

# Unsupervised Learning\*

Zoubin Ghahramani†

Gatsby Computational Neuroscience Unit  
University College London, UK  
zoubin@gatsby.ucl.ac.uk  
<http://www.gatsby.ucl.ac.uk/~zoubin>

September 16, 2004

## Abstract

We give a tutorial and overview of the field of unsupervised learning from the perspective of statistical modelling. Unsupervised learning can be motivated from information theoretic and Bayesian principles. We briefly review basic models in unsupervised learning, including factor analysis, PCA, mixtures of Gaussians, ICA, hidden Markov models, state-space models, and many variants and extensions. We derive the EM algorithm and give an overview of fundamental concepts in graphical models, and inference algorithms on graphs. This is followed by a quick tour of approximate Bayesian inference, including Markov chain Monte Carlo (MCMC), Laplace approximation, BIC, variational approximations, and expectation propagation (EP). The aim of this chapter is to provide a high-level view of the field. Along the way, many state-of-the-art ideas and future directions are also reviewed.

## Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	What is unsupervised learning? . . . . .	3
1.2	Machine learning, statistics, and information theory . . . . .	4
1.3	Bayes rule . . . . .	4
<b>2</b>	<b>Latent variable models</b>	<b>6</b>
2.1	Factor analysis . . . . .	6
2.2	Principal components analysis (PCA) . . . . .	7
2.3	Independent components analysis (ICA) . . . . .	7
2.4	Mixture of Gaussians . . . . .	7
2.5	K-means . . . . .	8
<b>3</b>	<b>The EM algorithm</b>	<b>8</b>
<b>4</b>	<b>Modelling time series and other structured data</b>	<b>9</b>
4.1	State-space models (SSMs) . . . . .	10
4.2	Hidden Markov models (HMMs) . . . . .	10
4.3	Modelling other structured data . . . . .	11
<b>5</b>	<b>Nonlinear, Factorial, and Hierarchical Models</b>	<b>11</b>
<b>6</b>	<b>Intractability</b>	<b>12</b>

---

\*This chapter will appear in Bousquet, O., Raetsch, G. and von Luxburg, U. (eds) *Advanced Lectures on Machine Learning* LNAI 3176. ©Springer-Verlag.

†The author is also at the Center for Automated Learning and Discovery, Carnegie Mellon University, USA.

<b>7</b>	<b>Graphical models</b>	<b>13</b>
7.1	Undirected graphs . . . . .	13
7.2	Factor graphs . . . . .	14
7.3	Directed graphs . . . . .	14
7.4	Expressive power . . . . .	15
<b>8</b>	<b>Exact inference in graphs</b>	<b>15</b>
8.1	Elimination . . . . .	16
8.2	Belief propagation . . . . .	16
8.3	Factor graph propagation . . . . .	18
8.4	Junction tree algorithm . . . . .	18
8.5	Cutest conditioning . . . . .	19
<b>9</b>	<b>Learning in graphical models</b>	<b>19</b>
9.1	Learning graph parameters . . . . .	20
9.1.1	The complete data case. . . . .	20
9.1.2	The incomplete data case. . . . .	20
9.2	Learning graph structure . . . . .	21
9.2.1	Scoring metrics. . . . .	21
9.2.2	Search algorithms. . . . .	21
<b>10</b>	<b>Bayesian model comparison and Occam’s Razor</b>	<b>21</b>
<b>11</b>	<b>Approximating posteriors and marginal likelihoods</b>	<b>22</b>
11.1	Laplace approximation . . . . .	23
11.2	The Bayesian information criterion (BIC) . . . . .	23
11.3	Markov chain Monte Carlo (MCMC) . . . . .	24
11.4	Variational approximations . . . . .	24
11.5	Expectation propagation (EP) . . . . .	25
<b>12</b>	<b>Conclusion</b>	<b>27</b>

# 1 Introduction

Machine learning is the field of research devoted to the formal study of learning systems. This is a highly interdisciplinary field which borrows and builds upon ideas from statistics, computer science, engineering, cognitive science, optimisation theory and many other disciplines of science and mathematics. The purpose of this chapter is to introduce in a fairly concise manner the key ideas underlying the sub-field of machine learning known as *unsupervised learning*. This introduction is necessarily incomplete given the enormous range of topics under the rubric of unsupervised learning. The hope is that interested readers can delve more deeply into the many topics covered here by following some of the cited references. The chapter starts at a highly tutorial level but will touch upon state-of-the-art research in later sections. It is assumed that the reader is familiar with elementary linear algebra, probability theory, and calculus, but not much else.

## 1.1 What is unsupervised learning?

Consider a machine (or living organism) which receives some sequence of inputs  $x_1, x_2, x_3, \dots$ , where  $x_t$  is the sensory input at time  $t$ . This input, which we will often call the *data*, could correspond to an image on the retina, the pixels in a camera, or a sound waveform. It could also correspond to less obviously sensory data, for example the words in a news story, or the list of items in a supermarket shopping basket.

One can distinguish between four different kinds of machine learning. In *supervised learning* the machine<sup>1</sup> is also given a sequence of desired outputs  $y_1, y_2, \dots$ , and the goal of the machine is to learn to produce the correct output given a new input. This output could be a class label (in classification) or a real number (in regression).

In *reinforcement learning* the machine interacts with its environment by producing actions  $a_1, a_2, \dots$ . These actions affect the state of the environment, which in turn results in the machine receiving some scalar rewards (or punishments)  $r_1, r_2, \dots$ . The goal of the machine is to learn to act in a way that maximises the future rewards it receives (or minimises the punishments) over its lifetime. Reinforcement learning is closely related to the fields of decision theory (in statistics and management science), and control theory (in engineering). The fundamental problems studied in these fields are often formally equivalent, and the solutions are the same, although different aspects of problem and solution are usually emphasised.

A third kind of machine learning is closely related to *game theory* and generalises reinforcement learning. Here again the machine gets inputs, produces actions, and receives rewards. However, the environment the machine interacts with is not some static world, but rather it can contain other machines which can also sense, act, receive rewards, and learn. Thus the goal of the machine is to act so as to maximise rewards in light of the other machines' current and future actions. Although there is a great deal of work in game theory for simple systems, the dynamic case with multiple adapting machines remains an active and challenging area of research.

Finally, in *unsupervised learning* the machine simply receives inputs  $x_1, x_2, \dots$ , but obtains neither supervised target outputs, nor rewards from its environment. It may seem somewhat mysterious to imagine what the machine could possibly learn given that it doesn't get any feedback from its environment. However, it is possible to develop a formal framework for unsupervised learning based on the notion that the machine's goal is to build representations of the input that can be used for decision making, predicting future inputs, efficiently communicating the inputs to another machine, etc. In a sense, unsupervised learning can be thought of as finding patterns in the data above and beyond what would be considered pure unstructured noise. Two very simple classic examples of unsupervised learning are clustering and dimensionality reduction. We discuss these in Section 2. The remainder of this chapter focuses on unsupervised learning, although many of the concepts discussed can be applied to supervised learning as well. But first, let us consider how unsupervised learning relates to statistics and information theory.

---

<sup>1</sup>Henceforth, for succinctness I'll use the term machine to refer both to machines and living organisms. Some people prefer to call this a system or agent. The same mathematical theory of learning applies regardless of what we choose to call the learner, whether it is artificial or biological.

## 1.2 Machine learning, statistics, and information theory

Almost all work in unsupervised learning can be viewed in terms of learning a probabilistic model of the data. Even when the machine is given no supervision or reward, it may make sense for the machine to estimate a model that represents the probability distribution for a new input  $x_t$  given previous inputs  $x_1, \dots, x_{t-1}$  (consider the obviously useful examples of stock prices, or the weather). That is, the learner models  $P(x_t|x_1, \dots, x_{t-1})$ . In simpler cases where the order in which the inputs arrive is irrelevant or unknown, the machine can build a model of the data which assumes that the data points  $x_1, x_2, \dots$  are independently and identically drawn from some distribution  $P(x)$ <sup>2</sup>.

Such a model can be used for *outlier detection* or *monitoring*. Let  $x$  represent patterns of sensor readings from a nuclear power plant and assume that  $P(x)$  is learned from data collected from a normally functioning plant. This model can be used to evaluate the probability of a new sensor reading; if this probability is abnormally low, then either the model is poor or the plant is behaving abnormally, in which case one may want to shut it down.

A probabilistic model can also be used for *classification*. Assume  $P_1(x)$  is a model of the attributes of credit card holders who paid on time, and  $P_2(x)$  is a model learned from credit card holders who defaulted on their payments. By evaluating the relative probabilities  $P_1(x')$  and  $P_2(x')$  on a new applicant  $x'$ , the machine can decide to classify her into one of these two categories.

With a probabilistic model one can also achieve efficient *communication* and *data compression*. Imagine that we want to transmit, over a digital communication line, symbols  $x$  randomly drawn from  $P(x)$ . For example,  $x$  may be letters of the alphabet, or images, and the communication line may be the internet. Intuitively, we should encode our data so that symbols which occur more frequently have code words with fewer bits in them, otherwise we are wasting bandwidth. Shannon's source coding theorem quantifies this by telling us that the optimal number of bits to use to encode a symbol with probability  $P(x)$  is  $-\log_2 P(x)$ . Using these number of bits for each symbol, the expected coding cost is the entropy of the distribution  $P$ .

$$H(P) \stackrel{\text{def}}{=} - \sum_x P(x) \log_2 P(x) \quad (1)$$

In general, the true distribution of the data is unknown, but we can learn a model of this distribution. Let's call this model  $Q(x)$ . The optimal code *with respect to this model* would use  $-\log_2 Q(x)$  bits for each symbol  $x$ . The expected coding cost, taking expectations with respect to the true distribution, is

$$- \sum_x P(x) \log_2 Q(x) \quad (2)$$

The difference between these two coding costs is called the Kullback-Leibler (KL) divergence

$$\text{KL}(P\|Q) \stackrel{\text{def}}{=} \sum_x P(x) \log \frac{P(x)}{Q(x)} \quad (3)$$

The KL divergence is non-negative and zero if and only if  $P=Q$ . It measures the coding inefficiency in bits from using a model  $Q$  to compress data when the true data distribution is  $P$ . *Therefore, the better our model of the data, the more efficiently we can compress and communicate new data.* This is an important link between machine learning, statistics, and information theory. An excellent text which elaborates on these relationships and many of the topics in this chapter is [48].

## 1.3 Bayes rule

Bayes rule,

$$P(y|x) = \frac{P(x|y)P(y)}{P(x)} \quad (4)$$

which follows from the equality  $P(x, y) = P(x)P(y|x) = P(y)P(x|y)$ , can be used to motivate a coherent statistical framework for machine learning. The basic idea is the following. Imagine we wish to design a

---

<sup>2</sup>We will use both  $P$  and  $p$  to denote probability distributions and probability densities. The meaning should be clear depending on whether the argument is discrete or continuous.

machine which has beliefs about the world, and updates these beliefs on the basis of observed data. The machine must somehow represent the strengths of its beliefs numerically. It has been shown that if you accept certain axioms of coherent inference, known as the *Cox axioms*, then a remarkable result follows [36]: If the machine is to represent the strength of its beliefs by real numbers, then the only reasonable and coherent way of manipulating these beliefs is to have them satisfy the rules of probability, such as Bayes rule. Therefore,  $P(X = x)$  can be used not only to represent the frequency with which the variable  $X$  takes on the value  $x$  (as in so-called frequentist statistics) but it can also be used to represent the degree of belief that  $X = x$ . Similarly,  $P(X = x|Y = y)$  can be used to represent the degree of belief that  $X = x$  given that one knows  $Y = y$ .<sup>3</sup>

From Bayes rule we derive the following simple framework for machine learning. Assume a universe of models  $\Omega$ ; let  $\Omega = \{1, \dots, M\}$  although it need not be finite or even countable. The machine starts with some prior beliefs over models  $m \in \Omega$  (we will see many examples of models later), such that  $\sum_{m=1}^M P(m) = 1$ . A model is simply some probability distribution over data points, i.e.  $P(x|m)$ . For simplicity, let us further assume that in all the models the data is taken to be independently and identically distributed (iid). After observing a data set  $\mathcal{D} = \{x_1, \dots, x_N\}$ , the beliefs over models is given by:

$$P(m|\mathcal{D}) = \frac{P(m)P(\mathcal{D}|m)}{P(\mathcal{D})} \propto P(m) \prod_{n=1}^N P(x_n|m) \quad (5)$$

which we read as the *posterior over models* is the *prior* multiplied by the *likelihood*, normalised.

The *predictive distribution* over new data, which would be used to encode new data efficiently, is

$$P(x|\mathcal{D}) = \sum_{m=1}^M P(x|m)P(m|\mathcal{D}) \quad (6)$$

Again this follows from the rules of probability theory, and the fact that the models are assumed to produce iid data.

Often models are defined by writing down a parametric probability distribution (again, we'll see many examples below). Thus, the model  $m$  might have parameters  $\theta$ , which are assumed to be unknown (this could in general be a vector of parameters). To be a well-defined model from the perspective of Bayesian learning, one has to define a prior over these model parameters  $P(\theta|m)$  which naturally has to satisfy the following equality

$$P(x|m) = \int P(x|\theta, m)P(\theta|m)d\theta \quad (7)$$

Given the model  $m$  it is also possible to infer the posterior over the parameters of the model, i.e.  $P(\theta|\mathcal{D}, m)$ , and to compute the predictive distribution,  $P(x|\mathcal{D}, m)$ . These quantities are derived in exact analogy to equations (5) and (6), except that instead of summing over possible models, we integrate over parameters of a particular model. All the key quantities in Bayesian machine learning follow directly from the basic rules of probability theory.

Certain approximate forms of Bayesian learning are worth mentioning. Let's focus on a particular model  $m$  with parameters  $\theta$ , and an observed data set  $\mathcal{D}$ . The predictive distribution averages over all possible parameters weighted by the posterior

$$P(x|\mathcal{D}, m) = \int P(x|\theta)P(\theta|\mathcal{D}, m)d\theta. \quad (8)$$

In certain cases, it may be cumbersome to represent the entire posterior distribution over parameters, so instead we will choose to find a *point-estimate* of the parameters  $\hat{\theta}$ . A natural choice is to pick the most

---

<sup>3</sup>Another way to motivate the use of the rules of probability to encode degrees of belief comes from game-theoretic arguments in the form of the *Dutch Book Theorem*. This theorem states that if you are willing to accept bets with odds based on your degrees of beliefs, then unless your beliefs are coherent in the sense that they satisfy the rules of probability theory, there exists a set of simultaneous bets (called a "Dutch Book") which you will accept and which is guaranteed to lose you money, no matter what the outcome. The only way to ensure that Dutch Books don't exist against you, is to have degrees of belief that satisfy Bayes rule and the other rules of probability theory.

probable parameter value given the data, which is known as the *maximum a posteriori* or MAP parameter estimate

$$\hat{\theta}_{\text{MAP}} = \arg \max_{\theta} P(\theta | \mathcal{D}, m) = \arg \max_{\theta} \left[ \log P(\theta | m) + \sum_n \log P(x_n | \theta, m) \right] \quad (9)$$

Another natural choice is the *maximum likelihood* or ML parameter estimate

$$\hat{\theta}_{\text{ML}} = \arg \max_{\theta} P(\mathcal{D} | \theta, m) = \arg \max_{\theta} \sum_n \log P(x_n | \theta, m) \quad (10)$$

Many learning algorithms can be seen as finding ML parameter estimates. The ML parameter estimate is also acceptable from a frequentist statistical modelling perspective since it does not require deciding on a prior over parameters. However, ML estimation does not protect against overfitting—more complex models will generally have higher maxima of the likelihood. In order to avoid problems with overfitting, frequentist procedures often maximise a *penalised* or *regularised* log likelihood (e.g. [26]). If the penalty or regularisation term is interpreted as a log prior, then maximising penalised likelihood appears identical to maximising a posterior. However, there are subtle issues that make a Bayesian MAP procedure and maximum penalised likelihood different [28]. One difference is that the MAP estimate is not invariant to reparameterisation, while the maximum of the penalised likelihood is invariant. The penalised likelihood is a function, not a density, and therefore does not increase or decrease depending on the Jacobian of the reparameterisation.

## 2 Latent variable models

The framework described above can be applied to a wide range of models. No single model is appropriate for all data sets. The art in machine learning is to develop models which are appropriate for the data set being analysed, and which have certain desired properties. For example, for high dimensional data sets it might be necessary to use models that perform dimensionality reduction. Of course, ultimately, the machine should be able to decide on the appropriate model without any human intervention, but to achieve this in full generality requires significant advances in artificial intelligence.

In this section, we will consider probabilistic models that are defined in terms of some latent or hidden variables. These models can be used to do dimensionality reduction and clustering, the two cornerstones of unsupervised learning.

### 2.1 Factor analysis

Let the data set  $\mathcal{D}$  consist of  $D$ -dimensional real valued vectors,  $\mathcal{D} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ . In factor analysis, the data is assumed to be generated from the following model

$$\mathbf{y} = \Lambda \mathbf{x} + \epsilon \quad (11)$$

where  $\mathbf{x}$  is a  $K$ -dimensional zero-mean unit-variance multivariate Gaussian vector with elements corresponding to hidden (or latent) factors,  $\Lambda$  is a  $D \times K$  matrix of parameters, known as the factor loading matrix, and  $\epsilon$  is a  $D$ -dimensional zero-mean multivariate Gaussian noise vector with diagonal covariance matrix  $\Psi$ . Defining the parameters of the model to be  $\theta = (\Psi, \Lambda)$ , by integrating out the factors, one can readily derive that

$$p(\mathbf{y} | \theta) = \int p(\mathbf{x} | \theta) p(\mathbf{y} | \mathbf{x}, \theta) d\mathbf{x} = \mathcal{N}(0, \Lambda \Lambda^\top + \Psi) \quad (12)$$

where  $\mathcal{N}(\mu, \Sigma)$  refers to a multivariate Gaussian density with mean  $\mu$  and covariance matrix  $\Sigma$ . For more details refer to [68].

Factor analysis is an interesting model for several reasons. If the data is very high dimensional ( $D$  is large) then even a simple model like the full-covariance multivariate Gaussian will have too many parameters to reliably estimate or infer from the data. By choosing  $K < D$ , factor analysis makes it possible to model a Gaussian density for high dimensional data without requiring  $\mathcal{O}(D^2)$  parameters. Moreover, given a new data point, one can compute the posterior over the hidden factors,  $p(\mathbf{x} | \mathbf{y}, \theta)$ ; since  $\mathbf{x}$  is lower dimensional than  $\mathbf{y}$  this provides a low-dimensional representation of the data (for example, one could pick the mean of  $p(\mathbf{x} | \mathbf{y}, \theta)$  as the representation for  $\mathbf{y}$ ).

## 2.2 Principal components analysis (PCA)

Principal components analysis (PCA) is an important limiting case of factor analysis (FA). One can derive PCA by making two modifications to FA. First, the noise is assumed to be isotropic, in other words each element of  $\epsilon$  has equal variance:  $\Psi = \sigma^2 I$ , where  $I$  is a  $D \times D$  identity matrix. This model is called *probabilistic PCA* [67, 78]. Second, if we take the limit of  $\sigma \rightarrow 0$  in probabilistic PCA, we obtain standard PCA (which also goes by the names Karhunen-Loève expansion, and singular value decomposition; SVD). Given a data set with covariance matrix  $\Sigma$ , for maximum likelihood factor analysis the goal is to find parameters  $\Lambda$ , and  $\Psi$  for which the model  $\Lambda\Lambda^\top + \Psi$  has highest likelihood. In PCA, the goal is to find  $\Lambda$  so that the likelihood is highest for  $\Lambda\Lambda^\top$ . Note that this matrix is singular unless  $K = D$ , so the standard PCA model is not a sensible model. However, taking the limiting case, and further constraining the columns of  $\Lambda$  to be orthogonal, it can be derived that the principal components correspond to the  $K$  eigenvectors with largest eigenvalue of  $\Sigma$ . PCA is thus attractive because the solution can be found immediately after eigendecomposition of the covariance. Taking the limit  $\sigma \rightarrow 0$  of  $p(\mathbf{x}|\mathbf{y}, \Lambda, \sigma)$  we find that it is a delta-function at  $\mathbf{x} = \Lambda^\top \mathbf{y}$ , which is the projection of  $\mathbf{y}$  onto the principal components.

## 2.3 Independent components analysis (ICA)

Independent components analysis (ICA) extends factor analysis to the case where the factors are non-Gaussian. This is an interesting extension because many real-world data sets have structure which can be modelled as linear combinations of sparse sources. This includes auditory data, images, biological signals such as EEG, etc. *Sparsity* simply corresponds to the assumption that the factors have distributions with higher kurtosis than the Gaussian. For example,  $p(x) = \frac{\lambda}{2} \exp\{-\lambda|x|\}$  has a higher peak at zero and heavier tails than a Gaussian with corresponding mean and variance, so it would be considered sparse (strictly speaking, one would like a distribution which had non-zero probability mass at 0 to get true sparsity).

Models like PCA, FA and ICA can all be implemented using neural networks (multilayer perceptrons) trained using various cost functions. It is not clear what advantage this implementation/interpretation has from a machine learning perspective, although it provides interesting ties to biological information processing.

Rather than ML estimation, one can also do Bayesian inference for the parameters of probabilistic PCA, FA, and ICA.

## 2.4 Mixture of Gaussians

The densities modelled by PCA, FA and ICA are all relatively simple in that they are unimodal and have fairly restricted parametric forms (Gaussian, in the case of PCA and FA). To model data with more complex structure such as clusters, it is very useful to consider mixture models. Although it is straightforward to consider mixtures of arbitrary densities, we will focus on Gaussians as a common special case. The density of each data point in a mixture model can be written:

$$p(\mathbf{y}|\theta) = \sum_{k=1}^K \pi_k p(\mathbf{y}|\theta_k) \quad (13)$$

where each of the  $K$  components of the mixture is, for example, a Gaussian with differing means and covariances  $\theta_k = (\mu_k, \Sigma_k)$  and  $\pi_k$  is the mixing proportion for component  $k$ , such that  $\sum_{k=1}^K \pi_k = 1$  and  $\pi_k > 0, \forall k$ .

A different way to think about mixture models is to consider them as latent variable models, where associated with each data point is a  $K$ -ary discrete latent (i.e. hidden) variable  $s$  which has the interpretation that  $s = k$  if the data point was generated by component  $k$ . This can be written

$$p(\mathbf{y}|\theta) = \sum_{k=1}^K P(s = k|\pi) p(\mathbf{y}|s = k, \theta) \quad (14)$$

where  $P(s = k|\pi) = \pi_k$  is the prior for the latent variable taking on value  $k$ , and  $p(\mathbf{y}|s = k, \theta) = p(\mathbf{y}|\theta_k)$  is the density under component  $k$ , recovering Equation (13).

## 2.5 K-means

The mixture of Gaussians model is closely related to an unsupervised clustering algorithm known as  $k$ -means as follows: Consider the special case where all the Gaussians have common covariance matrix proportional to the identity matrix:  $\Sigma_k = \sigma^2 I, \forall k$ , and let  $\pi_k = 1/K, \forall k$ . We can estimate the maximum likelihood parameters of this model using the iterative algorithm which we are about to describe, known as EM. The resulting algorithm, as we take the limit  $\sigma^2 \rightarrow 0$ , becomes exactly the  $k$ -means algorithm. Clearly the model underlying  $k$ -means has only singular Gaussians and is therefore an unreasonable model of the data; however,  $k$ -means is usually justified from the point of view of clustering to minimise a distortion measure, rather than fitting a probabilistic models.

## 3 The EM algorithm

The EM algorithm is an algorithm for estimating ML parameters of a model with latent variables. Consider a model with observed variables  $\mathbf{y}$ , hidden/latent variables  $\mathbf{x}$ , and parameters  $\theta$ . We can lower bound the log likelihood for any data point as follows

$$L(\theta) = \log p(\mathbf{y}|\theta) = \log \int p(\mathbf{x}, \mathbf{y}|\theta) d\mathbf{x} \quad (15)$$

$$= \log \int q(\mathbf{x}) \frac{p(\mathbf{x}, \mathbf{y}|\theta)}{q(\mathbf{x})} d\mathbf{x} \quad (16)$$

$$\geq \int q(\mathbf{x}) \log \frac{p(\mathbf{x}, \mathbf{y}|\theta)}{q(\mathbf{x})} d\mathbf{x} \stackrel{\text{def}}{=} F(q, \theta) \quad (17)$$

where  $q(\mathbf{x})$  is some arbitrary density over the hidden variables, and the lower bound holds due to the concavity of the log function (this inequality is known as Jensen's inequality). The lower bound  $F$  is a functional of both the density  $q(\mathbf{x})$  and the model parameters  $\theta$ . For a data set of  $N$  data points  $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)}$ , this lower bound is formed for the log likelihood term corresponding to each data point, thus there is a separate density  $q^{(n)}(\mathbf{x})$  for each point and  $F(q, \theta) = \sum_n F^{(n)}(q^{(n)}, \theta)$ .

The basic idea of the Expectation-Maximisation (EM) algorithm is to iterate between optimising this lower bound as a function of  $q$  and as a function of  $\theta$ . We can prove that this will never decrease the log likelihood. After initialising the parameters somehow, the  $k^{\text{th}}$  iteration of the algorithm consists of the following two steps:

**E step:** *optimise  $F$  with respect to the distribution  $q$  while holding the parameters fixed*

$$q_k(\mathbf{x}) = \arg \max_{q(\mathbf{x})} \int q(\mathbf{x}) \log \frac{p(\mathbf{x}, \mathbf{y}|\theta_{k-1})}{q(\mathbf{x})} d\mathbf{x} \quad (18)$$

$$q_k(\mathbf{x}) = p(\mathbf{x}|\mathbf{y}, \theta_{k-1}) \quad (19)$$

**M step:** *optimise  $F$  with respect to the parameters  $\theta$  while holding the distribution over hidden variables fixed*

$$\theta_k = \arg \max_{\theta} \int q_k(\mathbf{x}) \log \frac{p(\mathbf{x}, \mathbf{y}|\theta)}{q_k(\mathbf{x})} d\mathbf{x} \quad (20)$$

$$\theta_k = \arg \max_{\theta} \int q_k(\mathbf{x}) \log p(\mathbf{x}, \mathbf{y}|\theta) d\mathbf{x} \quad (21)$$

Let us be absolutely clear what happens for a data set of  $N$  data points: In the E step, for each data point, the distribution over the hidden variables is set to the posterior for that data point  $q_k^{(n)}(\mathbf{x}) = p(\mathbf{x}|\mathbf{y}^{(n)}, \theta_{k-1}), \forall n$ . In the M step the single set of parameters is re-estimated by maximising the sum of the expected log likelihoods:  $\theta_k = \arg \max_{\theta} \sum_n \int q_k^{(n)}(\mathbf{x}) \log p(\mathbf{x}, \mathbf{y}^{(n)}|\theta) d\mathbf{x}$ .

Two things are still unclear: how does (19) follow from (18), and how is this algorithm guaranteed to increase the likelihood? The optimisation in (18) can be written as follows since  $p(\mathbf{x}, \mathbf{y}|\theta_{k-1}) = p(\mathbf{y}|\theta_{k-1})p(\mathbf{x}|\mathbf{y}, \theta_{k-1})$ :

$$q_k(\mathbf{x}) = \arg \max_{q(\mathbf{x})} \left[ \log p(\mathbf{y}|\theta_{k-1}) + \int q(\mathbf{x}) \log \frac{p(\mathbf{x}|\mathbf{y}, \theta_{k-1})}{q(\mathbf{x})} d\mathbf{x} \right] \quad (22)$$



Now, the first term is a constant w.r.t.  $q(\mathbf{x})$  and the second term is the negative of the Kullback-Leibler divergence

$$\text{KL}(q(\mathbf{x})\|p(\mathbf{x}|\mathbf{y}, \theta_{k-1})) = \int q(\mathbf{x}) \log \frac{q(\mathbf{x})}{p(\mathbf{x}|\mathbf{y}, \theta_{k-1})} d\mathbf{x} \quad (23)$$

which we have seen in Equation (3) in its discrete form. This is minimised at  $q(\mathbf{x}) = p(\mathbf{x}|\mathbf{y}, \theta_{k-1})$ , where the KL divergence is zero. Intuitively, the interpretation of this is that in the E step of EM, the goal is to find the posterior distribution of the hidden variables given the observed variables and the current settings of the parameters. We also see that since the KL divergence is zero, at the end of the E step,  $F(q_k, \theta_{k-1}) = L(\theta_{k-1})$ .

In the M step,  $F$  is increased with respect to  $\theta$ . Therefore,  $F(q_k, \theta_k) \geq F(q_k, \theta_{k-1})$ . Moreover,  $L(\theta_k) = F(q_{k+1}, \theta_k) \geq F(q_k, \theta_k)$  after the next E step. We can put these steps together to establish that  $L(\theta_k) \geq L(\theta_{k-1})$ , establishing that the algorithm is guaranteed to increase the likelihood or keep it fixed (at convergence).

The EM algorithm can be applied to all the latent variable models described above, i.e. FA, probabilistic PCA, mixture models, and ICA. In the case of mixture models, the hidden variable is the discrete assignment  $s$  of data points to clusters; consequently the integrals turn into sums where appropriate. EM has wide applicability to latent variable models, although it is not always the fastest optimisation method [70]. Moreover, we should note that the likelihood often has many local optima and EM will converge to some local optimum which may not be the global one.

EM can also be used to estimate MAP parameters of a model, and as we will see in Section 11.4 there is a Bayesian generalization of EM as well.

## 4 Modelling time series and other structured data

So far we have assumed that the data is *unstructured*, that is, the observations are assumed to be independent and identically distributed. This assumption is unreasonable for many data sets in which the observations arrive in a sequence and subsequent observations are correlated. Sequential data can occur in time series modelling (as in financial data or the weather) and also in situations where the sequential nature of the data is not necessarily tied to time (as in protein data which consist of sequences of amino acids).

As the most basic level, time series modelling consists of building a probabilistic model of the present observation given all past observations  $p(\mathbf{y}_t|\mathbf{y}_{t-1}, \mathbf{y}_{t-2} \dots)$ . Because the history of observations grows arbitrarily large it is necessary to limit the complexity of such a model. There are essentially two ways of doing this.

The first approach is to limit the window of past observations. Thus one can simply model  $p(\mathbf{y}_t|\mathbf{y}_{t-1})$  and assume that this relation holds for all  $t$ . This is known as a first-order Markov model. A second-order Markov model would be  $p(\mathbf{y}_t|\mathbf{y}_{t-1}, \mathbf{y}_{t-2})$ , and so on. Such Markov models have two limitations: First, the influence of past observations on present observations vanishes outside this window, which can be unrealistic. Second, it may be unnatural and unwieldy to model directly the relationship between raw observations at one time step and raw observations at a subsequent time step. For example, if the observations are noisy images, it would make more sense to de-noise them, extract some description of the objects, motions, illuminations, and then try to predict from that.

The second approach is to make use of latent or hidden variables. Instead of modelling directly the effect of  $\mathbf{y}_{t-1}$  on  $\mathbf{y}_t$ , we assume that the observations were generated from some underlying hidden variable  $\mathbf{x}_t$  which captures the dynamics of the system. For example,  $\mathbf{y}$  might be noisy sonar readings of objects in a room, while  $\mathbf{x}$  might be the actual locations and sizes of these objects. We usually call this hidden variable  $\mathbf{x}$  the *state variable* since it is meant to capture all the aspects of the system relevant to predicting the future dynamical behaviour of the system.

In order to understand more complex time series models, it is essential that one be familiar with state-space models (SSMs) and hidden Markov models (HMMs). These two classes of models have played a historically important role in control engineering, visual tracking, speech recognition, protein sequence modelling, and error decoding. They form the simplest building blocks from which other richer time-series models can be developed, in a manner completely analogous to the role that FA and mixture models play in building more complex models for iid data.

## 4.1 State-space models (SSMs)

In a state-space model, the sequence of observed data  $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \dots$  is assumed to have been generated from some sequence of hidden state variables  $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots$ . Letting  $\mathbf{x}_{1:T}$  denote the sequence  $\mathbf{x}_1, \dots, \mathbf{x}_T$ , the basic assumption in an SSM is that the joint probability of the hidden states and observations factors in the following way:

$$p(\mathbf{x}_{1:T}, \mathbf{y}_{1:T} | \theta) = \prod_{t=1}^T p(\mathbf{x}_t | \mathbf{x}_{t-1}, \theta) p(\mathbf{y}_t | \mathbf{x}_t, \theta) \quad (24)$$

In other words, the observations are assumed to have been generated from the hidden states via  $p(\mathbf{y}_t | \mathbf{x}_t, \theta)$ , and the hidden states are assumed to have first-order Markov dynamics captured by  $p(\mathbf{x}_t | \mathbf{x}_{t-1}, \theta)$ . We can consider the first term  $p(\mathbf{x}_1 | \mathbf{x}_0, \theta)$  to be a prior on the initial state of the system  $\mathbf{x}_1$ .

The simplest kind of state-space model assumes that all variables are multivariate Gaussian distributed and all the relationships are linear. In such *linear-Gaussian state-space models*, we can write

$$\mathbf{y}_t = C\mathbf{x}_t + \mathbf{v}_t \quad (25)$$

$$\mathbf{x}_t = A\mathbf{x}_{t-1} + \mathbf{w}_t \quad (26)$$

where the matrices  $C$  and  $A$  define the linear relationships and  $\mathbf{v}$  and  $\mathbf{w}$  are zero-mean Gaussian noise vectors with covariance matrices  $R$  and  $Q$  respectively. If we assume that the prior on the initial state  $p(\mathbf{x}_1)$  is also Gaussian, then all subsequent  $\mathbf{x}$ s and  $\mathbf{y}$ s are also Gaussian due to the fact that Gaussian densities are closed under linear transformations. This model can be generalised in many ways, for example by augmenting it to include a sequence of observed inputs  $\mathbf{u}_1, \dots, \mathbf{u}_T$  as well as the observed model outputs  $\mathbf{y}_1, \dots, \mathbf{y}_T$ , but we will not discuss generalisations further.

By comparing equations (11) and (25) we see that linear-Gaussian SSMs can be thought of as a time-series generalisation of factor analysis where the factors are assumed to have linear-Gaussian dynamics over time.

The parameters of this model are  $\theta = (A, C, Q, R)$ . To learn ML settings of these parameters one can make use of the EM algorithm [73]. The E step of the algorithm involves computing  $q(\mathbf{x}_{1:T}) = p(\mathbf{x}_{1:T} | \mathbf{y}_{1:T}, \theta)$  which is the posterior over hidden state sequences. In fact, this whole posterior does not have to be computed or represented, all that is required are the marginals  $q(\mathbf{x}_t)$  and pairwise marginals  $q(\mathbf{x}_t, \mathbf{x}_{t+1})$ . These can be computed via the *Kalman smoothing algorithm*, which is an efficient algorithm for inferring the distribution over the hidden states of a linear-Gaussian SSM. Since the model is linear, the M step of the algorithm requires solving a pair of weighted linear regression problems to re-estimate  $A$  and  $C$ , while  $Q$  and  $R$  are estimated from the residuals of those regressions. This is analogous to the M step of factor analysis, which also involves solving a linear regression problem.

## 4.2 Hidden Markov models (HMMs)

Hidden Markov models are similar to state-space models in that the sequence of observations is assumed to have been generated from a sequence of underlying hidden states. The key difference is that in HMMs the state is assumed to be *discrete* rather than a continuous random vector. Let  $s_t$  denote the hidden state of an HMM at time  $t$ . We assume that  $s_t$  can take discrete values in  $\{1, \dots, K\}$ . The model can again be written as in (24):

$$P(s_{1:T}, \mathbf{y}_{1:T} | \theta) = \prod_{t=1}^T P(s_t | s_{t-1}, \theta) P(\mathbf{y}_t | s_t, \theta) \quad (27)$$

where  $P(s_1 | s_0, \theta)$  is simply some initial distribution over the  $K$  settings of the first hidden state; we can call this discrete distribution  $\boldsymbol{\pi}$ , represented by a  $K \times 1$  vector. The state-transition probabilities  $P(s_t | s_{t-1}, \theta)$  are captured by a  $K \times K$  transition matrix  $A$ , with elements  $A_{ij} = P(s_t = i | s_{t-1} = j, \theta)$ . The observations in an HMM can be either continuous or discrete. For continuous observations  $\mathbf{y}_t$  one can for example choose a Gaussian density; thus  $p(\mathbf{y}_t | s_t = i, \theta)$  would be a different Gaussian for each choice of  $i \in \{1, \dots, K\}$ . This model is the dynamical generalisation of a mixture of Gaussians. The marginal probability at each point in time is exactly a mixture of  $K$  Gaussians—the difference is that which component generates data point  $\mathbf{y}_t$

and which component generated  $\mathbf{y}_{t-1}$  are not independent random variables, but certain combinations are more and less probable depending on the entries in  $A$ . For  $\mathbf{y}_t$  a discrete observation, let us assume that it can take on values  $\{1, \dots, L\}$ . In that case the output probabilities  $P(\mathbf{y}_t | s_t, \theta)$  can be captured by an  $L \times K$  emission matrix,  $E$ .

The model parameters for a discrete-observation HMM are  $\theta = (\boldsymbol{\pi}, A, E)$ . Maximum likelihood learning of the model parameters can be approached using the EM algorithm, which in the case of HMMs is known as the *Baum-Welch algorithm*. The E step involves computing  $Q(s_t)$  and  $Q(s_t, s_{t+1})$  which are marginals of  $Q(s_{1:T}) = P(s_{1:T} | \mathbf{y}_{1:T}, \theta)$ . These marginals are computed as part of the *forward-backward algorithm* which as the name suggests sweeps forward and backward through the time series, and applies Bayes rule efficiently using the Markov conditional independence properties of the HMM, to compute the required marginals. The M step of HMM learning involves re-estimating  $\boldsymbol{\pi}$ ,  $A$ , and  $E$  by adding up and normalising expected counts for transitions and emissions that were computed in the E step.

### 4.3 Modelling other structured data

We have considered the case of iid data and time series data. The observations in real world data sets can have many other possible structures as well. Let us mention a few examples, although it is not possible to strive for completeness.

In spatial data, the points are assumed to live in some metric, often Euclidean, space. Three examples of spatial data include epidemiological data which can be modelled as a function of the spatial location of the measurement; data from computer vision where the observations are measurements of features on a 2D input to the camera; and functional neuroimaging where the data can be physiological measurements related to neural activity located in 3D voxels defining coordinates in the brain. Generalising HMMs, one can define Markov random field models where there are a set of hidden variables correlated to neighbours in some lattice, and related to the observed variables.

Hierarchical or tree-structured data contains known or unknown tree-like correlation structure between the data points or measured features. For example, the data points may be features of animals related through an evolutionary tree. A very different form of structured data is if each data point itself is tree-structured, for example if each point is a parse tree of a sentence in the English language.

Finally, one can take the structured dependencies between variables and consider the structure itself as an unknown part of the model. Such models are known as *probabilistic relational models* and are closely related to graphical models which we will discuss in Section 7.

## 5 Nonlinear, Factorial, and Hierarchical Models

The models we have described so far are attractive because they are relatively simple to understand and learn. However, their simplicity is also a limitation, since the intricacies of real-world data are unlikely to be well-captured by a simple statistical model. This motivates us to seek to describe and study learning in much more flexible models.

A simple combination of two of the ideas we have described for iid data is the *mixture of factor analysers* [23, 34, 77]. This model performs simultaneous clustering and dimensionality reduction on the data, by assuming that the covariance in each Gaussian cluster can be modelled by an FA model. Thus, it becomes possible to apply a mixture model to very high dimensional data while allowing each cluster to span a different sub-space of the data.

As their name implies linear-Gaussian SSMs are limited by assumptions of linearity and Gaussian noise. In many realistic dynamical systems there are significant nonlinear effects, which make it necessary to consider learning in *nonlinear state-space models*. Such models can also be learned using the EM algorithm, but the E step must deal with inference in non-Gaussian and potentially very complicated densities (since non-linearities will turn Gaussians into non-Gaussians), and the M step is nonlinear regression, rather than linear regression [25]. There are many methods of dealing with inference in non-linear SSMs, including methods such as particle filtering [29, 27, 40, 43, 35, 15], linearisation [2], the unscented filter [39, 80], the EP algorithm [52], and embedded HMMs [62].

Non-linear models are also important if we are to consider generalising simple dimensionality reduction models such as PCA and FA. These models are limited in that they can only find a linear subspace of the data to capture the correlations between the observed variables. There are many interesting and important nonlinear dimensionality reduction models, including generative topographic mappings (GTM) [11] (a probabilistic alternative to Kohonen maps), multi-dimensional scaling (MDS) [72, 45], principal curves [30], Isomap [76], and locally linear embedding (LLE) [69].

Hidden Markov models also have their limitations. Even though they can model nonlinear dynamics by discretising the hidden state space, an HMM with  $K$  hidden states can only capture  $\log_2 K$  bits of information in its state variable about the past of the sequence. HMMs can be extended by allowing a *vector* of discrete state variables, in an architecture known as a *factorial HMM* [24]. Thus a vector of  $M$  variables, each of which can take  $K$  states, can capture  $K^M$  possible states in total, and  $M \log_2 K$  bits of information about the past of the sequence. The problem is that such a model, if dealt with naively as an HMM would have exponentially many parameters and would take exponentially long to do inference in. Both the complexity in time and number of parameters can be alleviated by restricting the interactions between the hidden variables at one time step and at the next time step. A generalisation of these ideas is the notion of a *dynamical Bayesian network (DBN)* [56].

A relatively old but still quite powerful class of models for binary data is the *Boltzmann machine* (BM) [1]. This is a simple model inspired from Ising models in statistical physics. A BM is a multivariate model for capturing correlations and higher order statistics in vectors of binary data. Consider data consisting of vectors of  $M$  binary variables (the elements of the vector may, for example, be pixels in a black-and-white image). Clearly, each data point can be an instance of one of  $2^M$  possible patterns. An arbitrary distribution over such patterns would require a table with  $2^M - 1$  entries, again intractable in number of parameters, storage, and computation time. A BM allows one to define flexible distributions over the  $2^M$  entries of this table by using  $\mathcal{O}(M^2)$  parameters defining a symmetric matrix of weights connecting the variables. This can be augmented with hidden variables in order to enrich the model class, without adding exponentially many parameters. These hidden variables can be organised into layers of a hierarchy as in the Helmholtz machine [33]. Other hierarchical models include recent generalisations of ICA designed to capture higher order statistics in images [41].

## 6 Intractability

The problem with the models described in the previous section is that learning their parameters is in general computationally intractable. In a model with exponentially many settings for the hidden states, doing the E step of an EM algorithm would require computing appropriate marginals of a distribution over exponentially many possibilities.

Let us consider a simple example. Imagine we have a vector of  $N$  binary random variables  $\mathbf{s} = (s_1, \dots, s_N)$ , where  $s_i \in \{0, 1\}$  and a vector of  $N$  known integers  $(r_1, \dots, r_N)$  where  $r_i \in \{1, 2, 3, \dots, 10\}$ . Let the variable  $Y = \sum_{i=1}^N r_i s_i$ . Assume that the binary variables are all independent and identically distributed with  $P(s_i = 1) = 1/2, \forall i$ . Let  $N$  be 100. Now imagine that we are told  $Y = 430$ . How do we compute  $P(s_i = 1 | Y = 430)$ ? The problem is that even though the  $s_i$  were independent *before* we observed the value of  $Y$ , now that we know the value of  $Y$ , not all settings of  $\mathbf{s}$  are possible anymore. To figure out for some  $s_i$  the probability of  $P(s_i = 1 | Y = 430)$  requires that we enumerate all potentially exponentially many ways of achieving  $Y = 430$  and counting how many of those had  $s_i = 1$  vs  $s_i = 0$ .

This example illustrates the following ideas: Even if the prior is simple, the posterior can be very complicated. Whether two random variables are independent or not is a function of one's state of knowledge. Thus  $s_i$  and  $s_j$  may be independent if we are not told the value of  $Y$  but are certainly dependent given the value of  $Y$ . These type of phenomena are related to "explaining-away" which refers to the fact that if there are multiple potential causes for some effect, observing one, explains away the need for the others [64].

Intractability can thus occur if we have a model with discrete hidden variables which can take on exponentially many combinations. Intractability can also occur with continuous hidden variables if their density is not simply described, or if they interact with discrete hidden variables. Moreover, even for simple models, such as a mixture of Gaussians, intractability occurs when we consider the parameters to be unknown as




Figure 1: Three kinds of probabilistic graphical model: undirected graphs, factor graphs and directed graphs.

well, and we attempt to do Bayesian inference on them. To deal with intractability it is essential to have good tools for representing multivariate distributions, such as graphical models.

## 7 Graphical models

Graphical models are an important tool for representing the dependencies between random variables in a probabilistic model. They are important for two reasons. First, graphs are an intuitive way of visualising dependencies. We are used to graphical depictions of dependency, for example in circuit diagrams and in phylogenetic trees. Second, by exploiting the structure of the graph it is possible to devise efficient message passing algorithms for computing marginal and conditional probabilities in a complicated model. We discuss message passing algorithms for inference in Section 8.

The main statistical property represented explicitly by the graph is *conditional independence* between variables. We say that  $X$  and  $Y$  are conditionally independent given  $Z$ , if  $P(X, Y|Z) = P(X|Z)P(Y|Z)$  for all values of the variables  $X, Y$ , and  $Z$  where these quantities are defined (i.e. excepting settings  $z$  where  $P(Z = z) = 0$ ). We use the notation  $X \perp\!\!\!\perp Y|Z$  to denote the above conditional independence relation. Conditional independence generalises to sets of variables in the obvious way, and it is different from *marginal independence* which states that  $P(X, Y) = P(X)P(Y)$ , and is denoted  $X \perp\!\!\!\perp Y$ .

There are several different graphical formalisms for depicting conditional independence relationships. We focus on three of the main ones: undirected, factor, and directed graphs.

### 7.1 Undirected graphs

In an *undirected graphical model* each random variable is represented by a node, and the edges of the graph indicate conditional independence relationships. Specifically, let  $\mathcal{X}, \mathcal{Y}$ , and  $\mathcal{Z}$  be sets of random variables. Then  $\mathcal{X} \perp\!\!\!\perp \mathcal{Y}|\mathcal{Z}$  if every path on the graph from a node in  $\mathcal{X}$  to a node in  $\mathcal{Y}$  has to go through a node in  $\mathcal{Z}$ . Thus a variable  $X$  is conditionally independent of all other variables given the neighbours of  $X$ , and we say that the neighbours *separate*  $X$  from the rest of the graph. An example of an undirected graph is shown in Figure 1. In this graph  $A \perp\!\!\!\perp B|C$  and  $B \perp\!\!\!\perp E|\{C, D\}$ , for example, and the neighbours of  $D$  are  $B, C, E$ .

A *clique* is a fully connected subgraph of a graph. A *maximal clique* is not contained in any other clique of the graph. It turns out that the set of conditional independence relations implied by the separation properties in the graph are satisfied by probability distributions which can be written as a normalised product of non-negative functions over the variables in the maximal cliques of the graph (this is known as the Hammersley-Clifford Theorem [10]). In the example in Figure 1, this implies that the probability distribution over  $(A, B, C, D, E)$  can be written as:

$$P(A, B, C, D, E) = c g_1(A, C)g_2(B, C, D)g_3(C, D, E) \quad (28)$$

Here,  $c$  is the constant that ensures that the probability distribution sums to 1, and  $g_1, g_2$  and  $g_3$  are non-negative functions of their arguments. For example, if all the variables are binary the function  $g_2$  is a table with a non-negative number for each of the  $8 = 2 \times 2 \times 2$  possible settings of the variables  $B, C, D$ . These non-negative functions are supposed to represent how compatible these settings are with each other,

with a 0 encoding logical incompatibility. For this reason, the  $g$ 's are sometimes referred to as *compatibility functions*, other times as *potential functions*. Undirected graphical models are also sometimes referred to as *Markov networks*.

## 7.2 Factor graphs

In a *factor graph* there are two kinds of nodes, *variable nodes* and *factor nodes*, usually denoted as open circles and filled dots (Figure 1). Like an undirected model, the factor graph represents a factorisation of the joint probability distribution: each factor is a non-negative function of the variables connected to the corresponding factor node. Thus for the factor graph in Figure 1 we have:

$$P(A, B, C, D, E) = cg_1(A, C)g_2(B, C)g_3(B, D), g_4(C, D)g_5(C, E)g_6(D, E) \quad (29)$$

Factor nodes are also sometimes called function nodes. Again, as in an undirected graphical model, the variables in a set  $\mathcal{X}$  are conditionally independent of the variables in a set  $\mathcal{Y}$  given  $\mathcal{Z}$  if all paths from  $\mathcal{X}$  to  $\mathcal{Y}$  go through variables in  $\mathcal{Z}$ . Note that the factor graph in Figure 1 has exactly the same conditional independence relations as the undirected graph, even though the factors in the former are contained in the factors in the latter. Factor graphs are particularly elegant and simple when it comes to implementing message passing algorithms for inference (Section 8).

## 7.3 Directed graphs

In *directed graphical models*, also known as probabilistic directed acyclic graphs (DAGs), belief networks, and Bayesian networks, the nodes represent random variables and the directed edges represent statistical dependencies. If there exists an edge from  $A$  to  $B$  we say that  $A$  is a *parent* of  $B$ , and conversely  $B$  is a *child* of  $A$ . A directed graph corresponds to the factorisation of the joint probability into a product of the conditional probabilities of each node given its parents. For the example in Figure 1 we write:

$$P(A, B, C, D, E) = P(A)P(B)P(C|A, B)P(D|B, C)P(E|C, D) \quad (30)$$

In general we would write:

$$P(X_1, \dots, X_N) = \prod_{i=1}^N P(X_i | X_{\text{pa}_i}) \quad (31)$$

where  $X_{\text{pa}_i}$  denotes the variables that are parents of  $X_i$  in the graph.

Assessing the conditional independence relations in a directed graph is slightly less trivial than in undirected and factor graphs. Rather than simply looking at separation between sets of variables, one has to consider the directions of the edges. The graphical test for two sets of variables being conditionally independent given a third is called *d-separation* [64]. D-separation takes into account the following fact about *v-structures* of the graph, which consist of two (or more) parents of a child, as in the  $A \rightarrow C \leftarrow B$  sub-graph in Figure 1. In such a v-structure  $A \perp\!\!\!\perp B$ , but it is not true that  $A \perp\!\!\!\perp B | C$ . That is,  $A$  and  $B$  are marginally independent, but conditionally *dependent* given  $C$ . This can be easily checked by writing out  $P(A, B, C) = P(A)P(B)P(C|A, B)$ . Summing out  $C$  leads to  $P(A, B) = P(A)P(B)$ . However, given the value of  $C$ ,  $P(A, B|C) = P(A)P(B)P(C|A, B)/P(C)$  which does not factor into separate functions of  $A$  and  $B$ . As a consequence of this property of v-structures, in a directed graph a variable  $X$  is independent of all other variables given the parents of  $X$ , the children of  $X$ , and the parents of the children of  $X$ . This is the minimal set that d-separates  $X$  from the rest of the graph and is known as the Markov boundary for  $X$ .

It is possible, though not always appropriate, to interpret a directed graphical model as a causal generative model of the data. The following procedure would generate data from the probability distribution defined by a directed graph: draw a random value from the marginal distribution of all variables which do not have any parents (e.g.  $a \sim P(A)$ ,  $b \sim P(B)$ ), then sample from the conditional distribution of the children of these variables (e.g.  $c \sim P(C|A = a, B = b)$ ), and continue this procedure until all variables are assigned values. In the model,  $P(C|A, B)$  can capture the causal relationship between the causes  $A$  and  $B$  and the effect  $C$ . Such causal interpretations are much less natural for undirected and factor graphs, since even generating a sample from such models cannot easily be done in a hierarchical manner starting from “parents” to “children” except




Figure 2: No directed graph over 4 variables can represent the set of conditional independence relationships represented by this undirected graph.

in special cases. Moreover, the potential functions capture mutual compatibilities, rather than cause-effect relations.

A useful property of directed graphical models is that there is no global normalisation constant  $c$ . This global constant can be computationally intractable to compute in undirected and factor graphs. In directed graphs, each term is a conditional probability and is therefore already normalised  $\sum_x P(X_i = x | X_{\text{pa}_i}) = 1$ .

#### 7.4 Expressive power

Directed, undirected and factor graphs are complementary in their ability to express conditional independence relationships. Consider the directed graph consisting of a single v-structure  $A \rightarrow C \leftarrow B$ . This graph encodes  $A \perp\!\!\!\perp B$  but not  $A \perp\!\!\!\perp B | C$ . There exists no undirected graph or factor graph over these three variables which captures exactly these independencies. For example, in  $A - C - B$  it is not true that  $A \perp\!\!\!\perp B$  but it is true that  $A \perp\!\!\!\perp B | C$ . Conversely, if we consider the undirected graph in Figure 2, we see that some independence relationships are better captured by undirected models (and factor graphs).

## 8 Exact inference in graphs

Probabilistic *inference* in a graph usually refers to the problem of computing the conditional probability of some variable  $X_i$  given the observed values of some other variables  $X_{\text{obs}} = x_{\text{obs}}$  while marginalising out all other variables. Starting from a joint distribution  $P(X_1, \dots, X_N)$ , we can divide the set of all variables into three exhaustive and mutually exclusive sets  $\{X_1, \dots, X_N\} = \{X_i\} \cup X_{\text{obs}} \cup X_{\text{other}}$ . We wish to compute

$$P(X_i | X_{\text{obs}} = x_{\text{obs}}) = \frac{\sum_x P(X_i, X_{\text{other}} = x, X_{\text{obs}} = x_{\text{obs}})}{\sum_{x'} \sum_x P(X_i = x', X_{\text{other}} = x, X_{\text{obs}} = x_{\text{obs}})} \quad (32)$$

The problem is that the sum over  $x$  is exponential in the number of variables in  $X_{\text{other}}$ . For example, if there are  $M$  variables in  $X_{\text{other}}$  and each is binary, then there are  $2^M$  possible values for  $x$ . If the variables are continuous, then the desired conditional probability is the ratio of two high-dimensional integrals, which could be intractable to compute. Probabilistic inference is essentially a problem of computing large sums and integrals.

There are several algorithms for computing these sums and integrals which exploit the structure of the graph to get the solution efficiently for certain graph structures (namely trees and related graphs). For general graphs the problem is fundamentally hard [13].

## 8.1 Elimination

The simplest algorithm conceptually is *variable elimination*. It is easiest to explain with an example. Consider computing  $P(A = a|D = d)$  in the directed graph in Figure 1. This can be written

$$\begin{aligned}
P(A = a|D = d) &\propto \sum_c \sum_b \sum_e P(A = a, B = b, C = c, D = d, E = e) \\
&= \sum_c \sum_b \sum_e P(A = a)P(B = b)P(C = c|A = a, B = b) \\
&\quad P(D = d|C = c, B = b)P(E = e|C = c, D = d) \\
&= \sum_c \sum_b P(A = a)P(B = b)P(C = c|A = a, B = b) \\
&\quad P(D = d|C = c, B = b) \sum_e P(E = e|C = c, D = d) \\
&= \sum_c \sum_b P(A = a)P(B = b)P(C = c|A = a, B = b) \\
&\quad P(D = d|C = c, B = b)
\end{aligned}$$

What we did was (1) exploit the factorisation, (2) rearrange the sums, and (3) eliminate a variable,  $E$ . We could repeat this procedure and eliminate the variable  $C$ . When we do this we will need to compute a new function  $\phi(A = a, B = b, D = d) \stackrel{\text{def}}{=} \sum_c P(C = c|A = a, B = b)P(D = d|C = c, B = b)$ , resulting in:

$$P(A = a|D = d) \propto \sum_b P(A = a)P(B = b)\phi(A = a, B = b, D = d)$$

Finally, we eliminate  $B$  by computing  $\phi'(A = a, D = d) \stackrel{\text{def}}{=} \sum_b P(B = b)\phi(A = a, B = b, D = d)$  to get our final answer which can be written

$$P(A = a|D = d) \propto P(A = a)\phi'(A = a, D = d) = \frac{P(A = a)\phi'(A = a, D = d)}{\sum_a P(A = a)\phi'(A = a, D = d)}$$

The functions we get when we eliminate variables can be thought of as messages sent by that variable to its neighbours. Eliminating transforms the graph by removing the eliminated node and drawing (undirected) edges between all the nodes in the Markov boundary of the eliminated node.

The same answer is obtained no matter what order we eliminate variables in; however, the computational complexity can depend dramatically on the ordering used.

## 8.2 Belief propagation

The belief propagation (BP) algorithm is a message passing algorithm for computing conditional probabilities of any variable given the values of some set of other variables in a *singly-connected* directed acyclic graph [64]. The algorithm itself follows from the rules of probability and the conditional independence properties of the graph. Whereas variable elimination focuses on finding the conditional probability of a single variable  $X_i$  given  $X_{\text{obs}} = x_{\text{obs}}$ , belief propagation can compute at once all the conditionals  $p(X_i|X_{\text{obs}} = x_{\text{obs}})$  for all  $i$  not observed.

We first need to define singly-connected directed graphs. A directed graph is singly connected if between every pair of nodes there is only one undirected path. An *undirected path* is a path along the edges of the graph ignoring the direction of the edges: in other words the path can traverse edges both upstream and downstream. If there is more than one undirected path between any pair of nodes then the graph is said to be *multiply connected*, or *loopy* (since it has loops).

Singly connected graphs have an important property which BP exploits. Let us call the set of observed variables the *evidence*,  $e = X_{\text{obs}}$ . Every node in the graph divides the evidence into upstream  $e_X^+$  and downstream  $e_X^-$  parts. For example, in Figure 3 the variables  $U_1 \dots U_n$  their parents, ancestors, and children and descendents (not including  $X$ , its children and descendents) and anything else connected to  $X$  via an






Figure 3: Belief propagation in a directed graph.

edge directed toward  $X$  are all considered to be *upstream* of  $X$ ; anything connected to  $X$  via an edge away from  $X$  is considered *downstream* of  $X$  (e.g.  $Y_1$ , its children, the parents of its children, etc). Similarly, every edge  $X \rightarrow Y$  in a singly connected graph divides the evidence into upstream and downstream parts. This separation of the evidence into upstream and downstream components does not generally occur in multiply-connected graphs.

Belief propagation uses three key ideas to compute the probability of some variable given the evidence  $p(X|e)$ , which we can call the “belief” about  $X$ .<sup>4</sup> First, the belief about  $X$  can be found by combining upstream and downstream evidence:

$$P(X|e) = \frac{P(X, e)}{P(e)} \propto P(X, e_X^+, e_X^-) \propto P(X|e_X^+)P(e_X^-|X) \quad (33)$$

The last proportionality results from the fact that given  $X$  the downstream and upstream evidence are conditionally independent:  $P(e_X^-|X, e_X^+) = P(e_X^-|X)$ . Second, the effect of the upstream and downstream evidence on  $X$  can be computed via a local message passing algorithm between the nodes in the graph. Third, the message from  $X$  to  $Y$  has to be constructed carefully so that node  $X$  doesn’t send back to  $Y$  any information that  $Y$  sent to  $X$ , otherwise the message passing algorithm would reverberate information between nodes amplifying and distorting the final beliefs.

Using these ideas and the basic rules of probability we can arrive at the following equations, where  $\text{ch}(X)$  and  $\text{pa}(X)$  are children and parents of  $X$ , respectively:

$$\lambda(X) \stackrel{\text{def}}{=} P(e_X^-|X) = \prod_{j \in \text{ch}(X)} P(e_{XY_j}^-|X) \quad (34)$$

$$\pi(X) \stackrel{\text{def}}{=} P(X|e_X^+) = \sum_{U_1 \dots U_n} P(X|U_1, \dots, U_n) \prod_{i \in \text{pa}(X)} P(U_i|e_{U_i X}^+) \quad (35)$$

Finally, the messages from parents to children (e.g.  $X$  to  $Y_j$ ) and the messages from children to parents (e.g.  $X$  to  $U_i$ ) can be computed as follows:

$$\begin{aligned} \pi_{Y_j}(X) &\stackrel{\text{def}}{=} P(X|e_{XY_j}^+) \\ &\propto \left[ \prod_{k \neq j} P(e_{XY_k}^-|X) \right] \sum_{U_1, \dots, U_n} P(X|U_1 \dots U_n) \prod_i P(U_i|e_{U_i X}^+) \end{aligned} \quad (36)$$

$$\begin{aligned} \lambda_X(U_i) &\stackrel{\text{def}}{=} P(e_{U_i X}^-|U_i) \\ &= \sum_X P(e_X^-|X) \sum_{U_k: k \neq i} P(X|U_1 \dots U_n) \prod_{k \neq i} P(U_k|e_{U_k X}^+) \end{aligned} \quad (37)$$

<sup>4</sup>There is considerably variety in the field regarding the naming of algorithms. Belief propagation is also known as the sum-product algorithm, a name which some people prefer since beliefs seem subjective.

It is important to notice that in the computation of both the top-down message (36) and the bottom-up message (37) the recipient of the message is explicitly excluded. Pearl’s [64] mnemonic of calling these messages  $\lambda$  and  $\pi$  messages is meant to reflect their role in computing “likelihood” and “prior” terms.

BP includes as special cases two important algorithms: Kalman smoothing for linear-Gaussian state-space models, and the forward-backward algorithm for hidden Markov models. Although BP is only valid on singly connected graphs there is a large body of research on its application to multiply connected graphs—the use of BP on such graphs is called *loopy belief propagation* and has been analysed by several researchers [81, 82]. Interest in loopy belief propagation arose out of its impressive performance in decoding error correcting codes [21, 9, 50, 49]. Although the beliefs are not guaranteed to be correct on loopy graphs, interesting connections can be made to approximate inference procedures inspired by statistical physics known as the Bethe and Kikuchi free energies [84].

### 8.3 Factor graph propagation

In belief propagation, there is an asymmetry between the messages a child sends its parents and the messages a parent sends its children. Propagation in singly-connected factor graphs is conceptually much simpler and easier to implement. In a factor graph, the joint probability distribution is written as a product of factors. Consider a vector of variables  $\mathbf{x} = (x_1, \dots, x_n)$

$$p(\mathbf{x}) = p(x_1, \dots, x_n) = \frac{1}{Z} \prod_j f_j(\mathbf{x}_{S_j}) \quad (38)$$

where  $Z$  is the normalisation constant,  $S_j$  denotes the subset of  $\{1, \dots, n\}$  which participate in factor  $f_j$  and  $\mathbf{x}_{S_j} = \{x_i : i \in S_j\}$ .

Let  $n(x)$  denote the set of factor nodes that are neighbours of  $x$  and let  $n(f)$  denote the set of variable nodes that are neighbours of  $f$ . We can compute probabilities in a factor graph by propagating messages from variable nodes to factor nodes and vice-versa. The message from variable  $x$  to function  $f$  is:

$$\mu_{x \rightarrow f}(x) = \prod_{h \in n(x) \setminus \{f\}} \mu_{h \rightarrow x}(x) \quad (39)$$

while the message from function  $f$  to variable  $x$  is:

$$\mu_{f \rightarrow x}(x) = \sum_{\mathbf{x} \setminus x} \left( f(\mathbf{x}) \prod_{y \in n(f) \setminus \{x\}} \mu_{y \rightarrow f}(y) \right) \quad (40)$$

Once a variable has received all messages from its neighbouring factor nodes we can compute the probability of that variable by multiplying all the messages and renormalising:

$$p(x) \propto \prod_{h \in n(x)} \mu_{h \rightarrow x}(x) \quad (41)$$

Again, these equations can be derived by using Bayes rule and the conditional independence relations in a singly-connected factor graph. For multiply-connected factor graphs (where there is more than one path between at least one pair of variable nodes) one can apply a loopy version of factor graph propagation. Since the algorithms for directed graphs and factor graphs are essentially based on the same ideas, we also call the loopy version of factor graph propagation “loopy belief propagation”.

### 8.4 Junction tree algorithm

For multiply-connected graphs, the standard exact inference algorithms are based on the notion of a *junction tree* [46]. The basic idea of the junction tree algorithm is to group variables so as to convert the multiply-connected graph into a singly-connected undirected graph (tree) over sets of variables, and do inference in this tree.

We will not explain the algorithm in detail here, but rather give an overview of the steps involved. Starting from a directed graph, undirected edges are introduced between every pair of variables that share a child. This step is called “moralisation” in a tongue-in-cheek reference to the fact that it involves marrying the unmarried parents of every node. All the remaining edges are then changed from directed to undirected. We now have an undirected graph which does not imply any additional conditional or marginal independence relations which were not present in the original directed graph (although the undirected graph may easily have many fewer conditional or marginal independence relations than the directed graph). The next step of the algorithm is “triangulation” which introduces an edge cutting across every cycle of length 4. For example, the cycle  $A - B - C - D - A$  which would look like Figure 2 would be triangulated either by adding an edge  $A - C$  or an edge  $B - D$ . Once the graph has been triangulated, the maximal cliques of the graph are organised into a tree, where the nodes of the tree are cliques, by placing edges in the tree between some of the cliques with an overlap in variables (placing edges between all overlaps may not result in a tree). In general it may be possible to build several trees in this way, and triangulating the graph means that there exists a tree with the “running intersection property”. This property ensures that none of the variable is represented in disjoint parts of the tree, as this would cause the algorithm to come up with multiple possibly inconsistent beliefs about the variable. Finally, once the tree with the running intersection property is built (the junction tree) it is possible to introduce the evidence into the tree and apply what is essentially a variant of belief propagation to this junction tree. This BP algorithm is operating on sets of variables contained in the cliques of the junction tree, rather than on individual variables in the original graph. As such, the complexity of the algorithm scales exponentially with the size of the largest clique in the junction tree. For example, if moralisation and triangulation results in a clique containing  $K$  binary variables, the junction tree algorithm would have to store and manipulate tables of size  $2^K$ . Moreover, finding the optimal triangulation to get the most efficient junction tree for a particular graph is NP-complete [4, 44].

## 8.5 Cutset conditioning

In certain graphs the simplest inference algorithm is *cutset conditioning* which is related to the idea of “reasoning by assumptions”. The basic idea is very straightforward: find some small set of variables such that if they were given (i.e. you knew their values) it would make the remainder of the graph singly connected. For example, in the undirected graph in Figure 1, given  $C$  or  $D$ , the rest of the graph is singly connected. This set of variables is called the *cutset*. For each possible value of the variables in the cutset, run BP on the remainder of the graph to obtain the beliefs on the node of interest. These beliefs can be averaged with appropriate weights to obtain the true belief on the variable of interest. To make this more concrete, assume you want to find  $P(X|e)$  and you discover a cutset consisting of a single variable  $C$ . Then

$$P(X|e) = \sum_c P(X|C = c, e)P(C = c|e) \quad (42)$$

where the beliefs  $P(X|C = c, e)$  and corresponding weights  $P(C = c|e)$  are computed as part of BP, run once for each value of  $c$ .

## 9 Learning in graphical models

In Section 8 we described exact algorithms for inferring the value of variables in a graph with known parameters and structure. If the parameters and structure are unknown they can be learned from the data [31]. The learning problem can be divided into learning the graph parameters for a known structure, and learning the model structure (i.e. which edges should be present or absent).<sup>5</sup>

We focus here on directed graphs with discrete variables, although some of these issues become much more subtle for undirected and factor graphs [57]. The parameters of a directed graph with discrete variables parameterise the conditional probability tables  $P(X_i|X_{\text{pa}_i})$ . For each setting of  $X_{\text{pa}_i}$  this table contains a

<sup>5</sup>It should be noted that in Bayesian statistics there is no fundamental difference between parameters and variables, and therefore the learning and inference problems are really the same. All unknown quantities are treated as random variables, and learning is just inference about parameters and structure. It is however often useful to distinguish between parameters, which we assume to be fairly constant over the data, and variables, which we can assume to vary over each data point.

probability distribution over  $X_i$ . For example, if all variables are binary and  $X_i$  has  $K$  parents, then this *conditional probability table* has  $2^{K+1}$  entries; however, since the probability over  $X_i$  has to sum to 1 for each setting of its parents there are only  $2^K$  independent entries. The most general parameterisation would have a distinct parameter for each entry in this table, but this is often not a natural way to parameterise the dependency between variables. Alternatives (for binary data) are the noisy-or or sigmoid parameterisation of the dependencies [58]. Whatever the specific parameterisation, let  $\theta_i$  denote the parameters relating  $X_i$  to its parents, and let  $\theta$  denote all the parameters in the model. Let  $m$  denote the model structure, which corresponds to the set of edges in the graph. More generally the model structure can also contain the presence of additional hidden variables [16].

## 9.1 Learning graph parameters

We first consider the problem of learning graph parameters when the model structure is known and there are no missing or hidden variables. The presence of missing/hidden variables complicates the situation.

### 9.1.1 The complete data case.

Assume that the parameters controlling each *family* (a child and its parents) are distinct and that we observe  $N$  iid instances of all  $K$  variables in our graph. The data set is therefore  $\mathcal{D} = \{X^{(1)} \dots X^{(N)}\}$  and the likelihood can be written

$$P(\mathcal{D}|\theta) = \prod_{n=1}^N P(X^{(n)}|\theta) = \prod_{n=1}^N \prod_{i=1}^K P(X_i^{(n)}|X_{\text{pa}_i}^{(n)}, \theta_i) \quad (43)$$

Clearly, maximising the log likelihood with respect to the parameters results in  $K$  decoupled optimisation problems, one for each family, since the log likelihood can be written as a sum of  $K$  independent terms. Similarly, if the prior factors over the  $\theta_i$ , then the Bayesian posterior is also factored:  $P(\theta|\mathcal{D}) = \prod_i P(\theta_i|\mathcal{D})$ .

### 9.1.2 The incomplete data case.

When there is missing/hidden data, the likelihood no longer factors over the variables. Divide the variables in  $X^{(n)}$  into observed and missing components,  $X_{\text{obs}}^{(n)}$  and  $X_{\text{mis}}^{(n)}$ . The observed data is now  $\mathcal{D} = \{X_{\text{obs}}^{(1)} \dots X_{\text{obs}}^{(N)}\}$  and the likelihood is:

$$P(\mathcal{D}|\theta) = \prod_{n=1}^N P(X_{\text{obs}}^{(n)}|\theta) \quad (44)$$

$$= \prod_{n=1}^N \sum_{x_{\text{mis}}^{(n)}} P(X_{\text{mis}}^{(n)} = x_{\text{mis}}^{(n)}, X_{\text{obs}}^{(n)}|\theta) \quad (45)$$

$$= \prod_{n=1}^N \sum_{x_{\text{mis}}^{(n)}} \prod_{i=1}^K P(X_i^{(n)}|X_{\text{pa}_i}^{(n)}, \theta_i) \quad (46)$$

where in the last expression the missing variables are assumed to be set to the values  $x_{\text{mis}}^{(n)}$ . Because of the missing data, the cost function can no longer be written as a sum of  $K$  independent terms and the parameters are all coupled. Similarly, even if the prior factors over the  $\theta_i$ , the Bayesian posterior will couple all the  $\theta_i$ .

One can still optimise the likelihood by making use of the EM algorithm (Section 3). The E step of EM infers the distribution over the hidden variables given the current setting of the parameters. This can be done with BP for singly connected graphs or with the junction tree algorithm for multiply-connected graphs. In the M step, the objective function being optimised conveniently factors in exactly the same way as in the complete data case (c.f. Equation (21)). Whereas for the complete data case, the optimal ML parameters can often be computed in closed form, in the incomplete data case an iterative algorithm such as EM is usually required.

Bayesian parameter inference in the incomplete data case is also substantially more complicated. The parameters and missing data are coupled in the posterior distribution, as can be seen by multiplying (45) by the parameter prior and normalising. Inference can be achieved via approximate inference methods such as Markov chain Monte Carlo methods (Section 11.3, [59]) like Gibbs sampling, and variational approximations (Section 11.4, [6]).

## 9.2 Learning graph structure

There are two basic components to learning the structure of a graph from data: scoring and search. *Scoring* refers to computing a measure which can be used to compare different structures  $m$  and  $m'$  given a data set  $\mathcal{D}$ . *Search* refers to searching over the space of possible model structures, usually by proposing changes to the current model, so as to find the model with the highest score. This view of structure learning presupposes that the goal is to find a single structure with the highest score, although of course in the Bayesian inference framework it is desirable to infer the probability distribution over model structures given the data.

### 9.2.1 Scoring metrics.

Assume that you have a prior  $P(m)$  over model structures, which is ideally based on some domain knowledge. The natural score to use is the probability of the model given the data (although see [32]) or some monotonic function of this:

$$s(m, \mathcal{D}) = P(m|\mathcal{D}) \propto P(\mathcal{D}|m)P(m). \quad (47)$$

This score requires computing the *marginal likelihood*

$$P(\mathcal{D}|m) = \int P(\mathcal{D}|\boldsymbol{\theta}, m)P(\boldsymbol{\theta}|m)d\boldsymbol{\theta}. \quad (48)$$

We discuss the intuitions behind the marginal likelihood as a natural score for model comparison in Section 10.

For directed graphical models with fully-observed discrete variables and factored Dirichlet priors over the parameters of the conditional probability tables, the integral in (48) is analytically tractable. For models with missing/hidden data, alternative choices of priors and types of variables, the integral in (48) is often intractable and approximation methods are required. Some of the standard approximations that can be applied in this context and many other Bayesian inference problems are briefly reviewed in Section 11.

### 9.2.2 Search algorithms.

Given a way of scoring models, one can search over the space of all possible valid graphical models for the one with the highest score [19]. The space of all possible graphs is very large (exponential in the number of variables) and for directed graphs it can be expensive to check whether a particular change to the graph will result in a cycle being formed. Thus intelligent heuristics are needed to search the space efficiently [55]. An alternative to trying to find the most probable graph are methods that sample over the posterior distribution of graphs [20]. This has the advantage that it avoids the problem of overfitting which can occur for algorithms that select a single structure with highest score out of exponentially many.

## 10 Bayesian model comparison and Occam's Razor

So far in this chapter we have seen many different kinds of models. One of the most important problems in unsupervised learning is automatically determining which models are appropriate for a given data set. Model selection and comparison questions include all of the following:

- Are there clusters in the data and if so, how many? What are their shapes (e.g. Gaussian,  $t$ -distributed)?
- Does the data live on a low dimensional manifold? What dimensionality? Is this manifold flat or curved?
- Is the data discretised? If so, to what precision?

- Is the data a time series? If so, is it better modelled by an HMM, a state-space model? Linear or nonlinear? Gaussian or non-Gaussian noise? How many states should the HMM have? How many state variables should the SSM have?
- Can the data be modelled well by a directed graph? What is the structure of this graph? Does it have hidden variables? Are these continuous or discrete?

Clearly, this list could go on. A human may be able to answer these questions via careful use of visualisation, hypothesis testing, and guesswork. But ultimately, an intelligent unsupervised learning system should be able to answer all these questions automatically.

Fortunately, the framework of Bayesian inference can be used to provide a rational, coherent and automatic way of answering all of the above questions. This means that, *given a complete specification of the prior assumptions* there is an automatic procedure (based on Bayes rule) which provides a unique answer. Of course, as always, if the prior assumptions are very poor, the answers obtained could be useless. Therefore, it is essential to think carefully about the prior assumptions before turning the automatic Bayesian handle.

Let us go over this automatic procedure. Consider a model  $m_i$  coming from a set of possible models  $\{m_1, m_2, m_3, \dots\}$ . For instance, the model  $m_i$  might correspond to a Gaussian mixture model with  $i$  components. The models need not be nested, nor does the space of models need to be discrete (although we'll focus on that case). Given data  $\mathcal{D}$ , the natural way to compare models is via their probability:

$$P(m_i|\mathcal{D}) = \frac{P(\mathcal{D}|m_i)P(m_i)}{P(\mathcal{D})} \quad (49)$$

To compare models, the denominator, which sums over the potentially huge space of all possible models,  $P(\mathcal{D}) = \sum_j P(\mathcal{D}|m_j)P(m_j)$  is not required. Prior preference for models can be included in  $P(m_i)$ . However, it is interesting to look closely at the *marginal likelihood* term (sometimes called the *evidence* for model  $m_i$ ). Assume that model  $m_i$  has parameters  $\theta_i$  (e.g. the means and covariance matrices of the  $i$  Gaussians, along with the mixing proportions, c.f. Section 2.4). The marginal likelihood integrates over all possible parameter values

$$P(\mathcal{D}|m_i) = \int P(\mathcal{D}|\theta_i, m_i)P(\theta|m_i) d\theta_i \quad (50)$$

where  $P(\theta|m_i)$  is the prior over parameters, which is required for a complete specification of the model  $m_i$ .

The marginal likelihood has a very interesting interpretation. It is the probability of generating data set  $\mathcal{D}$  from parameters that are *randomly sampled* from under the prior for  $m_i$ . This should be contrasted with the maximum likelihood for  $m_i$  which is the probability of the data under the single setting of the parameters  $\hat{\theta}_i$  that maximises  $P(\mathcal{D}|\theta_i, m_i)$ . Clearly a more complicated model will have a higher maximum likelihood, which is the reason why maximising the likelihood results in *overfitting* — i.e. a preference for more complicated models than necessary. In contrast, the marginal likelihood can decrease as the model becomes more complicated. In a more complicated model sampling random parameter values can generate a wider range of possible data sets, but since the probability over data sets has to integrate to 1 (assuming a fixed number of data points) spreading the density to allow for more complicated data sets necessarily results in some simpler data sets having lower density under the model. This situation is diagrammed in Figure 4. The decrease in the marginal likelihood as additional parameters are added has been called the *automatic Occam's Razor* [37, 47, 66].

In theory all the questions posed at the beginning of this section could be addressed by defining appropriate priors and carefully computing marginal likelihoods of competing hypotheses. However, in practice the integral in (50) is usually very high dimensional and intractable. It is therefore necessary to approximate it.

## 11 Approximating posteriors and marginal likelihoods

There are many ways of approximating the marginal likelihood of a model, and the corresponding parameter posterior. In this section, we review some of the most frequently used methods.

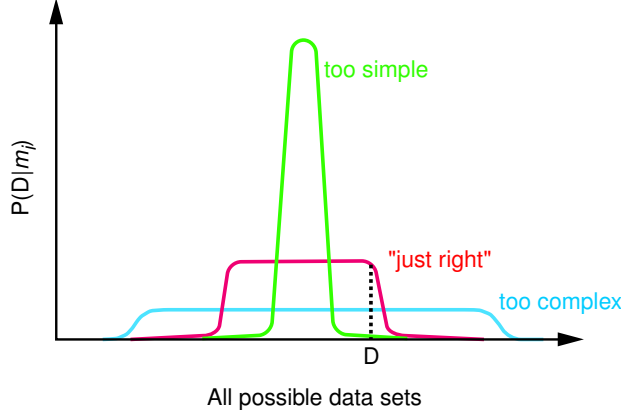


Figure 4: The marginal likelihood (evidence) as a function of an abstract one dimensional representation of “all possible” data sets of some size  $N$ . Because the evidence is a probability over data sets, it must normalise to one. Therefore very complex models which can account for many datasets only achieve modest evidence; simple models can reach high evidences, but only for a limited set of data. When a dataset  $\mathcal{D}$  is observed, the evidence can be used to select between model complexities.

### 11.1 Laplace approximation

It can be shown that under some regularity conditions, for large amounts of data  $N$  relative to the number of parameters in the model,  $d$ , the parameter posterior is approximately Gaussian around the MAP estimate,  $\hat{\theta}$ :

$$p(\theta|\mathcal{D}, m) \approx (2\pi)^{-\frac{d}{2}} |A|^{\frac{1}{2}} \exp \left\{ -\frac{1}{2} (\theta - \hat{\theta})^\top A (\theta - \hat{\theta}) \right\} \quad (51)$$

Here  $A$  is the  $d \times d$  negative of the Hessian matrix which measures the curvature of the log posterior at the MAP estimate:

$$A_{ij} = - \left. \frac{d^2}{d\theta_i d\theta_j} \log p(\theta|\mathcal{D}, m) \right|_{\theta=\hat{\theta}} \quad (52)$$

The matrix  $A$  is also referred to as the *observed information matrix*. Equation (51) is the *Laplace approximation* to the parameter posterior.

By Bayes rule, the marginal likelihood satisfies the following equality at any  $\theta$ :

$$p(\mathcal{D}|m) = \frac{p(\theta, \mathcal{D}|m)}{p(\theta|\mathcal{D}, m)} \quad (53)$$

The Laplace approximation to the marginal likelihood can be derived by evaluating the log of this expression at  $\hat{\theta}$ , using the Gaussian approximation to the posterior from equation (51) in the denominator:

$$\log p(\mathcal{D}|m) \approx \log p(\hat{\theta}|m) + \log p(\mathcal{D}|\hat{\theta}, m) + \frac{d}{2} \log 2\pi - \frac{1}{2} \log |A| \quad (54)$$

### 11.2 The Bayesian information criterion (BIC)

One of the disadvantages of the Laplace approximation is that it requires computing the determinant of the Hessian matrix. For models with many parameters, the Hessian matrix can be very large, and computing its determinant can be prohibitive.

The Bayesian Information Criterion (BIC) is a quick and easy way to compute an approximation to the marginal likelihood. BIC can be derived from the Laplace approximation by dropping all terms that do not depend on  $N$ , the number of data points. Starting from equation (54), we note that the first and third terms are constant with respect to the number of data points. Referring to the definition of the Hessian, we can see that its elements grow linearly with  $N$ . In the limit of large  $N$  we can therefore write  $A = N\tilde{A}$ , where

$\tilde{A}$  is a matrix independent of  $N$ . We use the fact that for any scalar  $c$  and  $d \times d$  matrix  $P$ , the determinant  $|cP| = c^d|P|$ , to get

$$\frac{1}{2} \log |A| \approx \frac{d}{2} \log N + \frac{1}{2} \log |\tilde{A}| \quad (55)$$

The last term does not grow with  $N$ , so by dropping it and substituting into Eq. (54) we get the BIC approximation:

$$\log p(\mathcal{D}|m) \approx \log p(\mathcal{D}|\hat{\boldsymbol{\theta}}, m) - \frac{d}{2} \log N \quad (56)$$

This expression is extremely easy to compute. Since the expression does not involve the prior it can be used either when  $\hat{\boldsymbol{\theta}}$  is the MAP or the ML parameter estimate, the latter choice making the entire procedure independent of a prior. The likelihood is penalised by a term that depends linearly on the number of parameters in the model; this term is referred to as the *BIC penalty*. This is how BIC approximates the Bayesian Occam's Razor effect which penalises overcomplex models. The BIC criterion can also be derived from within the *Minimum Description Length* (MDL) framework.

The BIC penalty is clearly attractive since it does not require any costly integrals or matrix inversions. However this simplicity comes at a cost in accuracy which can sometimes be catastrophic. One of the dangers of BIC is that it relies on the number of parameters. The basic assumption underlying BIC, that the Hessian converges to  $N$  times a full-rank matrix, only holds for models in which all parameters are identifiable and well-determined. This is often not true.

### 11.3 Markov chain Monte Carlo (MCMC)

Monte Carlo methods are a standard and often extremely effective way of computing complicated high dimensional integrals and sums. Many Bayesian inference problems can be seen as computing the integral (or sum) of some function  $f(\boldsymbol{\theta})$  under some probability density  $p(\boldsymbol{\theta})$ :

$$\bar{f} \stackrel{\text{def}}{=} \int f(\boldsymbol{\theta})p(\boldsymbol{\theta}) d\boldsymbol{\theta}. \quad (57)$$

For example, the marginal likelihood is the integral of the likelihood function under the prior. Simple Monte Carlo approximates (57) by sampling  $M$  independent draws  $\boldsymbol{\theta}_i \sim p(\boldsymbol{\theta})$  and computing the sample average of  $f$ :

$$\bar{f} \approx \frac{1}{M} \sum_{i=1}^M f(\boldsymbol{\theta}_i) \quad (58)$$

There are many limitations of simple Monte Carlo, for example it is often not possible to draw directly from  $p$ . Generalisations of simple Monte Carlo such as rejection sampling and importance sampling attempt to overcome some of these limitations.

An important family of generalisations of Monte Carlo methods are Markov chain Monte Carlo (MCMC) methods. These are commonly used and powerful methods for approximating the posterior over parameters and the marginal likelihood. Unlike simple Monte Carlo methods, the samples are not drawn independently but rather *dependently* in the form of a Markov chain  $\dots \boldsymbol{\theta}_i \rightarrow \boldsymbol{\theta}_{i+1} \rightarrow \boldsymbol{\theta}_{i+2} \dots$  where each sample depends on the value of the previous sample. MCMC estimates have the property that the asymptotic distribution of  $\boldsymbol{\theta}_i$  is the desired distribution. That is,  $\lim_{t \rightarrow \infty} p_t(\boldsymbol{\theta}_t) = p(\boldsymbol{\theta})$ . Creating MCMC methods is somewhat of an art, and there are many MCMC methods available, some of which are reviewed in [59]. Some notable examples are Gibbs sampling, the Metropolis algorithm, and Hybrid Monte Carlo.

### 11.4 Variational approximations

Variational methods can be used to derive a family of lower bounds on the marginal likelihood and to perform approximate Bayesian inference over the parameters of a probabilistic models [38, 83, 79]. Variational methods provide an alternative to the asymptotic and sampling-based approximations described above; they tend to be more accurate than the asymptotic approximations like BIC and faster than the MCMC approaches.



Let  $\mathbf{y}$  denote the observed variables,  $\mathbf{x}$  denote the latent variables, and  $\boldsymbol{\theta}$  denote the parameters. The log marginal likelihood of data  $\mathbf{y}$  can be lower bounded by introducing any distribution over both latent variables and parameters which has support where  $p(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y}, m)$  does, and then appealing to Jensen's inequality (due to the concavity of the logarithm function):

$$\ln p(\mathbf{y} | m) = \ln \int p(\mathbf{y}, \mathbf{x}, \boldsymbol{\theta} | m) d\mathbf{x} d\boldsymbol{\theta} = \ln \int q(\mathbf{x}, \boldsymbol{\theta}) \frac{p(\mathbf{y}, \mathbf{x}, \boldsymbol{\theta} | m)}{q(\mathbf{x}, \boldsymbol{\theta})} d\mathbf{x} d\boldsymbol{\theta} \quad (59)$$

$$\geq \int q(\mathbf{x}, \boldsymbol{\theta}) \ln \frac{p(\mathbf{y}, \mathbf{x}, \boldsymbol{\theta} | m)}{q(\mathbf{x}, \boldsymbol{\theta})} d\mathbf{x} d\boldsymbol{\theta}. \quad (60)$$

Maximising this lower bound with respect to the free distribution  $q(\mathbf{x}, \boldsymbol{\theta})$  results in  $q(\mathbf{x}, \boldsymbol{\theta}) = p(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y}, m)$  which when substituted above turns the inequality into an equality (c.f. Section 3). This does not simplify the problem since evaluating the true posterior distribution  $p(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y}, m)$  requires knowing its normalising constant, the marginal likelihood. Instead we use a simpler, factorised approximation  $q(\mathbf{x}, \boldsymbol{\theta}) = q_{\mathbf{x}}(\mathbf{x})q_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ :

$$\ln p(\mathbf{y} | m) \geq \int q_{\mathbf{x}}(\mathbf{x})q_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \ln \frac{p(\mathbf{y}, \mathbf{x}, \boldsymbol{\theta} | m)}{q_{\mathbf{x}}(\mathbf{x})q_{\boldsymbol{\theta}}(\boldsymbol{\theta})} d\mathbf{x} d\boldsymbol{\theta} \stackrel{\text{def}}{=} F_m(q_{\mathbf{x}}(\mathbf{x}), q_{\boldsymbol{\theta}}(\boldsymbol{\theta}), \mathbf{y}). \quad (61)$$

The quantity  $F_m$  is a functional of the free distributions,  $q_{\mathbf{x}}(\mathbf{x})$  and  $q_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ .

The variational Bayesian algorithm iteratively maximises  $F_m$  in equation (61) with respect to the free distributions,  $q_{\mathbf{x}}(\mathbf{x})$  and  $q_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ . We use elementary calculus of variations to take functional derivatives of the lower bound with respect to  $q_{\mathbf{x}}(\mathbf{x})$  and  $q_{\boldsymbol{\theta}}(\boldsymbol{\theta})$ , each while holding the other fixed. This results in the following update equations where the superscript  $(t)$  denotes the iteration number:

$$q_{\mathbf{x}}^{(t+1)}(\mathbf{x}) \propto \exp \left[ \int \ln p(\mathbf{x}, \mathbf{y} | \boldsymbol{\theta}, m) q_{\boldsymbol{\theta}}^{(t)}(\boldsymbol{\theta}) d\boldsymbol{\theta} \right] \quad (62)$$

$$q_{\boldsymbol{\theta}}^{(t+1)}(\boldsymbol{\theta}) \propto p(\boldsymbol{\theta} | m) \exp \left[ \int \ln p(\mathbf{x}, \mathbf{y} | \boldsymbol{\theta}, m) q_{\mathbf{x}}^{(t+1)}(\mathbf{x}) d\mathbf{x} \right] \quad (63)$$

When there is more than one data point then there are different hidden variables  $\mathbf{x}_i$  associated with each data point  $\mathbf{y}_i$  and the step in (62) has to be carried out for each  $i$ , where the distributions are  $q_{\mathbf{x}_i}^{(t)}(\mathbf{x}_i)$ .

Clearly  $q_{\boldsymbol{\theta}}(\boldsymbol{\theta})$  and  $q_{\mathbf{x}_i}(\mathbf{x}_i)$  are coupled, so we iterate these equations until convergence. Recalling the EM algorithm (Section 3 and [14, 63]) we note the similarity between EM and the iterative algorithm in (62) and (63). This procedure is called the *Variational Bayesian EM Algorithm* and generalises the usual EM algorithm; see also [5] and [22].

Re-writing (61), it is easy to see that maximising  $F_m$  is equivalent to minimising the KL divergence between  $q_{\mathbf{x}}(\mathbf{x}) q_{\boldsymbol{\theta}}(\boldsymbol{\theta})$  and the joint posterior  $p(\mathbf{x}, \boldsymbol{\theta} | \mathbf{y}, m)$ :

$$\ln p(\mathbf{y} | m) - F_m(q_{\mathbf{x}}(\mathbf{x}), q_{\boldsymbol{\theta}}(\boldsymbol{\theta}), \mathbf{y}) = \int q_{\mathbf{x}}(\mathbf{x}) q_{\boldsymbol{\theta}}(\boldsymbol{\theta}) \ln \frac{q_{\mathbf{x}}(\mathbf{x}) q_{\boldsymbol{\theta}}(\boldsymbol{\theta})}{p(\boldsymbol{\theta}, \mathbf{x} | \mathbf{y}, m)} d\mathbf{x} d\boldsymbol{\theta} = \text{KL}(q \| p) \quad (64)$$

Note that while this factorisation of the posterior distribution over latent variables and parameters may seem drastic, one can think of it as replacing stochastic dependencies between  $\mathbf{x}$  and  $\boldsymbol{\theta}$  with deterministic dependencies between relevant moments of the two sets of variables. To compare between models  $m$  and  $m'$  one can evaluate  $F_m$  and  $F_{m'}$ . This approach can, for example, be used to score graphical model structures [6].

Summarising, the variational Bayesian EM algorithm simultaneously computes an approximation to the marginal likelihood and to the parameter posterior by maximising a lower bound.

## 11.5 Expectation propagation (EP)

Expectation propagation (EP; [52, 53]) is another powerful method for approximate Bayesian inference. Consider a Bayesian inference problem in which you are given iid data  $\mathcal{D} = \{\mathbf{x}^{(1)} \dots, \mathbf{x}^{(N)}\}$  assumed to have come from a model  $p(\mathbf{x} | \boldsymbol{\theta})$  parameterised by  $\boldsymbol{\theta}$  with prior  $p(\boldsymbol{\theta})$ . The parameter posterior is:

$$p(\boldsymbol{\theta} | \mathcal{D}) = \frac{1}{p(\mathcal{D})} p(\boldsymbol{\theta}) \prod_{i=1}^N p(\mathbf{x}^{(i)} | \boldsymbol{\theta}). \quad (65)$$

To make the notation more general we can write the quantity we wish to approximate as a product of factors over  $\boldsymbol{\theta}$ ,

$$\prod_{i=0}^N f_i(\boldsymbol{\theta}) = p(\boldsymbol{\theta}) \prod_{i=1}^N p(\mathbf{x}^{(i)}|\boldsymbol{\theta}) \quad (66)$$

where  $f_0(\boldsymbol{\theta}) \stackrel{\text{def}}{=} p(\boldsymbol{\theta})$  and  $f_i(\boldsymbol{\theta}) \stackrel{\text{def}}{=} p(\mathbf{x}^{(i)}|\boldsymbol{\theta})$  and we will ignore the normalising constants. We wish to approximate this by a product of *simpler* terms

$$q(\boldsymbol{\theta}) \stackrel{\text{def}}{=} \prod_{i=0}^N \tilde{f}_i(\boldsymbol{\theta}). \quad (67)$$

For example, consider a binary linear classification problem where  $\boldsymbol{\theta}$  are the parameters of the classification hyperplane and  $p(\boldsymbol{\theta})$  is a Gaussian prior ([53], Chapter 5). The true posterior is the product of this Gaussian and  $N$  likelihood terms, each of which defines a half-plane consistent with the class label observed. This posterior has a complicated shape, but we can approximate it using EP by assuming that each of the approximate likelihood terms  $\tilde{f}_i$  is Gaussian in  $\boldsymbol{\theta}$ . Since the product of Gaussians is Gaussian,  $q(\boldsymbol{\theta})$  will be a Gaussian approximation to the posterior. In general, one makes the approximate terms  $\tilde{f}_i$  belong to some exponential family distribution so the overall approximation is in the same exponential family.

Having decided on the form of the approximation (67), let us consider how to tune this approximation so as to make it as accurate as possible. Ideally we would like to minimise the KL divergence between the true and the approximate distributions:

$$\min_{q(\boldsymbol{\theta})} \text{KL} \left( \prod_{i=0}^N f_i(\boldsymbol{\theta}) \parallel \prod_{i=0}^N \tilde{f}_i(\boldsymbol{\theta}) \right). \quad (68)$$

For example, if  $q(\boldsymbol{\theta})$  is a Gaussian density, minimising this KL divergence will result in finding the *exact* mean and covariance of the true posterior distribution over parameters. Unfortunately, this KL divergence involves averaging with respect to the true posterior distribution, which will generally be intractable. Note that the KL divergence in Equation (68) is different from the KL minimised by variational Bayesian methods (64); the former averages with respect to the true distribution and is therefore usually intractable, while the latter averages with respect to the approximate distribution and is often tractable. Moreover, for exponential family approximations the former KL has a unique global optimum, while the latter usually has multiple local optima.

Since we cannot minimise (68) we can instead consider minimising the KL divergence between each true term and the corresponding approximate term. That is, for each  $i$ :

$$\min_{\tilde{f}_i(\boldsymbol{\theta})} \text{KL} \left( f_i(\boldsymbol{\theta}) \parallel \tilde{f}_i(\boldsymbol{\theta}) \right). \quad (69)$$

This will usually be much easier to do, but each such approximate term will result in some error. Multiplying all the approximate terms together will probably result in an unacceptably inaccurate approximation. On the plus side, this approach is non-iterative in that once each term is approximated they are simply multiplied to get a final answer.

The Expectation Propagation (EP) algorithm is an iterative procedure which is as easy as the naive approach in (69) but which results in a much more accurate approximation. At each step of EP, one of the terms is optimised in the context of all the other approximate terms, i.e. for each  $i$ :

$$\min_{\tilde{f}_i(\boldsymbol{\theta})} \text{KL} \left( f_i(\boldsymbol{\theta}) \prod_{j \neq i} \tilde{f}_j(\boldsymbol{\theta}) \parallel \tilde{f}_i(\boldsymbol{\theta}) \prod_{j \neq i} \tilde{f}_j(\boldsymbol{\theta}) \right). \quad (70)$$

Since the approximate terms depend on each other, this procedure is iterated. On the left hand side of the KL divergence the  $i^{\text{th}}$  exact term is incorporated into  $\prod_{j \neq i} \tilde{f}_j(\boldsymbol{\theta})$ , which is assumed to be in the exponential family. The right hand side is an exponential-family approximation to this whole product. The minimisation is done by matching the appropriate moments (expectations) of  $f_i(\boldsymbol{\theta}) \prod_{j \neq i} \tilde{f}_j(\boldsymbol{\theta})$ . The name ‘‘Expectation

Propagation” comes from the fact that each step corresponds to computing certain expectations, and the effect of these expectations is propagated to subsequent terms in the approximation. In fact, the messages in belief propagation can be derived as a particular form of EP where the approximating distribution is assumed to be a fully factored product of marginals over the variables in  $\theta$ , i.e.  $q(\theta) = \prod_k q_k(\theta_k)$  [53].

In its simplest form, the EP algorithm can be summarised as in Figure 5. Although the algorithm as described here often converges, each step of the algorithm is not in fact decreasing any objective function so there is no guarantee of convergence. Convergent forms of EP can be derived by making use of the EP energy function [51] although these may not be as fast and simple to implement as the algorithm in Figure 5.

```

Input  $f_0(\theta) \dots f_N(\theta)$ 
Initialise  $\tilde{f}_0(\theta) = f_0(\theta)$ ,  $\tilde{f}_i(\theta) = 1$  for  $i > 0$ ,  $q(\theta) = \prod_i \tilde{f}_i(\theta)$ 
repeat
  for  $i = 0 \dots N$  do
    Deletion:  $q_{\setminus i}(\theta) \leftarrow \frac{q(\theta)}{\tilde{f}_i(\theta)} = \prod_{j \neq i} \tilde{f}_j(\theta)$ 
    Projection:  $\tilde{f}_i^{\text{new}}(\theta) \leftarrow \arg \min_{f(\theta)} \text{KL}(f_i(\theta)q_{\setminus i}(\theta) \| f(\theta)q_{\setminus i}(\theta))$ 
    Inclusion:  $q(\theta) \leftarrow \tilde{f}_i^{\text{new}}(\theta) q_{\setminus i}(\theta)$ 
  end for
until convergence

```

Figure 5: The EP algorithm. Some variations are possible: this assumes that  $f_0$  is in the exponential family, and updates sequentially over  $i$  rather than randomly. The names for the steps (deletion, projection, inclusion) are not the same as in [53]

## 12 Conclusion

In this chapter, we have seen that unsupervised learning can be viewed from the perspective of statistical modelling. Statistics provides a coherent framework for learning from data and for reasoning under uncertainty. Many interesting statistical models used for unsupervised learning can be cast as latent variable models and graphical models. These types of models have played an important role in defining unsupervised learning systems for a variety of different kinds of data. Graphical models have also played an important unifying framework for thinking about the role of conditional independence in inference in models with many variables. While for certain models exact inference is computationally tractable, for most of the models in this chapter we have seen that exact inference involves intractable sums and integrals. Thus, the study of unsupervised learning has lead us into focusing on ways of approximating high dimensional sums and integrals. We have reviewed many of the principal approximations, although of course in the limited space of this chapter one cannot hope to have a comprehensive review of approximation methods.

There are many interesting and relevant topics we did not get a chance to cover in this review of unsupervised learning. One of these is the interplay of unsupervised and supervised learning, in the form of semi-supervised learning. *Semi-supervised learning* refers to learning problems in which there is a small amount of labelled data and a large amount of unlabelled data. These problems are very natural, especially in domains where collecting data can be cheap (i.e. the internet) but labelling it can be expensive or time consuming. The key question in semi-supervised learning is how the data distribution from the unlabelled data should influence the supervised learning problem [71]. Many of the approaches to this problem attempt to infer a manifold, graph structure, or tree-structure from the unlabelled data and use spread in this structure to determine how labels will generalise to new unlabelled points [74, 85, 8, 42].

Another area of great interest which we did not have the space to cover are *nonparametric models*. The basic assumption of parametric statistical models is that the model is defined using a finite number of

parameters. The number of parameters is assumed fixed regardless of the number of data points. Thus the parameters provide a finite summary of the data. In nonparametric models, the number of “parameters” in the model is allowed to grow with the size of the data set. With more data, the model becomes more complex, with no a-priori limit on the complexity of the model. For this reason nonparametric models are also sometimes called *infinite models*. An important example of this are *infinite mixture models*, more formally known as *Dirichlet Process mixtures* [3, 18]. These correspond to mixture models (Section 2.4) where the number of components is assumed to be infinite. Inference can be done in these models using MCMC methods [17, 60, 65], variational methods [12], or the EP algorithm [54]. Just as hidden Markov models can be seen as an extension of finite mixture models to model time series data, it is possible to extend infinite mixture models to hidden Markov models with infinitely many states [7]. Infinite models based on Dirichlet processes have also been generalised to be hierarchical in several different ways [61, 75]. Bayesian inference in nonparametric models is one of the most active areas of research in unsupervised learning, and there still remain many open problems.

As we have seen, the field of unsupervised learning can be understood formally within the framework of information theory and statistics. However, it is important not to lose sight of the tremendous influence ideas from neuroscience and psychology have had on the field. Many of the models we have reviewed here started life as models of brain function. These models were inspired by the brain’s ability to extract statistical patterns from sensory data and to recognise complex visual scenes, sounds, and odours. Unsupervised learning theory and algorithms still have a long way to go to mimic some of the learning abilities of biological brains. As the boundaries of unsupervised learning get pushed forward, we will hopefully not only benefit from better learning machines and also improve our understanding of how the brain learns.

## References

- [1] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski. A learning algorithm for Boltzmann machines. *Cognitive Science*, 9:147–169, 1985.
- [2] B. D. O. Anderson and J. B. Moore. *Optimal Filtering*. Prentice-Hall, Englewood Cliffs, NJ, 1979.
- [3] C. E. Antoniak. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Annals of Statistics*, 2(6):1152–1174, 1974.
- [4] S. Arnborg, D. G. Corneil, and A. Proskurowski. Complexity of finding embeddings in a k-tree. *SIAM Journal of Algebraic and Discrete Methods*, 8:277–284, 1987.
- [5] H. Attias. Inferring parameters and structure of latent variable models by variational Bayes. In *Proc. 15th Conf. on Uncertainty in Artificial Intelligence*, 1999.
- [6] M. J. Beal and Z. Ghahramani. The variational Bayesian EM algorithm for incomplete data: With application to scoring graphical model structures. In J. M. Bernardo, A. P. Dawid, J. O. Berger, M. West, D. Heckerman, and M. J. Bayarri, editors, *Bayesian Statistics 7*. Oxford University Press, 2003.
- [7] M.J. Beal, Z. Ghahramani, and C.E. Rasmussen. The infinite hidden Markov model. In *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001.
- [8] M. Belkin and P. Niyogi. Semi-supervised learning on Riemannian manifolds. *Machine Learning*, 56(1–3):209–239, 2004.
- [9] C. Berrou, A. Glavieux, and P. Thitimajshima. Near shannon limit error-correcting coding and decoding: Turbo-codes (1). In *Proc. ICC ’93*, pages 1064–1070, 1993.
- [10] J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Ser. B*, 6:192–236, 1974.
- [11] C. M. Bishop, M. Svensen, and C. K. I. Williams. GTM: The generative topographic mapping. *Neural Computation*, 10(1):215–234, 1998.

- [12] D. Blei and M. I. Jordan. Variational methods for the Dirichlet process. In *Proceedings of the 21st International Conference on Machine Learning*, 2004.
- [13] G. F. Cooper. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2-3):393–405, 1990.
- [14] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Society Series B*, 39:1–38, 1977.
- [15] A. Doucet, J. F. G. de Freitas, and N.J. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, 2000.
- [16] G. Elidan, N. Lotner, N. Friedman, and D. Koller. Discovering hidden variables: A structure-based approach. In *Advances in Neural Information Processing Systems (NIPS)*, 2001.
- [17] M. D. Escobar and M. West. Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, 90(430):577–588, 1995.
- [18] T. S. Ferguson. Bayesian density estimation by mixtures of normal distributions. In *Recent Advances in Statistics*, pages 287–302, New York, 1983. Academic Press.
- [19] N. Friedman. The Bayesian structural EM algorithm. In *Proc. Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI '98)*, San Francisco, CA, 1998. Morgan Kaufmann.
- [20] N. Friedman and D. Koller. Being Bayesian about network structure: A Bayesian approach to structure discovery in Bayesian networks. *Machine Learning*, 50:95–126, 2003.
- [21] R. G. Gallager. *Low-Density Parity-Check Codes*. MIT Press, Cambridge, MA, 1963.
- [22] Z. Ghahramani and M. J. Beal. Propagation algorithms for variational Bayesian learning. In *Advances in Neural Information Processing Systems 13*, Cambridge, MA, 2001. MIT Press.
- [23] Z. Ghahramani and G. E. Hinton. The EM algorithm for mixtures of factor analyzers. *University of Toronto, Technical Report CRG-TR-96-1*, 1996.
- [24] Z. Ghahramani and M. I. Jordan. Factorial hidden Markov models. *Machine Learning*, 29:245–273, 1997.
- [25] Z. Ghahramani and S. T. Roweis. Learning nonlinear dynamical systems using an EM algorithm. In *NIPS 11*, pages 431–437, 1999.
- [26] F. Girosi, M. Jones, and T. Poggio. Regularization theory and neural networks architectures. *Neural Computation*, 7(2):219–269, 1995.
- [27] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. A novel approach to nonlinear/non-Gaussian Bayesian state space estimation. *IEE Proceedings F: Radar and Signal Processing*, 140:107–113, 1993.
- [28] P. J. Green. Penalized likelihood. In *Encyclopedia of Statistical Sciences, Update Volume 2*, 1998.
- [29] J. E. Handschin and D. Q. Mayne. Monte Carlo techniques to estimate the conditional expectation in multi-stage non-linear filtering. *International Journal of Control*, 9(5):547–559, 1969.
- [30] T. Hastie and W. Stuetzle. Principle curves. *Journal of the American Statistical Association*, 84(406):502–516, 1989.
- [31] D. Heckerman. A tutorial on learning with Bayesian networks. Technical Report MSR-TR-95-06, Microsoft Research, 1996.
- [32] D. Heckerman and D. M. Chickering. A comparison of scientific and engineering criteria for Bayesian model selection, 1996.

- [33] G. E. Hinton, P. Dayan, B. J. Frey, and R. M. Neal. The wake-sleep algorithm for unsupervised neural networks. *Science*, 268:1158–1161, 1995.
- [34] G. E. Hinton, P. Dayan, and M. Revow. Modeling the manifolds of images of handwritten digits. *IEEE Trans. Neural Networks*, 8(1):65–74, 1997.
- [35] M. Isard and A. Blake. Condensation – conditional density propagation for visual tracking, 1998.
- [36] E. T. Jaynes. *Probability Theory: The Logic of Science (Edited by G. Larry Bretthorst)*. Cambridge University Press, 2003.
- [37] W.H. Jefferys and J.O. Berger. Ockham’s razor and Bayesian analysis. *American Scientist*, 80:64–72, 1992.
- [38] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An introduction to variational methods in graphical models. *Machine Learning*, 37:183–233, 1999.
- [39] S. J. Julier and J. K. Uhlmann. A new extension of the Kalman filter to nonlinear systems. In *Int. Symp. Aerospace/Defense Sensing, Simulation and Controls*, 1997.
- [40] K. Kanazawa, D. Koller, and S. J. Russell. Stochastic simulation algorithms for dynamic probabilistic networks. In P. Besnard and S. Hanks, editors, *Uncertainty in Artificial Intelligence. Proceedings of the Eleventh Conference.*, pages 346–351. Morgan Kaufmann Publishers, San Francisco, CA, 1995.
- [41] Y. Karklin and M. S. Lewicki. Learning higher-order structures in natural images. *Network: Computation in Neural Systems*, 14:483–499, 2003.
- [42] C. Kemp, T. L. Griffiths, S. Stromsten, and J. B. Tenenbaum. Semi-supervised learning with trees. In *NIPS 16*, 2004.
- [43] G. Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *J. of Computational and Graphical Statistics*, 5:1–25, 1996.
- [44] U. Kjaerulff. Triangulation of graphs—algorithms giving small total state space, 1990.
- [45] J. B. Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29:1–27 and 115–129, 1964.
- [46] S. L. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *J. Royal Statistical Society B*, pages 157–224, 1988.
- [47] D. J. C. MacKay. Probable networks and plausible predictions—a review of practical Bayesian methods for supervised neural networks. *Network: Computation in Neural Systems*, 6:469–505, 1995.
- [48] D. J. C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
- [49] D. J. C. MacKay and R. M. Neal. Good error-correcting codes based on very sparse matrices. *IEEE Transactions on Information Theory*, 45(2):399–431, 1999.
- [50] R. J. McEliece, D. J. C. MacKay, and J.-F. Cheng. Turbo decoding as an instance of Pearl’s ‘Belief Propagation’ algorithm. *IEEE Journal on Selected Areas in Communications*, 16(2):140–152, February 1998.
- [51] T. P. Minka. The EP energy function and minimization schemes. Technical report, 2001.
- [52] T. P. Minka. Expectation propagation for approximate Bayesian inference. In *Uncertainty in Artificial Intelligence: Proceedings of the Seventeenth Conference (UAI-2001)*, pages 362–369, San Francisco, CA, 2001. Morgan Kaufmann Publishers.
- [53] T. P. Minka. *A family of algorithms for approximate Bayesian inference*. PhD thesis, MIT, 2001.

- [54] T. P. Minka and Z. Ghahramani. Expectation propagation for infinite mixtures. Technical report, Presented at NIPS'03 Workshop on Nonparametric Bayesian Methods and Infinite Models, 2003.
- [55] A. Moore and W-K. Wong. Optimal reinsertion: A new search operator for accelerated and more accurate Bayesian network structure learning. In T. Fawcett and N. Mishra, editors, *Proceedings of the 20th International Conference on Machine Learning (ICML '03)*, pages 552–559, Menlo Park, California, August 2003. AAAI Press.
- [56] K. P. Murphy. *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, UC Berkeley, Computer Science Division, 2002.
- [57] I. Murray and Z. Ghahramani. Bayesian learning in undirected graphical models: Approximate MCMC algorithms. In *Proceedings of UAI*, 2004.
- [58] R. M. Neal. Connectionist learning of belief networks. *Artificial Intelligence*, 56:71–113, 1992.
- [59] R. M. Neal. Probabilistic inference using Markov chain Monte Carlo methods. Technical report, Department of Computer Science, University of Toronto, September 1993.
- [60] R. M. Neal. Markov chain sampling methods for Dirichlet process mixture models. *Journal of Computational and Graphical Statistics*, 9:249–265, 2000.
- [61] R. M. Neal. Density modeling and clustering using Dirichlet diffusion trees. In J. M. Bernardo et al., editor, *Bayesian Statistics 7*, pages 619–629, 2003.
- [62] R. M. Neal, M. J. Beal, and S. T. Roweis. Inferring state sequences for non-linear systems with embedded hidden Markov models. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, Cambridge, MA, 2004. MIT Press.
- [63] R. M. Neal and G. E. Hinton. A new view of the EM algorithm that justifies incremental, sparse, and other variants. In M. I. Jordan, editor, *Learning in Graphical Models*. Kluwer Academic Press, 1998.
- [64] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- [65] C. E. Rasmussen. The infinite Gaussian mixture model. In *Adv. Neur. Inf. Proc. Sys. 12*, pages 554–560, 2000.
- [66] C. E. Rasmussen and Z. Ghahramani. Occam’s razor. In *Advances in Neural Information Processing Systems 13*, Cambridge, MA, 2001. MIT Press.
- [67] S. T. Roweis. EM algorithms for PCA and SPCA. In M. I. Jordan, M. J. Kearns, and S. A. Solla, editors, *Advances in Neural Information Processing Systems*, volume 10. The MIT Press, 1998.
- [68] S. T. Roweis and Z. Ghahramani. A unifying review of linear Gaussian models. *Neural Computation*, 11(2):305–345, 1999.
- [69] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, 2000.
- [70] R. Salakhutdinov, S. T. Roweis, and Z. Ghahramani. Optimization with EM and Expectation-Conjugate-Gradient. In *International Conference on Machine Learning (ICML-2003)*, pages 672–679, 2003.
- [71] M. Seeger. Learning with labeled and unlabeled data. Technical report, University of Edinburgh, 2001.
- [72] R. N. Shepard. The analysis of proximities: multidimensional scaling with an unknown distance function i and ii. *Psychometrika*, 27:125–139 and 219–246, 1962.
- [73] R. H. Shumway and D. S. Stoffer. An approach to time series smoothing and forecasting using the EM algorithm. *J. Time Series Analysis*, 3(4):253–264, 1982.

- [74] M. Szummer and T. S. Jaakkola. Partially labeled classification with Markov random walks. In *NIPS*, 2001.
- [75] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei. Hierarchical Dirichlet processes. Technical Report 653, Department of Statistics, University of California at Berkeley, 2004.
- [76] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- [77] M. E. Tipping and C. M. Bishop. Mixtures of probabilistic principal component analyzers. *Neural Computation*, 11(2):443–482, 1999.
- [78] M. E. Tipping and C. M. Bishop. Probabilistic principal component analysis. *Journal of the Royal Statistical Society, Series B*, 61(3):611–622, 1999.
- [79] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. Technical Report 649, UC Berkeley, Dept. of Statistics, 2003.
- [80] E. A. Wan, R. van der Merwe, and A. T. Nelson. Dual estimation and the unscented transformation. In *NIPS 12*, pages 666–672, 2000.
- [81] Y. Weiss. Correctness of local probability propagation in graphical models with loops. *Neural Computation*, 12(1):1–41, 2000.
- [82] Y. Weiss and W. T. Freeman. On the optimality of solutions of the max-product belief-propagation algorithm in arbitrary graphs. *IEEE Transactions on Information Theory, Special Issue on Codes on Graphs and Iterative Algorithms*, 47, 2001.
- [83] J. Winn. *Variational Message Passing and its Applications*. PhD thesis, Department of Physics, University of Cambridge, 2003.
- [84] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Generalized belief propagation. In *NIPS 13*, Cambridge, MA, 2001. MIT Press.
- [85] X. Zhu, Z. Ghahramani, and J. Lafferty. Semi-supervised learning using Gaussian fields and harmonic functions. In *The Twentieth International Conference on Machine Learning (ICML-2003)*, 2003.