensen's inequality}$$

8.2.3 Mutual Information

$$\mathcal{MJ}(X, Y) = \mathcal{KL}(p_{X,Y}, p_x p_y)$$

more detail
needed

8.3 Expectation Maximisation Algorithm

This algorithm is for probabilistic models with latent variables, observed variables, and parameters. We will discuss the algorithm in terms of the GMM. First, we can write out the likelihood from the graphical model in Figure 11, by marginalising the joint:

$$p(x_i|\theta) = \frac{p(x_i|z_i = c, \theta)p(z_i = c|\theta)}{p(z_i = c|x_i, \theta)} = \sum_{c=1}^K p(x_i|z_i = c, \theta)p(z_i = c|\theta)$$

and now we wish to compute the maximum marginal likelihood:

$$\sum_{i=1}^N \log \sum_{c=1}^K p(x_i|z_i = c, \theta)p(z_i = c|\theta)$$

we could maximise with a stochastic optimizer, but it's not very efficient. Instead, let's find a lower bound by applying Jensen's inequality:

$$\sum_{i=1}^N \log \sum_{c=1}^K \frac{q(z_i = c)}{q(z_i = c)} p(x_i, z_i = c|\theta) \geq \sum_{i=1}^N \sum_{c=1}^K q(z_i = c) \log \frac{p(x_i, z_i = c|\theta)}{q(z_i = c)}$$

This q is known as the variational distribution. In order to illustrate what maximising this lower bound does, we start with Bayes' rule:

$$p(x|\theta) = \frac{p(x|z, \theta)p(z|\theta)}{p(z|x, \theta)} = \frac{p(x|z, \theta)p(z|\theta)}{q(z)} \frac{q(z)}{p(z|x, \theta)}$$
$$\log p(x|\theta) = \log \frac{p(x|z, \theta)p(z|\theta)}{q(z)} + \log \frac{q(z)}{p(z|x, \theta)}$$

Finally we average both sides w.r.t. $q(z)$:

$$\log p(x|\theta) = \underbrace{\int q(z) \log \frac{p(x|z, \theta)p(z|\theta)}{q(z)} dz}_{=\mathcal{L}(x|\theta), \text{ lower-bound functional}} + \underbrace{\int q(z) \log \frac{q(z)}{p(z|x, \theta)} dz}_{\mathcal{KL}\text{-divergence}}$$

The lower-bound functional is a lower bound since the second term is always non-negative.

We now have the log-marginal-likelihood as a sum of two terms. But how do we select q ? We do this using the EM algorithm. The EM algorithm is as follows:

- Initialise randomly $\theta^{t=0}$, then for $t = 1..T$:
- **E-step**: for fixed θ^{t-1} maximise the lower-bound functional, $\mathcal{L}(x|\theta)$, wrt $q(z)$
 Since the marginal likelihood term $\log p(x|\theta)$ does not depend on $q(z)$, maximising the lower-bound consists of minimising the \mathcal{KL} -divergence, by setting $q^t(z) = p(z|x, \theta^{t-1})$
- **M-step**: for fixed $q^t(z)$, maximise the lower-bound functional, $\mathcal{L}(x|\theta)$, wrt θ

$$\arg \max_{\theta} \int q(z) \log \frac{p(x|z, \theta)p(z|\theta)}{q(z)} dz = \arg \max_{\theta} \int q(z) \log p(x|z, \theta)p(z|\theta) dz - \int q(z) \log q(z) dz$$

The second term is the entropy of $q(z)$ but is not dependent on θ so throw it away.

$$\theta^t = \arg \max_{\theta} \int q^t(z) \log p(x|z, \theta)p(z|\theta) dz = \arg \max_{\theta} \mathbb{E}_q [\log p(x, z|\theta)]$$

For the GMM algorithm, where the likelihood is Gaussian and the prior is π_c , we can work out the analytical expression of the M-step:

$$\begin{aligned} \theta^t &= \arg \max_{\theta} \int q^t(z) \log p(x|z, \theta)p(z|\theta) dz = \arg \max_{\theta} \mathbb{E}_q [\log p(x, z|\theta)] \\ &= \sum_{i=1}^N \sum_{c=1}^K q^t(z=c) \log \left(\frac{1}{Z} \exp \left(-\frac{(x_i - \mu_c)^2}{2\sigma_c^2} \right) \pi_c \right) \\ &= \sum_{i=1}^N \sum_{c=1}^K q^t(z=c) \left(\log \frac{\pi_c}{Z} - \frac{(x_i - \mu_c)^2}{2\sigma_c^2} \right) \end{aligned}$$

Using school-level maths, in order to maximise we take the derivative w.r.t. each μ_c , for example:

$$\begin{aligned} \frac{\partial}{\partial \mu_1} &= \sum_{i=1}^N q^t(z=1) \left(0 - \frac{(x_i - \mu_1)^2}{2\sigma_1^2} \right) = 0 \\ &= \sum_{i=1}^N q^t(z=1) (x_i - \mu_1) \\ \implies \mu_1 &= \frac{\sum_{i=1}^N q^t(z=1) x_i}{\sum_{i=1}^N q^t(z=1)} \end{aligned}$$

When can the EM algorithm be used?

The EM algorithm works when the posterior distribution is known and tractable.

8.4 Example: K-Means

K-means is a clustering algorithm: an unsupervised method of grouping together datapoints into groups (clusters). Datapoints are assigned to their closest cluster *centroid* by minimising squared Euclidean distance. It is important to note that this is **not** a probabilistic method, so one could argue it does not belong here. Nevertheless, we shall discover K-means by analysing its close similarity with GMM. We first fix the $\sigma_c = 1$ and the assignment weights π_c to be uniform ($\frac{1}{M}$ where M is the number of Gaussians). This means that the new likelihood is:

$$p(x|z = c, \theta) = \frac{1}{Z} \exp(-0.5\|x - \mu_c\|^2).$$

We modify the E-step and M-step in the following ways:

- **E-step:** We also choose our variational distribution $q(z = c) \in Q$ where Q is the set of delta functions. As we saw, the E-step consists only of minimising the \mathcal{KL} -divergence between the variational distribution and the true posterior. Thus the optimal q will be the delta function:

$$q(z) = \begin{cases} 1 & \text{if } z = c \\ 0 & \text{otherwise} \end{cases}$$

corresponding to the marginal likelihood maximisation:

$$\begin{aligned} c_i &= \arg \max_c p(z = c|x, \theta) = \arg \max_c \frac{1}{Z} p(x|z = c, \theta) p(z = c|\theta) \\ &= \arg \max_c \frac{1}{Z} \exp(-0.5\|x - \mu_c\|^2) \pi_c \\ &= \arg \min_c \|x - \mu_c\|^2 \end{aligned}$$

This result is easy to calculate by taking the derivative.

- **M-step:**

Recall that for fixed q^k , we maximise the lower bound w.r.t. θ . This is essentially the average of all datapoints assigned to a cluster.

$$\begin{aligned} \mu_1^k &= \frac{\sum_{i=1}^N q^t(z_i=1) x_i}{\sum_{i=1}^N q^t(z_i=1)} \\ \implies \mu_c^{k+1} &= \frac{\sum_{i:c_i=c} x_i}{|\{i:c_i=c\}|} \end{aligned}$$

8.5 Example: Topic Modelling


This section covers three models for modelling text documents, with increasing complexity. Starting with simple bag-of-words models, moving to the EM algorithm and finally Latent Dirichlet Allocation.

The general layout of the problem is as follows.

- There are D documents,
- with a vocab size of M words,
- N_d is the word count in doc d ,
- w_{nd} is the n -th word in document d , parameterised by β , the probabilities of each word. Our data is given a Categorical likelihood: $w_{nd}|\beta \sim \text{Cat}(w_{nd}|\beta)$

The simplest model is shown on the right. We will find an MLE solution, $\hat{\beta}$, which we constrain to ensure the β 's normalises to 1.

$$\begin{aligned}\hat{\beta} &= \arg \max_{\beta} \prod_{d=1}^D \prod_n^{N_d} p(\mathbf{w}|\beta) \text{ s.t. } \sum \beta_m = 1 \\ &= \arg \max_{\beta} \log \prod_{m=1}^M \beta_m^{c_m} = \arg \max_{\beta} \sum_{m=1}^M c_m \log \beta_m\end{aligned}$$



where c_m is total count of word m : $c_m = \sum_{d=1}^D \sum_n^{N_d} \mathbb{I}(w_{nd} = m)$. We use the Lagrange multiplier method to carry out this constrained optimisation problem. Setting derivatives to zero:

$$\begin{aligned}\mathcal{L} &= \sum_{m=1}^M c_m \log \beta_m + \lambda \left(1 - \sum_{m=1}^M \beta_m \right) \\ \frac{\partial \mathcal{L}}{\partial \beta_m} &= \frac{c_m}{\beta_m} - \lambda = 0 \implies c_m = \frac{\beta_m}{\lambda} \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= 1 - \sum_{m=1}^M \beta_m = 0 \implies \sum_{m=1}^M \beta_m = 1\end{aligned}$$

We can combine these two equations to yield our MLE:

$$\begin{aligned}\sum_{m=1}^M \frac{c_m}{\lambda} &= 1 \implies \lambda = \sum_{m=1}^M c_m \\ \implies \hat{\beta}_m &= \frac{c_m}{N}\end{aligned}$$

where N is total number of words (including repetitions) in the corpus.

This model is intuitive but limited since all documents are modelled by the same global word frequency distribution, but we want a model which accommodates different topics of documents.

In this new model (shown on the right), we add more latent variables: $z_d \in \{1, \dots, K\}$ assigns document d to one of K topics; and $\theta_k = p(z_d = k)$ is the probability of document d being assigned to k . They form the parameters of a categorical distribution, θ .




Figure 10: $z \sim \text{Cat}(\theta), w|z \sim \text{Cat}(\beta_{z_d})$

First, we write down the likelihood (try reading it off the graphical model!):

$$\begin{aligned}\log p(\mathbf{w}|\theta, \beta) &= \log \prod_{d=1}^D p(\mathbf{w}_d|\theta, \beta) \\ &= \log \prod_{d=1}^D \sum_{k=1}^K p(\mathbf{w}_d = \mathbf{1}, z_d = k|\theta, \beta) \quad (\text{marginalising the } z_k) \\ &= \log \prod_{d=1}^D \sum_{k=1}^K p(z_d = k|\theta) p(\mathbf{w}_d|z_d = k, \beta) \\ &= \sum_{d=1}^D \log \sum_{k=1}^K p(z_d = k|\theta) \prod_{n=1}^{N_d} p(w_{nd}|z_d = k, \beta)\end{aligned}$$

What does this look like? GMM! So we can use the EM algorithm! Recall that

$$\log p(x|\theta) = \int q(z) \log \frac{p(x|z, \theta)p(z|\theta)}{q(z)} dz + \int q(z) \log \frac{q(z)}{p(z|x, \theta)} dz$$

Our EM algorithm is then as follows:

- **E-step:** for fixed θ^{t-1} maximise the lower-bound wrt $q(z)$. This consists of minimising the \mathcal{KL} -divergence, by setting $q^t(z)$ to the posterior:

$$\begin{aligned} q^t(z_d = k) &= p(z_d = k | \theta^{t-1}) \prod_{n=1}^{N_d} p(w_{nd} | z_d = k, \beta) \\ &= \theta_k \text{Mult}(c_{1d}, \dots, c_{Md} | \beta_k, N_d) \\ &:= r_{kd} \end{aligned}$$

- **M-step:** for fixed $q^t(z)$, maximise the lower-bound wrt θ .

8.5.1 Latent Dirichlet Allocation

A *symmetric* Dirichlet simply has all parameters identical ($\forall i : \alpha_i = \alpha$).

9 Variational Inference

Recall our GMM model:




Figure 11: Graphical model for the GMM.

We seek the **posterior distribution** of the latent variable given the observations: $p(z|x)$

This is often **intractable**, why???

VI approximates the intractable distribution with a simpler one. The parameters of which are called **variational parameters** and are optimised. One such objective function is the ELBO which we came across in Section 8. Supposing we use a Gaussian posterior approximation, we would optimise the mean and deviation parameters **for each observation**. Variational parameters, therefore, scale with the dataset size.

9.1 Amortizing Variational Inference

Instead of optimising the free variational parameters directly, we can create a parameterised function that maps from observation space to variational parameter space. We end up with a constant number of variational parameters. The downside is that this results in less expressivity than free optimisation.

9.2 General Form

To give a more general comment on variational inference, we take:

$$\begin{aligned}\log p(x) &= \log \int p(x|z)p(z)dz = \log \int p(x, z) \frac{q(z)}{q(z)} dz \\ &= \log \mathbb{E} \left(\frac{p(x, z)}{q(z)} \right) \\ &\geq \mathbb{E} \left(\log \frac{p(x, z)}{q(z)} \right) = \int q(z) \log \frac{p(x|z)p(z)}{q(z)} dz\end{aligned}$$

which can be further decomposed into the lower bound and the \mathcal{KL} -divergence:

$$\begin{aligned}\int q(z) \log \frac{p(x|z)p(z)}{q(z)} dz &= \int q(z) \log p(x|z) dz - \int q(z) \log \frac{q(z)}{p(z)} \\ &= F(q) - \mathcal{KL}(q||p)\end{aligned}$$

9.2.1 Mean Field Approximation

This is a method of finding this \mathcal{KL} term. We select a family of distributions $Q = \{q|q(z) = \prod_{i=1}^d q_i(z_i)\}$

$$\begin{aligned}
\min_{q_k} \mathcal{KL}(\prod_{i=1}^d q_i || p) &= \int \prod_{i=1}^d q_i \log \frac{\prod_{i=1}^d q_i}{p} dz = \int \prod_{j=1}^d q_j \log \prod_{i=1}^d q_i dz - \int \prod_{i=1}^d q_i \log p dz \\
&= \sum_{i=1}^d \int \prod_{j=1}^d q_j \log q_i - \int \prod_{i=1}^d q_i \log p dz \\
&= \int \prod_{j=1}^d q_j \log q_k dz + \sum_{i \neq k} \int \prod_{j=1}^d q_j \log q_j - \int \prod_{i=1}^d q_i \log p dz
\end{aligned}$$

but $\int \prod_{j=1}^d q_j \log q_k dz = \int q_k \log q_k \underbrace{\int \prod_{j \neq k}^d q_j dz}_{1 \times \dots \times 1 = 1} dz_k$ and since we are optimising w.r.t q_k :

$$\implies \min_{q_k} \mathcal{KL}(\prod_{i=1}^d q_i || p) = \int q_k \log q_k dz_k - \int \prod_{i=1}^d q_i \log p dz$$

following a similar procedure for the second term
 $= t.b.c$

9.3 Examples

9.3.1 Inducing Point Approximation for Gaussian Processes

We will now go through the inducing point approximation which is commonly used with multi-output GPs. With D outputs, let y_d be the d th Gaussian process, such that $y_d = \mathbf{f}_d + \epsilon$ and $\epsilon \sim \mathcal{N}(0, \beta^2)$. The inducing points are \mathbf{u}_d . First, we apply the general form of VI from Section 9.2:

$$\log p(Y) = \log \int p(Y|X)p(X)dX \geq \int q(X) \log p(Y|X)dX - \int q(X) \log \frac{q(X)}{p(X)} dX \quad (10)$$

$$= F(q) - \mathcal{KL}(q||p) \quad (11)$$

We use the variational distribution $q(\mathbf{f}_d, \mathbf{u}_d) = p(\mathbf{f}_d|\mathbf{u}_d)\phi(\mathbf{u}_d)$. The marginal likelihood is now:

$$\begin{aligned}
\log p(y_d|X) &= \log \iint p(y_d|\mathbf{f}_d)p(\mathbf{f}_d|\mathbf{u}_d)p(\mathbf{u}_d) d\mathbf{f}_d d\mathbf{u}_d \\
&= \log \iint \frac{p(y_d|\mathbf{f}_d)p(\mathbf{f}_d|\mathbf{u}_d)p(\mathbf{u}_d)p(\mathbf{f}_d|\mathbf{u}_d)\phi(\mathbf{u}_d)}{p(\mathbf{f}_d|\mathbf{u}_d)\phi(\mathbf{u}_d)} d\mathbf{f}_d d\mathbf{u}_d
\end{aligned}$$

Now apply Jensen's inequality where \mathbb{E} is over variational distribution

$$\log p(y_d|X) \geq \iint p(\mathbf{f}_d|\mathbf{u}_d)\phi(\mathbf{u}_d) \log \frac{p(y_d|\mathbf{f}_d)p(\mathbf{f}_d|\mathbf{u}_d)p(\mathbf{u}_d)}{p(\mathbf{f}_d|\mathbf{u}_d)\phi(\mathbf{u}_d)} d\mathbf{f}_d d\mathbf{u}_d \quad (12)$$

$$= \int \phi(\mathbf{u}_d) \left[\int p(\mathbf{f}_d|\mathbf{u}_d) \log p(y_d|\mathbf{f}_d) + \underbrace{p(\mathbf{f}_d|\mathbf{u}_d) \log \frac{p(\mathbf{u}_d)}{\phi(\mathbf{u}_d)}}_{\text{integrates to 1}} d\mathbf{f}_d \right] d\mathbf{u}_d \quad (13)$$

$$= \int \phi(\mathbf{u}_d) \underbrace{\mathbb{E}_{p(\mathbf{f}_d|\mathbf{u}_d)}[\log p(y_d|\mathbf{f}_d)]}_{\leq \log \mathbb{E}_{p(\mathbf{f}_d|\mathbf{u}_d)}[p(y|\mathbf{f}_d)] = \log p(y|\mathbf{u}_d)} d\mathbf{u}_d + \underbrace{\int \phi(\mathbf{u}_d) \log \frac{p(\mathbf{u}_d)}{\phi(\mathbf{u}_d)} d\mathbf{u}_d}_{\mathcal{KL}(\phi||p)} \quad (14)$$

$$= \mathcal{L}_1 \quad (15)$$

- Note that $p(\mathbf{f}_d|\mathbf{u}_d) = \mathcal{N}(\mathbf{f}_d|\boldsymbol{\alpha}_d, K_{NN} - K_{NM}K_{MM}^{-1}K_{MN})$ and $p(y_d|\mathbf{f}_d) = \mathcal{N}(0, \sigma^2)$.
- Note that the brown expectation is a lower bound on the conditional $p(y|\mathbf{u}_d)$. With a Gaussian likelihood, $p(y|\mathbf{u}_d)$ is tractable, but in $O(N^3)$ compared with $O(M^3)$ for the lower bound.
- \mathcal{L}_1 is the lower bound from Titsias and Lawrence [2010], Hensman et al. [2013].

We take the trace of the log Gaussian since the trace of a scalar is the scalar.

$$\begin{aligned}
\mathbb{E}_{p(\mathbf{f}_d|\mathbf{u}_d)}[\log p(y_d|\mathbf{f}_d)] &= \mathbb{E}_{p(\mathbf{f}_d|\mathbf{u}_d)} \left[-\frac{N}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} (y_d^\top y_d - 2\mathbf{f}_d^\top y + \mathbf{f}_d^\top \mathbf{f}_d) \right] \\
&= -\frac{N}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \text{tr} \left(y_d^\top y_d - 2\mathbb{E}_{p(\mathbf{f}_d|\mathbf{u}_d)}[\mathbf{f}_d^\top y] + \underbrace{\mathbb{E}_{p(\mathbf{f}_d|\mathbf{u}_d)}[\mathbf{f}_d^\top \mathbf{f}_d]}_{\text{Cov}[f] + \boldsymbol{\alpha}\boldsymbol{\alpha}^\top} \right) \\
&= \underbrace{-\frac{N}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \text{tr}(y_d^\top y_d - 2\boldsymbol{\alpha}_d^\top y + \boldsymbol{\alpha}_d \boldsymbol{\alpha}_d^\top + K_{NN} - K_{NM}K_{MM}^{-1}K_{MN})}_{\log \mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2)} \\
&= \log \mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2) - \frac{1}{2\sigma^2} \text{tr} \left(K_{NN} - \underbrace{K_{NM}K_{MM}^{-1}K_{MN}}_{Q_{nn}} \right) \tag{16}
\end{aligned}$$

where $\boldsymbol{\alpha}_d = K_{NM}K_{MM}^{-1}\mathbf{u}_d$. We now plug Eq. 16 into Eq. 14, giving us Eq. 13 from Titsias and Lawrence [2010]:

$$\mathcal{L}_1 = \int \phi(\mathbf{u}_d) \log \frac{\mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2)p(\mathbf{u}_d)}{\phi(\mathbf{u}_d)} d\mathbf{u}_d - \frac{1}{2\sigma^2} \text{tr}(K_{NN} - Q_{nn}) \tag{17}$$

How do we acquire the optimal variational distribution?

1. We could take the (functional) derivative of \mathcal{L}_1 w.r.t. $\phi(\mathbf{u}_d)$, set it to zero, and plug that back into \mathcal{L}_1 .

$$0 = \frac{d\mathcal{L}_1}{d\phi} = \frac{d}{d\phi} \left[\int \phi(\mathbf{u}_d) \log \frac{\mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2)p(\mathbf{u}_d)}{\phi(\mathbf{u}_d)} d\mathbf{u}_d - \frac{1}{2\sigma^2} \text{tr}(K_{NN} - Q_{nn}) \right] \tag{18}$$

$$= \frac{d}{d\phi} \int \phi \log \frac{A}{\phi} d\mathbf{u}_d \tag{19}$$

$$= \int \frac{d}{d\phi} \phi \log \frac{A}{\phi} d\mathbf{u}_d = \int \frac{d}{d\phi} [\phi \log A - \phi \log \phi] d\mathbf{u}_d \tag{20}$$

$$= \int \log A - \left[\log \phi + \phi \frac{1}{\phi} \right] d\mathbf{u}_d = \int \log \frac{A}{\phi} - 1 d\mathbf{u}_d \tag{21}$$

$$= \int \log \frac{A}{\phi} - \log e d\mathbf{u}_d \tag{22}$$

$$\int \log \frac{A}{\phi} d\mathbf{u}_d = \int \log e d\mathbf{u}_d \tag{23}$$

$$\frac{A}{\phi} = e \implies \phi \propto A = \mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2)p(\mathbf{u}_d) \tag{24}$$

2. Another way is to reverse the Jensen's inequality, which we shall see now.

$$\mathcal{L}_1 = \int \underbrace{\phi(\mathbf{u}_d) \log \frac{\mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2)p(\mathbf{u}_d)}{\phi(\mathbf{u}_d)} d\mathbf{u}_d}_{\text{by reverse Jensen's inequality}} - \frac{1}{2\sigma^2} \text{tr}(K_{NN} - Q_{nn}) \quad (25)$$

$$\leq \log \int \phi(\mathbf{u}_d) \frac{\mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2)p(\mathbf{u}_d)}{\phi(\mathbf{u}_d)} d\mathbf{u}_d \quad \text{by reverse Jensen's inequality}$$

$$\Rightarrow \mathcal{L}_1 \leq \log \int \mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2)p(\mathbf{u}_d) d\mathbf{u}_d - \frac{1}{2\sigma^2} \text{tr}(K_{NN} - Q_{nn}) \quad (26)$$

$$= \log \int \mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2)\mathcal{N}(u|0, K_{MM}) d\mathbf{u}_d - \frac{1}{2\sigma^2} \text{tr}(K_{NN} - Q_{nn}) \quad (27)$$

$$= \log \int \mathcal{N}(y_d|\boldsymbol{\alpha}_d, \sigma^2) \underbrace{\mathcal{N}(K_{NM}K_{MM}^{-1}u|0, K_{NM}K_{MM}^{-1}K_{MN})}_{\text{scaled Gaussian: Cov}(BX)=BCov(X)B^T} d\mathbf{u}_d - \frac{1}{2\sigma^2} \text{tr}(K_{NN} - Q_{nn}) \quad (28)$$

$$= \log \int \mathcal{N}(y_d|0, Q_{NN} + \sigma^2 I) d\mathbf{u}_d - \frac{1}{2\sigma^2} \text{tr}(K_{NN} - Q_{NN}) \quad (29)$$

$$= \log \mathcal{N}(y|0, Q_{NN} + \sigma^2 I) - \frac{1}{2\sigma^2} \text{tr}(K_{NN} - Q_{nn}) \quad (30)$$

$$= \mathcal{L}_2 \quad (31)$$

where Eq. 29 comes from: $\mathcal{N}(a|\mu_1, \Sigma_1)\mathcal{N}(\mu_1|\mu_2, \Sigma_2) = \mathcal{N}(a|\mu_2, \Sigma_1 + \Sigma_2)$. This gives us \mathcal{L}_2 from Hensman et al. [2013].


Continuing with Hensman et al. [2013], we wish to derive a new lower bound which which we can apply stochastic variational inference (SVI). SVI scales to large datasets. Marginalising \mathbf{u}_d re-introduces dependencies on the observations, so Hensman et al. [2013] writes an explicit variational distribution $q(\mathbf{u}_d)$. This enables the derivation of **natural gradients** and therefore SVI.

SVI requires a set of global variables. For our purposes, \mathbf{u}_d will be the global variables and a new lower bound is defined:

$$\log p(y|\mathbf{X}) \geq \mathbb{E}_{q(\mathbf{u}_d)} [\mathcal{L}_1 + \log p(\mathbf{u}_d) - \log q(\mathbf{u}_d)] \triangleq \mathcal{L}_3$$

9.3.2 Deep Gaussian Processes

Imagine each layer is a multi-output GP, for example with Q inputs and D output units, $\mathbf{f} : \mathbb{R}^Q \rightarrow \mathbb{R}^D$. Each layer, therefore, would add many model parameters. Moreover, how do we pick the size of layers and number of layers? Deep GPs as defined in Damianou and Lawrence [2013], therefore, marginalise out the entire latent space. We start with a two-layer model:



First, the marginal likelihood:

$$\log p(\mathbf{Y}) = \log \iint p(\mathbf{Y}|\mathbf{X})p(\mathbf{X}|\mathbf{Z})p(\mathbf{Z})d\mathbf{X}d\mathbf{Y}$$

where $Z \sim \mathcal{N}(0, I)$

$$\log p(\mathbf{Y}) \geq \int \phi \log \frac{p(\mathbf{Y}, \mathbf{F}^{\mathbf{Y}}, \mathbf{F}^{\mathbf{X}}, \mathbf{X}, \mathbf{Z})}{\phi} d\mathbf{X}d\mathbf{Y} \quad (32)$$

hensman
2012 nips
paper: col-
lapsed stuff

natural gra-
dient intro

define global
variables

where we can factorise $p(\mathbf{Y}, \mathbf{F}^{\mathbf{Y}}, \mathbf{F}^{\mathbf{X}}, \mathbf{X}, \mathbf{Z}) = p(\mathbf{Y}|\mathbf{F}^{\mathbf{Y}})p(\mathbf{F}^{\mathbf{Y}}|\mathbf{X})p(\mathbf{X}|\mathbf{F}^{\mathbf{X}})p(\mathbf{F}^{\mathbf{X}}|\mathbf{Z})p(\mathbf{Z})$. This is intractable as \mathbf{X} and \mathbf{Z} appear in the integral non-linearly. Again, we use inducing points, yielding:

$$p(\mathbf{Y}, \mathbf{F}^{\mathbf{Y}}, \mathbf{F}^{\mathbf{X}}, \mathbf{X}, \mathbf{Z}, \mathbf{U}^{\mathbf{Y}}, \mathbf{U}^{\mathbf{X}}, \tilde{\mathbf{X}}, \tilde{\mathbf{Y}}) = p(\mathbf{Y}|\mathbf{F}^{\mathbf{Y}})p(\mathbf{F}^{\mathbf{Y}}|\mathbf{U}^{\mathbf{Y}}, \mathbf{X})p(\mathbf{U}^{\mathbf{Y}})p(\mathbf{X}|\mathbf{F}^{\mathbf{X}})p(\mathbf{F}^{\mathbf{X}}|\mathbf{U}^{\mathbf{X}}, \mathbf{Z})p(\mathbf{U}^{\mathbf{X}})p(\mathbf{Z})$$

Note that, for brevity, $p(\mathbf{U}^{\mathbf{Y}}) = p(\mathbf{U}^{\mathbf{Y}}|\tilde{\mathbf{X}})$ and similarly for $p(\mathbf{U}^{\mathbf{X}})$. We define ϕ as:

$$\phi(\mathbf{F}^{\mathbf{Y}}, \mathbf{U}^{\mathbf{Y}}, \mathbf{F}^{\mathbf{X}}, \mathbf{U}^{\mathbf{X}}) = p(\mathbf{F}^{\mathbf{Y}}|\mathbf{U}^{\mathbf{Y}}, \mathbf{X})p(\mathbf{F}^{\mathbf{X}}|\mathbf{U}^{\mathbf{X}}, \mathbf{Z})q(\mathbf{U}^{\mathbf{Y}})q(\mathbf{U}^{\mathbf{X}})q(\mathbf{X})q(\mathbf{Z})$$

We can now plug this into the unruly integral and separate terms (Eq. 13 in Damianou and Lawrence [2013]):

$$\log p(\mathbf{Y}) \geq \int \phi \log \frac{p(\mathbf{Y}|\mathbf{F}^{\mathbf{Y}})p(\mathbf{F}^{\mathbf{Y}}|\mathbf{U}^{\mathbf{Y}}, \mathbf{X})p(\mathbf{U}^{\mathbf{Y}})p(\mathbf{X}|\mathbf{F}^{\mathbf{X}})p(\mathbf{F}^{\mathbf{X}}|\mathbf{U}^{\mathbf{X}}, \mathbf{Z})p(\mathbf{U}^{\mathbf{X}})p(\mathbf{Z})}{p(\mathbf{F}^{\mathbf{Y}}|\mathbf{U}^{\mathbf{Y}}, \mathbf{X})p(\mathbf{F}^{\mathbf{X}}|\mathbf{U}^{\mathbf{X}}, \mathbf{Z})q(\mathbf{U}^{\mathbf{Y}})q(\mathbf{U}^{\mathbf{X}})q(\mathbf{X})q(\mathbf{Z})} d\{\mathbf{X}, \mathbf{Z}, \mathbf{F}^{\mathbf{X}}, \mathbf{F}^{\mathbf{Y}}, \mathbf{U}^{\mathbf{X}}, \mathbf{U}^{\mathbf{Y}}\} \quad (33)$$

$$= \int \phi \log \frac{p(\mathbf{Y}|\mathbf{F}^{\mathbf{Y}})p(\mathbf{U}^{\mathbf{Y}})p(\mathbf{X}|\mathbf{F}^{\mathbf{X}})p(\mathbf{U}^{\mathbf{X}})p(\mathbf{Z})}{q(\mathbf{U}^{\mathbf{Y}})q(\mathbf{U}^{\mathbf{X}})q(\mathbf{X})q(\mathbf{Z})} d\dots \quad (34)$$

$$= \mathbf{g}_Y + \mathbf{r}_X + \mathcal{H}_{q(\mathbf{X})} - \mathcal{KL}(q(\mathbf{Z})||p(\mathbf{Z})) \quad (35)$$

It's quite a large integral, but it is trivial algebraic rearrangements to derive \mathbf{g}_Y , \mathbf{r}_X , $\mathcal{H}_{q(\mathbf{X})}$, which I show below without derivation:

$$\mathbf{g}_Y = \mathbb{E}_{p(\mathbf{F}^{\mathbf{Y}}|\mathbf{U}^{\mathbf{Y}}, \mathbf{X})q(\mathbf{U}^{\mathbf{Y}})q(\mathbf{X})} \left[\log p(\mathbf{Y}|\mathbf{F}^{\mathbf{Y}}) + \log \frac{p(\mathbf{U}^{\mathbf{Y}})}{q(\mathbf{U}^{\mathbf{Y}})} \right] \quad (36)$$

$$\mathbf{r}_X = \mathbb{E}_{p(\mathbf{F}^{\mathbf{X}}|\mathbf{U}^{\mathbf{X}}, \mathbf{Z})q(\mathbf{U}^{\mathbf{X}})q(\mathbf{X})q(\mathbf{Z})} \left[\log p(\mathbf{X}|\mathbf{F}^{\mathbf{X}}) + \log \frac{p(\mathbf{U}^{\mathbf{X}})}{q(\mathbf{U}^{\mathbf{X}})} \right] \quad (37)$$

$$\mathcal{H}_{q(\mathbf{X})} = \int q(\mathbf{X}) \log q(\mathbf{X}) d\mathbf{X} \quad (38)$$

We can conclude:

- $q(\mathbf{U}^{\mathbf{Y}}), q(\mathbf{U}^{\mathbf{X}})$ are free-form variational distributions
- $q(\mathbf{X}), q(\mathbf{Z})$ are Gaussians factorised along the dimensions
- Both $\mathbf{g}_Y, \mathbf{r}_X$ are known Gaussian densities, so are tractable.
- Notice that \mathbf{g}_Y is equivalent to \mathcal{L}_1 from Eq. 11.

continue

9.4 Implementation Details

Now that we know some lower bounds, it is crucial that we understand how we might implement this. Suppose we have the lower bound:

$$\underbrace{\mathbb{E}_{q(\mathbf{X})} [\mathbb{E}_{q(F|m, S, X, Z)} [\log p(Y|F)]]}_{\mathbf{3}} - \underbrace{\mathcal{KL}[q(\mathbf{X})||p(\mathbf{X})]}_{\mathbf{2}} - \underbrace{\mathcal{KL}[q(U)||p(U)]}_{\mathbf{1}}$$

1: KL between the inducing points First, let's calculate $\mathcal{KL}[q(U)||p(U)]$. We need two variational parameters here: q_μ , the mean of the inducing points; and q_σ , the square-root of the inducing points covariance. The KL-divergence between two multivariate Gaussians is:

$$\frac{1}{2} \left[\log \frac{|\Sigma_2|}{|\Sigma_1|} - d + \text{tr}\{\Sigma_2^{-1}\Sigma_1\} + (\mu_2 - \mu_1)^\top \Sigma_2^{-1}(\mu_2 - \mu_1) \right]$$

where d is the number of dimensions plus the number of inducing points, $\Sigma_1 = \mathbf{q}_\sigma^2$, $\Sigma_2 = K_{MM}$, $\mu_1 = \mathbf{q}_\mu$, and $\mu_2 = 0$

Theorem 1. *The determinant of an upper triangular matrix is the product of the diagonal.*

Theorem 2. *The square-root of the determinant is equal to the determinant of the square-root. More generally, $|A^n| = |A|^n$*

Using these results, we can find the first term.

$$\begin{aligned} -\log |\mathbf{q}_\sigma^2| &= -\log |\mathbf{q}_\sigma|^2 = -\log \prod \text{diag}(\mathbf{q}_\sigma)^2 = -\sum \log \text{diag}(\mathbf{q}_\sigma)^2 \\ \log |K_{MM}| &= \log |LL^\top| \quad \text{where } L = \text{chol}(K_{MM}) \\ &= \log 2|L| = \log 2 \prod \text{diag}(L) \end{aligned}$$

Combining the two gives the first term. Now, for the trace term. To calculate $\text{tr}(X)$ where $X = K_{MM}^{-1}\mathbf{q}_\sigma^2$, notice that $X^{\frac{1}{2}} = L^{-1}\mathbf{q}_\sigma$. Therefore, since L is triangular, $\text{tr}(X) = \sum \text{triangular_solve}(L, \mathbf{q}_\sigma)^2$

Finally, we need the $\mathbf{q}_\mu^\top K_{MM}^{-1}\mathbf{q}_\mu$ term. $K_{MM}^{-1}\mathbf{q}_\mu$ can be calculated using `cholesky_solve(L, \mathbf{q}_μ)`, and therefore the whole thing is `$\mathbf{q}_\mu \text{cholesky_solve}(L, \mathbf{q}_\mu)$`

2: KL between the latents

Now, we wish to find $\mathcal{KL}[q(\mathbf{X})||p(\mathbf{X})]$.

3: Marginal likelihood

Finally, we can calculate the third term using Monte Carlo sampling of $q(\mathbf{X})$. First, we set

$$q(\mathbf{X}) = \prod_{n=1}^N \mathcal{N}(\boldsymbol{\mu}_n, S_n)$$

so we have $\{(\boldsymbol{\mu}_n, S_n)\}_n^N$ variational parameters where $\boldsymbol{\mu} \in \mathbb{R}^Q$.

$p(\mathbf{f}_d|\mathbf{u}_d) = \mathcal{N}(\mathbf{f}_d|\boldsymbol{\alpha}_d, K_{NN} - K_{NM}K_{MM}^{-1}K_{MN})$ where $\boldsymbol{\alpha}_d = K_{NM}K_{MM}^{-1}\mathbf{u}_d$.

Algorithm 1: Marginal likelihood.

Input : Current i

Output: Next i

First, we generate our latent samples:

$z \sim \mathcal{N}(0, I) \in \mathbb{R}^{N_s, N, Q}$

$z = \boldsymbol{\mu} + zS^{\frac{1}{2}}$ (reparameterisation trick)

Second, we sample from variational posterior:

$\text{temp} = L^{-1}K_{MN} = \text{triangular_solve}(L, K_{MN})$

$\boldsymbol{\alpha} = L^{-1\top}L^{-1} = \text{triangular_solve}(\text{temp},)$

10 Stochastic Calculus

When reading probabilistic machine learning literature, especially when applied to physical or dynamical systems, some foundational understanding of stochastic calculus is extremely useful. I cover only the main topics around stochastic processes and stochastic differential equations, so note that this only scratches the surface.

10.1 Probability Generating Function

We shall see that the probability generating function (pgf) is a useful **tool** when manipulating probabilities. It is defined as:

$$\mathcal{G}_X(s) = \mathbb{E} [s^X] = \sum_{x \geq 0} s^x p_X(x)$$

where s is a **complex variable in the unit circle** (see Fig. 12).

Note that the normalisation of the pmf p_x is $\mathcal{G}_X(1) = \sum_{x \geq 0} p_X(x)$. Therefore,,

$$\mathcal{G}_X(1) = \begin{cases} 1 & \text{if } p_X(\infty) = 0 \\ 1 - p_X(\infty) & \text{otherwise} \end{cases}$$

where $p_X(\infty)$ is the probability that X takes a value outside of $[0, \infty)$. For example, this could be the probability that the temperature will go below zero.

We can differentiate this function n times as follows:

$$\begin{aligned} \frac{d^n \mathcal{G}_X(s)}{ds^n} &= \sum_{x \geq n} x \cdot (x-1) \cdot \dots \cdot (x-n+1) \cdot s^{x-n} p_X(x) \\ &\quad \text{summand is zero for } 0 \leq x \leq n-1 \\ &= \sum_{x \geq n} \frac{x!}{(x-n)!} \cdot s^{x-n} p_X(x) \end{aligned}$$

Why is this useful? Notice that the first derivative at $s \rightarrow 1^-$ is the expectation:

$$\lim_{s \rightarrow 1^-} \frac{d\mathcal{G}_X(s)}{ds} = \sum_{x \geq 1} x \cdot s^{x-1} p_X(x) = \sum_{x \geq n} x p_X(x)$$

Similarly, we can work out the variance:

$$\begin{aligned} \text{Var}[X] &= \mathbb{E} [X^2] - \mathbb{E} [X]^2 \\ &= \mathcal{G}_X''(1) - \mathcal{G}_X'(1) + (\mathcal{G}_X'(1))^2 \end{aligned}$$

When $s \rightarrow 0$, the derivative of the pgf gives the probabilities:

$$\begin{aligned} \lim_{s \rightarrow 0} \sum_{x \geq n} \frac{x!}{(x-n)!} \cdot s^{x-1} p_X(x) &= n! \cdot p_X(x) + (n+1)! \cdot s p_X(x) + \frac{(n+1)!}{2!} \cdot s^2 p_X(x) + \dots \\ &= n! \cdot p_X(x) \end{aligned}$$

Let us derive the pgf for the Binomial distribution:

$$\mathcal{G}(s) = \sum_{x \geq 0} s^x \binom{n}{x} p^x q^{n-x} = (ps + q)^n \quad \text{by the Binomial Theorem}$$




Figure 12: Unit complex circle.

10.2 Stochastic Process

A stochastic process is an indexed collection of random variables, where the index set may be discrete or continuous. Some examples of stochastic processes include Brownian Motion (Section 10.3), Gaussian Processes (Section 4), and Markov chains (Section 6). Markov chains satisfy the **Markov property**, which states that the current state is sufficient for determining the future state. More formally:

$$p(X_t \leq x | X_{t-1}, X_{t-2}, \dots, X_0) = p(X_t \leq x | X_{t-1})$$

10.3 Brownian Motion

Brownian motion, B_t , is a stochastic process, in particular the continuous random walk that particles exhibit, for example in a gas. It satisfies the following properties:

- independent increments
- $\forall s < t : B_t - B_s \sim \mathcal{N}(0, t - s)$
- paths are continuous
- $B_0 = 0$

There are some interesting properties we can show. First, it is a discrete random walk taken to an infinite limit. Let X_n be zero-mean i.i.d. random variables with variance 1. The random walk starts at $S_0 = 0$ and proceeds with $S_n = S_0 + \sum_{i=1}^n X_i \quad \forall n \geq 1$. Suppose a continuous-time process:

$$B_t^n = \frac{S_{[nt]}}{\sqrt{n}} \quad \text{where } [nt] \text{ integer part of } nt$$

The Central Limit Theorem states that $Z = \frac{\hat{X}_n - \mu}{\sqrt{\sigma \cdot n}} = \mathcal{N}(0, 1)$ as $n \rightarrow \infty$.

$$\implies B_t^n = \frac{S_{[nt]}}{\sqrt{nt}} \frac{\sqrt{nt}}{\sqrt{n}} = \frac{S_{[nt]}}{\sqrt{nt}} \sqrt{t} = \sqrt{t} \cdot \mathcal{N}(0, 1) \quad \text{as } n \rightarrow \infty$$

Therefore, B_n^t converges to a scaled normal r.v., the variance of which will be:

$$\text{Var}[\sqrt{t}\mathcal{N}(0, 1)] = t \implies B_t^n \sim \mathcal{N}(0, t)$$

Furthermore,

$$B_t^n - B_s^n = \frac{S_{[nt]} - S_{[ns]}}{\sqrt{n}} = \frac{1}{\sqrt{n}} \sum_{i=[ns]+1}^{[nt]} X_i \stackrel{d}{=} \frac{S_{[nt]-[ns]}}{\sqrt{n(t-s)}} \sqrt{t-s} \sim \mathcal{N}(0, t-s)$$

10.3.1 Brownian Motion as a GP

Brownian motion can be interpreted as a zero-mean GP with covariance

$$\begin{aligned} \text{Cov}(B_s, B_t) &= \text{Cov}(B_s, B_s + B_t - B_s) = \underbrace{\text{Cov}(B_s, B_s)}_{\text{Var}(B_s)} + \underbrace{\text{Cov}(B_s, B_t - B_s)}_{0 \text{ (indep. increments)}} \\ &= s \end{aligned}$$

if $s < t$. Thus, in general, the kernel is $\kappa(t, s) = \min(s, t)$.

10.4 Stochastic Differential Equations

Modelling diffusive processes can be done using stochastic differential equations (SDEs). The typically look like this:

$$dX_t = \underbrace{\mu(X_t)}_{\text{drift}} dt + \underbrace{\sigma(X_t)}_{\text{diffusion}} dB_t$$

corresponding to (for $t > s$):

$$X_t = X_s + \int_s^t \mu(X_u) du + \int_s^t \sigma(X_u) dB_u$$

Some things to note:

- If there is no diffusion term then this reduces to an ODE.
- $\sigma(X_t) dB_t \sim \mathcal{N}(0, \sigma^2(X_t) dt)$

10.4.1 Itô vs. Stratonovich

The **Riemann-Stieltjes Integral** is the following:

$$\int_s^t f(u) dg(s) = \lim_{n \rightarrow \infty} \sum_{i=1}^n f(\underbrace{\tau_i}_{\in [t_i, t_{i+1}]}) (g(t_{i+1}) - g(t_i))$$

We wish to use this to solve the diffusion term. There are two ways of picking τ_i :

- **Itô:** $\tau_i = t_i$
- **Stratonovich:** $\tau_i = \frac{t_i + t_{i+1}}{2}$

11 Advanced Generative Models

11.1 Variational Autoencoders

The autoencoder is a method of compressing data through a bottleneck such that it can be reconstructed with minimal error. A simple example architecture is shown in Figure 13. The \mathbf{z} vector in the bottleneck layer is referred to as the *latent state*. The variational autoencoder (VAE) [Kingma and Welling, 2013] is an extension which combines the autoencoder and latent variable approach, resulting in a continuous latent state that has meaningful samples. This means that the latent space can be explored by decoding these samples. You might have guessed from the “variational” in the name that a lower bound is imminent. We will derive it now.

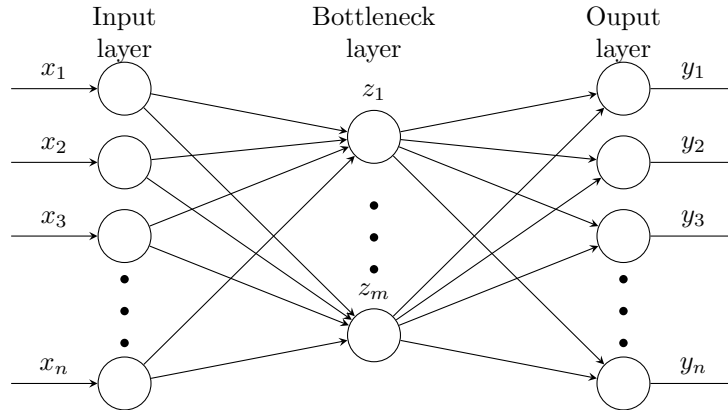


Figure 13: General autoencoder architecture.

The fundamental model assumption is that our dataset $\mathbf{X} = \{\mathbf{x}^{(i)}\}_{i=1}^N$ are samples of some random variable \mathbf{x} generated by a continuous random variable \mathbf{z} with prior $p(\mathbf{z})$. In our case, the *decoder* represents the conditional $p(\mathbf{x}|\mathbf{z})$. Moreover, the aforementioned pdfs are differentiable w.r.t. parameters and latent variables. Our *encoder* is also our variational posterior: $q(\mathbf{z}|\mathbf{x})$.

When are VAEs used?

- Probabilistically, the VAE is used for cases where the true posterior distribution is intractable.
- Practically, the VAE is used in an unsupervised learning setting where we wish to generate new data that mimics the same generative process as the real data.

Our log marginal likelihood is:

$$\begin{aligned}
 \log p(\mathbf{x}^{(i)}) &= \int q(\mathbf{z}|\mathbf{x}^{(i)}) \log \frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{q(\mathbf{z}|\mathbf{x}^{(i)})} dz + \int q(\mathbf{z}|\mathbf{x}^{(i)}) \log \frac{q(\mathbf{z}|\mathbf{x}^{(i)})}{p(\mathbf{z}|\mathbf{x}^{(i)})} dz \\
 &= \mathcal{L}(\phi; \mathbf{x}^{(i)}) + \mathcal{KL} \left[q(\mathbf{z}|\mathbf{x}^{(i)}) || p(\mathbf{z}|\mathbf{x}^{(i)}) \right] \\
 &\geq \mathcal{L}(\phi; \mathbf{x}^{(i)}) = \mathbb{E}_{q(\mathbf{z}|\mathbf{x}^{(i)})} \left[\log p(\mathbf{x}|\mathbf{z})p(\mathbf{z}) - \log q(\mathbf{z}|\mathbf{x}^{(i)}) \right] \\
 \implies \mathcal{L}(\phi; \mathbf{x}^{(i)}) &= \mathbb{E}_{q(\mathbf{z}|\mathbf{x}^{(i)})} \left[\log p(\mathbf{x}^{(i)}|\mathbf{z}) \right] - \mathcal{KL} \left[q(\mathbf{z}|\mathbf{x}^{(i)}) || p(\mathbf{z}) \right]
 \end{aligned}$$

Note how we have removed dependency on the true posterior by using this lower bound. Note also that our variational posterior is conditioned on the data. One way of optimising such a bound is

using Monte Carlo gradient sampling, where we seek the gradient w.r.t. the parameters:

$$\nabla_{\phi} \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x}^{(i)})} [f(\mathbf{z})] \approx \frac{1}{T} \sum_{t=1}^T \nabla_{q_{\phi}(\mathbf{z}^{(i,t)}|\mathbf{x}^{(i)})} f(\mathbf{z}^{(t)})$$

where $\mathbf{z}^{(t)} \sim q_{\phi}(\mathbf{z}|\mathbf{x}^{(i)})$ and $f(\mathbf{z}) = \log p(\mathbf{x}^{(i)}|\mathbf{z})$. See Section 6 for intuition of the above.

However, this method “exhibits high variance” [Kingma and Welling, 2013]. In order to find another method of inferring a posterior, we must introduce the **reparameterisation trick**.

11.1.1 Reparameterisation Trick

Let $\mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{x})$ (our posterior). We represent \mathbf{z} as a non-random variable $\mathbf{z} = g(\boldsymbol{\epsilon}, \mathbf{x})$ where $\boldsymbol{\epsilon} \sim p(\boldsymbol{\epsilon})$ and g is a vector-valued function with parameters ϕ . An expectation w.r.t. $q_{\phi}(\mathbf{z}|\mathbf{x})$ can be reparameterised as an expectation that is differentiable w.r.t. ϕ

As we saw in Section 1.2.3, under the conservation of probability, we have:

$$q_{\phi}(\mathbf{z}|\mathbf{x}) d\mathbf{z} = p(\boldsymbol{\epsilon}) d\boldsymbol{\epsilon}$$

Substituting both this equation and our function g into the Monte Carlo sampler yields:

$$\begin{aligned} \nabla_{\phi} \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x}^{(i)})} [f(\mathbf{z})] &= \nabla_{\phi} \int f(\mathbf{z}) q_{\phi}(\mathbf{z}|\mathbf{x}) d\mathbf{z} \\ &= \nabla_{\phi} \int f(g(\boldsymbol{\epsilon}, \mathbf{x})) p(\boldsymbol{\epsilon}) d\boldsymbol{\epsilon} \\ &= \nabla_{\phi} \mathbb{E}_{p(\boldsymbol{\epsilon})} [f(g(\boldsymbol{\epsilon}, \mathbf{x}))] \\ &= \nabla_{\phi} \mathbb{E}_{p(\boldsymbol{\epsilon})} [\log p(\mathbf{x}^{(i)}|\mathbf{z})] \\ &= \mathbb{E}_{p(\boldsymbol{\epsilon})} [\nabla_{\phi} \log p(\mathbf{x}^{(i)}|\mathbf{z})] \end{aligned}$$

where $\mathbf{z} = g(\boldsymbol{\epsilon}, \mathbf{x})$.

Key Concept

Note that in the last step the gradient can pass through the expectation since it is unrelated to the differential, $p(\boldsymbol{\epsilon})$. This enables us to propagate a gradient through the sampling procedure!

In order to define g for the VAE, we should specify the variational posterior, which we pick to be a multivariate Gaussian:

$$q_{\phi}(\mathbf{z}|\mathbf{x}^{(i)}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}^{(i)}, \boldsymbol{\sigma}^{2(i)} I)$$

where $\boldsymbol{\mu}^{(i)}$ and $\boldsymbol{\sigma}^{(i)}$ are output from our encoder neural network. A valid g , therefore, could be:

$$\mathbf{z}^{(i,t)} \sim q_{\phi}(\mathbf{z}^{(i,t)}|\mathbf{x}^{(i)}) = g(\mathbf{z}, \boldsymbol{\epsilon}) = \boldsymbol{\mu}^{(i)} + \boldsymbol{\sigma}^{(i)} \odot \boldsymbol{\epsilon}$$

Our final loss function, therefore, is:

$$\mathcal{L}(\phi; \mathbf{x}^{(i)}) \approx -\mathcal{KL} [q(\mathbf{z}|\mathbf{x}^{(i)})||p(\mathbf{z})] + \frac{1}{T} \sum_{t=1}^T \log p(\mathbf{x}^{(i)}|\mathbf{z}^{(i,t)})$$

where $\mathbf{z}^{(i,t)} = \boldsymbol{\mu}^{(i)} + \boldsymbol{\sigma}^{(i)} \odot \boldsymbol{\epsilon}$ and $\boldsymbol{\epsilon} \sim \mathcal{N}$ and $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$.

11.2 Generative Adversarial Networks

write

11.3 Normalising Flows

Recall that for two distributions X and Z , and function $f : X \rightarrow Z$,

$$\begin{aligned} dx \cdot p_X(x) &= dz \cdot p_Z(z) \\ \implies p_X(x) &= \det\left(\frac{df(x)}{dx}\right) \cdot p_Z(f(x)) \end{aligned}$$

where the determinant is to generalise the result to multidimensional distributions. Taking the logarithm:

$$\log p_X(x) = \log \det\left(\frac{df}{dx}\right) + \log p_Z(f(x))$$

finish

References

- Andreas Damianou and Neil Lawrence. Deep gaussian processes. In *Artificial Intelligence and Statistics*, pages 207–215, 2013.
- James Hensman, Nicolo Fusi, and Neil D Lawrence. Gaussian processes for big data. *arXiv preprint arXiv:1309.6835*, 2013.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- Michalis Titsias and Neil D Lawrence. Bayesian gaussian process latent variable model. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 844–851, 2010.

A Derivations

A.1 Double Integral of a Mean Function

$$\begin{aligned}
\int_b \int_a (ax + b)p(a)p(b) &= \int_b \int_a axp(a)p(b) + bp(a)p(b) \\
&= \int_b \int_a axp(a)p(b) + \int_b \int_a bp(a)p(b) \\
&= \int_a axp(a) \int_b p(b) + \int_a p(a) \int_b bp(b) \quad \text{but } \int_a p(a) = 1 \\
&= \int_a axp(a) + \int_b bp(b)
\end{aligned}$$

A.2 Expansion of Strange Gaussian Identity

$$\begin{aligned}
p(t_g | w_{I_g}, w_{J_g}) &\propto \exp\left(-\frac{1}{2}(w_{I_g} - w_{J_g} - t_g)^2\right) \\
&\propto \mathcal{N}\left(-\frac{1}{2} \begin{bmatrix} w_{I_g} - \mu_1 & w_{J_g} - \mu_2 \end{bmatrix} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} w_{I_g} - \mu_1 \\ w_{J_g} - \mu_2 \end{bmatrix}, 1\right) \\
&\propto \mathcal{N}\left(-\frac{1}{2} \left[(w_{I_g} - \mu_1) - (w_{J_g} - \mu_2) \quad -(w_{I_g} - \mu_1) + w_{J_g} - \mu_2 \right] \begin{bmatrix} w_{I_g} - \mu_1 \\ w_{J_g} - \mu_2 \end{bmatrix}\right) \\
&\propto \mathcal{N}\left(-\frac{1}{2} \left[w_{I_g} - w_{J_g} - \mu_1 + \mu_2 \quad -w_{I_g} + w_{J_g} + \mu_1 - \mu_2 \right] \begin{bmatrix} w_{I_g} - \mu_1 \\ w_{J_g} - \mu_2 \end{bmatrix}\right) \\
&\propto \mathcal{N}\left(-\frac{1}{2} \left[w_{I_g} - w_{J_g} - t_g \quad -w_{I_g} + w_{J_g} + t_g \right] \begin{bmatrix} w_{I_g} - \mu_1 \\ w_{J_g} - \mu_2 \end{bmatrix}\right) \\
&\propto \mathcal{N}\left(-\frac{1}{2} (w_{I_g} - w_{J_g} - t_g)(w_{I_g} - \mu_1) + (-w_{I_g} + w_{J_g} + t_g)(w_{J_g} - \mu_2)\right) \\
&\propto \mathcal{N}\left(-\frac{1}{2} (w_{I_g}^2 - w_{J_g} w_{I_g} - t_g w_{I_g} - \mu_1 (w_{I_g} - w_{J_g} - t_g) - w_{I_g} w_{J_g} + w_{J_g}^2 + t_g w_{J_g} + \mu_2 (w_{I_g} - w_{J_g} - t_g))\right) \\
&\propto \mathcal{N}\left(-\frac{1}{2} (w_{I_g}^2 - 2w_{J_g} w_{I_g} - t_g w_{I_g} + (\mu_2 - \mu_1)(w_{I_g} - w_{J_g} - t_g) + w_{J_g}^2 + t_g w_{J_g})\right) \\
&\propto \mathcal{N}\left(-\frac{1}{2} (w_{I_g}^2 - 2w_{J_g} w_{I_g} - t_g w_{I_g} + t_g (w_{J_g} - w_{I_g} + t_g) + w_{J_g}^2 + t_g w_{J_g})\right) \\
&\propto \mathcal{N}\left(-\frac{1}{2} (w_{I_g}^2 + w_{J_g}^2 + t_g^2 - 2w_{J_g} w_{I_g} - 2t_g w_{I_g} + 2t_g w_{J_g})\right)
\end{aligned}$$

A.3 Derivation of Biased Variance Estimator

$$\begin{aligned}
\mathbb{E}[s^2] &= \mathbb{E}\left[\frac{1}{N}\sum_{i=1}^N\left(X_i - \frac{1}{N}\sum_{j=1}^N X_j\right)^2\right] = \frac{1}{N}\sum_{i=1}^N \mathbb{E}\left[\underbrace{\left(X_i - \frac{1}{N}\sum_{j=1}^N X_j\right)^2}_{X_i^2 - \frac{2}{N}\sum_{j=1}^N X_i X_j + \frac{1}{N^2}\left(\sum_{j=1}^N X_j\right)^2}\right] \\
&= \frac{1}{N}\sum_{i=1}^N \left[\mathbb{E}[X_i^2] - \frac{2}{N}\mathbb{E}\left[\sum_{j=1}^N X_i X_j\right] + \frac{1}{N^2}\mathbb{E}\left[\underbrace{\left(\sum_{j=1}^N X_j\right)^2}_{\sum_{j=1}^N X_j^2 + 2\sum_{j=1}^N \sum_{k=j+1}^N X_j X_k}\right] \right] \\
&= \frac{1}{N}\sum_{i=1}^N \left[\mathbb{E}[X_i^2] - \frac{2}{N}\left(\sum_{j \neq i}^N \underbrace{\mathbb{E}[X_i X_j]}_{\mathbb{E}[X_i]\mathbb{E}[X_j]} + \mathbb{E}[X_i^2]\right) + \frac{1}{N^2}\left(\sum_{j=1}^N \mathbb{E}[X_j^2] + 2\sum_{j=1}^N \sum_{k=j+1}^N \mathbb{E}[X_j X_k]\right) \right] \\
&= \frac{1}{N}\sum_{i=1}^N \left[\left(1 - \frac{2}{N} + \frac{1}{N}\right)\mathbb{E}[X_i^2] - \frac{2}{N}(N-1)\mathbb{E}[X_i]^2 + \frac{2}{N^2}N(N-1)\mathbb{E}[X_i]^2 \right] \\
&= \frac{1}{N}\sum_{i=1}^N \left[\frac{N-1}{N}\left(\mathbb{E}[X_i^2] - \mathbb{E}[X_i]^2\right) \right] \\
&= \frac{1}{N}\sum_{i=1}^N \left[\frac{N-1}{N}\text{Var}[X_i] \right] = \frac{N-1}{N}\text{Var}[X_i]
\end{aligned}$$

A.4 Derivations of Means and Variances

We go through derivations for some common distributions. Recall $\text{Var}[X] = \mathbb{E}[(X - \mu)^2] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$

| DISTRIBUTION | DERIVATION |
|--|---|
| $X \sim \text{Bernoulli}(p)$ | $\mathbb{E}[X] = \sum xp_X(x) = 1p + 0(1-p) = p$ $\text{Var}[X] = \sum x^2 p_X(x) - p^2$ $= p - p^2 = p(1-p)$ |
| $X \sim \text{Binomial}(N, p)$ | <p>Note $k \binom{n}{k-1} = n \binom{n-1}{k-1}$</p> $\mathbb{E}[X] = \sum_{k=0}^n k \binom{n}{k} p^k q^{n-k}$ $= n \sum_{k=0}^n \binom{n-1}{k-1} p p^{k-1} q^{(n-1)-(k-1)}$ $= np$ $\text{Var}[X] = \sum k^2 \binom{n}{k} p^k q^{n-k} - (np)^2$ $= np \sum (1+k-1) \binom{n-1}{k-1} p^{k-1} q^{n-k} - (np)^2$ $= np \sum \binom{n-1}{k-1} p^{k-1} q^{n-k} + np \sum (k-1) \binom{n-1}{k-1} p^{k-1} q^{n-k} - (np)^2$ $= np + np(p(n-1)) \sum \binom{n-2}{k-2} p^{k-2} q^{n-k} - (np)^2$ $= np(np-p+1) - (np)^2 = np(1-p)$ |
| $X \sim \text{Geometric}(p)$ $p_X(x) = (1-p)^{x-1} p$ | $\mathbb{E}[X] = \sum_{x=1}^{\infty} x q^{x-1} p = p \sum_{x=1}^{\infty} \frac{d}{dq} q^x$ $= p \frac{d}{dq} (1-q)^{-1}$ (power series) $= \frac{p}{(1-q)^2} = \frac{1}{p}$ <p>Note $\frac{1}{1-q} = \sum_{x=1}^{\infty} q^x$</p> $\left(\frac{1}{1-q}\right)'' = \sum_{x=1}^{\infty} q^{-1} \sum x^2 q^{x-1} - q^{-1} \sum x q^{x-1}$ $\sum x^2 q^{x-1} = q \left(\frac{1}{1-q}\right)'' + \left(\frac{1}{1-q}\right)'$ $\text{Var}[X] = p \sum_{x=1}^{\infty} x^2 q^{x-1} - \frac{1}{p^2}$ $= \frac{2(1-p)+p-1}{p^2} = \frac{1-p}{p^2}$ |

B Kullback Leibler divergence

To minimise the KL divergence $\int q(x) \log \frac{q(x)}{p(x)} dx$, we add a Lagrange multiplier to normalise the $q(z)$ to 1. We will work with $\mathcal{KL}(q(x)||p(x))$

$$\frac{\delta}{\delta q(x)} \left[\int q(x) \log \frac{q(x)}{p(x)} dx + \lambda(1 - \int q(x) dx) \right] = \log \frac{q(x)}{p(x)} + 1 - \lambda$$

Since $q(x) \log \frac{q(x)}{p(x)} = q(x) \log q(x) - q(x) \log p(x)$

$$\implies \frac{\delta}{\delta q(x)} = q(x)/q(x) + \log q(x) - \log p(x) = 1 + \log \frac{q(x)}{p(x)}$$

$q(x) = \exp(\lambda - 1)p(x)$ and we set $\lambda = 1$ for normalisation, so $q(x) = p(x)$ at the minimum.

C Backup

We can now work out the bound from Hensman et al. [2013]. Let $\mathcal{L}_1 = \mathbb{E}_{p(\mathbf{f}_d|\mathbf{u}_d)}[\log p(\mathbf{y}_d|\mathbf{f}_d)]:$

$$\exp(\mathcal{L}_1) = \prod_{i=1}^N \mathcal{N}(y_d^{(i)} | \boldsymbol{\alpha}_d^{(i)}, \sigma^2) \exp\left(\frac{1}{2\sigma^2} \tilde{k}_{ii}\right)$$

where \tilde{k}_{ii} is the i th element of $\text{tr}(K_{NN} - Q_{nn})$. Note that this assumes that $p(\mathbf{y}_d|\mathbf{f}_d)$ factorises over the data. Next, by plugging in the exponentiated \mathcal{L}_1 into the likelihood, we get:

$$\log p(y_d|X) \geq \log \int \exp(\mathcal{L}_1) p(\mathbf{u}_d) d\mathbf{u}_d = \mathcal{L}_2$$

This is from Titsias, it may not be necessary. Plugging into the lower bound from Eq. 11:

$$\begin{aligned}
F_d(q) &\geq \int q(X) \left[\int \phi(\mathbf{u}_d) \left[\log \frac{\mathcal{N}(y_d | \boldsymbol{\alpha}_d, \sigma^2) p(\mathbf{u}_d)}{\phi(\mathbf{u}_d)} d\mathbf{u}_d - \frac{1}{2\sigma^2} \text{tr}(K_{NN} - Q_{nn}) \right] dX \\
&= \int \phi(\mathbf{u}_d) \left[\underbrace{\langle \log \mathcal{N}(y_d | \boldsymbol{\alpha}_d, \sigma^2) \rangle_{q(X)} + \log \frac{p(\mathbf{u}_d)}{\phi(\mathbf{u}_d)}}_{\text{now reverse Jensen's inequality: } \mathbb{E}(\cdot) \rightarrow \log \mathbb{E} \exp(\cdot)} \right] d\mathbf{u}_d - \frac{1}{2\sigma^2} \langle \text{tr}(K_{NN} - Q_{nn}) \rangle_{q(X)} \\
&= \log \int \phi(\mathbf{u}_d) \exp(\langle \log \mathcal{N}(y_d | \boldsymbol{\alpha}_d, \sigma^2) \rangle_{q(X)}) \frac{p(\mathbf{u}_d)}{\phi(\mathbf{u}_d)} d\mathbf{u}_d - \frac{1}{2\sigma^2} \langle \text{tr}(K_{NN} - Q_{nn}) \rangle_{q(X)} \\
&= \log \int \exp(\langle \log \mathcal{N}(y_d | \boldsymbol{\alpha}_d, \sigma^2) \rangle_{q(X)}) p(\mathbf{u}_d) d\mathbf{u}_d - \frac{1}{2\sigma^2} \langle \text{tr}(K_{NN} - Q_{nn}) \rangle_{q(X)}
\end{aligned}$$