

Bayesian Interpolation

David J.C. MacKay
Computation and Neural Systems*
California Institute of Technology 139-74
Pasadena CA 91125
`mackay@hope.caltech.edu`

Abstract

Although Bayesian analysis has been in use since Laplace, the Bayesian method of *model-comparison* has only recently been developed in depth. In this paper, the Bayesian approach to regularisation and model-comparison is demonstrated by studying the inference problem of interpolating noisy data. The concepts and methods described are quite general and can be applied to many other data modelling problems.

Regularising constants are set by examining their posterior probability distribution. Alternative regularisers (priors) and alternative basis sets are objectively compared by evaluating the *evidence* for them. ‘Occam’s razor’ is automatically embodied by this process.

The way in which Bayes infers the values of regularising constants and noise levels has an elegant interpretation in terms of the effective number of parameters determined by the data set. This framework is due to Gull and Skilling.

1 Data modelling and Occam’s razor

In science, a central task is to develop and compare models to account for the data that are gathered. In particular this is true in the problems of learning, pattern classification, interpolation and clustering. Two levels of **inference** are involved in the task of data modelling (figure 1). At the first level of inference, we assume that one of the models that we invented is true, and we fit that model to the data. Typically a model includes some free parameters; fitting the model to the data involves inferring what values those parameters should probably take, given the data. The results of this inference are often summarised by the most probable parameter values and error bars on those parameters. This is repeated for each model. The second level of inference is the task of model comparison. Here, we wish to compare the models in the light of the data, and assign some sort of preference or ranking to the alternatives.¹

For example, consider the task of interpolating a noisy data set. The data set could be interpolated using a splines model, using radial basis functions, using polynomials, or using

*Address from January 1st 1992: Darwin College, Cambridge CB2 9EU, U.K. `mackay@eng.cam.ac.uk`

¹Note that both levels of *inference* are distinct from *decision theory*. The goal of inference is, given a defined hypothesis space and a particular data set, to assign probabilities to hypotheses. Decision theory typically chooses between alternative actions on the basis of these probabilities so as to minimise the expectation of a ‘loss function.’ This paper concerns inference alone and no loss functions or utilities are involved.

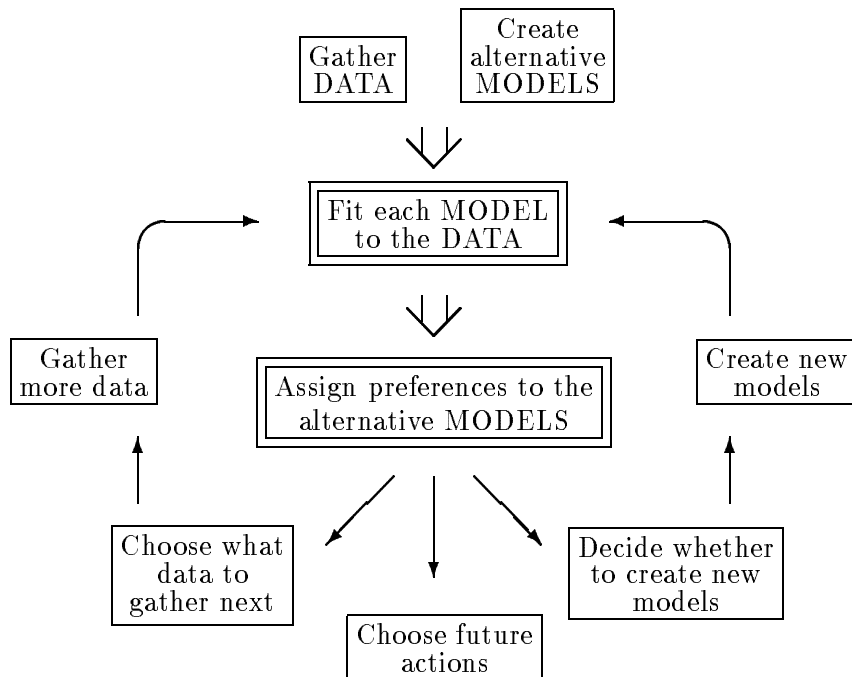


Figure 1: **Where Bayesian inference fits into the data modelling process.**

This figure illustrates an abstraction of the part of the scientific process, in which data is collected and modelled. In particular, this figure applies to pattern classification, learning, interpolation, etc.. The two double-framed boxes denote the two steps which involve *inference*. It is only in those two steps that Bayes' rule can be used. Bayes does not tell you how to invent models, for example.

The first box, 'fitting each model to the data,' is the task of inferring what the model parameters might be given the model and the data. Bayes may be used to find the most probable parameter values, and error bars on those parameters. The result of applying Bayes to this problem is often little different from the result of using orthodox statistics.

The second inference task, model comparison in the light of the data, is where Bayes is in a class of its own. This second inference problem requires a quantitative Occam's razor to penalise over-complex models. Bayes can assign objective preferences to the alternative models in a way that automatically embodies Occam's razor.

feedforward neural networks. At the first level of inference, we take each model individually and find the best fit interpolant for that model. At the second level of inference we want to rank the alternative models and state for our particular data set that, for example, 'splines are probably the best interpolation model,' or 'if the interpolant is modelled as a polynomial, it should probably be a cubic.'

Bayesian methods are able consistently and quantitatively to solve both these inference tasks. There is a popular myth that states that Bayesian methods only differ from orthodox (also known as 'frequentist' or 'sampling theory') statistical methods by the inclusion of subjective priors which are arbitrary and difficult to assign, and usually don't make much difference to the conclusions. It is true that at the first level of inference, a Bayesian's results will often differ little from the outcome of an orthodox attack. What is not widely appreciated is how Bayes performs the second level of inference. It is here that Bayesian methods are totally different from orthodox methods. Indeed, when regression and density estimation are discussed in most statistics texts, the task of model comparison is virtually ignored; no general orthodox method exists for solving this problem.




Figure 2: **Why Bayes embodies Occam’s razor**

This figure gives the basic intuition for why complex models are penalised. The horizontal axis represents the space of possible data sets D . Bayes rule rewards models in proportion to how much they *predicted* the data that occurred. These predictions are quantified by a normalised probability distribution on D . In this paper, this probability of the data given model \mathcal{H}_i , $P(D|\mathcal{H}_i)$, is called the evidence for \mathcal{H}_i .

A simple model \mathcal{H}_1 makes only a limited range of predictions, shown by $P(D|\mathcal{H}_1)$; a more powerful model \mathcal{H}_2 , that has, for example, more free parameters than \mathcal{H}_1 , is able to predict a greater variety of data sets. This means however that \mathcal{H}_2 does not predict the data sets in region \mathcal{C}_1 as strongly as \mathcal{H}_1 . Assume that equal prior probabilities have been assigned to the two models. Then if the data set falls in region \mathcal{C}_1 , the *less powerful* model \mathcal{H}_1 will be the *more probable* model.

Model comparison is a difficult task because it is not possible simply to choose the model that fits the data best: more complex models can always fit the data better, so the maximum likelihood model choice would lead us inevitably to implausible over-parameterised models which generalise poorly. ‘Occam’s razor’ is the principle that states that unnecessarily complex models should not be preferred to simpler ones. Bayesian methods automatically and quantitatively embody Occam’s razor (Gull, 1988, Jeffreys, 1939), without the introduction of ad hoc penalty terms. Complex models are automatically self-penalising under Bayes’ rule. Figure 2 gives the basic intuition for why this should be expected; the rest of this paper will explore this property in depth.

Bayesian methods were first laid out in depth by the Cambridge geophysicist Sir Harold Jeffreys (1939). The logical basis for the Bayesian use of probabilities as measures of plausibility was subsequently established by Cox (1964), who proved that consistent inference in a closed hypothesis space can be mapped onto probabilities. For a general review of Bayesian philosophy the reader is encouraged to read the excellent papers by Jaynes (1986) and Loredo (1989). Since Jeffreys the emphasis of most Bayesian probability theory has been ‘to formally utilize prior information’ (Berger, 1985), *i.e.* to perform inference in a way that makes explicit the prior knowledge and ignorance that we have, which orthodox methods omit. However, Jeffreys’ work also laid the foundation for Bayesian model comparison, which does not involve an emphasis on prior information. Jeffreys applied this theory to simple model comparison problems in geophysics, for example testing whether a single additional parameter is justified by the data. Since the 1960s, Jeffreys’ model comparison methods have been applied and extended in the economics literature (Zellner, 1984). Only recently has this aspect of Bayesian analysis been further developed and applied to more complex problems in other fields.

This paper will review Bayesian model comparison, ‘regularisation,’ and noise estimation, by studying the problem of interpolating noisy data. The Bayesian framework I will describe for these tasks is due to Gull (1988, 1989a, 1991) and Skilling (1991), who have used Bayesian methods to achieve the state of the art in image reconstruction. The same approach to regularisation has also been developed in part by Szeliski (1989). Bayesian model comparison is also discussed by Bretthorst (1990), who has used Bayesian methods to push back the limits of NMR signal detection. The same Bayesian theory underlies the unsupervised classification system, AutoClass (Hanson *et. al.*, 1991). The fact that Bayesian model comparison embodies Occam’s razor has been rediscovered by Kashyap (1977) in the context of modelling time series; his paper includes a thorough discussion of how Bayesian model comparison is different from orthodox ‘Hypothesis testing.’ One of the earliest applications of these sophisticated Bayesian methods of model comparison to real data is by Patrick and Wallace (1982); in this fascinating paper, competing models accounting for megalithic stone circle geometry are compared within the description length framework, which is equivalent to Bayes.

As the quantities of data collected throughout science and engineering continue to increase, and the computational power and techniques available to model that data also multiply, I believe Bayesian methods will prove an ever more important tool for refining our modelling abilities. I hope that this review will help to introduce these techniques to the ‘neural’ modelling community. A companion paper (MacKay, 1991b) will demonstrate how these techniques can be fruitfully applied to backpropagation neural networks. Another paper will show how this framework relates to the task of selecting where next to gather data so as to gain maximal information about our models (MacKay, 1991c).

2 The evidence and the Occam factor

Let us write down Bayes’ rule for the two levels of inference described above, so as to see explicitly how Bayesian model comparison works. Each model \mathcal{H}_i (\mathcal{H} stands for ‘hypothesis’) is assumed to have a vector of parameters \mathbf{w} . A model is defined by its functional form and two probability distributions: a ‘prior’ distribution $P(\mathbf{w}|\mathcal{H}_i)$ which states what values the model’s parameters might plausibly take; and the predictions $P(D|\mathbf{w}, \mathcal{H}_i)$ that the model makes about the data D when its parameters have a particular value \mathbf{w} . Note that models with the same parameterisation but different priors over the parameters are therefore defined to be different models.

1. **Model fitting.** At the first level of inference, we assume that one model \mathcal{H}_i is true, and we infer what the model’s parameters \mathbf{w} might be given the data D . Using Bayes’ rule, the **posterior probability** of the parameters \mathbf{w} is:

$$P(\mathbf{w}|D, \mathcal{H}_i) = \frac{P(D|\mathbf{w}, \mathcal{H}_i)P(\mathbf{w}|\mathcal{H}_i)}{P(D|\mathcal{H}_i)} \quad (1)$$

In words:

$$\text{Posterior} = \frac{\text{Likelihood} \times \text{Prior}}{\text{Evidence}}$$

The normalising constant $P(D|\mathcal{H}_i)$ is commonly ignored, since it is irrelevant to the first level of inference, *i.e.* the choice of \mathbf{w} ; but it will be important in the second level of inference, and we name it the **evidence** for \mathcal{H}_i . It is common to use gradient-based methods to find the maximum of the posterior, which defines

the most probable value for the parameters, \mathbf{w}_{MP} ; it is then common to summarise the posterior distribution by the value of \mathbf{w}_{MP} , and error bars on these best fit parameters. The error bars are obtained from the curvature of the posterior; writing the Hessian $\mathbf{A} = -\nabla\nabla \log P(\mathbf{w}|D, \mathcal{H}_i)$ and Taylor-expanding the log posterior with $\Delta\mathbf{w} = \mathbf{w} - \mathbf{w}_{\text{MP}}$,

$$P(\mathbf{w}|D, \mathcal{H}_i) \simeq P(\mathbf{w}_{\text{MP}}|D, \mathcal{H}_i) \exp\left(-\frac{1}{2}\Delta\mathbf{w}^T \mathbf{A} \Delta\mathbf{w}\right) \quad (2)$$

we see that the posterior can be locally approximated as a gaussian with covariance matrix (error bars) \mathbf{A}^{-1} .²

2. **Model comparison.** At the second level of inference, we wish to infer which model is most plausible given the data. The posterior probability of each model is:

$$P(\mathcal{H}_i|D) \propto P(D|\mathcal{H}_i)P(\mathcal{H}_i) \quad (3)$$

Notice that the data-dependent term $P(D|\mathcal{H}_i)$ is the evidence for \mathcal{H}_i , which appeared as the normalising constant in (1). The second term, $P(\mathcal{H}_i)$, is a ‘subjective’ prior over our hypothesis space which expresses how plausible we thought the alternative models were before the data arrived. We will see later that this subjective part of the inference will typically be overwhelmed by the objective term, the evidence. Assuming that we have no reason to assign strongly differing priors $P(\mathcal{H}_i)$ to the alternative models, **models \mathcal{H}_i are ranked by evaluating the evidence.** Equation (3) has not been normalised because in the data modelling process we may develop new models after the data have arrived (figure 1), when an inadequacy of the first models is detected, for example. So we do not start with a completely defined hypothesis space. Inference is open-ended: we continually seek more probable models to account for the data we gather. New models are compared with previous models by evaluating the evidence for them.

The key concept of this paper is this: to assign a preference to alternative models \mathcal{H}_i , a Bayesian evaluates the evidence $P(D|\mathcal{H}_i)$. This concept is very general: the evidence can be evaluated for parametric and ‘non-parametric’ models alike; whether our data modelling task is a regression problem, a classification problem, or a density estimation problem, the evidence is the Bayesian’s transportable quantity for comparing alternative models. In all these cases the evidence naturally embodies Occam’s razor; we will examine how this works shortly.

Of course, the evidence is not the whole story if we have good reason to assign unequal priors to the alternative models \mathcal{H} . (To only use the evidence for model comparison is equivalent to using maximum likelihood for parameter estimation.) The classic example is the ‘Sure Thing’ hypothesis, © E.T Jaynes, which is the hypothesis that the data set will be D , the precise data set that actually occurred; the evidence for the Sure Thing hypothesis is huge. But Sure Thing belongs to an immense class of similar hypotheses which should all be assigned correspondingly tiny prior probabilities; so the posterior probability for Sure

²Whether this approximation is a good one or not will depend on the problem we are solving. For the interpolation models discussed in this paper, there is only a single maximum in the posterior distribution, and the gaussian approximation is exact. For more general statistical models we still expect the posterior to be dominated by locally gaussian peaks on account of the central limit theorem (Walker, 1967). Multiple maxima which arise in more complex models complicate the analysis, but Bayesian methods can still successfully be applied (Hanson *et. al.*, 1991, MacKay, 1991b, Neal, 1991).

Thing is negligible alongside any sensible model. Models like Sure Thing are rarely seriously proposed in real life, but if such models are developed then clearly we need to think about precisely what priors are appropriate. Patrick and Wallace (1982), studying the geometry of ancient stone circles (for which some people have proposed extremely elaborate theories!), discuss a practical method of assigning relative prior probabilities to alternative models by evaluating the lengths of the computer programs that decode data previously encoded under each model. This procedure introduces a second sort of Occam’s razor into the inference, namely a *prior* bias against complex models. However, this paper will not include such prior biases; we will address only the data’s preference for the alternative models, *i.e.* the evidence, and the Occam’s razor that it embodies. In the limit of large quantities of data this objective Occam’s razor will always be the more important of the two.

A modern Bayesian approach to priors

It should be pointed out that the emphasis of this modern Bayesian approach is not on the inclusion of priors into inference. There is not one significant ‘subjective prior’ in this entire paper. (If you are interested to see problems where significant subjective priors do arise see (Gull, 1989b, Skilling, 1989b).) The emphasis is that consistent degrees of preference for alternative hypotheses are represented by probabilities, and relative preferences for models are assigned by evaluating those probabilities. Historically, Bayesian analysis has been accompanied by methods to work out the ‘right’ prior $P(\mathbf{w}|\mathcal{H})$ for a problem, for example, the principles of insufficient reason and maximum entropy. The modern Bayesian however does not take a fundamentalist attitude to assigning the ‘right’ priors — many different priors can be tried; each particular prior corresponds to a different hypothesis about the way the world is. We can compare these alternative hypotheses in the light of the data by evaluating the evidence. This is the way in which alternative regularisers are compared, for example. If we try one model and obtain awful predictions, we have *learned* something. ‘A failure of Bayesian prediction is an opportunity to learn’ (Jaynes, 1986), and we are able to come back to the same data set with new models, using new priors for example.

Evaluating the evidence

Let us now explicitly study the evidence to gain insight into how the Bayesian Occam’s razor works. The evidence is the normalising constant for equation (1):

$$P(D|\mathcal{H}_i) = \int P(D|\mathbf{w}, \mathcal{H}_i) P(\mathbf{w}|\mathcal{H}_i) d\mathbf{w} \quad (4)$$

For many problems, including interpolation, it is common for the posterior $P(\mathbf{w}|D, \mathcal{H}_i) \propto P(D|\mathbf{w}, \mathcal{H}_i)P(\mathbf{w}|\mathcal{H}_i)$ to have a strong peak at the most probable parameters \mathbf{w}_{MP} (figure 3). Then the evidence can be approximated by the height of the peak of the integrand $P(D|\mathbf{w}, \mathcal{H}_i)P(\mathbf{w}|\mathcal{H}_i)$ times its width, $\Delta\mathbf{w}$:

$$\begin{aligned} P(D|\mathcal{H}_i) &\simeq \underbrace{P(D|\mathbf{w}_{\text{MP}}, \mathcal{H}_i)}_{\text{Best fit likelihood}} \underbrace{P(\mathbf{w}_{\text{MP}}|\mathcal{H}_i) \Delta\mathbf{w}}_{\text{Occam factor}} \end{aligned} \quad (5)$$

Thus the evidence is found by taking the best fit likelihood that the model can achieve and multiplying it by an ‘Occam factor’ (Gull, 1988), which is a term with magnitude less than one that penalises \mathcal{H}_i for having the parameter \mathbf{w} .




Figure 3: **The Occam factor**

This figure shows the quantities that determine the Occam factor for a hypothesis \mathcal{H}_i having a single parameter \mathbf{w} . The prior distribution (dotted line) for the parameter has width $\Delta^0 \mathbf{w}$. The posterior distribution (solid line) has a single peak at \mathbf{w}_{MP} with characteristic width $\Delta \mathbf{w}$. The Occam factor is $\frac{\Delta \mathbf{w}}{\Delta^0 \mathbf{w}}$.

Interpretation of the Occam factor

The quantity $\Delta \mathbf{w}$ is the posterior uncertainty in \mathbf{w} . Imagine for simplicity that the prior $P(\mathbf{w}|\mathcal{H}_i)$ is uniform on some large interval $\Delta^0 \mathbf{w}$, representing the range of values of \mathbf{w} that \mathcal{H}_i thought possible before the data arrived (figure 3). Then $P(\mathbf{w}_{\text{MP}}|\mathcal{H}_i) = \frac{1}{\Delta^0 \mathbf{w}}$, and

$$\text{Occam factor} = \frac{\Delta \mathbf{w}}{\Delta^0 \mathbf{w}},$$

i.e. the ratio of the posterior accessible volume of \mathcal{H}_i 's parameter space to the prior accessible volume, or the factor by which \mathcal{H}_i 's hypothesis space collapses when the data arrive (Gull, 1988, Jeffreys, 1939). The model \mathcal{H}_i can be viewed as being composed of a certain number of equivalent submodels, of which only one survives when the data arrive. The Occam factor is the inverse of that number. The log of the Occam factor can be interpreted as the amount of information we gain about the model when the data arrive.

Typically, a complex model with many parameters, each of which is free to vary over a large range $\Delta^0 \mathbf{w}$, will be penalised with a larger Occam factor than a simpler model. The Occam factor also provides a penalty for models which have to be finely tuned to fit the data; the Occam factor promotes models for which the required precision of the parameters $\Delta \mathbf{w}$ is coarse. The Occam factor is thus a measure of complexity of the model, but unlike the V-C dimension or algorithmic complexity, it relates to the complexity of the predictions that the model makes in data space; therefore it depends on the number of data points and other properties of the data set. Which model achieves the greatest evidence is determined by a trade-off between minimising this natural complexity measure and minimising the data misfit.

Occam factor for several parameters

If \mathbf{w} is k -dimensional, and if the posterior is well approximated by a gaussian, the Occam factor is given by the determinant of the gaussian's covariance matrix:

$$\begin{aligned} P(D|\mathcal{H}_i) &\simeq \underbrace{P(D|\mathbf{w}_{\text{MP}}, \mathcal{H}_i)}_{\text{Best fit likelihood}} \underbrace{P(\mathbf{w}_{\text{MP}}|\mathcal{H}_i) (2\pi)^{k/2} \det^{-\frac{1}{2}} \mathbf{A}}_{\text{Occam factor}} \end{aligned} \quad (6)$$

where $\mathbf{A} = -\nabla\nabla \log P(\mathbf{w}|D, \mathcal{H}_i)$, the Hessian which we already evaluated when we calculated the error bars on \mathbf{w}_{MP} . As the amount of data collected, N , increases, this gaussian approximation is expected to become increasingly accurate on account of the central limit theorem (Walker, 1967). For the linear interpolation models discussed in this paper, this gaussian expression is exact for any N .

Comments

- Bayesian model selection is a simple extension of maximum likelihood model selection: **the evidence is obtained by multiplying the best fit likelihood by the Occam factor.**

To evaluate the Occam factor all we need is the Hessian \mathbf{A} , if the gaussian approximation is good. Thus the Bayesian method of model comparison by evaluating the evidence is computationally no more demanding than the task of finding for each model the best fit parameters and their error bars.

- It is common for there to be degeneracies in models with many parameters, *i.e.* several equivalent parameters could be relabelled without affecting the likelihood. In these cases, the right hand side of equation (6) should be multiplied by the degeneracy of \mathbf{w}_{MP} to give the correct estimate of the evidence.
- ‘Minimum description length’ (MDL) methods are closely related to this Bayesian framework (Rissanen, 1978, Wallace and Boulton, 1968, Wallace and Freeman, 1987). The log evidence $\log_2 P(D|\mathcal{H}_i)$ is the number of bits in the ideal shortest message that encodes the data D using model \mathcal{H}_i . Akaike’s (1970) criterion can be viewed as an approximation to MDL (Schwarz, 1978, Zellner, 1984). Any implementation of MDL necessitates approximations in evaluating the length of the ideal shortest message. Although some of the earliest work on complex model comparison involved the MDL framework (Patrick and Wallace, 1982), I can see no advantage in MDL, and recommend that the evidence should be approximated directly.
- It should be emphasised that the Occam factor has nothing to do with how computationally complex it is to *use* a model. The evidence is a measure of *plausibility* of a model. How much CPU time it takes to use each model is certainly an interesting issue which might bias our decisions towards simpler models, but Bayes’ rule does not address that. Choosing between models on the basis of how many function calls they need is an exercise in *decision theory*, which is not addressed in this paper. Once the probabilities described above have been inferred, optimal actions can be chosen using standard decision theory with a suitable utility function.

3 The noisy interpolation problem

Bayesian interpolation through *noise-free* data has been studied by Skilling and Sibisi (1991). In this paper I study the problem of interpolating through data where the dependent variables are assumed to be noisy (a task also known as ‘regression,’ ‘curve-fitting,’ ‘signal estimation,’ or, in the neural networks community, ‘learning’). I am not examining the case where the independent variables are also noisy. This different and more difficult problem has been studied for the case of straight line-fitting by Gull (1989b).

Let us assume that the data set to be interpolated is a set of pairs $D = \{x_m, t_m\}$, where $m = 1 \dots N$ is a label running over the pairs. For simplicity I will treat x and t as scalars, but the method generalises to the multidimensional case. To define a linear interpolation model, a set of k fixed basis functions³ $\mathcal{A} = \{\phi_h(x)\}$ is chosen, and the interpolated function is assumed to have the form:

$$y(x) = \sum_{h=1}^k w_h \phi_h(x) \quad (7)$$

where the parameters w_h are to be inferred from the data. The data set is modelled as deviating from this mapping under some additive noise process:

$$t_m = y(x_m) + \nu_m \quad (8)$$

If ν is modelled as zero-mean gaussian noise with standard deviation σ_ν , then the probability of the data⁴ given the parameters \mathbf{w} is:

$$P(D|\mathbf{w}, \beta, \mathcal{A}) = \frac{\exp(-\beta E_D(D|\mathbf{w}, \mathcal{A}))}{Z_D(\beta)} \quad (9)$$

where $\beta = 1/\sigma_\nu^2$, $E_D = \sum_m \frac{1}{2}(y(x_m) - t_m)^2$, and $Z_D = (2\pi/\beta)^{N/2}$. $P(D|\mathbf{w}, \beta, \mathcal{A})$ is called the likelihood. It is well known that finding the maximum likelihood parameters \mathbf{w}_{ML} may be an ‘ill-posed’ problem. That is, the \mathbf{w} that minimises E_D is underdetermined and/or depends sensitively on the details of the noise in the data; the maximum likelihood interpolant in such cases oscillates wildly so as to fit the noise. Thus it is clear that to complete an interpolation model we need a prior \mathcal{R} that expresses the sort of smoothness we expect the interpolant $y(x)$ to have. A model may have a prior of the form

$$P(y|\mathcal{R}, \alpha) = \frac{\exp(-\alpha E_y(y|\mathcal{R}))}{Z_y(\alpha)} \quad (10)$$

where E_y might be for example the functional $E_y = \int y''(x)^2 dx$ (which is the regulariser for cubic spline interpolation⁵). The parameter α is a measure of how smooth $f(x)$ is expected to be. Such a prior can also be written as a prior on the parameters \mathbf{w} :

$$P(\mathbf{w}|\mathcal{A}, \mathcal{R}, \alpha) = \frac{\exp(-\alpha E_W(\mathbf{w}|\mathcal{A}, \mathcal{R}))}{Z_W(\alpha)} \quad (11)$$

where $Z_W = \int d^k \mathbf{w} \exp(-\alpha E_W)$. E_W (or E_y) is commonly referred to as a regularising function.

The interpolation model is now complete, consisting of a choice of basis functions \mathcal{A} , a noise model with parameter β , and a prior (regulariser) \mathcal{R} , with regularising constant α .

³the case of *adaptive* basis functions, also known as feedforward neural networks, is examined in a companion paper.

⁴Strictly, this probability should be written $P(\{t_m\}|\{x_m\}, \mathbf{w}, \beta, \mathcal{A})$, since these interpolation models do not predict the distribution of input variables $\{x_m\}$; this liberty of notation will be taken throughout this paper and its companion.

⁵Strictly, this particular prior may be improper because a $y(x)$ of the form $w_1 x + w_0$ is not constrained by this prior.

The first level of inference

If α and β are known, then the posterior probability of the parameters \mathbf{w} is:⁶

$$P(\mathbf{w}|D, \alpha, \beta, \mathcal{A}, \mathcal{R}) = \frac{P(D|\mathbf{w}, \beta, \mathcal{A})P(\mathbf{w}|\alpha, \mathcal{A}, \mathcal{R})}{P(D|\alpha, \beta, \mathcal{A}, \mathcal{R})} \quad (12)$$

Writing⁷

$$M(\mathbf{w}) = \alpha E_W + \beta E_D, \quad (13)$$

the posterior is

$$P(\mathbf{w}|D, \alpha, \beta, \mathcal{A}, \mathcal{R}) = \frac{\exp(-M(\mathbf{w}))}{Z_M(\alpha, \beta)} \quad (14)$$

where $Z_M(\alpha, \beta) = \int d^k \mathbf{w} \exp(-M)$. We see that minimising the combined objective function M corresponds to finding the *most probable interpolant*, \mathbf{w}_{MP} . Error bars on the best fit interpolant⁸ can be obtained from the Hessian of M , $\mathbf{A} = \nabla \nabla M$, evaluated at \mathbf{w}_{MP} .

This is the well known Bayesian view of regularisation (Poggio *et. al.*, 1985, Titterton, 1985).

Bayes can do a lot more than just provide an interpretation for regularisation. What we have described so far is just the first of three levels of inference. (The second level described in sections 1 and 2, ‘model comparison,’ splits into a second and a third level for this problem, because each interpolation model is made up of a continuum of sub-models with different values of α and β .) At the second level, Bayes allows us to objectively assign values to α and β , which are commonly unknown *a priori*. At the third, Bayes enables us to quantitatively rank alternative basis sets \mathcal{A} , alternative regularisers (priors) \mathcal{R} , and, in principle, alternative noise models. Furthermore, we can quantitatively compare interpolation under any model \mathcal{A}, \mathcal{R} with other interpolation and learning models such as neural networks, if a similar Bayesian approach is applied to them. Neither the second nor the third level of inference can be successfully executed without Occam’s razor.

The Bayesian theory of the second and third levels of inference has only recently been worked out (Gull, 1989a), and this paper’s goal is to review that framework. Section 4 will describe the Bayesian method of inferring α and β ; section 5 will describe Bayesian model comparison for the interpolation problem. Both these inference problems are solved by evaluating the appropriate evidence.

⁶The regulariser α, \mathcal{R} has been omitted from the conditioning variables in the likelihood because the data distribution does not depend on the prior once \mathbf{w} is known. Similarly the prior does not depend on β .

⁷The name M stands for ‘misfit’; it will be demonstrated later that M is the natural measure of misfit, rather than $\chi_D^2 = 2\beta E_D$.

⁸These error bars represent the uncertainty of the interpolant, and should not be confused with the typical scatter of noisy data points relative to the interpolant.

4 Selection of parameters α and β

Typically, α is not known *a priori*, and often β is also unknown. As α is varied, the properties of the best fit (most probable) interpolant vary. Assume that we are using a prior that encourages smoothness, and imagine that we interpolate at a very large value of α ; then this will constrain the interpolant to be very smooth and flat, and it will not fit the data at all well (figure 4a). As α is decreased, the interpolant starts to fit the data better (figure 4b). If α is made even smaller, the interpolant oscillates wildly so as to overfit the noise in the data (figure 4c). The choice of the ‘best’ value of α is our first ‘Occam’s razor’ problem: large values of α correspond to simple models which make constrained and precise predictions, saying ‘the interpolant is expected to not have extreme curvature anywhere’; a tiny value of α corresponds to the more powerful and flexible model that says ‘the interpolant could be anything at all, our prior belief in smoothness is very weak.’ The task is to find a value of α which is small enough that the data are fitted but not so small that they are overfitted. For more severely ill-posed problems such as deconvolution, the precise value of the regularising parameter is increasingly important. Orthodox statistics has ways of assigning values to such parameters, based for example on misfit criteria, the use of test data, and cross-validation. Gull (1989a) has demonstrated why the popular use of misfit criteria is incorrect and how Bayes sets these parameters. The use of test data may be an unreliable technique unless large quantities of data are available. Cross-validation, the orthodox ‘method of choice’ (Eubank, 1988), will be discussed more in section 6 and (MacKay, 1991b). I will explain the Bayesian method of inferring α and β after first reviewing some statistics of misfit.

Misfit, χ^2 , and the effect of parameter measurements

For N independent gaussian variables with mean μ and standard deviation σ , the statistic $\chi^2 = \sum (x - \mu)^2 / \sigma^2$ is a measure of misfit. If μ is known *a priori*, χ^2 has expectation $N \pm \sqrt{N}$. However, if μ is fitted from the data by setting $\mu = \bar{x}$, we ‘use up a degree of freedom,’ and χ^2 has expectation $N - 1$. In the second case μ is a ‘well-measured parameter.’ When a parameter is determined by the data in this way it is unavoidable that the parameter fits some of the noise in the data as well. That is why the expectation of χ^2 is reduced by one. This is the basis of the distinction between the σ_N and σ_{N-1} buttons on your calculator. It is common for this distinction to be ignored, but in cases such as interpolation where the number of free parameters is similar to the number of data points, it is essential to find and make the analogous distinction. It will be demonstrated that the Bayesian choices of both α and β are most simply expressed in terms of the effective number of well-measured parameters, γ , to be derived below.

Misfit criteria are ‘principles’ which set parameters like α and β by requiring that χ^2 should have a particular value. The discrepancy principle requires $\chi^2 = N$. Another principle requires $\chi^2 = N - k$, where k is the number of free parameters. We will find that an intuitive misfit criterion arises for the most probable value of β ; on the other hand, the Bayesian choice of α will be unrelated to the value of the misfit.




Figure 4: **How the best interpolant depends on α**

These figures introduce a data set, 'X,' which is interpolated with a variety of models in this paper. Notice that the density of data points is not uniform on the x -axis. In the three figures the data set is interpolated using a radial basis function model with a basis of 60 equally spaced Cauchy functions, all with radius 0.2975. The regulariser is $E_W = \frac{1}{2} \sum w^2$, where w are the coefficients of the basis functions. Each figure shows the most probable interpolant for a different value of α : a) 6000; b) 2.5; c) 10^{-7} . Note at the extreme values how the data are oversmoothed and overfitted respectively. Assuming a flat prior, $\alpha = 2.5$ is the *most probable* value of α . In b), the most probable interpolant is displayed with its 1σ error bars, which represent how uncertain we are about the interpolant at each point, under the assumption that the interpolation model and the value of α are correct. Notice how the error bars increase in magnitude where the data are sparse. The error bars do not get bigger near the datapoint close to (1,0), because the radial basis function model does not expect sharp discontinuities; the error bars are obtained *assuming the model is correct*, so that point is interpreted as an improbable outlier.

Bayesian choice of α and β

To infer from the data what value α and β should have,⁹ Bayesians evaluate the posterior probability distribution:

$$P(\alpha, \beta | D, \mathcal{A}, \mathcal{R}) = \frac{P(D | \alpha, \beta, \mathcal{A}, \mathcal{R}) P(\alpha, \beta)}{P(D | \mathcal{A}, \mathcal{R})} \quad (15)$$

The data dependent term $P(D | \alpha, \beta, \mathcal{A}, \mathcal{R})$ has already appeared earlier as the normalising constant in equation (12), and it is called the evidence for α and β . Similarly the normalising constant of (15) is called the evidence for \mathcal{A}, \mathcal{R} , and it will turn up later when we compare alternative models \mathcal{A}, \mathcal{R} in the light of the data.

If $P(\alpha, \beta)$ is a flat prior¹⁰ (which corresponds to the statement that we don't know what value α and β should have), the evidence is the function that we use to assign a preference to alternative values of α and β . It is given in terms of the normalising constants defined earlier by

$$P(D | \alpha, \beta, \mathcal{A}, \mathcal{R}) = \frac{Z_M(\alpha, \beta)}{Z_W(\alpha) Z_D(\beta)}. \quad (16)$$

Occam's razor is implicit in this formula: if α is small, the large freedom in the range of possible values of \mathbf{w} is automatically penalised by the consequent large value of Z_W ; models that fit the data well achieve a large value of Z_M . The optimum value of α achieves a compromise between fitting the data well and being a simple model.

Now to assign a preference to (α, β) , our computational task is to evaluate the three integrals Z_M , Z_W and Z_D . We will come back to this task in a moment.

But that sounds like determining your prior after the data have arrived!

This is an aside which can be omitted on a first reading. When I first heard the preceding explanation of Bayesian regularisation I was discontent because it seemed that the prior is being chosen from an ensemble of possible priors *after* the data have arrived. To be precise, as described above, the most probable value of α is selected; then the prior corresponding to that value of α alone is used to infer what the interpolant might be. This is not how Bayes would have us infer what the interpolant is. It is the combined ensemble of priors that define our prior, and we should integrate over this ensemble when we do inference.¹¹ Let us work out what happens if we follow this proper approach. The preceding method of using only the most probable prior will emerge as a good approximation.

The true posterior $P(\mathbf{w} | D, \mathcal{A}, \mathcal{R})$ is obtained by integrating over α and β :

$$P(\mathbf{w} | D, \mathcal{A}, \mathcal{R}) = \int P(\mathbf{w} | D, \alpha, \beta, \mathcal{A}, \mathcal{R}) P(\alpha, \beta | D, \mathcal{A}, \mathcal{R}) d\alpha d\beta \quad (17)$$

⁹Note that it is not satisfactory to simply maximise the likelihood simultaneously over \mathbf{w} , α and β ; the likelihood has a skew peak such that the maximum likelihood value for the parameters is not in the same place as most of the posterior probability (Gull, 1989a). To get a feeling for this here is a more familiar problem: examine the posterior probability for the parameters of a gaussian (μ, σ) given N samples: the maximum likelihood value for σ is σ_N , but the most probable value for σ (found by integrating over μ) is σ_{N-1} . It should be emphasised that this distinction has nothing to do with the prior over the parameters, which is flat here. It is the process of marginalisation that corrects the bias of maximum likelihood.

¹⁰Since α and β are scale parameters, this prior should be understood as a flat prior over $\log \alpha$ and $\log \beta$.

¹¹It is remarkable that Laplace almost got this right in 1774 (Stigler, 1986); when inferring the mean of a Laplacian distribution, he both inferred the posterior probability of a nuisance parameter like β in (15), and then attempted to integrate out the nuisance parameter as in equation (17).




Figure 5: **Choosing α**

a) **The evidence as a function of α :** Using the same radial basis function model as in figure 4, this graph shows the log evidence as a function of α , and shows the functions which make up the log evidence, namely the data misfit $\chi^2_D = 2\beta E_D$, the weight penalty term $\chi^2_w = 2\alpha E_W$, and the log of the Occam factor $(2\pi)^{k/2} \det^{-\frac{1}{2}} \mathbf{A} / Z_W(\alpha)$.

b) **Criteria for optimising α :** This graph shows the log evidence as a function of α , and the functions whose intersection locates the evidence maximum: the number of good parameter measurements γ , and χ^2_w . Also shown is the test error (rescaled) on two test sets; finding the test error minimum is an alternative criterion for setting α . Both test sets were more than twice as large in size as the interpolated data set. Note how the point at which $\chi^2_w = \gamma$ is clear and unambiguous, which cannot be said for the minima of the test energies. The evidence gives α a 1- σ confidence interval of [1.3, 5.0]. The test error minima are more widely distributed because of finite sample noise.

In words, the posterior probability over \mathbf{w} can be written as a linear combination of the posteriors for all values of α, β . Each posterior density is weighted by the probability of α, β given the data, which appeared in (15). This means that if $P(\alpha, \beta | D, \mathcal{A}, \mathcal{R})$ has a dominant peak at $\hat{\alpha}, \hat{\beta}$, then the true posterior $P(\mathbf{w} | D, \mathcal{A}, \mathcal{R})$ will be dominated by the density $P(\mathbf{w} | D, \hat{\alpha}, \hat{\beta}, \mathcal{A}, \mathcal{R})$. As long as the properties of the posterior $P(\mathbf{w} | D, \alpha, \beta, \mathcal{A}, \mathcal{R})$ do not change rapidly with α, β near $\hat{\alpha}, \hat{\beta}$ and the peak in $P(\alpha, \beta | D, \mathcal{A}, \mathcal{R})$ is strong, we are justified in using the approximation:

$$P(\mathbf{w} | D, \mathcal{A}, \mathcal{R}) \simeq P(\mathbf{w} | D, \hat{\alpha}, \hat{\beta}, \mathcal{A}, \mathcal{R}) \quad (18)$$

This approximation is valid under the same conditions as in footnote 12. It is a matter of ongoing research to develop computational methods for cases where this approximation is invalid (Sibisi and Skilling, personal communication).

Evaluating the evidence

Let us return to our train of thought at equation (16). To evaluate the evidence for α, β , we want to find the integrals Z_M , Z_W and Z_D . Typically the most difficult integral to evaluate is Z_M .

$$Z_M(\alpha, \beta) = \int d^k \mathbf{w} \exp(-M(\mathbf{w}, \alpha, \beta))$$

If the regulariser \mathcal{R} is a quadratic functional (and the favourites are), then E_D and E_W are quadratic functions of \mathbf{w} , and we can evaluate Z_M exactly. Letting $\nabla\nabla E_W = \mathbf{C}$ and $\nabla\nabla E_D = \mathbf{B}$ then using $\mathbf{A} = \alpha\mathbf{C} + \beta\mathbf{B}$, we have:

$$M = M(\mathbf{w}_{\text{MP}}) + \frac{1}{2}(\mathbf{w} - \mathbf{w}_{\text{MP}})^T \mathbf{A}(\mathbf{w} - \mathbf{w}_{\text{MP}})$$

where $\mathbf{w}_{\text{MP}} = \mathbf{A}^{-1}\mathbf{B}\mathbf{w}_{\text{ML}}$. This means that Z_M is the gaussian integral:

$$Z_M = e^{-M_{\text{MP}}} (2\pi)^{k/2} \det^{-\frac{1}{2}} \mathbf{A} \quad (19)$$

In many cases where the regulariser is not quadratic (for example, entropy-based), this gaussian approximation is still servicable (Gull, 1989a). Thus we can write the log evidence for α and β as:

$$\log P(D|\alpha, \beta, \mathcal{A}, \mathcal{R}) = -\alpha E_W^{\text{MP}} - \beta E_D^{\text{MP}} - \frac{1}{2} \log \det \mathbf{A} - \log Z_W(\alpha) - \log Z_D(\beta) + \frac{k}{2} \log 2\pi \quad (20)$$

The term βE_D^{MP} represents the misfit of the interpolant to the data. The three terms $-\alpha E_W^{\text{MP}} - \frac{1}{2} \log \det \mathbf{A} - \log Z_W(\alpha)$ constitute the log of the ‘Occam factor’ penalising over-powerful values of α : the ratio $(2\pi)^{k/2} \det^{-\frac{1}{2}} \mathbf{A} / Z_W(\alpha)$ is the ratio of the posterior accessible volume in parameter space to the prior accessible volume, and the term αE_W^{MP} measures how far \mathbf{w}_{MP} is from its null value. Figure 5a illustrates the behaviour of these various terms as a function of α for the same radial basis function model as illustrated in figure 4.

Now we could just proceed to evaluate the evidence numerically as a function of α and β , but a more deep and fruitful understanding of this problem is possible.

Properties of the evidence maximum

The maximum over α, β of $P(D|\alpha, \beta, \mathcal{A}, \mathcal{R}) = Z_M(\alpha, \beta) / (Z_W(\alpha) Z_D(\beta))$ has some remarkable properties which give deeper insight into this Bayesian approach. The results of this section are useful both numerically and intuitively.

Following Gull (1989a), we transform to the basis in which the Hessian of E_W is the identity, $\nabla\nabla E_W = \mathbf{I}$. This transformation is simple in the case of quadratic E_W : rotate into the eigenvector basis of \mathbf{C} and stretch the axes so that the quadratic form E_W becomes homogeneous. This is the natural basis for the prior. I will continue to refer to the parameter vector in this basis as \mathbf{w} , so from here on $E_W = \frac{1}{2} \sum w_i^2$. Using $\nabla\nabla M = \mathbf{A}$ and $\nabla\nabla E_D = \mathbf{B}$ as above, we differentiate the log evidence with respect to α and β so as to find the condition that is satisfied at the maximum. The log evidence, from (20), is:

$$\log P(D|\alpha, \beta, \mathcal{A}, \mathcal{R}) = -\alpha E_W^{\text{MP}} - \beta E_D^{\text{MP}} - \frac{1}{2} \log \det \mathbf{A} + \frac{k}{2} \log \alpha + \frac{N}{2} \log \beta - \frac{N}{2} \log 2\pi \quad (21)$$

First, differentiating with respect to α , we need to evaluate $\frac{d}{d\alpha} \log \det \mathbf{A}$. Using $\mathbf{A} = \alpha\mathbf{I} + \beta\mathbf{B}$,

$$\begin{aligned} \frac{d}{d\alpha} \log \det \mathbf{A} &= \text{Trace} \left(\mathbf{A}^{-1} \frac{d\mathbf{A}}{d\alpha} \right) \\ &= \text{Trace} \left(\mathbf{A}^{-1} \mathbf{I} \right) = \text{Trace} \mathbf{A}^{-1} \end{aligned}$$

This result is exact if E_W and E_D are quadratic. Otherwise this result is an approximation, omitting terms in $\partial\mathbf{B}/\partial\alpha$. Now, differentiating (21) and setting the derivative to zero, we obtain the following condition for the most probable value of α :

$$2\alpha E_W^{\text{MP}} = k - \alpha \text{Trace} \mathbf{A}^{-1} \quad (22)$$




Figure 6: **Good and bad parameter measurements**

Let w_1 and w_2 be the components in parameter space in two directions parallel to eigenvectors of the data matrix \mathbf{B} . The circle represents the characteristic prior distribution for \mathbf{w} . The ellipse represents a characteristic contour of the likelihood, centred on the maximum likelihood solution \mathbf{w}_{ML} . \mathbf{w}_{MP} represents the most probable parameter vector. w_1 is a direction in which λ_1 is small compared to α , *i.e.* the data have no strong preference about the value of w_1 ; w_1 is a poorly measured parameter, and the term $\frac{\lambda_1}{\lambda_1 + \alpha}$ is close to zero. w_2 is a direction in which λ_1 is large; w_2 is well determined by the data, and the term $\frac{\lambda_2}{\lambda_2 + \alpha}$ is close to one.

The quantity on the left is the dimensionless measure of the amount of structure introduced into the parameters by the data, *i.e.* how much the fitted parameters differ from their null value. It can be interpreted as the χ^2 of the parameters, since it is equal to $\chi_W^2 = \sum w_i^2 / \sigma_w^2$, with $\alpha = 1/\sigma_w^2$.

The quantity on the right of (22) is called the number of good parameter measurements, γ , and has value between 0 and k . It can be written in terms of the eigenvalues of $\beta\mathbf{B}$, λ_a , where the subscript a runs over the k eigenvectors. The eigenvalues of \mathbf{A} are $\lambda_a + \alpha$, so we have:

$$\gamma = k - \alpha \text{Trace} \mathbf{A}^{-1} = k - \sum_{a=1}^k \frac{\alpha}{\lambda_a + \alpha} = \sum_{a=1}^k \frac{\lambda_a}{\lambda_a + \alpha} \quad (23)$$

Each eigenvalue λ_a measures how strongly one parameter is determined by the data. The constant α measures how strongly the parameters are determined by the prior. The a th term $\gamma_a = \lambda_a / (\lambda_a + \alpha)$ is a number between 0 and 1 which measures the strength of the data relative to the prior (figure 6): the components of \mathbf{w}_{MP} are given by $\mathbf{w}_{\text{MP}a} = \gamma_a \mathbf{w}_{\text{ML}a}$.

A direction in parameter space for which λ_a is small compared to α does not contribute to the number of good parameter measurements. γ is thus a measure of the effective number of parameters which are well determined by the data. As $\alpha/\beta \rightarrow 0$, γ increases from 0 to k . The condition (22) for the most probable value of α can therefore be interpreted as an estimation of the variance σ_w^2 of the gaussian distribution from which the weights are drawn, based on γ effective samples from that distribution: $\sigma_w^2 = \sum w_i^2 / \gamma$.

This concept is not only important for locating the optimum value of α : it is only the γ good parameter measurements which are expected to contribute to the reduction of the data misfit that occurs when a model is fitted to noisy data. In the process of fitting \mathbf{w} to the data, it is unavoidable that some fitting of the model to noise will occur, because some components of the noise are indistinguishable from real data. Typically, one unit (χ^2) of noise will be fitted for every well-determined parameter. Poorly determined parameters are determined by the regulariser only, so they do not reduce χ_D^2 in this way. We will now examine how this concept enters into the Bayesian choice of β .

Recall that the expectation of the χ^2 misfit between the true interpolant and the data is N . However we do not know the true interpolant, and the only misfit measure we have access

to is the χ^2 between the *inferred* interpolant and the data, $\chi_D^2 = 2\beta E_D$. The ‘discrepancy principle’ of orthodox statistics states that the model parameters should be adjusted so as to make $\chi_D^2 = N$. Work on un-regularised least squares regression suggests that we should estimate the noise level so as to set $\chi_D^2 = N - k$, where k is the number of free parameters. Let us find out the opinion of Bayes’ rule on this matter.

We differentiate the log evidence (21) with respect to β and obtain, setting the derivative to zero:

$$2\beta E_D = N - \gamma \quad (24)$$

Thus the most probable noise estimate, $\hat{\beta}$, does not satisfy $\chi_D^2 = N$ or $\chi_D^2 = N - k$; rather, $\chi_D^2 = N - \gamma$. This Bayesian estimate of noise level naturally takes into account the fact that the parameters which have been determined by the data inevitably suppress some of the noise in the data, while the poorly measured parameters do not. Note that the value of χ_D^2 only enters into the determination of β : misfit criteria have no role in the Bayesian choice of α (Gull, 1989a).

In summary, at the optimum value of α and β , $\chi_W^2 = \gamma$, $\chi_D^2 = N - \gamma$. Notice that this implies that the total misfit $M = \alpha E_W + \beta E_D$ satisfies the simple equation $2M = N$.

The interpolant resulting from the Bayesian choice of α is illustrated by figure 4b. Figure 5b illustrates the functions involved with the Bayesian choice of α , and compares them with the ‘test error’ approach. Demonstration of the Bayesian choice of β is omitted, since it is straightforward; β is fixed to its true value for the demonstrations in this paper. Inference of an input-dependent noise level $\beta(x)$ will be demonstrated in a future publication.

These results generalise to the case where there are two or more separate regularisers with independent regularising constants $\{\alpha_c\}$ (Gull, 1989a). In this case, each regulariser has a number of good parameter measurements γ_c associated with it. Multiple regularisers will be used in the companion paper on neural networks.

Finding the evidence maximum with a head on approach would involve evaluating $\det \mathbf{A}$ while searching over α, β ; the above results (22,24) enable us to speed up this search (for example by the use of re-estimation formulae like $\alpha := \gamma/2E_W$) and replace the evaluation of $\det \mathbf{A}$ by the evaluation of $\text{Trace} \mathbf{A}^{-1}$. For large dimensional problems where this task is demanding, Skilling (1989b) has developed methods for estimating $\text{Trace} \mathbf{A}^{-1}$ statistically in k^2 time.

5 Model comparison

To rank alternative basis sets \mathcal{A} and regularisers (priors) \mathcal{R} in the light of the data, we examine the posterior probabilities:

$$P(\mathcal{A}, \mathcal{R} | D) \propto P(D | \mathcal{A}, \mathcal{R}) P(\mathcal{A}, \mathcal{R}) \quad (25)$$

The data-dependent term, the evidence for \mathcal{A}, \mathcal{R} , appeared earlier as the normalising constant in (15), and it is evaluated by integrating the evidence for (α, β) :

$$P(D | \mathcal{A}, \mathcal{R}) = \int P(D | \mathcal{A}, \mathcal{R}, \alpha, \beta) P(\alpha, \beta) d\alpha d\beta \quad (26)$$

Assuming that we have no reason to assign strongly differing priors $P(\mathcal{A}, \mathcal{R})$, alternative models \mathcal{A}, \mathcal{R} are ranked just by examining the evidence. The evidence can also be compared with the evidence found by an equivalent Bayesian analysis of other learning and interpolation models so as to allow the data to assign a preference to the alternative models.

Notice as pointed out earlier that this modern Bayesian framework includes no emphasis on defining the ‘right’ prior \mathcal{R} with which we ought to interpolate. Rather, we invent as many priors (regularisers) as we want, and allow the data to tell us which prior is *most probable*. Having said this, I would still recommend that the ‘maximum entropy principle’ and other respected guides should be consulted when inventing these priors (see Gull, 1988, for example).

Evaluating the evidence for \mathcal{A}, \mathcal{R}

As α and β vary, a single evidence maximum is obtained, at $\hat{\alpha}, \hat{\beta}$ (at least for quadratic E_D and E_W). The evidence maximum is often well approximated¹² by a separable gaussian, and differentiating (21) twice we obtain gaussian error bars for $\log \alpha$ and $\log \beta$:

$$\begin{aligned} (\Delta \log \alpha)^2 &\simeq 2/\gamma \\ (\Delta \log \beta)^2 &\simeq 2/(N - \gamma) \end{aligned}$$

Putting these error bars into (26), we obtain the evidence.¹³

$$P(D|\mathcal{A}, \mathcal{R}) \simeq P(D|\hat{\alpha}, \hat{\beta}, \mathcal{A}, \mathcal{R}) P(\hat{\alpha}, \hat{\beta}) 2\pi \Delta \log \alpha \Delta \log \beta \quad (27)$$

How is the prior $P(\hat{\alpha}, \hat{\beta})$ assigned? This is the first time in this paper that we have met one of the infamous ‘subjective priors’ which are supposed to plague Bayesian methods. Here are some answers to this question. (a) Any other method of assigning a preference to alternatives must implicitly assign such priors. Bayesians adopt the healthy attitude of not sweeping them under the carpet. (b) With some thought, reasonable values can usually be assigned to subjective priors, and the degree of reasonable subjectivity in these assignments can be quantified. For example, a reasonable prior on an unknown standard deviation states that σ is unknown over a range of (3 ± 2) orders of magnitude. This prior contributes a subjectivity of about ± 1 to the value of the log evidence. This degree of subjectivity is often negligible compared to the log evidence differences. (c) In the noisy interpolation example, all models considered include the free parameters α and β . So in this paper I do not need to assign a value to $P(\hat{\alpha}, \hat{\beta})$; I assume that it is a flat prior (flat over $\log \alpha$ and $\log \beta$, since α and β are scale parameters) which cancels out when we compare alternative interpolation models.

6 Demonstration

These demonstrations will use two one-dimensional data sets, in imitation of (Sibisi, 1991). The first data set, ‘X,’ has discontinuities in derivative (figure 4), and the second is a smoother data set, ‘Y’ (figure 8). In all the demonstrations, β was not left as a free parameter, but was fixed to its known true value.

The Bayesian method of setting α , assuming a single model is correct, has already been demonstrated, and quantified error bars have been placed on the most probable interpolant (figure 4). The method of evaluating the error bars is to use the posterior covariance matrix of the parameters w_h , \mathbf{A}^{-1} , to get the variance on $y(x)$, which for any x is a linear function

¹²This approximation is valid when, in the spectrum of eigenvalues of $\beta \mathbf{B}$, the number of eigenvalues within ϵ -fold of $\hat{\alpha}$ is $O(1)$.

¹³There are analytic methods for performing such integrals over β (Bretthorst, 1990).

of the parameters, $y(x) = \sum_h \phi_h(x)w_h$. The error bars at a single point x are given by $\text{var } y(x) = \phi^T \mathbf{A}^{-1} \phi$. Actually we have access to the full covariance information for the entire interpolant, not just the pointwise error bars. It is possible to visualise the *joint* error bars on the interpolant by making typical samples from the posterior distribution, performing a random walk around the posterior ‘bubble’ in parameter space (Sibisi, 1991, Skilling *et. al.*, 1991). Figure 8 shows data set Y interpolated by three typical interpolants found by random sampling from the posterior distribution. These error bar properties are found under the assumption that the model is correct; so it is possible for the true interpolant to lie significantly outside the error bars of a poor model.

In this section Bayesian model comparison will be demonstrated first with models differing only in the number of free parameters (for example polynomials of different degrees), then with comparisons between models as disparate as splines, radial basis functions and feedforward neural networks. For each individual model, the value of α is optimised, and the evidence is evaluated by integrating over α using the gaussian approximation. All logarithms are to base e .

Legendre polynomials: Occam’s razor for the number of basis functions

Figure 7a shows the evidence for Legendre polynomials of different degrees for data set X. The basis functions were chosen to be orthonormal on an interval enclosing the data, and a regulariser of the form $E_W = \sum \frac{1}{2} w_h^2$ was used.

Notice that an evidence maximum is obtained: beyond a certain number of terms, the evidence starts to decrease. This is the Bayesian Occam’s razor at work. The additional terms make the model more powerful, able to make more predictions. This power is automatically penalised. Notice the characteristic shape of the ‘Occam hill.’ On the left, the hill is steep as the over-simple models fail to fit the data; the penalty for misfitting the data scales as N , the number of data measurements. The other side of the hill is much less steep; the log Occam factors here only scale as $k \log N$, where k is the number of parameters. We note in table 1 the value of the maximum evidence achieved by these models, and move on to alternative models.

The choice of orthonormal Legendre polynomials described above was motivated by a maximum entropy argument (Gull, 1988). Models using other polynomial basis sets have also been tried. For less well motivated basis sets such as Hermite polynomials, it was found that the Occam factors were far bigger and the evidence was substantially smaller. If the size of the Occam factor increases rapidly with over-parameterisation, it is generally a sign that the space of alternative models is poorly matched to the problem.

Fixed radial basis functions

For a radial basis function or ‘kernel’ model, the basis functions are $\phi_h(x) = g((x - x_h)/r)/r$; here the x_h are equally spaced over the range of interest. I examine two choices of g : a gaussian and a Cauchy function, $1/(1 + x^2)$. We can quantitatively compare these alternative models of spatial correlation for any data set by evaluating the evidence. The regulariser is $E_W = \sum \frac{1}{2} w_h^2$. Note that this model includes one new free parameter, r ; in these demonstrations this parameter has been set to its most probable value (*i.e.* the value which maximises the evidence). To penalise this free parameter an Occam factor is included, $\sqrt{2\pi} P(\log r) \Delta \log r$, where $\Delta \log r$ = posterior uncertainty in $\log r$, and $P(\log r)$ is the prior on $\log r$, which is subjective to a small degree (I used $P(\log r) = 1/(4 \pm 2)$).

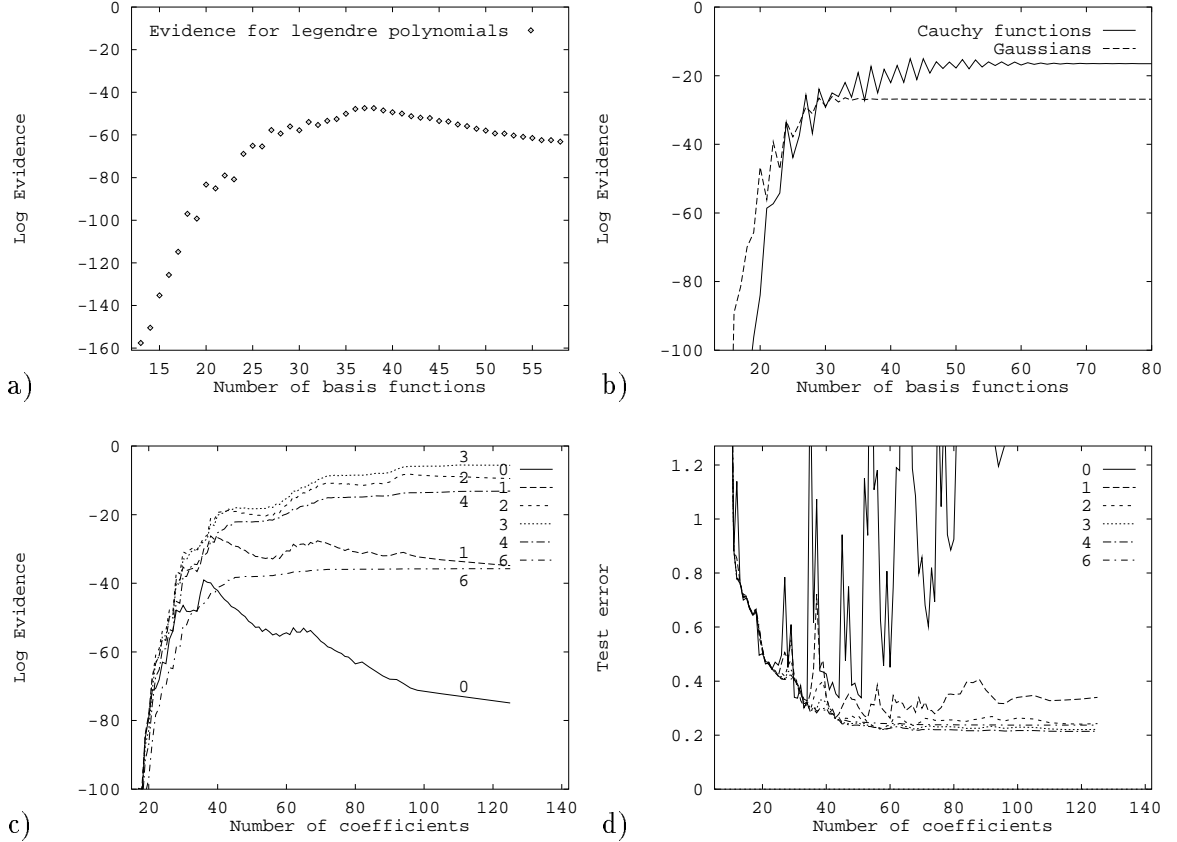


Figure 7: **The Evidence for data set X** (see also table 1)

a) Log Evidence for Legendre polynomials. Notice the evidence maximum. The gentle slope to the right is due to the ‘Occam factors’ which penalise the increasing complexity of the model.
b) Log Evidence for radial basis function models. Notice that there is no Occam penalty for the additional coefficients in these models, because increased density of radial basis functions does not make the model more powerful. The oscillations in the evidence are due to the details of the pixellation of the basis functions relative to the data points.
c) Log Evidence for splines. The evidence is shown for the alternative splines regularisers $p = 0 \dots 6$ (see text). In the representation used, each spline model is obtained in the limit of an infinite number of coefficients. For example, $p = 4$ yields the cubic splines model.
d) Test error for splines. The number of data points in the test set was 90, *c.f.* number of data points in training set = 37. The y axis shows E_D ; the value of E_D for the true interpolant has expectation 0.225 ± 0.02 .

Table 1: **Evidence for models interpolating data sets X and Y**

All logs are natural. The evidence $P(D|\mathcal{H})$ is a density over D space, so the absolute value of the log evidence is arbitrary within an additive constant. Only differences in values of log evidences are relevant, relating directly to probability ratios.

Model	Data Set X		Data Set Y	
	Best parameter values	Log evidence	Best parameter values	Log evidence
Legendre polynomials	$k = 38$	-47	$k = 11$	23.8
Gaussian radial basis functions	$k > 40$, $r = .25$	-28.8 ± 1.0	$k > 50$, $r = .77$	27.1 ± 1.0
Cauchy radial basis functions	$k > 50$, $r = .27$	-18.9 ± 1.0	$k > 50$, $r = 1.1$	25.7 ± 1.0
Splines, $p = 2$	$k > 80$	-9.5	$k > 50$	8.2
Splines, $p = 3$	$k > 80$	-5.6	$k > 50$	19.8
Splines, $p = 4$	$k > 80$	-13.2	$k > 50$	22.1
Splines, $p = 5$	$k > 80$	-24.9	$k > 50$	21.8
Splines, $p = 6$	$k > 80$	-35.8	$k > 50$	20.4
Hermite functions	$k = 18$	-66	$k = 3$	42.2
Neural networks	8 neurons, $k = 25$	-12.6	6 neurons, $k = 19$	25.7

This radial basis function model is the same as the ‘intrinsic correlation’ model of Charter (1991), Gull (1989a), Skilling and Sibisi (1991). Figure 7b shows the evidence as a function of the number of basis functions, k . Note that for these models there is *not* an increasing Occam penalty for large numbers of parameters. The reason for this is that these extra parameters do not make the model any more powerful (for fixed α and r). The increased density of basis functions does not enable the model to make any significant new predictions because the kernel g band-limits the possible interpolants.

Splines: Occam’s razor for the choice of regulariser

The splines model was implemented as follows: let the basis functions be a Fourier set $\cos hx, \sin hx$, $h = 0, 1, 2, \dots$. Use the regulariser $E_W = \sum \frac{1}{2} h^p w_{h(\cos)}^2 + \sum \frac{1}{2} h^p w_{h(\sin)}^2$. If $p = 4$ then in the limit $k \rightarrow \infty$ we have the cubic splines regulariser $E_y^{(4)} = \int y''(x)^2 dx$; if $p = 2$ we have the regulariser $E_y^{(2)} = \int y'(x)^2 dx$, etc. Notice that the ‘non-parametric’ splines model can easily be put in an explicit parameterised representation.

Figure 7c shows the evidence for data set X as a function of the number of terms, for $p = 0, 1, 2, 3, 4, 6$. Notice that in terms of Occam’s razor, both cases discussed above occur: for $p = 0, 1$, as k increases, the model becomes more powerful and there is an Occam penalty. For $p = 3, 4, 6$, increasing k gives rise to no penalty. The case $p = 2$ seems to be on the fence between the two.

As p increases, the regulariser becomes more opposed to strong curvature. Once we reach $p = 6$, the model becomes improbable because the data demand sharp discontinuities. The evidence can choose the order of our splines regulariser for us. For this data set, it

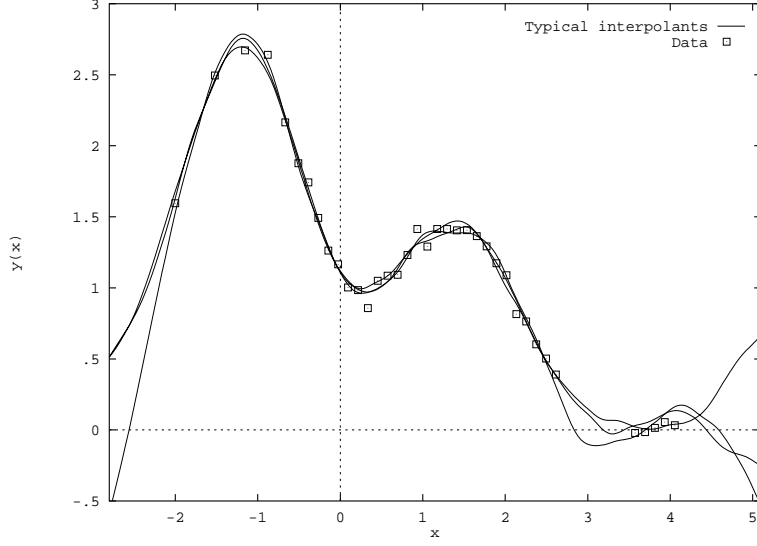


Figure 8: **Data set ‘Y,’ interpolated with splines, $p = 5$.**

The data set is shown with three typical interpolants drawn from the posterior probability distribution. Contrast this with figure 4b, in which the most probable interpolant is shown with its pointwise error bars.

turns out that $p=3$ is the most probable value of p , by a few multiples of e .

In passing, the radial basis function models described above can be transformed into the splines models’ Fourier representation. If the radial basis function kernel is $g(x)$ then the regulariser in the splines representation is $E_W = \sum \frac{1}{2}(w_{h(\cos)}^2 + w_{h(\sin)}^2)G_h^{-2}$, where G_h is the discrete Fourier transform of g .

Results for a smoother data set

Figure 8 shows data set Y, which comes from a much smoother interpolant than data set X. Table 1 summarises the evidence for the alternative models. We can confirm that the evidence behaves in a reasonable manner by noting the following differences between data sets X and Y:

In the splines family, the most probable value of p has shifted upwards to the stiffer splines with $p=4-5$, as we would intuitively expect.

Legendre polynomials: an observant reader may have noticed that when data set X was modelled with Legendre polynomials, the most probable number of coefficients $k=38$ was suspiciously similar to the number of data points $N=37$. For data set Y, however, the most probable number of coefficients is 11, which confirms that the evidence does not always prefer the polynomial with $k=N$! Data set X behaved in this way because it is very poorly modelled by polynomials.

The Hermite function model, which was a poor model for data set X, is now the most probable, by a long way (over a million times more probable). The reason for this is that actually the data *were* generated from a Hermite function!

Why Bayes can't systematically reject the truth

Let us ask a sampling theory question: if one of the models we offer to Bayes is actually true, *i.e.* it is the model from which the data were generated, then is it possible for Bayes to systematically (over the ensemble of possible data sets) prefer a false model? Clearly under a worst case analysis, a Bayesian's posterior may favour a false model. Furthermore, Skilling (1991) demonstrated that with some data sets a free form (maximum entropy) model can have greater evidence than the truth; but is it possible for this to happen in the *typical* case, as Skilling seems to claim? I will show that the answer is no, the effect that Skilling demonstrated cannot be systematic. To be precise, the expectation over possible data sets of the log evidence for the true model is greater than the expectation of the log evidence for any other fixed model (Osteyee and Good, 1974).¹⁴

Proof. Suppose that the truth is actually \mathcal{H}_1 . A single data set arrives and we compare the evidences for \mathcal{H}_1 and \mathcal{H}_2 , a different fixed model. Both models may have free parameters, but this will be irrelevant to the argument. Intuitively we expect that the evidence for \mathcal{H}_1 , $P(D|\mathcal{H}_1)$, should usually be greatest. Let us examine the difference in log evidence between \mathcal{H}_1 and \mathcal{H}_2 . The expectation of this difference, given that \mathcal{H}_1 is true, is

$$\left\langle \log \frac{P(D|\mathcal{H}_1)}{P(D|\mathcal{H}_2)} \right\rangle = \int d^N D P(D|\mathcal{H}_1) \log \frac{P(D|\mathcal{H}_1)}{P(D|\mathcal{H}_2)}.$$

(Note that this integral implicitly integrates over all \mathcal{H}_1 's parameters according to their prior distribution under \mathcal{H}_1 .) Now it is well known that for normalised p and q , $\int p \log \frac{p}{q}$ is minimised by setting $q = p$ (Gibb's theorem). Therefore a distinct model \mathcal{H}_2 is never expected to *systematically* defeat the true model, for just the same reason that it is not wise to bet differently from the true odds. •

This has two important implications. First, it gives us frequentist confidence in the ability of Bayesian methods on the average to identify the true model. Secondly, it provides a stringent test of numerical implementations of Bayesian model comparison: imagine that we have written a program that evaluates the evidence for models \mathcal{H}_1 and \mathcal{H}_2 ; then we can generate mock data from sources simulating \mathcal{H}_1 and \mathcal{H}_2 and evaluate the evidences; if there is any systematic bias, averaged over several mock data sets, for the estimated evidence to favour the false model, then we can be sure that our numerical implementation is not evaluating the evidence correctly.

This issue is illustrated using data set Y. The 'truth' is that this data set was actually generated from a quadratic Hermite function, $1.1(1-x+2x^2)e^{-x^2/2}$. By the above argument the evidence ought probably to favour the model 'the interpolant is a 3-coefficient Hermite function' over our other models. Table 1 shows the evidence for the true Hermite function model, and for other models. As already stated, the truth is indeed considerably more probable than the alternatives.

Having demonstrated that Bayes cannot systematically fail when one of the models is true, we now examine the way in which this framework can fail, if none of the models offered to Bayes is any good.

¹⁴Skilling's result presumably occurred because the particular parameter values of the true model that generated the data were not typical of the prior used when evaluating the evidence for that model. In such a case, the log evidence difference can show a transient bias against the true model, for small quantities of data; such biases are usually corrected by greater quantities of data.

Comparison with ‘generalisation error’

It is a popular and intuitive criterion for choosing between alternative interpolants (found using different models) to compare their errors on a test set that was not used to derive the interpolants. ‘Cross-validation’ is a more refined and more computationally expensive version of this same idea. How does this method relate to the evaluation of the evidence described in this paper?

Figure 7c displayed the evidence for the family of spline interpolants. Figure 7d shows the corresponding test error, measured on a test set with size over twice as big (90) as the ‘training’ data set (37) used to determine the interpolant. A similar comparison was made in figure 5b. Note that the overall trends shown by the evidence are matched by trends in the test error (if you flip one graph upside down). Also, for this particular problem, the ranks of the alternative spline models under the evidence are similar to their ranks under the test error. And in figure 5b, the evidence maximum over α was surrounded by the test error minima. Thus this suggests that the evidence might be a reliable predictor of generalisation ability. However, this is not necessarily the case. There are five reasons why the evidence and the test error might not be correlated.

First, the test error is a noisy quantity. It is necessary to devote large quantities of data to the test set to obtain a reasonable signal to noise ratio. In figure 5b more than twice as much data is in each test set but the difference in $\log \alpha$ between the two test error minima exceeds the size of the Bayesian confidence interval for $\log \alpha$.

Second, the model with greatest evidence is not expected to be the best model all the time — Bayesian inferences are uncertain. The whole point of Bayes is that it quantifies precisely those uncertainties: the relative values of the evidence for alternative models express the plausibility of the models, given the data and the underlying assumptions.

Third, there is more to the evidence than there is to the generalisation error. For example, imagine that for two models, the most probable interpolants happen to be identical. In this case, the two solutions will have the same generalisation error, but the evidence will not in general be the same: typically, the model that was *a priori* more complex will suffer a larger Occam factor and will have a smaller evidence.

Fourth, the test error is a measure of performance only of the single most probable interpolant: the evidence is a measure of plausibility of the entire posterior ensemble around the best fit interpolant. Probably a stronger correlation between the evidence and the test statistic would be obtained if the test statistic used were the average of the test error over the posterior ensemble of solutions. This ensemble test error is not so easy to compute.

The fifth and most interesting reason why the evidence might not be correlated with the generalisation error is that there might be a flaw in the underlying assumptions such that the models being compared might all be poor models. If a poor regulariser is used, for example, one that is ill-matched to the statistics of the world, then the Bayesian choice of α will often not be the best in terms of generalisation error (Davies and Anderssen, 1986, Gull, 1989a, Haussler *et. al.*, 1991). Such a failure occurs in the companion paper on neural networks. What is our attitude to such a failure of Bayesian prediction? The failure of the evidence does not mean that we should discard Bayes’ rule and use the generalisation error as our criterion for choosing α . A failure is an opportunity to learn; a healthy scientist actively searches for such failures, because they yield insights into the defects of the current model. The detection of such a failure (by evaluating the generalisation error for example) motivates the search for new models which do not fail in this way; for example alternative regularisers can be tried until a model is found that makes the data more probable.

If one only uses the generalisation error as a criterion for model comparison, one is denied this mechanism for learning. The development of maximum entropy image deconvolution was held up for years because no-one used the Bayesian choice of α ; once the Bayesian choice of α was used (Gull, 1989a), the results obtained were most dissatisfactory, making clear what a poor regulariser was being used; this motivated an immediate search for alternative priors; the new, more probable priors discovered by this search are now at the heart of the state of the art in image deconvolution (Weir, 1991).

The similarity between regularisation and ‘early stopping’

While an over-parameterised model is fitted to a data set using gradient descent on the data error, it is sometimes noted that the model’s generalisation error passes through a minimum, rather than decreasing monotonically. This is known as ‘over-learning’ in the neural networks community, and some researchers advocate the use of ‘early stopping,’ that is, stopping gradient descent before the data error minimum is reached, so as to try to obtain solutions with smaller generalisation error.

This author believes that ‘over-learning’ should be viewed as a symptom of a model ill-matched to the data set, and that the appropriate response is not to patch up a bad model, but rather to search for models which are well matched to our data. In particular, the use of models incorporating simple regularisers is expected to give results qualitatively similar to early stopping. This can be seen by examining figure 6. The regulariser moves the minimum of the objective function from \mathbf{w}_{ML} to \mathbf{w}_{MP} ; as the strength of the regulariser α is increased, \mathbf{w}_{MP} follows a knee-shaped trajectory from \mathbf{w}_{ML} to the origin; a typical solution \mathbf{w}_{MP} is shown in figure 6. If on the other hand gradient descent on the likelihood (data error) is used, and if the typical initial condition is close to the origin, then gradient descent will follow a similar knee-shaped trajectory. Thus qualitatively similar solutions are expected from increasingly early stopping and increasingly strong regularisation with complete minimisation. Regularisation is to be preferred as a more robust, repeatable and comprehensible procedure.

Admitting neural networks into the canon of Bayesian interpolation models

A second paper will discuss how to apply this Bayesian framework to feedforward neural networks. Preliminary results using these methods are included in table 1. Assuming that the approximations used were valid, it is interesting that the evidence for neural nets is actually good for both the spiky and the smooth data sets. Furthermore, neural nets, in spite of their arbitrariness, yield a relatively compact model, with fewer parameters needed than to specify the splines and radial basis function solutions.

7 Conclusions

The recently developed methods of Bayesian model comparison and regularisation have been presented. Models can be ranked by evaluating the evidence, a solely data-dependent measure which intuitively and consistently combines a model’s ability to fit the data with its complexity. The precise posterior probabilities of the models also depend on the subjective priors that we assign to them, but these terms are typically overwhelmed by the evidence.

Regularising constants are set by maximising the evidence. For many regularisation problems, the theory of the number of well-measured parameters makes it possible to perform this optimisation on-line.

In the interpolation examples discussed, the evidence was used to set the number of basis functions k in a polynomial model; to set the characteristic size r in a radial basis function model; to choose the order p of the regulariser for a spline model; and to rank all these different models in the light of the data.

Further work is needed to formalise the relationship of this framework to the pragmatic model comparison technique of cross-validation. By using the two techniques in parallel it is possible to detect flaws in the underlying assumptions implicit in the data models being used. Such failures direct us in our search for superior models, providing a powerful tool for human learning.

There are thousands of data modelling tasks waiting for the evidence to be evaluated. It will be exciting to see how much we can learn when this is done.

References

- H. Akaike (1970). ‘Statistical predictor identification’, *Ann. Inst. Statist. Math.* **22**, 203–217.
- J. Berger (1985). *Statistical decision theory and Bayesian analysis*, Springer.
- G.L. Bretthorst (1990). ‘Bayesian Analysis. I. Parameter Estimation Using Quadrature NMR Models. II. Signal Detection and Model Selection. III. Applications to NMR.’, *J. Magnetic Resonance* **88** 3, 533–595.
- M.K. Charter (1991). ‘Quantifying drug absorption’, in Grandy and Schick, eds. (1991), 245–252.
- R.T. Cox (1964). ‘Probability, frequency, and reasonable expectation’, *Am. J. Physics* **14**, 1–13.
- A.R. Davies and R.S. Anderssen (1986). ‘Optimization in the regularization of ill-posed problems’, *J. Austral. Mat. Soc. Ser. B* **28**, 114–133.
- R.L. Eubank (1988). ‘Spline smoothing and non-parametric regression’, Marcel Dekker.
- W.T. Grandy, Jr. and L.H. Schick, eds. (1991). *Maximum entropy and Bayesian methods*, Laramie 1990, Kluwer.
- S.F. Gull (1988). ‘Bayesian inductive inference and maximum entropy’, in *Maximum entropy and Bayesian methods in science and engineering, vol. 1: Foundations*, G.J. Erickson and C.R. Smith, eds., Kluwer.
- S.F. Gull (1989). ‘Developments in maximum entropy data analysis’, in J. Skilling, ed. (1989a), 53–71.
- S.F. Gull (1989). ‘Bayesian data analysis: straight-line fitting’, in J. Skilling, ed. (1989a), 511–518.
- S.F. Gull and J. Skilling (1991). *Quantified maximum entropy. MemSys5 user’s manual*. M.E.D.C., 33 North End, Royston, SG8 6NR, England.
- R. Hanson, J. Stutz and P. Cheeseman (1991). ‘Bayesian classification theory’, NASA Ames TR FIA-90-12-7-01.
- D. Haussler, M. Kearns and R. Schapire (1991). ‘Bounds on the sample complexity of Bayesian learning using information theory and the VC dimension’, preprint.
- E.T. Jaynes (1986). ‘Bayesian methods: general background’, in *Maximum entropy and Bayesian methods in applied statistics*, ed. J.H. Justice, C.U.P.

- H. Jeffreys (1939). *Theory of probability*, Oxford Univ. Press.
- R.L. Kashyap (1977). ‘A Bayesian comparison of different classes of dynamic models using empirical data’, *IEEE Transactions on Automatic Control* **AC-22** 5, 715–727.
- T.J. Loredo (1989). ‘From Laplace to supernova SN 1987A: Bayesian inference in astrophysics’, in *Maximum entropy and Bayesian methods*, ed. P. Fougere, Kluwer.
- D.J.C. MacKay (1991b) ‘A practical Bayesian framework for backprop networks’, *Neural Computation*, this volume.
- D.J.C. MacKay (1991c) ‘Information-based objective functions for active data selection’, *Neural Computation*, this volume.
- R.M. Neal (1991). ‘Bayesian mixture modeling by Monte Carlo simulation’, Preprint.
- D.B. Osteyee and I.J. Good (1974). *Information, weight of evidence, the singularity between probability measures and signal detection*, Springer.
- J.D. Patrick and C.S. Wallace (1982). ‘Stone circle geometries: an information theory approach’, in *Archaeoastronomy in the Old World*, D.C. Heggie, editor, Cambridge Univ. Press.
- T. Poggio, V. Torre and C. Koch (1985). ‘Computational vision and regularization theory’, *Nature* **317** 6035, 314–319.
- J. Rissanen (1978). ‘Modeling by shortest data description’, *Automatica* **14**, 465–471.
- G. Schwarz (1978). ‘Estimating the dimension of a model’, *Ann. Stat.* **6** 2, 461–464.
- S. Sibisi (1991). ‘Bayesian interpolation’, in Grandy and Schick, eds. (1991), 349–355.
- J. Skilling, editor (1989a). *Maximum entropy and Bayesian methods, Cambridge 1988*, Kluwer.
- J. Skilling (1989b). ‘The eigenvalues of mega-dimensional matrices’, in J. Skilling, ed. (1989a), 455–466.
- J. Skilling (1991). ‘On parameter estimation and quantified MaxEnt’, in Grandy and Schick, eds. (1991), 267–273.
- J. Skilling, D.R.T. Robinson, and S.F. Gull (1991). ‘Probabilistic displays’, in Grandy and Schick, eds. (1991), 365–368.
- S.M. Stigler (1986). ‘Laplace’s 1774 memoir on inverse probability’, *Stat. Sci.* **1** 3, 359–378.
- R. Szeliski (1989). ‘Bayesian modeling of uncertainty in low level vision’, Kluwer.
- D. Titterton (1985). ‘Common structure of smoothing techniques in statistics’, *Int. Statist. Rev.* **53**, 141–170.
- A.M. Walker (1967). ‘On the asymptotic behaviour of posterior distributions’, *J. R. Stat. Soc. B* **31**, 80–88.
- C. S. Wallace and D. M. Boulton (1968). ‘An information measure for classification’, *Comput. J.* **11** 2, 185–194.
- C. S. Wallace and P. R. Freeman (1987). ‘Estimation and Inference by Compact Coding’, *J. R. Statist. Soc. B* **49** 3, 240–265.
- N. Weir (1991). ‘Applications of maximum entropy techniques to HST data’, *Proceedings of the ESO/ST-ECF data analysis workshop*, April 1991.
- A. Zellner (1984). *Basic issues in econometrics*, Chicago.

I thank Mike Lewicki, Nick Weir and David R.T. Robinson for helpful conversations, and Andreas Herz for comments on the manuscript. I am grateful to Dr. R. Goodman and Dr. P. Smyth for funding my trip to Maxent 90.

This work was supported by a Caltech Fellowship and a Studentship from SERC, UK.