



Cite this article: Ghahramani Z. 2013
Bayesian non-parametrics and the
probabilistic approach to modelling. *Phil Trans
R Soc A* 371: 20110553.
<http://dx.doi.org/10.1098/rsta.2011.0553>

One contribution of 17 to a Discussion Meeting
Issue ‘Signal processing and inference for the
physical sciences’.

Subject Areas:

pattern recognition, computer modelling and
simulation, statistics, mathematical modelling

Keywords:

probabilistic modelling, Bayesian statistics,
non-parametrics, machine learning

Author for correspondence:

Zoubin Ghahramani
e-mail: zoubin@eng.cam.ac.uk

Bayesian non-parametrics and the probabilistic approach to modelling

Zoubin Ghahramani

Department of Engineering, University of Cambridge,
Cambridge CB2 1PZ, UK

Modelling is fundamental to many fields of science and engineering. A model can be thought of as a representation of possible data one could predict from a system. The probabilistic approach to modelling uses probability theory to express all aspects of uncertainty in the model. The probabilistic approach is synonymous with Bayesian modelling, which simply uses the rules of probability theory in order to make predictions, compare alternative models, and learn model parameters and structure from data. This simple and elegant framework is most powerful when coupled with flexible probabilistic models. Flexibility is achieved through the use of Bayesian non-parametrics. This article provides an overview of probabilistic modelling and an accessible survey of some of the main tools in Bayesian non-parametrics. The survey covers the use of Bayesian non-parametrics for modelling unknown functions, density estimation, clustering, time-series modelling, and representing sparsity, hierarchies, and covariance structure. More specifically, it gives brief non-technical overviews of Gaussian processes, Dirichlet processes, infinite hidden Markov models, Indian buffet processes, Kingman’s coalescent, Dirichlet diffusion trees and Wishart processes.

1. Introduction

Modelling is central to the sciences. Models allow one to make predictions, to understand phenomena, and to quantify, compare and falsify hypotheses.

Modelling is also at the core of intelligence. Both artificial and biological systems that exhibit intelligence must be able to make predictions, anticipate outcomes of

© 2012 The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License <http://creativecommons.org/licenses/by/3.0/>, which permits unrestricted use, provided the original author and source are credited.

their actions and update their ability to make predictions in light of new data. It is hard to imagine how a system could do this without building models of the environment that the system interacts with. It is thus not surprising that many theories in cognitive science are based around the idea of building internal models [1–3].

A model is simply a compact representation of possible data one could observe. As such, it may be more interpretable than the observed data itself, providing a useful representation of data. A model must be able to make forecasts of possible future data; otherwise, it seems impossible to falsify a model in light of new data. I will use the term *forecast* in a very general way to refer to the process of making any claims about unobserved data on the basis of observed data; I will also use *predict* interchangeably with forecast. For all non-trivial phenomena, forecasts have to include some representation of the forecasting uncertainty. Deterministic forecasts (e.g. tomorrow's high temperature *will be* 17°C) are too brittle and therefore easy to falsify.

Ideally, a model should also be adaptive. By *adaptive*, I mean that the forecasts of the model should change, depending on the data observed so far. Such adaptation, or learning, should hopefully have the effect of making the model's forecasts be better aligned with actual data. For example, if the forecasts are in the form of probability distributions, adaptation should have the effect that the forecast probability assigned to what actually happens should increase after the model has seen more data, although this cannot be generally guaranteed.

It is clear that all forecasts need to represent uncertainty. Observable data are generally corrupted by noise in the measurement process, and this noise needs to be incorporated in the forecast uncertainty. But uncertainty lurks at many other levels in any model. The amount of noise in the measurement process may itself be unknown. The model may have a number of parameters that are unknown. The structure of the model itself may be uncertain, or there may be multiple plausible competing models for the data. The forecasting system should ideally produce forecasts that incorporate all reasonable sources of uncertainty. As we will see, the Bayesian framework provides a natural and coherent approach for representing and manipulating all forms of uncertainty in modelling.

2. The Bayesian framework

The fundamental idea in Bayesian modelling is to use the mathematics of probability theory to represent and manipulate all forms of uncertainty in the model. This is a surprisingly simple yet powerful idea.

The good news is that there are only two rules of probability theory one needs to remember: the sum rule and the product rule.¹ Consider a pair of random variables x and y taking on values in some spaces \mathcal{X} and \mathcal{Y} , respectively. The sum rule states that if I know the joint probability of two random variables, x and y , I can obtain the marginal probability of x by summing over all possible values of y ,

$$P(x) = \sum_{y \in \mathcal{Y}} P(x, y).$$

If y is continuous, we simply replace the sum with an integral. If I have a model for the joint probability distribution of the high temperature in London and Cambridge, I can obtain the marginal distribution for the high temperature in Cambridge summing out London's temperature.

The product rule states that the joint probability of x and y can be decomposed into the product of the marginal probability of x and the conditional probability of y given x (or the other way around),

$$P(x, y) = P(x)P(y|x) = P(y)P(x|y).$$

¹With apologies to probabilists and statisticians, my presentation and notation will eschew formality in favour of didactic clarity.

Combining the sum and product rule and rearranging a bit, we obtain Bayes rule as a corollary,

$$P(y|x) = \frac{P(x|y)P(y)}{P(x)} = \frac{P(x|y)P(y)}{\sum_{y' \in \mathcal{Y}} P(x|y')P(y')}.$$

Let us now use these equations in a canonical modelling context. We will use \mathcal{D} to represent the observable data. Often we talk about the data being a *set* consisting of a number of ‘data points’ or measurements, $\mathcal{D} = \{x_1, \dots, x_N\}$, but this is in no way a necessary assumption. For example, the data could be a single image, an observed graph or an ordered *sequence* of measurements rather than a set. Our model will be indexed by m and we may want to consider multiple alternative models, e.g. m, m' , etc. Each model usually has a number of free parameters, which we will denote with θ , which can be a vector if needed.

First we need to make sure that the model m is well defined, in the sense that it can predict or forecast data. As previously discussed, we use probability theory to represent the model forecasts. For any given setting of the model parameters, the model must be able to produce a forecast of the form

$$P(\mathcal{D}|\theta, m).$$

This probability of the data as a function of the parameters is the *likelihood* of the parameters. With the likelihood, for any given setting of the parameters, we are in the position of making forecasts. However, the model m is not fully defined until we specify a ‘range’ for the parameter values, θ . A linear regression model that states that the slope can take values between -1 and $+1$ is a very different model from one that specifies that the slope can take values between -100 and $+100$. In fact, to fully specify a model, we need to do a bit more than specify the range of parameters, we need to define a distribution over this range. Only then will our model m be able to make forecasts. We see this by writing the forecast probability of the model using the sum and product rules,

$$P(\mathcal{D}|m) \stackrel{\text{sum}}{=} \int P(\mathcal{D}, \theta|m) d\theta \stackrel{\text{prod}}{=} \int P(\mathcal{D}|\theta, m)P(\theta|m) d\theta. \quad (2.1)$$

The expression $P(\mathcal{D}|m)$ is variously called the *marginal likelihood*, *model evidence* or *integrated likelihood*. The *prior* over the parameters, $P(\theta|m)$, plays the role of specifying the range as described above (e.g. it could be uniform on $[-1, +1]$) in the form of a distribution over the allowable values of the parameters. Without a prior, our model is not well defined: we cannot generate or forecast data until we know how to choose values for θ .² Once the prior and likelihood are defined, and only then, is the model m fully specified in the sense that it can *generate possible datasets*.

People often object to Bayesian methods on the basis that it forces one to define priors on the parameters. This, in my opinion, is completely misguided. All models make assumptions; without assumptions, it is impossible to make any forecasts or predictions from observed data. The first stage of the Bayesian modelling framework is to explicitly state all assumptions using the language of probability theory. Specifying both the prior and likelihood is a necessary requirement so that the model is well defined. In fact, the distinction between prior and likelihood is arbitrary³; both are essential parts of the *model*.

People also object to the use of priors on the parameters on the grounds that they do not think of the parameters as being ‘random’ variables. For example, if one is estimating the mass of a planet from astronomical data, the mass of the planet is not ‘random’ in the colloquial sense of the word.⁴ This is a misunderstanding of the semantics of probabilities in Bayesian

²Of course, our model may have a fixed value of θ , e.g. 7.213, which corresponds to a delta function prior.

³Consider, for example, a model that defines the probability of the data (i.e. the likelihood) using a student *t*-distribution. For a single data point, this can be equivalently expressed as a model with a Gaussian likelihood and a Gamma prior on the precision of that Gaussian.

⁴I do not like the term ‘random variable’ because it suggests that there is some true source of randomness driving some process. To give another example, people would naturally and understandably flinch at the idea of calling something like the millionth digit of π or the number of people alive today a random variable. The notion that Bayesians consider true parameters to be ‘random variables’ makes the framework less intuitive than it actually is. It is much more natural to use the term *uncertain*. Clearly I can be uncertain about the millionth digit of π , even though its value is not random at all. My uncertainty about this digit will change after I look it up on the web, but presumably will not go to zero as there is always the chance of someone having made a mistake somewhere.

modelling. Probabilities are used to represent our *uncertainty* about unknown quantities. They are just as good at modelling uncertainty for repeatable experiments such as the outcome of a roll of a die, as they are for modelling uncertainty in the mass of a planet. Incidentally, both forms of uncertainty are fundamentally subjective; one's uncertainty about a roll of a die depends on one's knowledge of the exact initial conditions of the roll, in just the same way as the uncertainty about the mass of the planet depends on knowledge of the observational data on the planet's orbit.

Finally, people trained in the sciences are uncomfortable with the very notion of subjectivity in data analysis and modelling. This again is deeply misguided: all models involve assumptions, and all conclusions drawn from data analysis are conditional on assumptions. The probabilistic framework *forces* the scientist to be completely transparent about the assumptions, by expressing all assumptions as distributions on unknown quantities. These assumptions can then be easily contested, and the same data can be reanalysed under different modelling assumptions (and priors). The fact that conclusions could change depending on the assumptions is essential to good scientific practice. Fortunately, given enough data, the effect of the prior is generally overcome by the likelihood, and posterior conclusions will converge [4–6]. This is directly analogous to the progress of science, where of many possible hypotheses, only the ones consistent with the data will survive. Bayesian modelling is *subjective* but not *arbitrary*: given a full specification of the model, and the data, there is only one way to reason about any quantity of interest.

Modelling thus becomes a very simple procedure:

- write down your assumptions (possible models, parameters, noise processes, etc.), representing all forms of uncertainty using the language of probability theory and
- given the data, use probability theory to make inferences about any unknown quantities in your model, or to make predictions from the model.

This process lends itself very naturally to a sequential processing of data. Your posterior after observing some data \mathcal{D}_{old} , $P(\theta|\mathcal{D}_{\text{old}}, m)$ is your prior before observing new data, \mathcal{D}_{new} ,

$$P(\theta|\mathcal{D}_{\text{new}}, \mathcal{D}_{\text{old}}, m) \propto P(\mathcal{D}_{\text{new}}|\mathcal{D}_{\text{old}}, \theta, m)P(\theta|\mathcal{D}_{\text{old}}, m).$$

The sum and product rule also tell us how to make predictions from a model. Consider predicting some unknown quantity x (e.g. the next data point) given observed data \mathcal{D} and model m ,

$$P(x|\mathcal{D}, m) = \int P(x|\theta, \mathcal{D}, m)P(\theta|\mathcal{D}, m) d\theta. \quad (2.2)$$

This is a very intuitively satisfying expression which tells us that our predictions or forecasts are a weighted average of the forecasts from different parameter values, weighted by the posterior probability of each parameter value given the data observed so far. For parametric models, this simplifies because *given* the parameters forecasts are independent of the observed data: $P(x|\theta, \mathcal{D}, m) = P(x|\theta, m)$. We will revisit this point as we discuss parametric versus non-parametric models in §3. If we are considering a number of models m, m' , etc., then by the sum and product rules, our forecasts are an average over models weighted by their posteriors.

The probabilistic modelling framework also provides intuitive answers to problems in model comparison. Assuming a set of competing probabilistic models \mathcal{M} , given some observed data, we can evaluate the posterior probability of a particular model m ,

$$P(m|\mathcal{D}) = \frac{P(\mathcal{D}|m)P(m)}{P(\mathcal{D})} = \frac{P(\mathcal{D}|m)P(m)}{\sum_{m' \in \mathcal{M}} P(\mathcal{D}|m')P(m')}.$$

Note the prominent role played by the marginal likelihood, $P(\mathcal{D}|m)$.




Figure 1. Marginal likelihoods, Occam's razor and overfitting: consider modelling a function $y = f(x) + \epsilon$ describing the relationship between some input variable x , and some output or response variable y . (a) The red dots in the plots on the left-hand side are a dataset of eight (x, y) pairs of points. There are many possible f that could model this given data. Let us consider polynomials of different order, ranging from constant ($M = 0$), linear ($M = 1$), quadratic ($M = 2$), etc., to seventh order ($M = 7$). The blue curves depict maximum-likelihood polynomials fit to the data under Gaussian noise assumptions (i.e. least-squares fits). Clearly, the $M = 7$ polynomial can fit the data perfectly, but it seems to be overfitting wildly, predicting that the function will shoot off up or down between neighbouring observed data points. By contrast, the constant polynomial may be underfitting, in the sense that it might not pick up some of the structure in the data. The green curves indicate 20 random samples from the Bayesian posterior of polynomials of different order given this data. A Gaussian prior was used for the coefficients, and an inverse gamma prior on the noise variance (these conjugate choices mean that the posterior can be analytically integrated). The samples show that there is considerable posterior uncertainty given the data, and also that the maximum-likelihood estimate can be very different from the typical sample from the posterior. (b) The normalized model evidence or marginal likelihood for this model is plotted as a function of the model order, $P(Y|M)$, where the dataset Y are the eight observed output y values. Note that given the data, model orders ranging from $M = 0$ to $M = 3$ have considerably higher marginal likelihood than other model orders, which seems plausible given the data. Higher-order models, $M > 3$, have relatively much smaller marginal likelihood, which is not visible on this scale. The decrease in marginal likelihood as a function of model order is a reflection of the automatic Occam razor that results from Bayesian marginalization.

Importantly, this marginal likelihood captures a preference for simpler models known as Bayesian Occam's razor [7–9]. Consider a set of nested models, for example, different order polynomials (e.g. constant, linear, quadratic, cubic, etc.) used to fit some regression relationship (figure 1). Clearly, a higher-order model such as the cubic polynomial is strictly more complex than a lower-order model such as the linear polynomial. Model fitting procedures based on optimization (such as maximum-likelihood methods or penalized-likelihood methods) need to take great care not to *overfit* the data by fitting the parameters of an overly complex model to a relatively small dataset. Overfitting is not a problem for fully Bayesian methods, as there is no 'fitting' of the model to the data. We have only the sum rule and the product rule to work with, there is no 'optimize' rule in probability theory. A more complex model, say one that has more parameters, simply spreads its predictive probability mass over more possible datasets than a simpler model. If all models are specified as probability distributions over datasets, because probability distributions have to sum to one, all models have the same amount of probability mass to spread over possible data (figure 2). Given a particular dataset, it therefore becomes possible to reject both models that are too simple or too complex simply by using the rules of probability theory.

This approach to Bayesian model comparison can be used to solve a vast range of problems in learning the structure of complex models. For example, it has been used to learn the number of clusters in a mixture model [12,13], finding relevant variables or features in a




Figure 2. An illustration of Occam's razor. Consider all possible datasets of some fixed size n . Competing probabilistic models correspond to alternative distributions over the datasets. Here, we have illustrated three possible models that spread their probability mass in different ways over these possible datasets. A *complex* model (shown in blue) spreads its mass over many more possible datasets, whereas a *simple* model (shown in green) concentrates its mass on a smaller fraction of possible data. Because probabilities have to sum to one, the complex model spreads its mass at the cost of not being able to model simple datasets as well as a simple model—this normalization is what results in an automatic Occam razor. Given any particular dataset, here indicated by the dotted line, we can use the marginal likelihood to reject both overly simple models, and overly complex models. This figure is inspired by a figure from Mackay [10], and an actual realization of this figure on a toy classification problem is discussed in Murray & Ghahramani [11].

prediction problem [14,15], discovering the number of states in a hidden Markov model (HMM) [16] and learning the dependency structure between variables in a probabilistic graphical model [17].

The above approach to model comparison relies on the ability to enumerate a set of models \mathcal{M} to be compared. This is often reasonable in scientific settings where there are a number of competing hypotheses to explain some phenomenon. Bayesian Occam's razor ensures that overly complex models are adequately penalized when doing model comparison. However, the complexity of real-world phenomena often requires us to consider complex models that are flexible enough to capture the structure in real data. Flexible models are not only more realistic, but will also generally result in more reasonable forecasts than simpler models. Bayesian non-parametrics provides a natural framework for defining flexible models.

3. Non-parametric models

The best way to understand non-parametric models is by first reminding ourselves of what parametric models are. A parametric model has some finite set of parameters θ . Given these parameters, future predictions, x , are independent of the observed data,

$$p(x|\theta, \mathcal{D}) = p(x|\theta).$$

The parameters therefore capture everything there is to know about the data that is relevant for predicting future data.

We can think of all models as *information channels* from past data \mathcal{D} to future predictions x . The parameters in a parametric model constitute a bottleneck in this information channel. The complexity of the model, and the capacity of the channel, is bounded, even if the amount of observed data becomes unbounded. Parametric models are therefore not generally very flexible.

By contrast, a *non-parametric* model assumes that the data distribution cannot be defined in terms of such a finite set of parameters. However, often we can think of non-parametric models as being defined in terms of an infinite-dimensional θ . More formally, the infinite-dimensional

θ is often represented as a function. The term *non-parametric* is therefore a bit of a misnomer; it is not that non-parametric models do not have parameters; in fact, they have infinitely many parameters. Because of this, non-parametric models cannot be explicitly represented in terms of their parameters.

From the information channel viewpoint, we have removed the bottleneck. The amount of information that θ can capture about the data \mathcal{D} grows as the amount of data grows. This makes non-parametric models more flexible than parametric models.⁵

There is another way to view the difference between parametric and non-parametric models. Predictions from a parametric model are explicitly and compactly summarized through the parameters θ , $P(x|\theta)$. Non-parametric models, by contrast, cannot be summarized in this way. Because of this, predictions from a non-parametric are necessarily *memory-based*, $P(x|\mathcal{D})$; to make predictions, we need to store or remember a growing amount of information about the training data, \mathcal{D} .

Non-parametric models are inextricably tied to the notion of *exchangeability*. A sequence is exchangeable if its joint distribution is invariant under arbitrary permutation of the indices. Consider modelling a dataset $\{x_1, \dots, x_N\}$ under the assumption that the ordering of the elements is uninformative. This dataset may be a collection of documents, images or any other sort of object collected in a manner such that either the ordering is irrelevant or completely unknown.⁶ De Finetti's theorem states that a sequence is exchangeable if and only if there exists some θ such that the elements x_n are independently and identically distributed (iid) from some unknown distribution indexed by θ [19]. Importantly, θ may need to be infinite dimensional as it needs to index the space of probability measures (non-negative functions that normalize to one).

The consequence of de Finetti's theorem is that if we want to model exchangeable data in full generality, we need to consider putting distributions on unknown probability measures.

Distributions on measures, functions and other infinite-dimensional objects are thus central to Bayesian non-parametric modelling. Many of these distributions are infinite-dimensional versions of their finite-dimensional counterparts, and in fact a good strategy for deriving Bayesian non-parametric models is to start from a parametric model and 'take the infinite limit' [20]. Distributions on infinite-dimensional objects are the main subject of study in stochastic process theory, and therefore much of the terminology used in Bayesian non-parametrics is borrowed from this field.⁷

Two of the classical building blocks for Bayesian non-parametric models are the Gaussian process (GP) and the Dirichlet process (DP). I will give an overview of these models in §§4 and 5, with an emphasis on their applications to general problems in machine learning and statistics, including regression, classification, clustering and density estimation (these problems will also be described in those sections). I will also cover newer building blocks that can represent non-parametric distributions over time series, sparse matrices, hierarchies, and covariances (§§6–9). In order to give a broad overview, I will necessarily have to avoid going in much depth into any of the specific topics, providing instead pointers to the relevant literature. Admittedly, my overview is based on my personal view of the field, and is therefore biased towards areas of Bayesian non-parametrics to which my colleagues and I have contributed, and misses other important areas.

⁵To the best of my knowledge, this information channel view of non-parametric modelling is not explicitly articulated in the literature, although to be sure, there are a number of strong links between Bayesian inference and information theory [8,18]. These ideas could certainly be further formalized by explicitly attempting to measure or estimate the channel capacity of different models. This would help provide a further definition of the *complexity* of a model, measured in bits per data point, which in many ways is more satisfactory than definitions that rely on counting parameters.

⁶There are generalizations of exchangeability to handle time series and arrays.

⁷The word 'process' usually evokes the idea of something evolving temporally, but it is important to note that stochastic processes can be indexed by time, space or any other index space. Many of the uses of stochastic processes in Bayesian non-parametrics do not correspond to indexing by time.

4. Modelling functions, classification and regression: Gaussian processes

GPs are a distribution over functions that can be used in numerous contexts where one's model requires one to represent an unknown function [21]. One-dimensional GPs indexed by time are familiar to many fields: Brownian motion, Wiener processes, Ornstein–Uhlenbeck processes, linear Gaussian state-space models, and many random walks are all examples of GPs. For historical reasons, GPs are also sometimes associated with spatial statistics, for example, modelling temperature as a function of spatial location. Within machine learning, GPs are often used for nonlinear regression and classification.

Consider the following simple model for nonlinear regression:

$$y_n = f(x_n) + \epsilon_n.$$

Here, we wish to model the relationship between an input (or covariate) x and an output (or response variable) y . The subscript n indexes data points in our dataset \mathcal{D} , ϵ_n is some additive noise and crucially, f is the unknown regression function we wish to learn about from data.

Assume we have a distribution over functions f and we evaluate the marginal distribution it assigns to the vector $\mathbf{f} = (f(x_1), \dots, f(x_N))$. If, for any choice of input points, (x_1, \dots, x_N) , the marginal distribution over \mathbf{f} is multi-variate Gaussian, then the distribution over the function f is said to be a GP. In other words, a GP is an infinite-dimensional generalization of the multi-variate Gaussian. Analogous to the Gaussian, a GP is parametrized by a mean function, and a covariance function.

The application of GPs to regression is straightforward. Starting with a prior on functions $p(f)$, we condition on the data to obtain a posterior $p(f|\mathcal{D})$. When the noise ϵ is assumed to be Gaussian, all computations reduce to operations on N -dimensional Gaussians. The infinitely many other dimensions of f can be marginalized out analytically.

Classification problems correspond to predicting categorical response or output variables, e.g. $y \in \{\text{cat}, \text{dog}\}$. GP regression can be modified to do classification simply by introducing a link function that maps the real values $f(x)$ into probabilities over the classes. Computing $p(f|\mathcal{D})$ exactly becomes intractable, but many good methods exist for approximating the required integrals (§10).

5. Density estimation and clustering: Dirichlet processes and Chinese restaurant processes

We now consider two distinct problems—density estimation and clustering—and describe a close link between the two when approached from a Bayesian non-parametric modelling approach.

Density estimation refers to the problem of inferring an unknown density p from data $\mathcal{D} = \{x_1, \dots, x_N\}$. Let us first consider a very simple situation where the data points belong to a discrete and finite space with K possible values, $x \in \mathcal{X} = \{1, \dots, K\}$. Any distribution on \mathcal{X} can be represented by a K -dimensional non-negative vector \mathbf{p} that sums to one. To infer \mathbf{p} from \mathcal{D} , we need a reasonable prior on finite distributions. The Dirichlet distribution is a natural choice that takes the form

$$P(\mathbf{p}|\alpha, \mathbf{h}) = \frac{1}{Z} \prod_{k=1}^K p_k^{\alpha h_k - 1}.$$

Here, Z is a normalizing constant, \mathbf{h} is the mean of \mathbf{p} and $\alpha > 0$ controls the dispersion around the mean. A very attractive property of the Dirichlet distribution is that it is *conjugate* in the sense that the posterior $P(\mathbf{p}|\mathcal{D}, \alpha, \mathbf{h})$ is also Dirichlet. Another nice property is that for all but pathological choices of the parameters, it has *good coverage* in the sense that it puts some non-zero probability mass near all possible values of \mathbf{p} .

The DP extends the Dirichlet distribution to countable or uncountably infinite spaces \mathcal{X} . It has the property that for any *finite* partition of $\mathcal{X} = (\mathcal{X}_1, \dots, \mathcal{X}_K)$, it marginalizes to a K -dimensional Dirichlet distribution on the measure (i.e. mass) assigned to each element of the partition [22]. We

write that the probability measure G is drawn from a DP prior as $G \sim \text{DP}(\alpha, H)$ analogous to our notation for the Dirichlet distribution. An excellent introduction to the DP is provided by Teh [23].

Conjugacy and good coverage suggest that the DP could be a very good general purpose non-parametric density estimator. Unfortunately, the distributions drawn from a DP prior are, with probability one, discrete so they do not have a density. In fact, a draw from a DP prior can be represented in the following way:

$$G = \sum_{k=1}^{\infty} \pi_k \delta_{x_k}, \quad (5.1)$$

where the sum is an infinite sum, the π_k are masses that sum to one, δ is the Dirac-delta function and the x_k are locations for those probability masses drawn iid from the base measure H , which controls the mean of G . To alleviate the problem of discreteness, the DP is often used as a prior on the parameters of a mixture model, with the whole model now called a Dirichlet process mixture (DPM) [20,24–26].

The particular example of a DPM of Gaussians, also known as an infinite Gaussian mixture model [27], is widely used for both density estimation and clustering. *Clustering* refers to the problem of finding groupings of objects or data points such that similar objects belong to the same group and dissimilar objects belong to different groups. An example application of clustering is finding groups of similar celestial objects in a large astronomical survey. Posed abstractly in this way, clustering is not a well-defined problem (how many groups should we find? What does ‘similar’ mean?). However, if we restrict ourselves to assuming that each cluster can be captured by some parametrized probability distribution over data points, such as a Gaussian, then clustering becomes a well-defined probabilistic inference problem. Bayesian non-parametric clustering using DPMs is a natural extension of classical clustering algorithms such as the Expectation Maximization algorithm for finite mixture models or k-means [28]. In a finite mixture model, the data are assumed to come from a distribution composed of K components,

$$p(x|\boldsymbol{\pi}, \boldsymbol{\theta}) = \sum_{k=1}^K \pi_k p(x|\theta_k), \quad (5.2)$$

with mixing weights $\boldsymbol{\pi}$ that sum to one, and parameters θ_k for each component. An infinite mixture model considers the limit of $K \rightarrow \infty$, and has the property that it allows the number of observed clusters to grow with the number of observed data points. A DPM can be obtained as an infinite limit of a finite mixture model in many ways, but let us consider the following construction:

$$x_n \sim p(x|\theta_n),$$

$$\theta_n \sim G$$

and

$$G \sim \text{DP}(\alpha, H).$$

As G is discrete, the values of θ_n will repeat, which results in a clustering of the data points. By equation (5.1), the π_k correspond to the mixing weights of the infinitely many clusters, which can be compared with the finite counterpart (equation (5.2)). The distribution over partitions of the data points induced by the DPM is known as a *Chinese restaurant process* (CRP; Aldous [29]).

The DPM, apart from being mathematically elegant, has some practical advantages over traditional clustering algorithms. There can be computational advantages to running a single instance of a DPM inference algorithm that automatically infers the number of clusters, rather than having to run multiple instances of an algorithm to compare different hypotheses on the number of clusters. Moreover, at test time, the DPM always allows for the possibility that a new test point (e.g. a new astronomical object) belongs to a cluster that was not observed in the training data. This makes DPM predictions somewhat more robust to outliers.

6. Discrete state time series: infinite hidden Markov models

Many processes of interest produce sequential data for which models that assume exchangeability are inadequate. For example, natural language text, protein amino acid sequences, financial time series and natural sounds all have sequential structure that is important to model.

One of the most widely used models for times series is the HMM. An HMM assumes that a series of observations (x_1, \dots, x_T) was generated by a process that can be in one of K different discrete states at each point in time, $s_t \in \{1, \dots, K\}$. Moreover, in an HMM, s_t fully captures the state of the system at time t in the sense that given s_t the future evolution of the system does not depend on the state at times previous to t ; this is the *Markov property*: $P(s_{t+1}|s_1, \dots, s_t) = P(s_{t+1}|s_t)$. Finally, the observations are assumed to be drawn independently given the hidden states, through an emission process $P(x_t|s_t)$.

An example of an interesting application of HMMs in the physical and biological sciences is the modelling of single ion channel kinetics [30]. Here, the measured time series of currents from an ion channel are modelled by assuming that at each point in time the channel macromolecule can be in one of many different conformational states, and that there is a transition matrix defining the probability of transitioning between each of these states.

Learning an HMM involves inferring both parameters of the transition process $P(s_{t+1}|s_t)$, which is in general a $K \times K$ transition matrix, and parameters of the emission process, $P(x_t|s_t)$ [31]. The problem of learning the structure of an HMM corresponds to inferring the number of hidden states, K , from data. Rather than doing model selection over varying K , we would like to develop a non-parametric approach to HMMs where the model has countably infinitely many hidden states at its disposal. This can be useful both when we do not believe that any finite HMM can capture the data generating process well, and in situations where we believe that the data was actually generated from a finite-state process, but we simply do not know how many states this process should have.⁸

The key insight that allows one to develop a non-parametric HMM is that the finite HMM is a time-series generalization of finite mixture models. At each time step, the HMM assumes that the observation x_t was generated by one of K mixture components, where s_t indicated the component (or cluster) used. The only difference between HMMs and mixture models is that the mixture indicators in an HMM depend on the ones at the previous time step.

Using this insight, Beal *et al.* [32] developed the infinite HMM (iHMM). The basic idea was to consider a Bayesian HMM with countably infinitely many states. The main difficulty was to define a sensible prior on the parameters of the $\infty \times \infty$ transition matrix. In the finite K -dimensional case, one would typically use independent symmetric Dirichlet prior distributions for each row of the transition matrix (where a k th row corresponds to the vector of all outgoing transition probabilities from state k). In the infinite limit, the independent Dirichlet prior does not result in a sensible model, as under this prior, with probability one, the HMM will keep transitioning to new states rather than revisiting previous states. The solution developed in Beal *et al.* [32] was to couple the rows by using a hierarchical DP, a solution analogous to a reinforced urn process in probability theory [33]. This work was followed up in the elegant paper by Teh *et al.* [34], which further developed the hierarchical DP and proposed an improved Markov chain Monte Carlo (MCMC) sampler for the iHMM.

Since the original paper, there have been a number of conceptual and algorithmic developments of the iHMM. The beam sampler provides an efficient way of sampling the iHMM by using dynamic programming forward-backward style message passing [35].⁹ Parallel and distributed implementations of the iHMM allow larger scale deployments [36]. The block diagonal iHMM is an extension that groups the hidden states into clusters of states, effectively

⁸There is a subtle but important distinction between assuming an infinite model, and assuming a finite but potentially unbounded model. The difference comes to light when we consider what we expect to happen in the limit of infinitely large datasets. In the former case, the posterior will have visited infinitely many distinct states, while in the latter case, the posterior should converge on some finite number of states.

⁹The beam sampler is available in the software package <http://mloss.org/software/view/205/>.

hierarchically partitioning the state space [37]. The iHMM can be extended to have a power-law structure on the hidden states by using the Pitman–Yor process [38,39] and has been successfully applied to diverse problems such as language modelling [39] and speaker diarization [40].

7. Sparse matrices and overlapping clusters: Indian buffet processes

One limitation of DPMs, the iHMM and clustering models in general is that each data point is modelled as belonging to one of a set of mutually exclusive clusters. Although the non-parametric variants are flexible, in that they allow countably infinitely many clusters, they do not allow a data point to belong to multiple clusters at the same time—they fundamentally define distributions over *partitions* of the data.

We would like building blocks for our models that can allow overlapping cluster membership. For example, to understand a person’s network of friendships, we really need models that can capture the idea that a person can belong *simultaneously* to many possible social groupings based on workplace, family, housing location, high school, hobbies, etc. This type of hidden structure in data is sometimes called *factorial* structure [41–43].

The *Indian buffet process* (IBP; [44,45]) is a probabilistic object that can be used to represent non-parametric factorial structure. There are a number of ways of understanding and deriving the IBP. Let us start with a sparse binary matrix view of IBPs.

Consider an $N \times K$ binary matrix, Z , where we can think of the rows of the matrix as corresponding to objects or data points, and the columns as features or clusters. An element of this matrix $z_{nk} = 1$ may denote that object n possesses hidden feature k (or belongs to cluster k), and features (clusters) are not mutually exclusive in that an object can have multiple features. We wish to define a very simple distribution that is exchangeable over the objects (rows). A very simple choice is to assume that the columns are independent. We can therefore model each column through a single unknown parameter, θ_k , representing the frequency of feature k , $p(z_{nk} = 1 | \theta_k) = \theta_k$. Full specification for this model requires some choice for the prior distribution of θ_k . A natural choice is the beta distribution (the special case of a Dirichlet when there are only two outcomes) that happens to be conjugate to the Bernoulli likelihood, allowing θ_k to be integrated out.

Like in the previous cases, we wish to consider models with infinitely many possible features or clusters, so we therefore have to examine the limit $K \rightarrow \infty$. The beta distribution has two parameters, α, β , and the mean of the beta distribution is $\alpha / (\alpha + \beta)$. For fixed α, β and N in the limit $K \rightarrow \infty$, the matrix Z will have infinitely many ones in it, which makes it computationally and statistically of limited interest. However, consider scaling the first parameter of the beta distribution and setting the second parameter to one, i.e. using a $\text{beta}(\alpha/K, 1)$ prior for each θ_k .¹⁰ In this case, the limit on the distribution of Z has a number of nice properties: (i) the number of ones in each row is distributed as $\text{Poisson}(\alpha)$, (ii) the total expected number of ones is αN , (iii) the number of non-zero columns grown as $O(\alpha \log N)$, and (iv) the rows are exchangeable. This distribution is the IBP.

Note that the distribution described earlier has infinitely many columns of zeros. Because the columns are all iid, once we sample a matrix, we can reorder its columns into a more usable representation such as the one shown in figure 3. In many applications, we are specifically interested in sparse binary matrices; for example, to represent which factors are non-zero in sparse latent factor models [46], for binary matrix factorization [47] or to represent the adjacency matrix in a graph [48]. However, it is sometimes useful to view the IBP as a more abstract probabilistic object. Whereas the CRP is an infinitely exchangeable distribution over *partitions* of a set of objects, the IBP is an infinitely exchangeable distribution over *subsets* of a set of objects. Each object belongs to $\text{Poisson}(\alpha)$ subsets (or clusters), and as N increases, both the number of subsets grows (logarithmically) and the size of each subset grows (linearly) with N .

¹⁰It is not hard to reintroduce the second parameter to get the two-parameter IBP.




Figure 3. A sample from an IBP matrix, with columns reordered. Each row has, on average, 10 ones. Note the logarithmic growth of non-zero columns with rows. For the ‘restaurant’ analogy where customers enter a buffet with infinitely many dishes, you can refer to the original IBP papers.

Since its introduction, many interesting properties and extensions of the IBP have been uncovered. Because the IBP defines an exchangeable distribution, it has a de Finetti mixing distribution; for the IBP, Thibaux & Jordan [49] showed that this is the beta process [50]. An important extension of the IBP is the three-parameter model that exhibits a power-law growth in the number of clusters [51].

Non-parametric models that use the IBP to define sparse latent variables have been applied to a number of different problems, as reviewed in Griffiths & Ghahramani [45]. A very interesting application of the IBP is to the problem of *network modelling*: modelling the connections between objects or entities in social, biological and physical networks [52]. While many models of networks are based on the idea of discovering communities or clusters of the nodes, the IBP allows each node to belong to multiple overlapping communities or clusters, a property that can be exploited to obtain improved predictive performance in tasks such as link prediction [53–55].

Figure 4 shows a useful visualization of the relationship between a number of models. Here, we can see that the DPM (§5), the iHMM (§6) and the IBP (§7) can all be related to each other. Combining the key features of all models results in the infinite factorial HMM, a Bayesian non-parametric time-series model with factorial hidden structure [56].

8. Hierarchies: Kingman’s coalescent, and Dirichlet and Pitman–Yor diffusion trees

It is useful to distinguish between *flat* clustering models and *hierarchical* clustering models. In the former, data are partitioned into clusters, while in the latter, these clusters in turn are also partitioned into superclusters in a hierarchical manner. This results in a *tree* (or dendrogram)




Figure 4. A diagram representing how some models relate to each other. We start from finite mixture models and consider three different ways of extending them. Orange arrows correspond to time-series versions of static (iid) models. Blue arrows correspond to Bayesian non-parametric versions of finite parametric models. Green arrows correspond to factorial (overlapping subset) versions of clustering (non-overlapping) models. ifHMM, infinite factorial hidden Markov model.

where at the top level, we have a single root node corresponding to one coarse cluster with all the data points, while at the bottom, we have leaves corresponding to fine clusters with a single data point in each cluster [28].

Such hierarchical clustering models are useful for many reasons. First, many natural phenomena are well modelled through hierarchies. For example, the evolution of biological organisms results in phylogenetic trees that are well modelled by a hierarchy, and real-world objects can often be decomposed into parts, subparts, etc. Second, hierarchies allow one to tie together parameters of complex models so as to improve generalization in learning. For example, if one is building statistical models of patient outcome across a country, it might be natural to group together parameters of regions, cities, hospitals and individual doctors, corresponding to multiple levels of a hierarchy. Third, hierarchies provide good abstractions for interpretability of complex models. For example, rather than trying to understand what the hundreds of states in an HMM each do, it would be useful to have an HMM that partitions the states in a coarse to fine hierarchy so that one can start out by interpreting the coarse states and gradually move down to the fine states.

Note that while hierarchical clustering models are often described in terms of hierarchies over data points, in the above examples, we have seen that hierarchies can be useful more generally over internal components of models, such as hidden states of an HMM, or parameters of a model.

We have seen that the DPM model (and the associated CRP) can be used to define Bayesian non-parametric models for flat clustering. Are there equivalent Bayesian non-parametric models that result in hierarchical clusterings?

The answer is yes. Here, we very briefly touch upon two frameworks for generating hierarchies that can be used in non-parametric models. For the non-parametric setting, the key requirement is that the models define an infinitely exchangeable distribution over the data points. For this to hold, the models must be projective in the sense that marginalizing over the $(N + 1)$ th point should result in a coherent model over N points [57].

The first solution is given by *Kingman's coalescent* [58]. This model has been widely used for modelling population genetics [59] and more recently in the machine learning community for hierarchical clustering [60]. The coalescent defines distributions over trees by starting with every data point in its own cluster and considering a process that merges clusters *backwards* in time until

only one cluster remains. The key property of the coalescent is that for each pair of clusters the event that they merge is independent of the event that any other pair merges, and the time to this event is drawn from an exponential distribution with constant rate (which we can set to one without loss of generality). This process is well defined in the limit of $N \rightarrow \infty$ data points, and defines an infinitely exchangeable distribution over points.

A second solution is given by the Dirichlet diffusion tree (DDT; Neal [61]). Like Kingman's coalescent, the DDT also defines a tree by considering a Markovian process evolving in time; however, the DDT starts at time $t=0$ and evolves *forward* for one unit of time.¹¹ We denote the evolution of the n th point in time via $x_n(t)$. The first data point starts at a location $x_1(0)=0$ and follows a Brownian motion process, ending up at some point drawn marginally from a unit variance Gaussian, $x_1(1) \sim N(0,1)$. The second data point exactly follows the path of the first data point until a divergence event occurs, after which point its path is independent of the path of the first point. The time to divergence is parametrized through a *hazard function*, an object commonly used in survival analysis. Subsequent data points follow the paths of previous data points, diverging according to a scaled form of the hazard function, and when reaching a branch point choosing a branch with probability proportional to the number of points that chose that branch before. This process defines an exchangeable distribution over data points, parametrized by the unknown tree. Using the DDT prior, the problem of hierarchical clustering becomes one of inferring the unknown tree given some data.

Both Kingman's coalescent and the DDT generate binary trees with probability one.¹² A generalization of the DDT that allows arbitrary branching of the trees is given by the Pitman–Yor diffusion tree (PYDT; [62]). The process is generalized to allow, at each branch point, for the new data point either to follow the path of one of the previous points, or to create a new branch. Like the Pitman–Yor process [63], the PYDT has two parameters controlling its branching structure. Certain settings of these parameters result in the DDT, while other settings recover the distribution over trees induced by Kingman's coalescent. The tree distribution induced by the PYDT is the multi-furcating Gibbs fragmentation tree [64], the most general Markovian exchangeable distribution over trees. General distributions over hierarchies and other clustering structures can also be derived through the elegant theory of fragmentation–coagulation processes [65].

9. Covariance matrices: generalized Wishart processes

Often we are interested in modelling the relationship between a number of variables. One can express relationships between variables in many different ways, but perhaps the simplest representation of dependence is the *covariance matrix*, Σ , a symmetric positive (semi)definite matrix representing second-order statistics of the variables. The covariance matrix plays a key role in parametrizing many distributions, most notably the multi-variate Gaussian, but also extensions such as the multi-variate t-distribution or more generally elliptical distributions [66].

In certain applications, we wish to model how such a covariance matrix might depend on some other variables. For example, in econometrics and finance, one is often interested in modelling a time-varying covariance matrix $\Sigma(t)$ —this is the key object of interest in the field of multi-variate stochastic volatility [67]. More generally, we would like to place distributions on covariance matrices that can depend on arbitrary variables, $\Sigma(x)$, not just scalar time. Viewed as a *function* of x , we want to be able to define distributions on matrix-valued functions restricted to the space of symmetric-positive-definite (s.p.d.) matrix values. Is there a convenient and simple way to define such a stochastic process?

Indeed there is, and the key insight comes from the observation that one can generate s.p.d. matrices by taking the outer products of random vectors. Consider the following construction,

¹¹Note that both in the coalescent and in the DDT, 'time' is an indexing variable used to generate a tree structure, and need not correspond to a real notion of time.

¹²In the case of the DDT, binary trees result in all but pathological cases where the hazard function diverges before $t=1$.

where we draw D -dimensional vectors independently from a multi-variate Gaussian distribution $\mathbf{u}_i \sim \mathcal{N}(0, V)$ with covariance matrix V , and we define

$$\Sigma = \sum_{i=1}^{\nu} \mathbf{u}_i \mathbf{u}_i^{\top}.$$

Such a matrix Σ is Wishart distributed with ν degrees of freedom and is s.p.d. with probability one as long as $\nu \geq D$. The mean of Σ is proportional to V .

We can generalize the Wishart distribution to a stochastic process indexed by x in any space \mathcal{X} by replacing the elements of each \mathbf{u}_i with draws from a GP: $\mathbf{u}_i(x) = (u_{i1}(x), u_{i2}(x), \dots, u_{iD}(x))$, where

$$u_{id} \sim \text{GP}(0, K),$$

where K is the covariance function or kernel of the GP [21].

The desired stochastic process is obtained by the same construction

$$\Sigma(x) = \sum_{i=1}^{\nu} \mathbf{u}_i(x) \mathbf{u}_i(x)^{\top}.$$

A special case of such a construction where the GPs are assumed to be Brownian has been studied in probability theory and is known as a *Wishart process* [68]; these processes have recently been applied to problems in econometrics [69,70]. The more general case for arbitrary GPs is developed in Wilson & Ghahramani [71] where applications to multi-variate stochastic volatility are also explored. In this general framework, the model parameters controlling V , K and ν can be learned from data.

10. Inference

So far we have focused on describing a general probabilistic approach to modelling and some non-parametric distributions that can be used in models of complex datasets. The three key ingredients in any approach to modelling are the data, the model and an *algorithm for doing probabilistic inference* in the model. The problem of probabilistic inference corresponds to computing the conditional probability of some variables of interest given some observed variables, while marginalizing out all other variables. Thus, both the computation of a model's marginal likelihood (equation (2.1)) and prediction (equation (2.2)) are inference problems. Filling in missing data given observed data, computing the parameter posterior, or evaluating expectations of various quantities can all also be phrased as instances of the inference problem.

Generally, the inference problem boils down to a problem of numerical integration or summation over large state spaces. For most models of interest, especially non-parametric models, exact inference is computationally *intractable*, in the sense that all known algorithms for computing the exact probabilities of interest scale exponentially with some aspect of the problem, such as the number of data points or variables. In many problems, even approximating these exact probabilities to within some small error tolerance is intractable in the worst case.¹³

A wide variety of approximation methods have been developed to solve Bayesian inference problems. These can be roughly divided into stochastic approximations (which make extensive use of random numbers) and deterministic approximations. Some examples of widely used stochastic approximate inference methods include MCMC methods (for an excellent review, see Neal [72]), exact sampling methods [73] and particle filtering methods [74]. Some examples of deterministic algorithms include the Laplace approximation, variational methods [75] and expectation propagation [76]. Both deterministic and stochastic algorithms for inference can often exploit the conditional independence relationships that exist between the variables in a model to perform the relevant computations efficiently using local messages passed between nodes of a graphical model [77–79].

¹³Most theoretical results on intractability focus on the worst case of a problem instance. The real-world inference problem may in fact be much easier to approximate.

A complete review of approximate inference methods is beyond the scope of this paper, but a couple of points are worth making. All approximate inference methods can be characterized in terms of a speed–accuracy trade-off. Some methods are fast but often inaccurate, while other methods are slow or, such as MCMC, can be run for increasing amounts of time to give increasingly accurate results. There is no general rule of thumb for which approximate inference method is best—different models and problems tend to favour different methods. However, for a particular problem, the difference between a good choice of inference algorithm and a poor choice can be orders of magnitude of computation. Thus, it is well worth being familiar with a number of inference algorithms and being willing to try several methods on a particular problem.

The field of Bayesian statistics has thrived in recent years, both due to the availability of better inference algorithms and the dramatic growth in computing power. Ironically, the most widely used inference method in Bayesian statistics is MCMC, which itself is a thoroughly frequentist (non-Bayesian) method for approximating integrals. From a Bayesian perspective, numerical computation problems are also inference problems. In numerical integration, one is trying to infer the value of an integral by computing the integrand at a limited number of locations. The values of the integrand are the *data*, which, combined with a prior on the integrand, can result in a posterior on the value of the integral. The choice of where to evaluate the integrand can be made using Bayesian decision theory, clearly the random evaluations of MCMC are not an optimally efficient method for evaluating integrals. The Bayesian approach to numerical integration is known variously as Bayesian quadrature or Bayesian Monte Carlo, and although it is not as widely used as MCMC, it can be dramatically more efficient in situations where evaluating the integrand is computationally costly [80–83].

Deriving the approximate inference equations for each new model can be tedious and error prone, thereby inhibiting the researcher’s ability to explore many variations on any given model. The field of *probabilistic programming* offers an exciting alternative approach to building and evaluating models. Three notable examples of probabilistic programming frameworks are *BUGS*, *Church* and *Infer.NET* [84–86]. The basic idea is to write down a computer program defining the generative model in a programming language that is augmented to have random variables. This computer program that defines the generative model can then be *automatically* transformed or manipulated into a program that performs approximate posterior inference in the model. Thus, the process of deriving the approximate inference equations is automated, reducing the risk of human error. One potential disadvantage of this approach is that the automatically derived inference code may not be as efficient as an expertly designed inference method for a particular model. However, the rapid model development cycle that can result from the use of probabilistic programming languages offers tremendous advantages for the future of probabilistic modelling.

11. Conclusions

Modelling is fundamentally a process of quantifying uncertainty about possible predictions given available information. Probability theory provides an elegant framework for representing all sources of uncertainty in a model. Probability theory also provides the simple rules with which to manipulate models so as to obtain predictions, to compare models and to learn models from data. Bayesian statistics is simply the application of the rules of probability theory to modelling from data.

The probabilistic approach to modelling is most effective when the models are flexible enough to capture relevant aspects of the problem domain. Flexibility can be achieved by allowing models either to have many parameters, or in the limiting case, to have infinitely many parameters, in other words to be *non-parametric*. Learning models with infinitely many parameters may seem statistically and computationally daunting. However, statistically we have seen in §2 that the process of Bayesian averaging or marginalization avoids overfitting and therefore allows one to use models with infinitely many parameters (§3). Computationally, we have seen that there are a wide range of approximation methods that can be used to perform the relevant marginalizations required for inference (§10).

This article has attempted to give an overview of the basics of Bayesian modelling, the motivation for Bayesian non-parametrics and some of the more widely used non-parametric models. We have touched upon some computational issues, and briefly alluded to some successful applications of Bayesian non-parametrics. Of course, there is a great deal of current work in this rapidly advancing field that was not covered in this study.

There are four areas of future work worth highlighting: theory, new models, scalable algorithms and new applications. Regarding *theory*, developing frequentist consistency and convergence rate results for Bayesian non-parametric models is challenging but important. In particular, the effect of the prior on an infinite-dimensional space can result in inconsistent models in certain cases, so careful study of consistency is required.

While we have many building blocks for models, *new models* are an exciting area for work as they can lead to novel applications. General results on the construction of Bayesian non-parametric models, such as the recent work of Orbanz [57], are particularly useful for developing new models.

Many modern applications of statistical modelling and machine learning require the analysis of very large datasets (this is sometimes called ‘big data’). While a number of very good approximate inference methods have been developed in the last few decades, making these highly *scalable* is a challenge that requires extensive tools from computer science including efficient data structures, and parallel and distributed computing paradigms.

Finally, novel *applications* of Bayesian non-parametric modelling will advance the field in many ways. First, many new models are derived in response to the challenges arising from new applications. Second, applications can motivate general scalable inference solutions. Finally, successful applications will help in widening the adoption of a Bayesian non-parametric approach to modelling.

References

1. Wolpert DM, Ghahramani Z, Jordan MI. 1995 An internal model for sensorimotor integration. *Science* **269**, 1880–1882. (doi:10.1126/science.7569931)
2. Knill D, Richards W. 1996 *Perception as Bayesian inference*. Cambridge, UK: Cambridge University Press.
3. Griffiths TL, Tenenbaum JB. 2006 Optimal predictions in everyday cognition. *Psychol. Sci.* **17**, 767–773. (doi:10.1111/j.1467-9280.2006.01780.x)
4. Doob JL. 1949 Application of the theory of martingales. *Coll. Int. Centre Nat. Res. Sci.* **13**, 23–27.
5. Le Cam L. 1986 *Asymptotic methods in statistical decision theory*. Berlin, Germany: Springer.
6. Van der Vaart A. 2000 *Asymptotic statistics*. Cambridge, UK: Cambridge University Press.
7. Jefferys WH, Berger JO. 1992 Ockham’s razor and Bayesian analysis. *Am. Sci.* **80**, 64–72.
8. MacKay DJC. 2003 *Information theory, inference, and learning algorithms*. Cambridge, UK: Cambridge University Press.
9. Rasmussen CE, Ghahramani Z. 2001 Occam’s razor. In *Advances in Neural Information Processing Systems* (eds TK Leen, TG Dietterich and V Tresp), vol. 13, pp. 294–300. Cambridge, MA: MIT Press.
10. MacKay DJC. 1991 Bayesian methods for adaptive models. PhD thesis, California Institute of Technology.
11. Murray IA, Ghahramani, Z. 2005 A note on the evidence and Bayesian Occam’s razor. Technical Report GCNU-TR 2005-003, Gatsby Unit, University College London, London, UK.
12. Roberts SJ, Husmeier D, Rezek I, Penny W. 1998 Bayesian approaches to Gaussian mixture modeling. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**, 1133–1142. (doi:10.1109/34.730550)
13. Ghahramani Z, Beal MJ. 2000 Variational inference for Bayesian mixtures of factor analysers. In *Advances in Neural Information Processing Systems*, vol. 12, pp. 449–455.
14. MacKay DJC. 1994 Bayesian nonlinear modeling for the prediction competition. *ASHRAE Trans.* **100**, 1053–1062.
15. Neal RM. 1998 Assessing relevance determination methods using DELVE. In *Neural networks and machine learning* (ed. CM Bishop), pp. 97–129. Berlin, Germany: Springer.
16. MacKay DJC. 1997 Ensemble learning for hidden Markov models. Technical report, Cavendish Laboratory, University of Cambridge, Cambridge, UK.

17. Friedman N, Koller D. 2003 Being Bayesian about network structure: a Bayesian approach to structure discovery in Bayesian networks. *Mach. Learn.* **50**, 95–126. (doi:10.1023/A:1020249912095)
18. Zellner A. 1988 Optimal information processing and Bayes's theorem. *Am. Stat.* **42**, 278–280. (doi:10.2307/2685148)
19. Kallenberg O. 2005 *Probabilistic symmetries and invariance principles*. Berlin, Germany: Springer.
20. Neal RM. 2000 Markov chain sampling methods for Dirichlet process mixture models. *J. Comput. Graph. Stat.* **9**, 249–265. (doi:10.1080/10618600.2000.10474879)
21. Rasmussen CE, Williams CKI. 2006 *Gaussian processes for machine learning*. Cambridge, MA: MIT Press.
22. Ferguson TS. 1973 A Bayesian analysis of some nonparametric problems. *Ann. Stat.* **1**, 209–230. (doi:10.1214/aos/1176342360)
23. Teh YW. 2010 Dirichlet processes. In *Encyclopedia of machine learning* (eds C Sammut, GI Webb), pp. 280–287. Berlin, Germany: Springer.
24. Antoniak CE. 1974 Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Ann. Stat.* **2**, 1152–1174. (doi:10.1214/aos/1176342871)
25. Ferguson TS. 1983 Bayesian density estimation by mixtures of normal distributions. *Recent Adv. Stat.* 287–303.
26. Lo A. 1984 On a class of Bayesian nonparametric estimates. I. Density estimates. *Ann. Stat.* **12**, 351–357. (doi:10.1214/aos/1176346412)
27. Rasmussen, CE. 2000 The infinite Gaussian mixture model. In *Advances in Neural Information Processing Systems*, vol. 12, pp. 554–560.
28. Duda RO, Hart PE. 1973 *Pattern classification and scene analysis*. London, UK: Wiley.
29. Aldous D. 1985 Exchangeability and related topics. In *École d'été de probabilités de Saint-Flour XIII–1983*, pp. 1–198. Berlin, Germany: Springer.
30. Chung SH, Moore JB, Xia L, Premkumar LS, Gage PW. 1990 Characterization of single channel currents using digital signal processing techniques based on Hidden Markov models. *Phil. Trans. R. Soc. Lond. B* **329**, 265–285. (doi:10.1098/rstb.1990.0170)
31. Rabiner LR, Juang, BH. 1986 An introduction to hidden Markov models. *IEEE Acoust. Speech Signal Process. Mag.* **3**, 4–16. (doi:10.1109/MASSP.1986.1165342)
32. Beal MJ, Ghahramani Z, Rasmussen CE. 2002 The infinite hidden Markov model. In *Advances in Neural Information Processing Systems*, vol. 14, pp. 577–584.
33. Fortini S, Petrone S. 2012 Hierarchical reinforced urn processes. *Stat. Probab. Lett.* **82**, 1521–1529. (doi:10.1016/j.spl.2012.04.012)
34. Teh YW, Jordan MI, Beal MJ, Blei DM. 2006 Hierarchical Dirichlet processes. *J. Am. Stat. Assoc.* **101**, 1566–1581. (doi:10.1198/016214506000000302)
35. Van Gael J, Saatici Y, Teh YW, Ghahramani Z. 2008 Beam sampling for the infinite hidden Markov model. In *Proc. 25th Int. Conf. Machine Learning, Helsinki, Finland, 5–9 June 2008* (eds WW Cohen, A McCallum, ST Roweis), pp. 1088–1095.
36. Bratieres S, Van Gael J, Vlachos A, Ghahramani Z. 2010 Scaling the iHMM: parallelization versus Hadoop. In *IEEE Int. Conf. Computing and Information Technology, Bradford, UK, 29 June–1 July 2010*.
37. Stepleton T, Ghahramani Z, Gordon G, Lee T-S. 2009 The block diagonal infinite hidden Markov model. In *12th Int. Conf. Artificial Intelligence Statistics*, vol. 5 (eds D van Dyk, M Welling), pp. 552–559. Brookline, MA: Microtome Publishing.
38. Van Gael J, Ghahramani Z. 2011 Nonparametric hidden Markov models. In *Bayesian time series models* (eds D Barber, A Cemgil, S Chiappa), pp. 317–340. Cambridge, UK: Cambridge University Press.
39. Blunsom P, Cohn T. 2011 A hierarchical Pitman–Yor process HMM for unsupervised part of speech induction. In *Proc. 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, 19–24 June 2011, Portland, OR*, vol. 1, pp. 865–874. Association for Computational Linguistics.
40. Fox EB, Sudderth EB, Jordan MI, Willsky AS. 2008 An HDP–HMM for systems with state persistence. In *Proc. 25th Int. Conf. Machine Learning, Helsinki, Finland*, vol. 25. ACM Press.
41. Hinton GE, Zemel RS. 1994 Autoencoders, minimum description length, and Helmholtz free energy. In *Advances in Neural Information Processing Systems*, vol. 6 (eds J Cowan, G Tesauro, J Alspecter), pp. 3–10. San Francisco, CA: Morgan Kaufmann Publishers.

42. Saund E. 1994 Unsupervised learning of mixtures of multiple causes in binary data. In *Advances in neural information processing systems*, vol. 6 (eds J Cowan, G Tesauro, J Alspector), pp. 27–34. San Francisco, CA: Morgan Kaufmann.
43. Ghahramani Z. 1995 Factorial learning and the EM algorithm. In *Advances in neural information processing systems*, vol. 7 (eds G Tesauro, DS Touretzky, TK Leen), pp. 617–624. Cambridge, MA: MIT Press.
44. Griffiths TL, Ghahramani Z. 2006 Infinite latent feature models and the Indian buffet process. In *Advances in neural information processing systems*, vol. 18 (eds B Schölkopf, JC Platt, T Hoffman), pp. 475–482. Cambridge, MA: MIT Press.
45. Griffiths TL, Ghahramani Z. 2011 The Indian buffet process: an introduction and review. *J. Mach. Learn. Res.* **12**, 1185–1224.
46. Knowles DA, Ghahramani Z. 2007 Infinite sparse factor analysis and infinite independent components analysis. In *7th Int. Conf. Independent Component Analysis and Signal Separation* (eds ME Davies, CJ James, SA Abdallah, MD Plumbley), pp. 381–388. London, UK: Springer.
47. Meeds E, Ghahramani Z, Neal R, Roweis S. 2007 Modelling dyadic data with binary latent factors. In *Advances in Neural Information Processing Systems*, vol. 19 (eds B Schölkopf, J Platt, T Hofmann), pp. 977–984. Cambridge, MA: MIT Press.
48. Adams RP, Wallach H, Ghahramani Z. 2010 Learning the structure of deep sparse graphical models. In *13th Int. Conf. Artificial Intelligence and Statistics, Sardinia, Italy* (eds YW Teh, M Titterton), pp. 1–8.
49. Thibaux R, Jordan MI. 2007 Hierarchical beta processes and the Indian buffet process. In *11th Int. Conf. Artificial Intelligence and Statistics (AISTATS 2007), San Juan, Puerto Rico, 21–24 March 2007* (eds M. Meila, X Shen).
50. Hjort NL. 1990 Nonparametric Bayes estimators based on Beta processes in models for life history data. *Ann. Stat.* **18**, 1259–1294. (doi:10.1214/aos/1176347749)
51. Teh YW, Görür D. 2009 Indian buffet processes with power-law behavior. In *Advances in Neural Information Processing Systems*, vol. 22, pp. 1838–1846.
52. Newman, M. 2010 *Networks: an introduction*. Oxford, UK: Oxford University Press.
53. Miller K, Griffiths T, Jordan M. 2009 Nonparametric latent feature models for link prediction. In *Advances in Neural Information Processing Systems*, vol. 22, pp. 1276–1284.
54. Mørup M, Schmidt M, Hansen L. 2011 Infinite multiple membership relational modeling for complex networks. In *IEEE Int. Workshop on Machine Learning for Signal Processing (MLSP), Beijing, China, 18–21 September 2011*, pp. 1–6. IEEE.
55. Palla K, Knowles DA, Ghahramani Z. 2012 An infinite latent attribute model for network data. In *Proc. 29th Int. Conf. Machine Learning, Edinburgh, UK, July 2012* (eds J Langford, J Pineau), pp. 1607–1614. Omnipress.
56. Van Gael J, Teh YW, Ghahramani Z. 2008 The infinite factorial hidden Markov model. In *Advances in Neural Information Processing Systems*, vol. 21, pp. 1697–1704.
57. Orbanz P. 2009 Construction of nonparametric Bayesian models from parametric Bayes equations. In *Advances in Neural Information Processing Systems*, vol. 22, pp. 1392–1400.
58. Kingman JFC. 1982 The coalescent. *Stoch. Process. Appl.* **13**, 235–248. (doi:10.1016/0304-4149(82)90011-4)
59. Donnelly P, Tavaré S. 1995 Coalescents and genealogical structure under neutrality. *Annu. Rev. Genet.* **29**, 401–421. (doi:10.1146/annurev.ge.29.120195.002153)
60. Teh YW, Daumé H, Roy DM. 2008 Bayesian agglomerative clustering with coalescents. In *Advances in Neural Information Processing Systems*, vol. 22, pp. 1473–1480.
61. Neal RM. 2003. Density modeling and clustering using Dirichlet diffusion trees. In *Bayesian statistics*, vol. 7 (ed. Bernardo JM), pp. 619–629. Oxford University Press.
62. Knowles DA, Ghahramani Z. 2011 Pitman–Yor diffusion trees. In *Uncertainty in Artificial Intelligence (UAI 2011), Barcelona, Spain, 14–17 July 2011* (eds FG Cozman, A Pfeffer), pp. 410–418. AUAI Press.
63. Pitman J, Yor M. 1997 The two-parameter Poisson–Dirichlet distribution derived from a stable subordinator. *Ann. Probab.* **25**, 855–900. (doi:10.1214/aop/1024404422)
64. McCullagh P, Pitman J, Winkel M. 2008 Gibbs fragmentation trees. *Bernoulli* **14**, 988–1002. (doi:10.3150/08-BEJ134)
65. Teh YW, Blundell C, Elliott LT. 2011 Modelling genetic variations with fragmentation–coagulation processes. In *Advances in Neural Information Processing Systems*, vol. 23, pp. 819–827.

66. Muirhead R. 1982 *Aspects of multivariate statistical theory*, vol. 42. London, UK: Wiley.
67. Asai M, McAleer M, Yu J. 2006 Multivariate stochastic volatility: a review. *Econ. Rev.* **25**, 145–175. (doi:10.1080/07474930600713564)
68. Bru M. 1991 Wishart processes. *J. Theor. Probab.* **4**, 725–751. (doi:10.1007/BF01259552)
69. Philipov A, Glickman M. 2006 Multivariate stochastic volatility via Wishart processes. *J. Bus. Econ. Stat.* **24**, 313–328. (doi:10.1198/073500105000000306)
70. Gouriéroux C, Jasiak J, Sufana R. 2009 The Wishart autoregressive process of multivariate stochastic volatility. *J. Econ.* **150**, 167–181. (doi:10.1016/j.jeconom.2008.12.016)
71. Wilson AG, Ghahramani Z. 2011 Generalised Wishart processes. In *Uncertainty in Artificial Intelligence, Barcelona, Spain, 14–17 July 2011* (eds FG Cozman, A Pfeffer), pp. 736–744. AUAI Press.
72. Neal RM. 1993 Probabilistic inference using Markov chain Monte Carlo methods. Technical Report CRG-TR-93-1, Department of Computer Science, University of Toronto.
73. Propp JG, Wilson DB. 1996 Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct. Algorithms* **9**, 223–252. (doi:10.1002/(SICI)1098-2418(199608/09)9:1/2<223::AID-RSA14>3.0.CO;2-O)
74. Doucet A, de Freitas JFG, Gordon NJ. 2000 *Sequential Monte Carlo methods in practice*. New York, NY: Springer.
75. Jordan MI, Ghahramani Z, Jaakkola TS, Saul LK. 1999 An introduction to variational methods for graphical models. *Mach. Learn.* **37**, 183–233. (doi:10.1023/A:1007665907178)
76. Minka TP. 2001 Expectation propagation for approximate Bayesian inference. In *Uncertainty in artificial intelligence*, vol. 17, pp. 362–369. Morgan Kaufmann.
77. Winn J, Bishop CM. 2005 Variational message passing. *J. Mach. Learn. Res.* **6**, 661–694.
78. Minka, TP. 2005 Divergence measures and message passing. Microsoft Research Technical Report MSR-TR-2005-173.
79. Koller D, Friedman N. 2009 *Probabilistic graphical models: principles and techniques*. Cambridge, MA: MIT Press.
80. Diaconis P. 1988 Bayesian numerical analysis. In *Statistical decision theory and related topics IV*, vol. 1, pp. 163–175. Berlin, Germany: Springer.
81. O’Hagan A. 1991 Bayes–Hermite quadrature. *J. Stat. Plan. Inference* **29**, 245–260. (doi:10.1016/0378-3758(91)90002-V)
82. Rasmussen CE, Ghahramani Z. 2003 Bayesian Monte Carlo. In *Advances in Neural Information Processing Systems*, vol. 15, pp. 489–496. Cambridge, MA: MIT Press.
83. Osborne MA, Garnett R, Roberts S, Hart C, Aigrain S, Gibson N. 2012 Bayesian quadrature for ratios. In *Proc. 15th Int. Conf. Artificial Intelligence and Statistics (AISTATS 2012), La Palma, Canary Islands, 21–23 April 2012* (eds N Lawrence, M Girolami), pp. 832–840.
84. Lunn DJ, Thomas A, Best N, Spiegelhalter D. 2000 WinBUGS: a Bayesian modelling framework: concepts, structure, and extensibility. *Stat. Comput.* **10**, 325–337. (doi:10.1023/A:1008929526011)
85. Goodman N, Mansinghka V, Roy D, Bonawitz K, Tenenbaum J. 2008 Church: a language for generative models. In *Uncertainty in artificial intelligence, Helsinki, Finland, 9–12 July 2008*, vol. 22, p. 23.
86. Minka TP, Winn JM, Guiver JP, Knowles DA. 2010 Infer.NET 2.4. Microsoft Research Cambridge. See <http://research.microsoft.com/infernet>.