apts.ac.uk
Academy for PhD Training in Statistics

Adrian Bowman, Ludger Evers
School of Mathematics & Statistics
The University of Glasgow

Nonparametric Smoothing

Lecture Notes

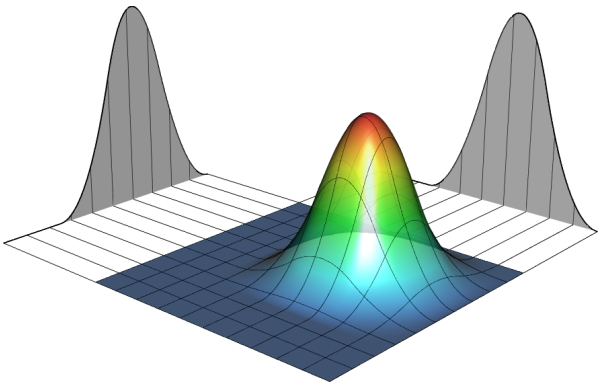

1

Introduction

1.1 What this course is about

This APTS course will cover a variety of methods which enable data to be modelled in a flexible manner. It will use and extend a variety of topics covered in earlier APTS courses, including

– linear models, including the Bayesian version;
– generalised linear models;
– R programming;
– Taylor series expansions and standard asymptotic methods.

The main emphasis will be on regression settings, because of the widespread use and application of this kind of data structure. However, the material of the first chapter will include the simple case of density estimation, also covered in the preliminary material, to introduce some of the main ideas of nonparametric smoothing and to highlight some of the main issues involved.

As with any statistical topic, a rounded treatment involves a variety of approaches, including

– clear understanding of the underlying concepts;
– technical understanding of methods, with an exploration of their properties;
– appreciation of the practical computational issues;
– some knowledge of the tools available to fit relevant models in R;
– understanding of how these models can bring insight into datasets and applications.

The aim is to reflect all of these aspects in the course, but to varying degrees in different sections. There will not be time to cover all the material in the notes and some of the material is intended to provide pointers to topics which it might be of interest to explore in the context of your own research.

1.2 Broad concepts and issues of smoothing

The simple case of density estimation highlights features and issues which are common to a wide range of problems involving the estimation of functions, relationships or patterns which are *nonparametric* but *smooth*. The term nonparametric is used in this context to mean that the relationships or patterns of interest cannot be expressed in specific formulae which involved a fixed number of unknown parameters. This means that the parameter space is the space of functions, whose dimensionality is infinite. This takes us outside of the standard framework for parametric models and the main theme of the course will be to discuss how we can do this while producing tools which are highly effective for modelling and analysing data from a wide variety of contexts and exhibiting a wide variety of structures.

On a side note, the term nonparametric is sometimes used in the narrower setting of simple statistical methods based on the ranks of the data, rather than the original measurements. This is not the sense in which it will be used here.

Further details on density estimation are given in Silverman (1986), Scott (1992), Wand and Jones (1995) & Simonoff (1996).

The issues raised by our brief discussion of density estimation include

- how to construct estimators which match the type of data we are dealing with;
- how to find a suitable balance between being faithful to the observed data and incorporating the underlying regularity or smoothness which we believe to be present;
- how to construct and make use of suitable inferential tools which will allow the models to weigh the evidence for effects of interest, in a setting which takes us outside of standard parametric methods.

These broad issues will be explored in a variety of contexts in the remainder of the course.

1.3 Nonparametric regression

Regression is one of the most widely used model paradigms and this will be the main focus in the remainder of the course. Here is an example which will be used to illustrate the initial discussion.

Example 1.1 (Great Barrier Reef data). A survey of the fauna on the sea bed lying between the coast of northern Queensland and the Great Barrier Reef was carried out. The sampling region covered a zone which was closed to commercial fishing, as well as neighbouring zones where fishing was permitted. The variables are:

Zone	an indicator for the closed (1) and open (0) zones
Year	an indicator of 1992 (0) or 1993 (1)
Latitude	latitude of the sampling position
Longitude	longitude of the sampling position
Depth	bottom depth
Score1	catch score 1
Score2	catch score 2

The details of the survey and an analysis of the data are provided by Poiner et al. (1997), *The effects of prawn trawling in the far northern section of the Great Barrier Reef*, CSIRO Division of Marine Research, Queensland Dept. of Primary Industries. ◁

The relationship between catch score (Score1) and longitude is of particular interest because, at this geographical location, the coast runs roughly north-south and so longitude is a proxy for distance offshore. We might therefore reasonably expect the abundance of marine life to change with longitude. The first of the three panels below shows that there is indeed a strong underlying negative relationship, with considerable variability also present. The middle panel summarises this in a simple linear regression which captures much of this relationship. However, if we allow our regression model to be more flexible then a more complex relationship is suggested in the right hand panel, with a broadly similar mean level for some distance offshore followed by a marked decline, possibly followed by some levelling off thereafter. This gives valuable informal and graphical insight into the data but how can flexible regression models can be constructed and how can we use them to evaluate whether there is really evidence of non-linear behaviour in the data?

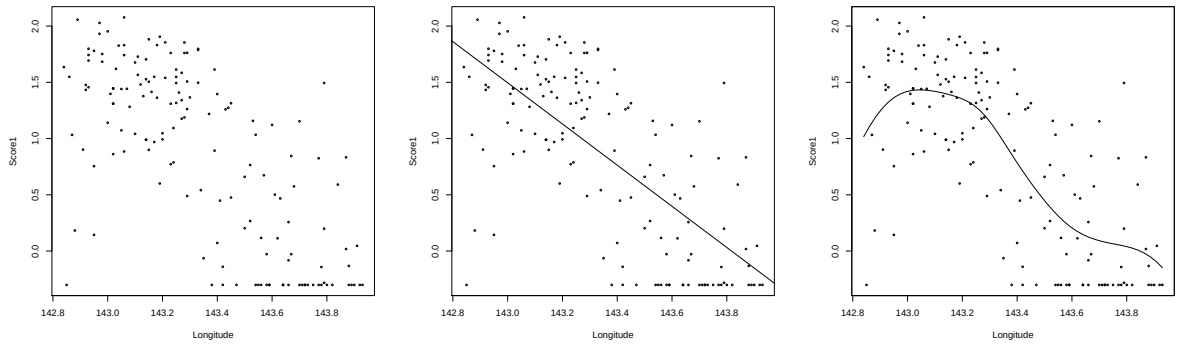

1.3.1 A local fitting approach

A simple nonparametric model has the form

$$y_i = m(x_i) + \varepsilon_i,$$

where the data (x_i, y_i) are described by a smooth curve m plus independent errors ε_i. One approach to fitting this is to take a model we know and fit it locally. For example, we can

construct a *local linear regression*. This involves solving the least squares problem

$$\min_{\alpha,\beta} \sum_{i=1}^{n} \{y_i - \alpha - \beta(x_i - x)\}^2 \, w(x_i - x \, ; h)$$

and taking as the estimate at x the value of $\hat{\alpha}$, as this defines the position of the local regression line at the point x. This has an appealing simplicity and it can be generalised quite easily to other situations. This was the approach used to produce the nonparametric regression of the Reef data in the plot above.

There is a variety of other ways in which smooth curve estimates can be produced and a further approach is outlined in the next section. It can sometimes reasonably be argued that the precise mechanism usually isn't too important and can be chosen for convenience.

1.3.2 Basis function approaches

Basis approaches to function approximation have a very long history. One of the first was Fourier series, which is based on the expansion

$$m(x) \approx \frac{a_0}{2} + \sum_{j=1}^{p} a_j cos\left(\frac{2\pi jx}{P}\right) + b_j sin\left(\frac{2\pi jx}{P}\right)$$

where P is the range of the covariate. The plots below show the effects of fitting a Fourier series with $p = 1, 2, 4, 6$ frequencies.

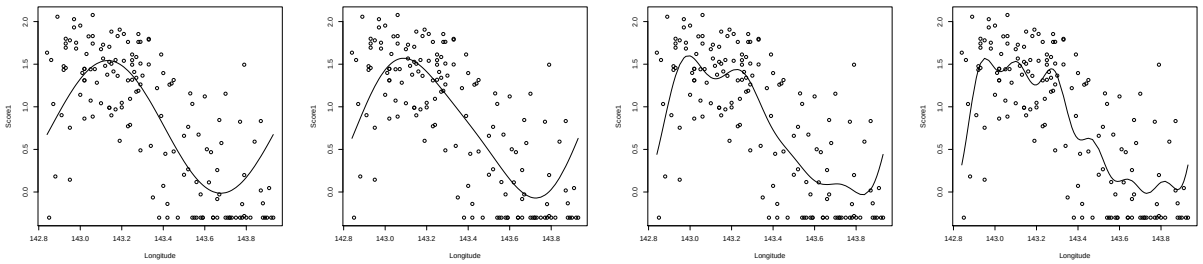

Each basis function in a Fourier expansion has effects across the entire range of the data, with the result that the effect of fitting data well in one part of the sample space can create artefacts elsewhere. An alternative approach is to use a set of basis functions which are more local in their effects. The figure below shows a set of *b-spline* functions, $b_1(x), \ldots, b_p(x)$ which can be used for this.

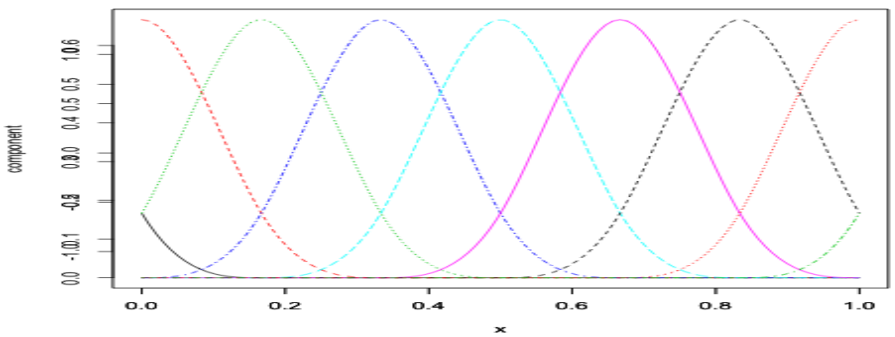

A curve estimate can then be produced simply by fitting the regression

$$y_i = \beta_0 + \beta_1 b_1(x_i) + \beta_p b_2(x_i) + \ldots + \beta_p b_p(x_i) + \varepsilon_i.$$

A popular approach, known as *p-splines*, uses a large number of basis functions but includes a penalty to control the smoothness of the resulting estimator. Specifically, the coefficients β_i are chosen to minimise

$$\sum_{i=1}^{n} \{y_i - \beta_0 - \beta_1 b_1(x_i) - \beta_p b_2(x_i) - \ldots - \beta_p b_p(x_i)\}^2 + \lambda P(\beta),$$

where the penalty function P might, for example, have the form

$$P(\beta) = \sum_{j=2}^{p} (\beta_j - \beta_{j-1})^2$$

or some other measure of roughness along the sequence of coefficients. Here the parameter λ controls the degree of smoothness of the resulting estimator. This idea of using a penalty to find a good balance between fit to the data and smoothness of the estimator is a general theme which will reappear at various points during the course.

A useful connection can be made here with wavelets, where local basis functions are also used to construct sophisticated function estimators.

1.4 Further illustrations of smoothing

This section gives a few brief illustrations of how the ideas outlined in simple settings can be extended to much more complex situations and data structures. Some of these will be revisited in greater detail later in the course.

1.4.1 Regression with more than one covariate

It is rare to have problems which involve only a single covariate. For the Reef data a natural extension is to look at the relationship between the catch score and both latitude (x_1) and longitude (x_2), in a model

$$y_i = m(x_{1i}, x_{2i}) + \varepsilon_i.$$

the top left hand panel of the plot below shows the effect of this. The effect of longitude dominates, as we see from the earlier nonparametric regression. However, a small effect of latitude is also suggested. The methods by which surfaces of this type can be constructed will be discussed in the next two chapters.

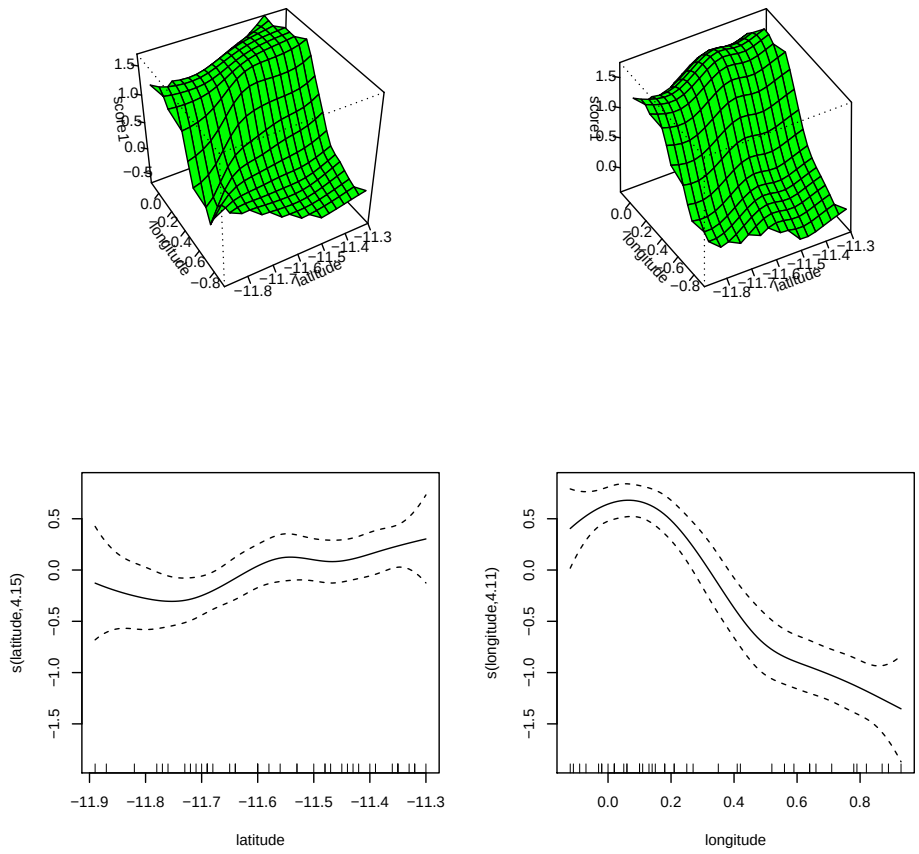

It would be unrealistic to generalise this much further, by modelling additional covariates through functions of ever-increasing dimension. A very powerful approach is to construct *additive models* of the form

$$y_i = m_1(x_{1i}) + m_2(x_{2i}) + \varepsilon_i,$$

where the component functions m_1 and m_2 describe the separate and additive effects of the two covariates. Estimated additive model components, and their combined effects as an additive surface, are displayed in the other panels of the figure above. Methods for fitting models of this type will also be discussed later.

1.4.2 Areal data

Data quantifying population-level summaries of disease prevalence for n non-overlapping areal units are available from both the English and Scottish Neighbourhood Statistics databases. They are used in many different applications, including quantifying the effect of an exposure on health, and identifying clusters of areal units that exhibit elevated risks of disease. The health data are denoted by $\mathbf{Y} = (Y_1, \ldots, Y_n)$ and $\mathbf{E} = (E_1, \ldots, E_n)$, which are the observed and expected numbers of disease cases in each areal unit. The covariates are denoted by an $n \times p$ matrix $X = (\mathbf{x}_1, \ldots, \mathbf{x}_n)$, and could include environmental exposures or measures of socio-economic deprivation.

One example concerns the prevalence of respiratory disease in Glasgow in 2010, where interest focusses on the hospitalisation risk $\text{SIR}_k = Y_k/E_k$. The data are displayed in the plot below.

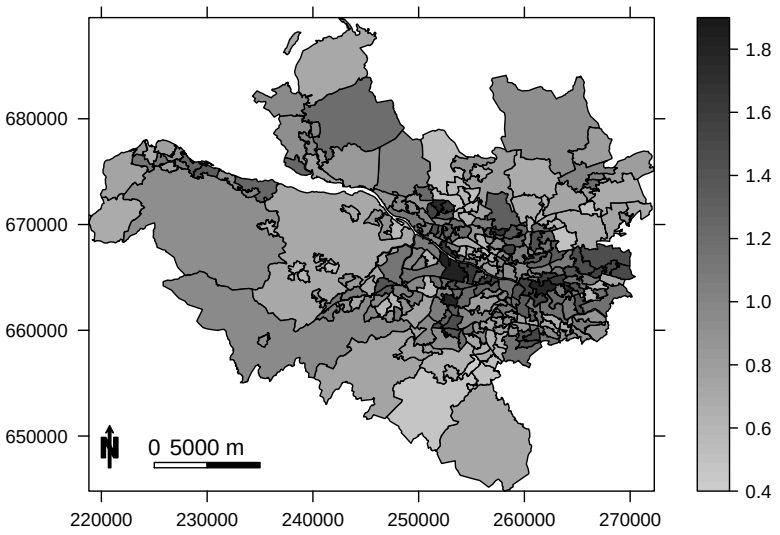

A suitable statistical model is

$$\begin{aligned} Y_k &\sim \text{Poisson}(E_k R_k), \\ \log(R_k) &= \mathbf{x}_k^{\mathrm{T}}\boldsymbol{\beta} + \phi_k, \\ \phi_k|\boldsymbol{\phi}_{-k}, \tau^2, W &\sim \text{N}\left(\frac{\sum_{i=1}^n w_{ki}\phi_i}{\sum_{i=1}^n w_{ki}}, \frac{\tau^2}{\sum_{i=1}^n w_{ki}}\right), \end{aligned}$$

where

– R_k quantifies disease risk in area k.

– $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_n)$ are random effects that model residual spatial autocorrelation and improve estimation of disease risk by smoothing the risks across neighbouring areas.

– Commonly, conditional autoregressive (CAR) models are used for this smoothing, where $W = (w_{ki})$ is a binary $n \times n$ neighbourhood matrix.

Material kindly provided by Duncan Lee, University of Glasgow.

1.4.3 Network data

Models for spatial data often involve some form of smoothing, either implicitly or explicitly. An interesting forms of spatial data arises from river networks, where models for the variation in measurements should respect both the spatial pattern of the interconnecting water channels as well as the mixing effect of confluence points where channels meet. The data below refer to nitrate pollution in the River Tweed. The point measurements on the left can be used to construct a spatial model for the whole network, whose predictions are shown in the panel on

the right.

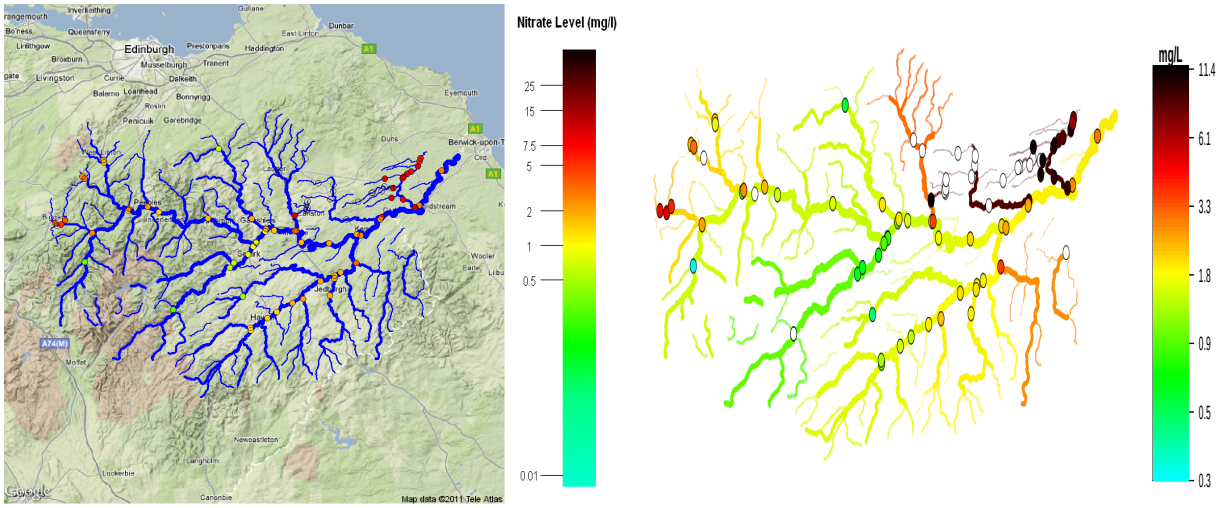

This example arises from joint work with Alastair Rushworth, David O'Donnell, Marian Scott, Mark Hallard (SEPA).

2

Density estimation

2.1 Motivating example

A probability density function is a key concept through which variability can be expressed precisely. In statistical modelling its role is often to capture variation sufficiently well, within a model where the main interest lies in structural terms such as regression coefficients. However, there are some situations where the shape of the density function itself is the focus of attention. The example below illustrates this.

Example 2.1 (Aircraft data). These data record six characteristics of aircraft designs which appeared during the twentieth century. The variables are:

Yr: year of first manufacture
Period: a code to indicate one of three broad time periods
Power: total engine power (kW)
Span: wing span (m)
Length: length (m)
Weight: maximum take-off weight (kg)
Speed: maximum speed (km/h)
Range: range (km)

A brief look at the data suggests that the six measurements on each aircraft should be expressed on the log scale to reduce skewness. Span is displayed on a log scale below, for Period 3 which corresponds to the years after the second World War. ◁

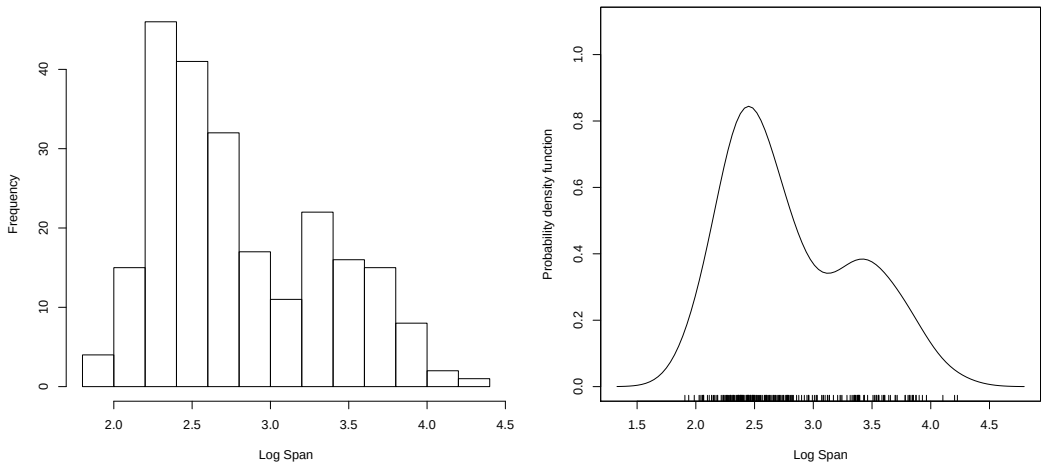

The pattern of variability shown in both the histogram and the density estimate exhibits some skewness. There is perhaps even a suggestion of a subsidiary mode at high values of log span, although this is difficult to evaluate.

2.2 Kernel density estimation

The histogram is a very familiar object. It can be written as

$$\tilde{f}(y) = \sum_{i=1}^{n} I(y - \tilde{y}_i; h),$$

where $\{y_1, \ldots, y_n\}$ denote the observed data, $\tilde{y}_i$ denotes the centre of the interval in which y_i falls and $I(z; h)$ is the indicator function of the interval $[-h, h]$. The form of the construction of $\tilde{f}$ highlights some feature which are open to criticism if we view the histogram as an estimator of the underlying density function. Firstly the histogram is not smooth, when we expect that the underlying density usually will be. Secondly, some information is lost when we replace each observation y_i by the bin mid-point $\tilde{y}_i$. Both of the issues can be addressed by using a density estimator in the form

$$\hat{f}(y) = \frac{1}{n} \sum_{i=1}^{n} w(y - y_i; h),$$

where w is a probability density, called here a *kernel function*, whose variance is controlled by the *smoothing parameter* h.

Large changes in the value of the smoothing parameter have large effects on the smoothness of the resulting estimates, as the plots below illustrate.

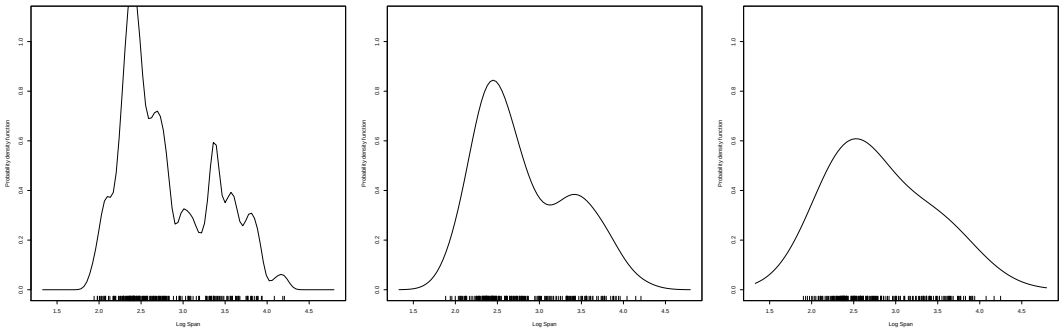

One advantage of density estimates is that it is a simple matter to superimpose these to allow different groups to be compared. Here the groups for the three different time periods are compared. It is interesting that the 'shoulder' appears in all three time periods.

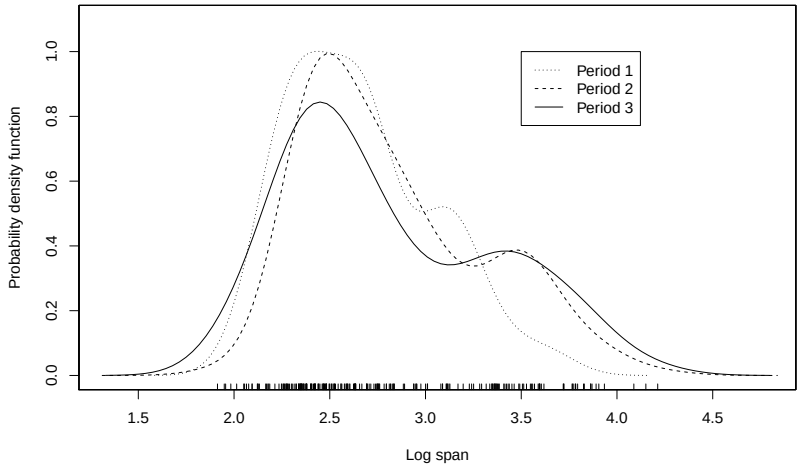

2.2.1 Simple asymptotic properties

Without any real restriction, we can assume that the kernel function can be written in the simple form $w(y - y_i; h) = \frac{1}{h} w\left(\frac{y-y_i}{h}\right)$. The preliminary material showed that the mean and variance of a density estimator can then be expressed as

$$\begin{aligned} \mathbb{E}\left\{\hat{f}(y)\right\} &= f(y) + \frac{h^2}{2}\sigma_w^2 f''(y) + o(h^2), \\ \mathsf{var}\left\{\hat{f}(y)\right\} &= \frac{1}{nh} f(y)\,\alpha(w) + o\left(\frac{1}{nh}\right), \end{aligned}$$

where we assume that the kernel function is symmetric so that $\int uw(u)du = 0$, and where σ_w^2 denotes the variance of the kernel, namely $\int u^2 w(u)du$, and $\alpha(w) = \int w^2(u)du$.

These expressions capture the essential features of smoothing. In particular, bias is incurred and we can see that this is controlled by f'', which means that where the density has peaks and valleys the density estimate will underestimate and overestimate respectively. This makes intuitive sense.

If we need it, a useful global measure of performance is the *mean integrated squared error* (MISE) which balances squared bias and variance.

$$\text{MISE}(\hat{f}) = \frac{1}{4}\, h^4 \sigma_w^4 \int f''(y)^2 dy + \frac{1}{nh}\, \alpha(w) + o\left(h^4 + \frac{1}{nh}\right).$$

2.2.2 Extension to other sample spaces

The simple idea of density estimation is to place a kernel function, which in fact is itself a density function, on top of each observation and average these functions. This extends very naturally to a wide variety of other types of data and sample spaces.

For example, a two-dimensional density estimate can be constructed from bivariate data $\{(y_{1i}, y_{2i}) : i = 1, \ldots, n\}$ by employing a two-dimensional kernel function in the form

$$\hat{f}(y_1, y_2) = \frac{1}{n} \sum_{i=1}^{n} w(y_1 - y_{1i}; h_1)\, w(y_1 - y_{2i}; h_2).$$

Notice that there are now two smoothing parameters, (h_1, h_2). A more general two-dimensional kernel function could be used, but the simple product form is very convenient and usually very effective.

Here is an example which uses the scores from the first two principal components of the aircraft data, again focussing on the third time period. The left hand scatterplot shows the individual scores while the right hand plot shows a density estimate, from which suggests three separate modes. This feature is not so easily seen from the raw scatterplot.

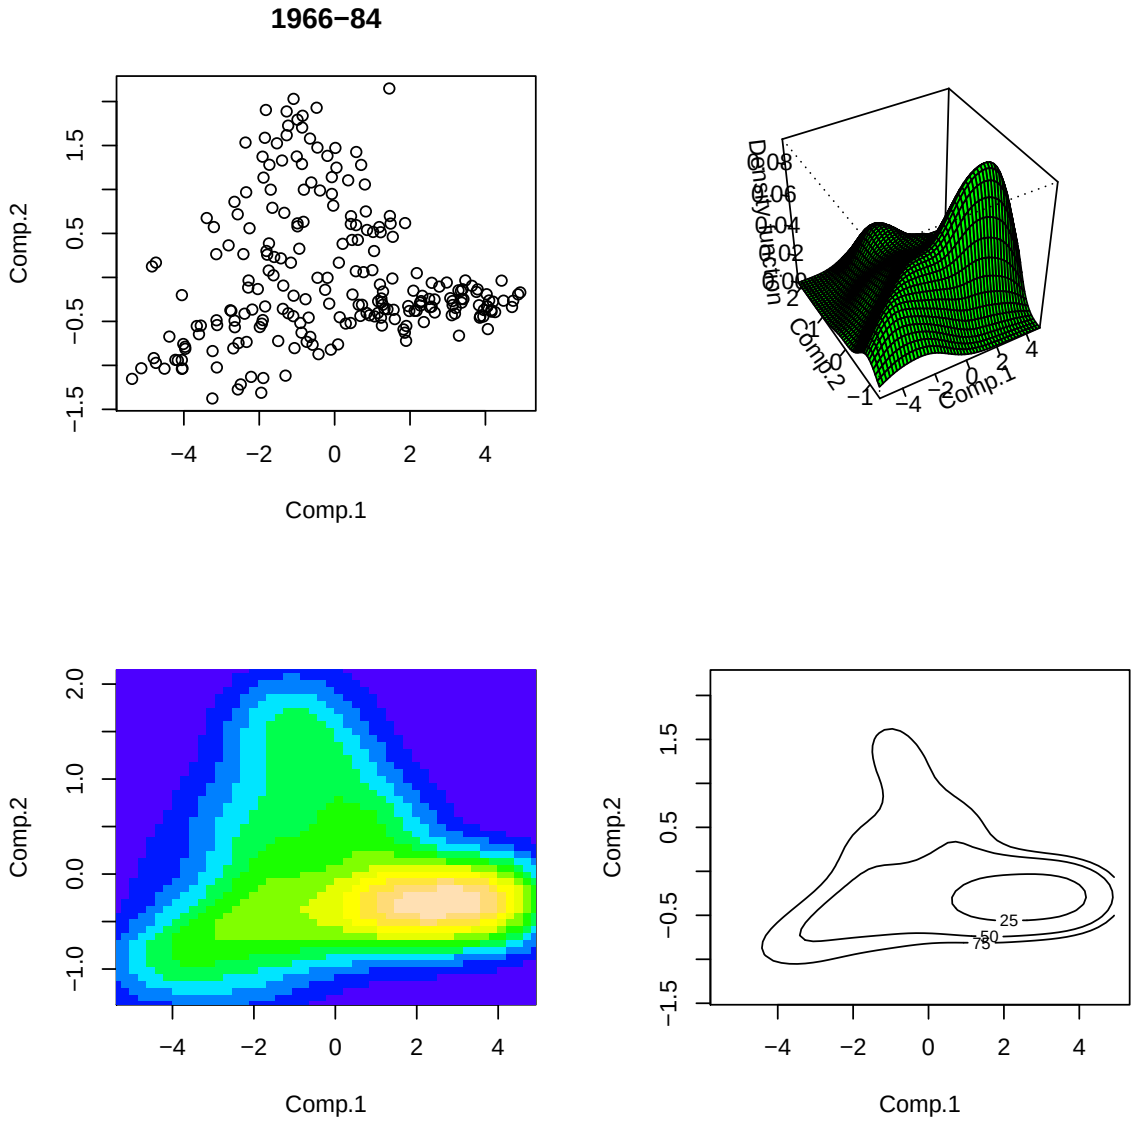

The lower two plots show alternative ways of presenting a two-dimensional estimate, using a coloured image on the left and contour lines on the right. Notice that the contours on the right have been chosen carefully to contain the quarters of the data with successively higher density, in a manner which has some similarities with a box plot.

This principle extends to all kinds of other data structures and sample spaces by suitable choice of an appropriate kernel function.

2.2.3 Deciding how much to smooth

It is not too hard to show that the value of h which minimizes MISE in an asymptotic sense is

$$h_{\text{opt}} = \left\{ \frac{\gamma(w)}{\beta(f)n} \right\}^{1/5},$$

where $\gamma(w) = \alpha(w)/\sigma_w^4$, and $\beta(f) = \int f''(y)^2 dy$. Of course, this is of rather limited use because it is a function of the unknown density. However, there are two practical approaches

which can be taken to deciding on a suitable smoothing parameter to use. One is to construct an estimate of MISE and minimise this. Another is to estimate the optimal smoothing parameter. These two approaches are outlined below.

Cross-validation

The integrated squared error (ISE) of a density estimate is

$$\int \{\hat{f}(y) - f(y)\}^2 dy = \int \hat{f}(y)^2 dy - 2\int f(y)\hat{f}(y)dy + \int f(y)^2 dy.$$

Only the first two of these terms involve h and these terms can be estimated by

$$\frac{1}{n}\sum_{i=1}^{n} \int \hat{f}_{-i}^2(y)dy - \frac{2}{n}\sum_{i=1}^{n} \hat{f}_{-i}(y_i),$$

where $\hat{f}_{-i}(y)$ denotes the estimator constructed from the data without the observation y_i. The value of h which minimises this expression is known as the *cross-validatory* smoothing parameter.

Plug-in methods

By inserting suitable estimates of the unknown quantities in the formula for the optimal smoothing parameter, a *plug-in* choice can be constructed. The difficult part is the estimation of $\beta(f)$ as this involves the second derivative of the density function. Sheather & Jones (JRSSB 53, 683–90) came up with a good, stable way of doing this. The Sheather-Jones method remains one of the most effective strategies for choosing the smoothing parameter.

A very simple plug-in approach is to use the normal density function in the expression for the optimal smoothing parameter. This yields the simple formula

$$h = \left(\frac{4}{3n}\right)^{1/5} \sigma,$$

where σ denotes the standard deviation of the distribution. This is a surprisingly effective means of smoothing data, in large part because it is very stable.

2.2.4 Some simple inferential tools

Once an estimate has been constructed, a natural next step is to find its standard error. The earlier result on the variance of $\hat{f}$ is a natural starting point, but this expression involves the unknown density. A helpful route is to consider a 'variance stabilising' transformation. For any transformation $t(\cdot)$, a Taylor series argument shows that

$$\mathsf{var}\left\{t(\hat{f}(y))\right\} \approx \mathsf{var}\left\{\hat{f}(y)\right\} \left[t'\left(\mathbb{E}\left\{\hat{f}(y)\right\}\right)\right]^2.$$

When $t(\cdot)$ is the square root transformation, the principal term of this expression becomes

$$\mathsf{var}\left\{\sqrt{\hat{f}(y)}\right\} \approx \frac{1}{4}\frac{1}{nh}\,\alpha(w),$$

which does not depend on the unknown density f. This forms the basis of a useful *variability band*. We cannot easily produce proper confidence intervals because of the bias present in the estimate. However, if the standard error is constructed and the intervals corresponding to two s.e.'s on the square root scale are transformed back to the origin scale, then a very useful indication of the variability of the density estimate can be produced. This is shown below for the aircraft span data from period 3.

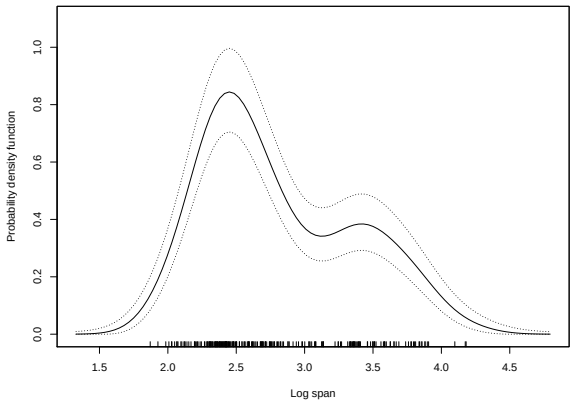

A useful variation on this arises when the true density function is assumed to be normal with mean μ and variance σ^2, and the kernel function w is also normal. If the standard normal density function is denoted by ϕ, then the mean and variance of the density estimate at the point y are then

$$\begin{aligned}
\mathbb{E}\left\{\hat{f}(y)\right\} &= \phi\left(y-\mu; \sqrt{h^2+\sigma^2}\right) \\
\mathsf{var}\left\{\hat{f}(y)\right\} &= \frac{1}{n}\phi\left(0; \sqrt{2}\,h\right)\ \phi\left(y-\mu; \sqrt{\sigma^2+\frac{1}{2}h^2}\right) \\
&\quad -\frac{1}{n}\phi\left(y-\mu; \sqrt{\sigma^2+h^2}\right)^2
\end{aligned}$$

These expressions allow the likely range of values of the density estimate to be calculated, under the assumption that the data are normally distributed. This can be expressed graphically through a *reference band*.

Example 2.2 (Icelandic tephra layer). Data on the percentages of aluminium oxide found in samples from a tephra layer resulting from a volcanic eruption in Iceland around 3500 years ago are available in the `tephra` dataset in the `sm` package. To deal with the percentage scale, apply the logit transformation

```
logit <- log(tephra$Al2O3/(100-tephra$Al2O3))
```

.

Can the variation in the tephra data be adequately modelled by a normal distribution on this scale? ◁

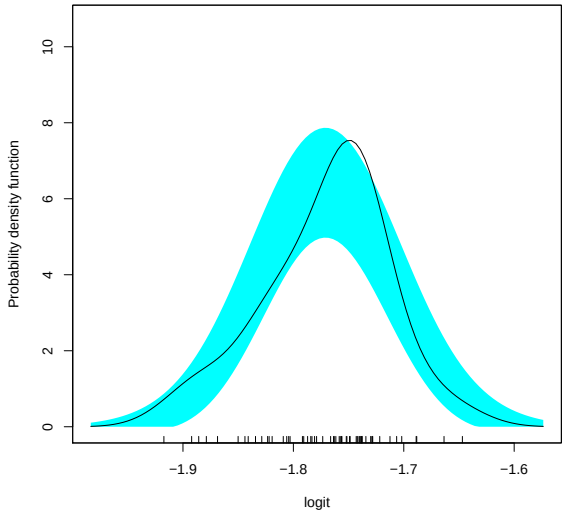

The density estimate does not fit comfortably within the reference band at all points and this effect persists across a wide range of smoothing parameters. A global test of normality could be developed but the graphical device of a reference band offers very useful informal insight.

The preliminary material also discussed the role of the bootstrap in capturing the variability, but not immediately the bias, of a density estimate.

2.3 Density estimation using mixtures

2.3.1 Introduction

Kernel density estimation is not the only way of estimating a density. When using a Gaussian kernel for density estimation we have constructed the estimate by adding up n Gaussian densities (where n is the sample size). When using mixtures we try to achieve the same using fewer Gaussian densities. However, in order for this to work, we need to estimate the mean and dispersion parameter of each Gaussian rather than considering Gaussian densities of the same dispersion anchored at each observation. Mixture-based methods for density estimation typically fare better than kernel-based methods for medium-dimensional data (none of the approaches work particularly well for high-dimensional data).

The key idea of mixture models is that whilst the density of interest might not necessarily be of a simple form, we can split the data into groups such that within each group the density has a simple form (e.g. Gaussian). This idea is illustrated by the example below.

Example 2.3 (Fisher's iris data). Consider the measurements of the petal width in the iris dataset. Figure 2.1 shows a plot of the (estimated) densities of the petal width for the three different species and a plot of the density for the whole dataset. The plots suggest that it is reasonable

to assume that the petal width in each group is from a Gaussian distribution. Denote with Y_i the petal width and with S_i the species name of the i-th flower. Then Y_i is (assumed to be) Gaussian in each group and the mean and the variance can be estimated by the empirical mean and variance in the different populations:

$(Y_i|S_i = se) \sim N(\mu_{se}, \sigma^2_{se})$ for *iris setosa* (estimates: $\hat{\mu}_{se} = 0.246$, $\hat{\sigma}^2_{se} = 0.01133$)
$(Y_i|S_i = ve) \sim N(\mu_{ve}, \sigma^2_{ve})$ for *iris versicolor* (estimates: $\hat{\mu}_{ve} = 1.326$, $\hat{\sigma}^2_{ve} = 0.03990$)
$(Y_i|S_i = vi) \sim N(\mu_{vi}, \sigma^2_{vi})$ for *iris virginica* (estimates: $\hat{\mu}_{vi} = 2.026$, $\hat{\sigma}^2_{vi} = 0.07697$)

The density of Y (species confounded) is then

$$f(y) = \pi_{se} f_{\mu_{se},\sigma^2_{se}}(y) + \pi_{ve} f_{\mu_{ve},\sigma^2_{ve}}(y) + \pi_{vi} f_{\mu_{vi},\sigma^2_{vi}}(y)$$

The density of Y is thus a *mixture* of three normal distributions. The π_{se}, π_{ve} and π_{vi} denote the probability that a randomly picked iris is of the corresponding species. These probabilities are sometimes referred to as *mixing weights*.

In the typical context of density estimation we would not have the grouping information provided by the species labels, i.e. we now assume we only know the petal widths. Nonetheless, we can stick to our idea that the density of Y is a mixture of Gaussians. Thus we assume

$$f(y) = \pi_1 f_{\mu_1,\sigma^2_1}(y) + \pi_2 f_{\mu_2,\sigma^2_2}(y) + \pi_3 f_{\mu_3,\sigma^2_3}(y).$$

Note that we had to replace the species names with indexes, because we don't possess the species information any more. This makes the estimation of the means μ_1, μ_2, μ_3 and variances $\sigma^2_1, \sigma^2_2, \sigma^2_3$ far more difficult. The species name S_i is now a *latent* variable, as we are unable to observe it. Once estimated, the S_i however provide potentially interesting information: they show how the data can be clustered into three groups. ◁

2.3.2 Model

We will for now assume that the component densities are Gaussian distributions, as this is by far the most common model. However any other distribution could be used. If we want the mixture to be able to approximate an arbitrary distribution we would have to choose a component density that can, in the limit, be reduced to a point mass (just like a Gaussian is reduced to a point mass if $\sigma^2 \to 0$.

We assume that the data comes from K groups ("clusters"). For the moment K is assumed to be fixed and known. We assume that the data in cluster k is from a (multivariate) Gaussian distribution with mean $\boldsymbol{\mu}_k$ and covariance matrix $\boldsymbol{\Sigma}_k$. Thus the density of $\mathbf{y}$ in cluster k is

$$f_{\boldsymbol{\mu}_k,\boldsymbol{\Sigma}_k}(\mathbf{y}) = \frac{1}{\sqrt{(2\pi)^p \cdot |\boldsymbol{\Sigma}_k|}} e^{-\frac{1}{2}(\mathbf{y}-\boldsymbol{\mu}_k)'\boldsymbol{\Sigma}_k^{-1}(\mathbf{y}-\boldsymbol{\mu}_k)}.$$

Often one assumes that all clusters have the same orientation and spread, i.e. $\boldsymbol{\Sigma}_k = \boldsymbol{\Sigma}$ or that all clusters are spherical, i.e. $\boldsymbol{\Sigma} = \sigma^2 \mathbf{I}$. Another popular assumption is that the covariance

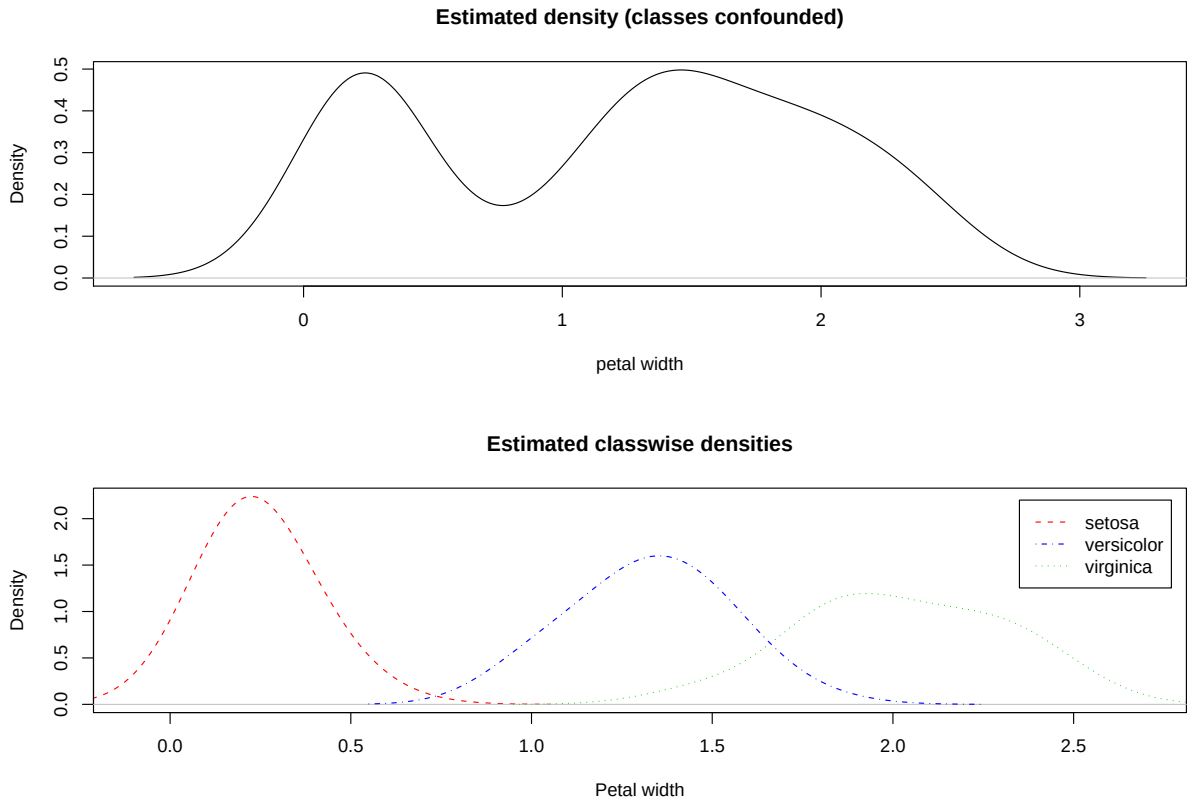

Figure 2.1. Estimated density of the petal width for all species confounded (above) and by species (below)

matrices are diagonal, i.e. $\mathbf{\Sigma}_k = \begin{bmatrix} {\sigma_1^{(k)}}^2 & \cdots 0 & \\ \vdots & \ddots & \vdots \\ 0 & \cdots & {\sigma_p^{(k)}}^2 \end{bmatrix}$. In this case the clusters are ellipses with axes parallel to to coordinate axes. Figure 2.2 visualises these different assumptions.

The density of $\mathbf{y}$ (all clusters confounded) is the *mixture* of the distributions in the clusters:

$$f(\mathbf{y}) = \sum_{k=1}^{K} \pi_k f_{\boldsymbol{\mu}_k, \mathbf{\Sigma}_k}(\mathbf{y}).$$

with $\pi_k \geq 0$ such that $\sum_{k=1}^{K} \pi_k = 1$. This yields the likelihood

$$\ell(\pi_1, \ldots, \pi_K, \boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_K, \mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_K) = \prod_{i=1}^{n} \left(\sum_{k=1}^{K} \pi_k f_{\boldsymbol{\mu}_k, \mathbf{\Sigma}_k}(\mathbf{y}_i) \right). \tag{2.1}$$

2.3.3 Inference

Estimating the parameters of the mixture model appears difficult at first sight. Each contribution to the likelihood (2.1) is a sum rather than a product, which is usually the case. Often the loglikelihood is a simpler expression than the likelihood itself (think of a Gaussian), however for mixtures this is not the case: each contribution to the loglikelihood is the logarithm of a sum, which cannot be simplified any further.

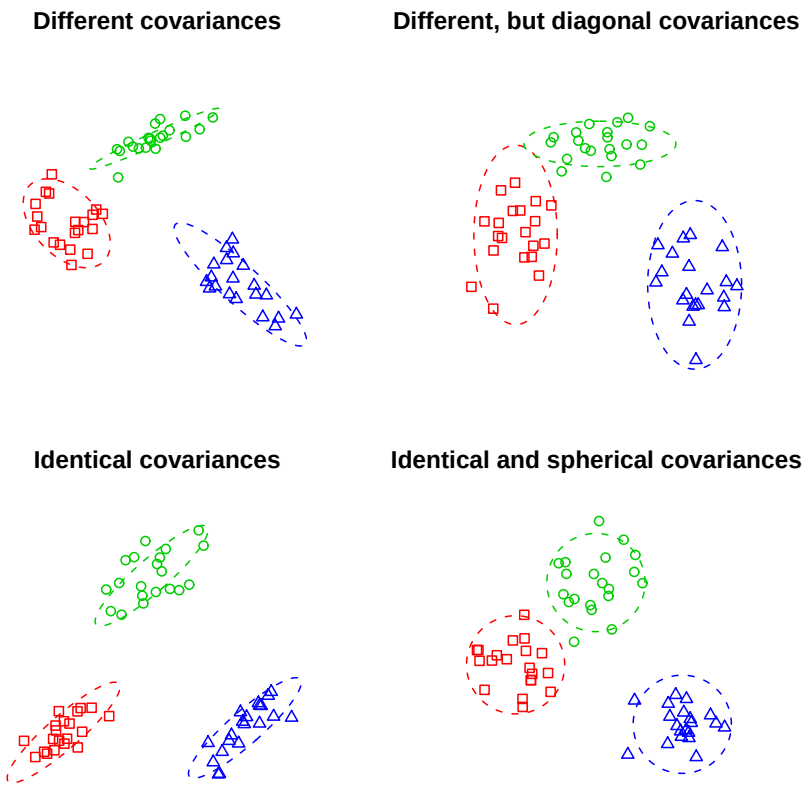

Figure 2.2. Examples for the different assumptions on the covariance structure

The problem can be solved by considering an EM algorithm, which introduces auxiliary variables S_i which indicate the cluster allocation. The EM algorithm now consists of the following iterative process.

For $h = 1, 2, \ldots$ iterate the E-Step and the M-Step:

E-Step Estimate the distribution of S_i given the data and the previous values of the parameters $\hat{\pi}_k^{(h-1)}$, $\hat{\boldsymbol{\mu}}_k^{(h-1)}$ and $\hat{\boldsymbol{\Sigma}}_k^{(h-1)}$ $(k = 1, \ldots, K)$. This yields the "responsibilities" $w_{ik}^{(h)}$:

$$w_{ik}^{(h)} := \frac{\pi_k^{(h-1)} f_{\hat{\boldsymbol{\mu}}_k^{(h-1)}, \hat{\boldsymbol{\Sigma}}_k^{(h-1)}}(\mathbf{y}_i)}{\sum_{\zeta=1}^{K} \pi_\zeta^{(h-1)} f_{\hat{\boldsymbol{\mu}}_\zeta^{(h-1)}, \hat{\boldsymbol{\Sigma}}_\zeta^{(h-1)}}(\mathbf{y}_i)}$$

M-Step Compute the updated $\hat{\pi}_k^{(h)}$, $\hat{\boldsymbol{\mu}}_k^{(h)}$, and $\hat{\boldsymbol{\Sigma}}_k^{(h)}$ $(k = 1, \ldots, K)$ by estimating them from the data using the vector $\mathbf{w}_k^{(h)} = (w_{1k}^{(h)}, \ldots, w_{nk}^{(h)})$ of responsibilities as weights: This leads to the updated estimators

$$\hat{\pi}_k^{(h)} = \frac{1}{n} \sum_{i=1}^{n} w_{ik}^{(h)},$$

$$\hat{\boldsymbol{\mu}}_k^{(h)} := \frac{1}{\sum_{i=1}^{n} w_{ik}^{(h)}} \sum_{i=1}^{n} w_{ik}^{(h)} \mathbf{y}_i,$$

and

$$\hat{\boldsymbol{\Sigma}}_k^{(h)} := \frac{1}{\sum_{i=1}^{n} w_{ik}^{(h)}} \sum_{i=1}^{n} w_{ik}^{(h)} (\mathbf{y}_i - \hat{\boldsymbol{\mu}}_k^{(h)})(\mathbf{y}_i - \hat{\boldsymbol{\mu}}_k^{(h)})^\top$$

for the covariances $\boldsymbol{\Sigma}_k$ (assuming different covariances in the classes).

Figure 2.3 visualises the flow of the algorithm for an easy toy example.

Derivation of the EM algorithm for mixtures

The EM algorithm is based on the likelihood that is obtained when assuming the class labels S_i are known. In this case the likelihood is

$$\prod_{i=1}^{n} \pi_{S_i} f_{\boldsymbol{\mu}_{S_i}, \boldsymbol{\Sigma}_{S_i}}(\mathbf{y})$$

Note that the sum has now disappeared. This is due to the fact that we know the cluster labels now, so we just multiply the densities of the distribution in the corresponding clusters instead of multiplying the unpleasant mixture density. For simplification of the notation denote with $\boldsymbol{\theta}$ the whole parameter vector, i.e.

$$\boldsymbol{\theta} := (\pi_1, \ldots, \pi_K, \boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_K, \boldsymbol{\Sigma}_1, \ldots, \boldsymbol{\Sigma}_K)$$

The log-likelihood is thus

$$L(\boldsymbol{\theta}) = \sum_{i=1}^{n} (\log \pi_{S_i} + \log f_{\boldsymbol{\mu}_{S_i}, \boldsymbol{\Sigma}_{S_i}}(\mathbf{y})) \tag{2.2}$$

In reality however, the S_i are unknown. So the EM-algorithm uses the trick of taking the (conditional) expectation of the likelihood.

In the E-step we have to compute the conditional expectation of the log-likelihood (2.2) given the data using the previous values of the parameters $\hat{\boldsymbol{\theta}}^{(h-1)}$. Thus we have to compute the conditional distribution of S_i (using the Bayes formula) yielding the responsibilities

$$w_{ik}^{(h)} := \mathbb{P}_{\hat{\boldsymbol{\theta}}^{(h-1)}}(S_i = k) = \frac{\hat{\pi}_k^{(h-1)} f_{\hat{\boldsymbol{\mu}}_k^{(h-1)}, \hat{\boldsymbol{\Sigma}}_k^{(h-1)}}(\mathbf{y}_i)}{\sum_{\zeta=1}^{K} \hat{\pi}_\zeta^{(h-1)} f_{\hat{\boldsymbol{\mu}}_\zeta^{(h-1)}, \hat{\boldsymbol{\Sigma}}_\zeta^{(h-1)}}(\mathbf{y}_i)}$$

The conditional expectation is thus

$$\mathbb{E}_{\hat{\boldsymbol{\theta}}^{(h-1)}}\left(L(\boldsymbol{\theta})|\mathbf{y}_1, \ldots, \mathbf{y}_n\right) = \sum_{i=1}^{n}\sum_{k=1}^{K}(w_{ik}^{(h)} \log \pi_k + w_{ik}^{(h)} \log f_{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k}(\mathbf{y})). \tag{2.3}$$

In the M-step we have to find $\hat{\boldsymbol{\theta}}^{(h)}$ by maximising the conditional expectation (2.3) ober $\boldsymbol{\theta}$. One can easily see that the updated parameter estimates $\hat{\pi}_k^{(h)}$, $\hat{\boldsymbol{\mu}}_k^{(h)}$, and $\boldsymbol{\Sigma}_k^{(h)}$ maximise (2.3).

2.3.4 Estimating the number of components

Estimating the number of components is, at least from a theoretical point of view, difficult. We cannot use classical statistical tests for this. To illustrate this consider a mixture model with two univariate Gaussian components

$$f(y) = \pi f_{\mu_1, \sigma_1^2}(y) + (1 - \pi) f_{\mu_2, \sigma_2^2}(y).$$

If we want to test whether we need one or two components, we could test the null hypothesis that $\pi = 1$. The first obvious issue is that $\pi = 1$ is on the boundary of the parameter space for π, which is $[0, 1]$, so standard approaches for likelihood ratio tests would fail. The even bigger problem however is that if $\pi = 1$, the parameters μ_2 and σ_2^2 do not influence the likleihood any more. In other words, the effective dimension of the problem collapses on the boundary, which makes deriving theoretical results for likelihood ratio tests very challenging.

For this reason, it is common practice to use model selection criteria such as BIC to select the number of components.

However, if our only intention is to estimate the density rather than interpret the clustering identified, getting the number of components exactly right is less of an issue.

2.3.5 Example

Example 2.4 (Aircraft data (ctd.)). Figure 2.4 shows the results obtained when using mixtures of Gaussians to estimate the density of the first two principal components of the aircrcraft data for

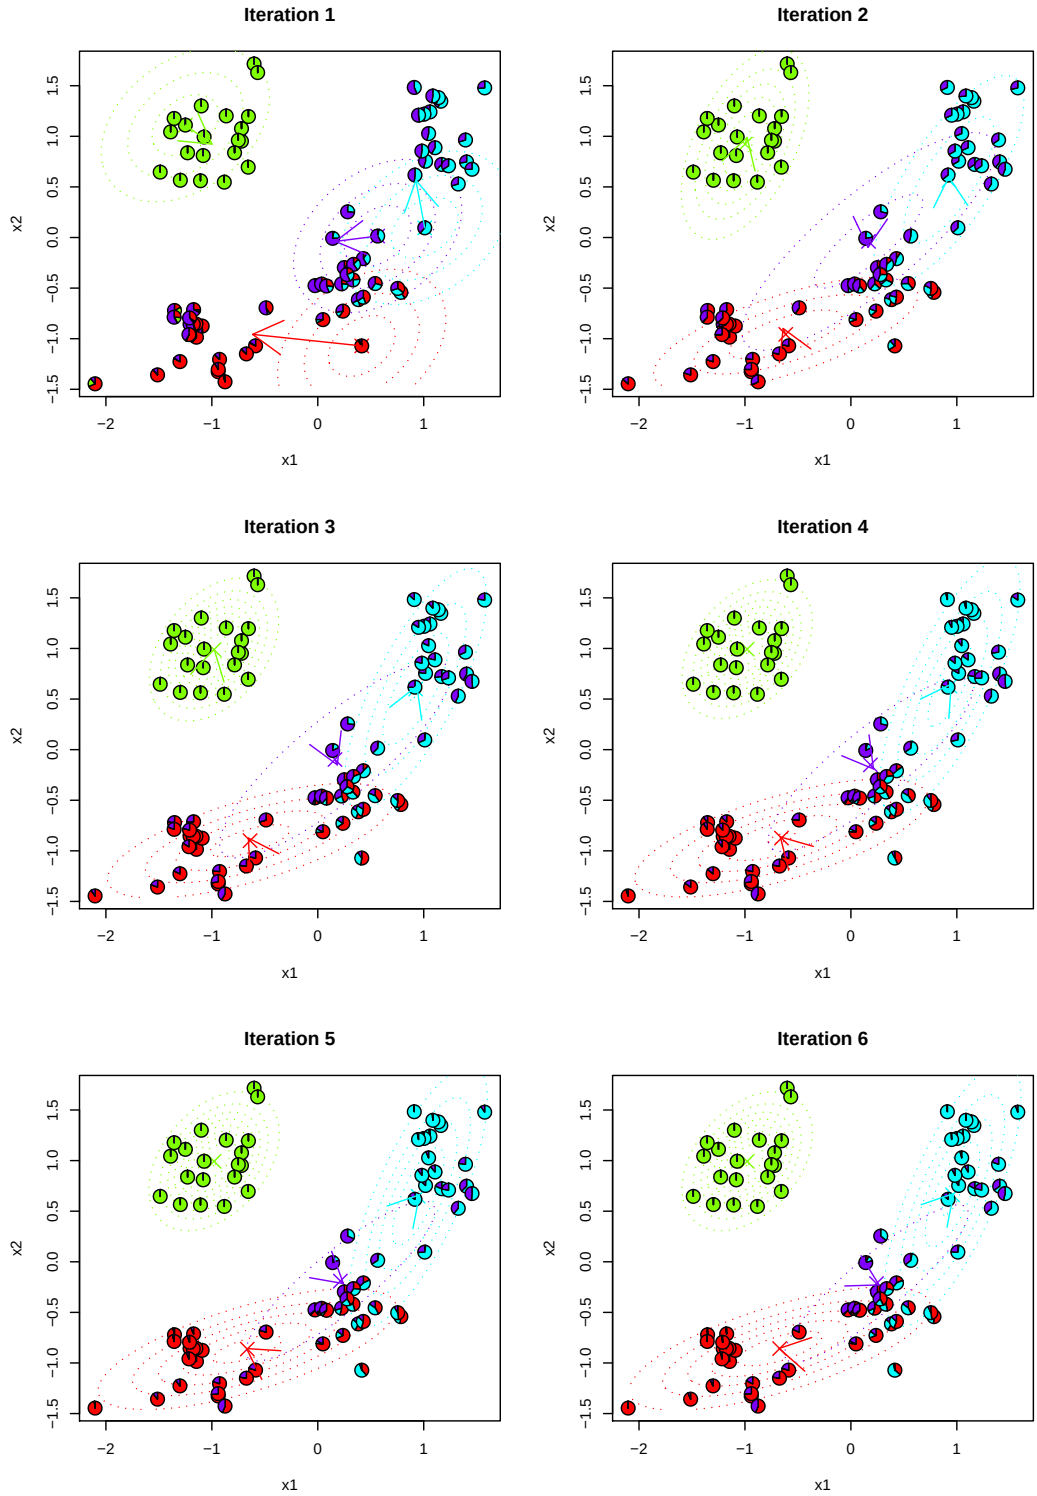

Figure 2.3. First 6 iterations of the EM algorithm for fitting a Gaussian mixture model to a toy example. The colouring of the data points corresponds to the responsibilities, the dotted lines show the isodensites of the Gaussians, and the arrow indicates where the means of the clusters are moved.

the third period. Due to the nature of the mixture model, the peaks are much more pronounced and most importantly there is a very concentrated fourth peak that was not apparent in the results obtained using kernel density estimation. Note that there would have been no need to compute principal components before the density estimation using mixtures of normals. There would not have been an issue using the raw data of the six measurements per aircraft. The principal components were used to allow better comparison to the results obtained using a kernel density estimate, which would not be able to cope well with data of that dimension. ◁

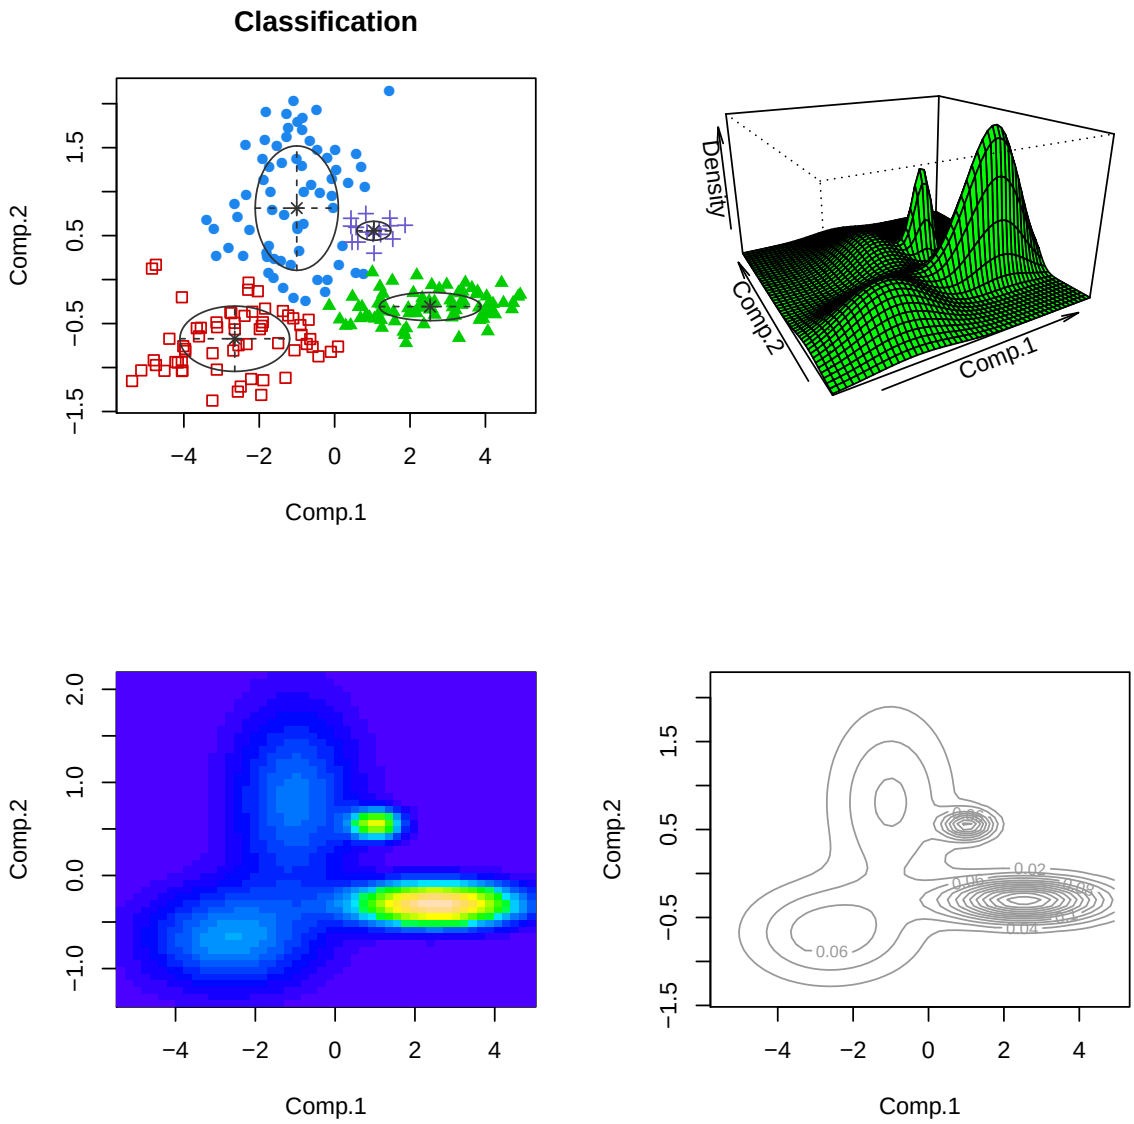

Figure 2.4. Estimated density using mixtures of normals for the first two principal components of aircraft data (using the R package `mclust`)

3

Splines

3.1 Introduction

This chapter covers splines, which are one of the most popular tools for flexible modelling. This section discusses a number of more philosophical concepts, some of which we have already touched upon.

In parametric modelling (*e.g.* estimating the rate of a Poisson distribution, linear regression) we assume we know the data generating process up to a finite number of parameters. In flexible modelling we want to fit a function to data, without making such a strict parametric assumption. All we are willing to assume is typically that the function of interest is sufficiently smooth. More formally speaking, this corresponds to working with an infinite-dimensional parameter space. This flexible approach has a number of advantages, most importantly it is less likely that the model is mis-specified. However there is a price to pay. Estimation becomes more difficult.

Example 3.1. Figure 3.1 shows two smooth functions describing the relationship between the response Y_i and the covariate x_i. In this example both functions yield the same fitted values $\hat{y}_i = \hat{m}(x_i)$. This also implies that the least-squares loss $\sum_{i=1}^{n}(y_i - \hat{m}(x_i))^2$ is the same for both functions, i.e. the data alone does not tell us which function does a better job. There is no global answer to this question.

Which of the two functions appears better suited to us depends on the context and also to some extent our subjective choice. In most circumstances we would prefer the function in the left-hand panel as it is the "simpler" function. However, if we expect the signal to have a periodic component (say we are expecting a day-of-the-week effect) then we might prefer the function shown in the right-hand panel. ◁

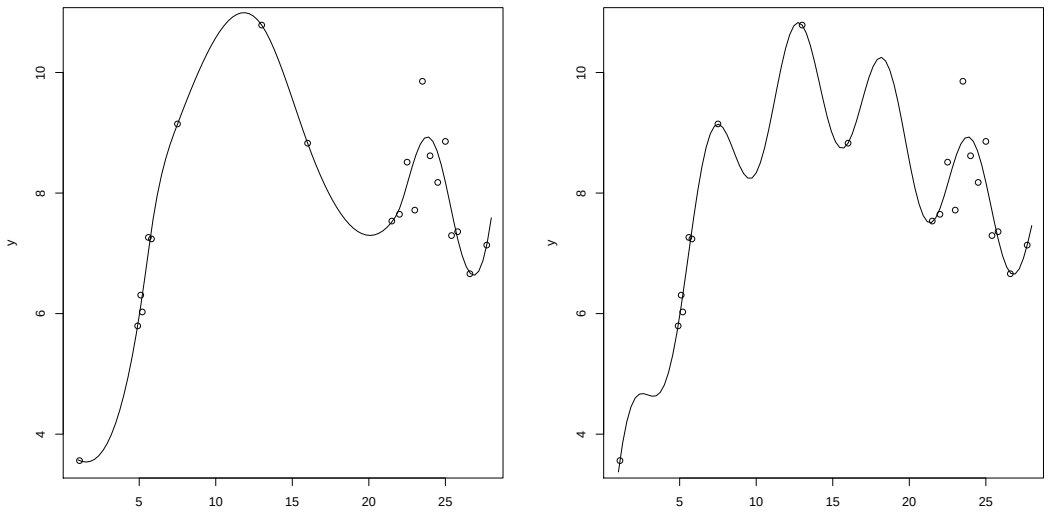

Figure 3.1. Two possible smooth functions modelling the relationship between the response Y_i and the covariate x_i. Note that both functions yield the same fitted values $\hat{y}_i = \hat{m}(x_i)$ and thus the same least-squares loss $\sum_{i=1}^{n}(y_i - \hat{m}(x_i))^2$.

What we have seen in the example is simply that the family of smooth functions is so large that observing a finite sample alone will not tell us enough to learn the function of interest $m(\cdot)$.

We need to provide additional information, which can be of different types:

– We can assume that the function of interest $m(\cdot)$ comes from a more restricted family of functions. We might even assume a rich class of parametric models. We will use this idea when we are looking at splines based on truncated power series and B-splines in section 3.2.4.
– We express a preference for some functions over others (without looking at the data) and use this in the model fitting procedure. Typically we prefer a smooth function to a more wiggly function. In a frequentist setting, this leads to penalty-based approach, or can be viewed as a Bayesian prior over the space of functions. We will discuss this in sections 3.2.3 and 3.3.

3.2 Univariate splines

3.2.1 Polynomial regression

We will start by revising polynomial regression. To fix notation, we quickly state the simple linear regression model

$$\mathbb{E}(Y_i) = \beta_0 + \beta_1 x_i \qquad \text{for } i = 1, \ldots, n,$$

or equivalently, in matrix-vector notation,

$$\mathbb{E}(\mathbf{y}) = \mathbf{B}\boldsymbol{\beta} \qquad \text{with } \mathbf{y} = (Y_1, \ldots, Y_n)^\top \text{ and } \mathbf{B} = \begin{pmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_n \end{pmatrix}.$$

The simple linear regression model can be extended into a polynomial regression model by including powers of the covariates x_i in the design matrix. The polynomial regression model

$$\mathbb{E}(Y_i) = \beta_0 + \beta_1 x_i + \ldots + \beta_r x_i^r \qquad \text{for } i = 1, \ldots, n,$$

just corresponds to linear regression using the expanded design matrix

$$\mathbf{B} = \begin{pmatrix} 1 & x_1 & \ldots & x_1^r \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_n & \ldots & x_n^r \end{pmatrix}.$$

We can then estimate $\boldsymbol{\beta}$ using the same techniques as used in multiple linear regression, i.e. the least-squares estimator is

$$\hat{\boldsymbol{\beta}} = (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{y}$$

Polynomial regression is a very simple example of a basis expansion technique. We have simply replaced the design matrix of simple linear regression by an augmented design matrix. In the case of polynomial regression we have simply added powers of the x_i's. Many of the techniques covered in this chapter will be based in this idea of basis expansions.

Polynomial regression can be a useful tool if a polynomial of very low order yields a sufficient fit to the data.

Example 3.2 (Glucose levels in potatoes). Figure 3.2 shows a quadratic regression model fitted to a simple data set from an experiment in which the glucose level in potatoes was measured over the course of several weeks. Given the small number of observations there is little need to go beyond a simple quadratic regression model. ◁

However, polynomial regression is not very well suited for modelling more complex relationships, as the following example shows.

Example 3.3. Consider the data set simulated using the model

$$y_i = 1 - x_i^3 - 2\exp(-100x_i^2) + \varepsilon_i$$

with $\mathbf{x} = (-1, -0.98, \ldots, 0.98, 1)$ and $\varepsilon_i \sim \mathsf{N}(0, 0.1^2)$. Figure 3.3(a) shows the data together with the fitted function obtained for a polynomial regression model of degree 10. The polynomial model of degree 10 is not flexible enough to capture the sharp dip around 0. If we

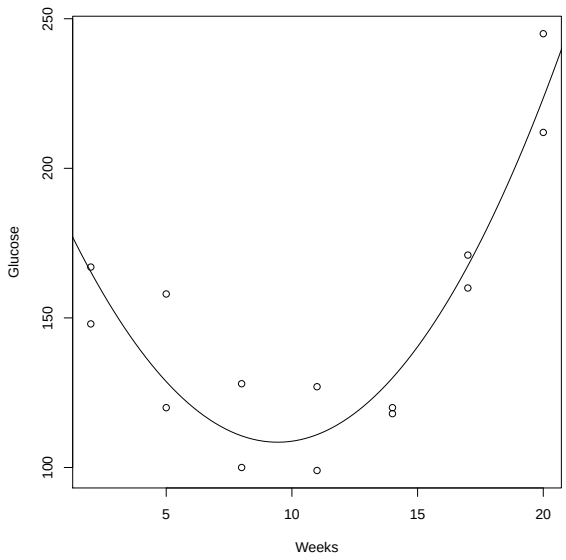

Figure 3.2. Glucose level in potatoes. The solid line is the fitted regression function obtained from quadratic regression.

increase the degree to 17 we can capture the dip better (panel (b)). However, the polynomial fit of degree 17 shows strong oscillations which are not supported by the data. Panel (c) shows the fitted regression function using a spline based model, which we will discuss later on in this chapter. The spline-based approach can capture the sharp dip much better and without yielding any oscillations.

Figure 3.4 allows some insight into why the polynomial model struggles. It shows `image` plots of the hat matrix $\mathbf{S} = \mathbf{B}(\mathbf{B}^\top\mathbf{B})^{-1}\mathbf{B}^\top$ for the three models under consideration. The hat matrix maps the observed response to the fitted response, i.e.

$$\hat{\mathbf{y}} = \mathbf{B}\hat{\boldsymbol{\beta}} = \mathbf{B}(\mathbf{B}^\top\mathbf{B})^{-1}\mathbf{B}^\top\mathbf{y} = \mathbf{S}\mathbf{y}$$

When performing flexible regression we would expect the prediction at x_i to almost only depend on observations close to x_i, i.e. we would expect the hat matrix $\mathbf{S}$ to be largely band-diagonal with a rather narrow band width. However, polynomials are not "local". As one can see from a Taylor series expansion, the coefficients of the polynomial can be learnt from higher order derivatives observed at a single point. The problem is that sharp dip provides more information than the data on either side of it, yielding to a poor fit on both sides. This is known as Runge's phenomenon in Numerical Analysis.

Figure 3.3(a) and figure 3.3(b) shows another drawback of polynomial regression. As $x \to \pm\infty$ the polynomial must go to $\pm\infty$ as well. This often leads to very high curvature at both ends of the range, which is typically not supported by the data.

Yet another reason for avoiding polynomial regression is that it is highly likely to be numerically unstable. Due to the large correlations between the powers of the x_i, which make up the columns of the design matrix, the design matrix $\mathbf{B}$ and the matrix of cross-products

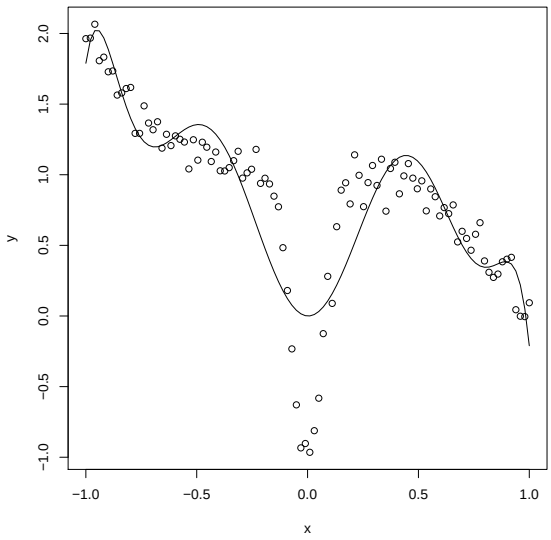

(a) Polynomial regression of degree 10

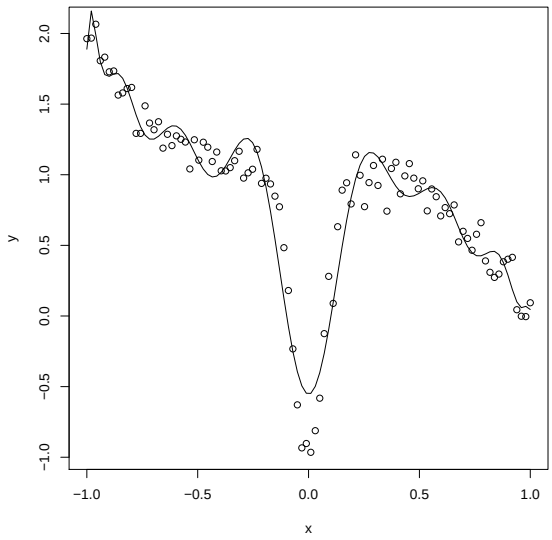

(b) Polynomial regression of degree 17

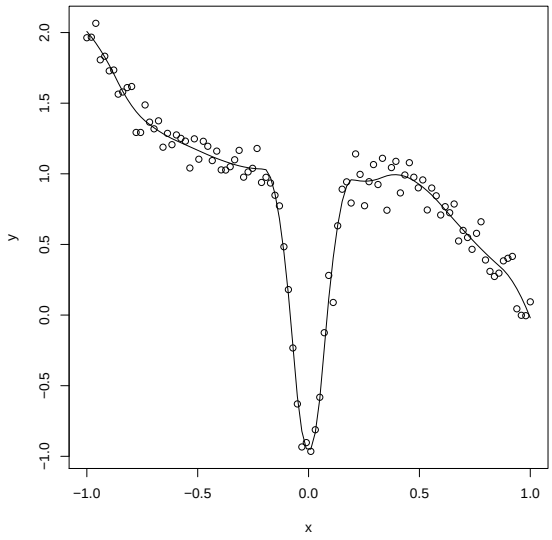

(c) Quadratic-spline-based regression

Figure 3.3. Data and fitted function for the simulated data from example 3.3 for polynomial regression of degrees 10 and 17 as well as for a spline-based model.

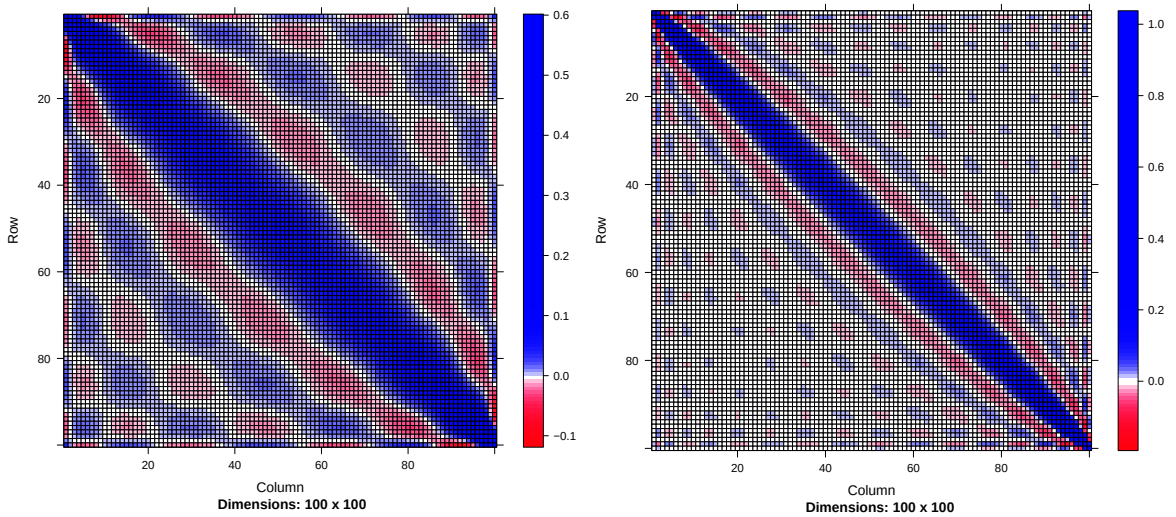

(a) Polynomial regression of degree 10

(b) Polynomial regression of degree 17

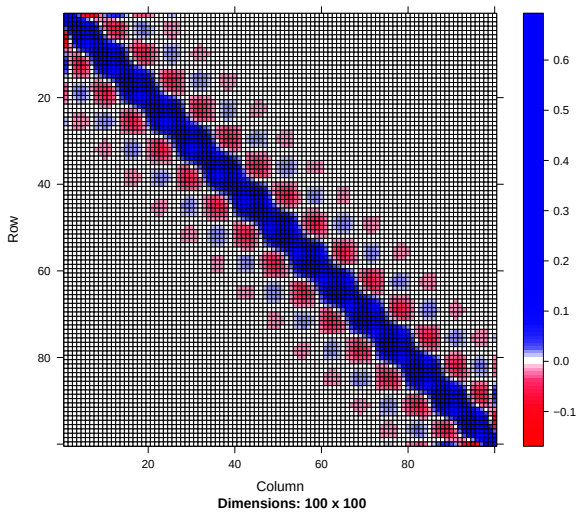

(c) Quadratic-spline-based regression

Figure 3.4. Hat matrix $\mathbf{S} = \mathbf{B}(\mathbf{B}^\top\mathbf{B})^{-1}\mathbf{B}^\top$ for polynomial regression of degrees 10 and 17 as well as for splines applied to the simulated data from example 3.3.

$\mathbf{B}^\top\mathbf{B}$ is very likely to be ill-conditioned. The condition number[1] of $\mathbf{B}^\top\mathbf{B}$ for the polynomial regression model of degree 17 is 1.56×10^{12}, i.e. $\mathbf{B}^\top\mathbf{B}$ is barely invertible. For comparison, the corresponding condition number for the spline-based model is 32.49. ◁

As we have seen in the example above, polynomial regression is, unless modelling very simple relationships, not a suitable tool for flexible regression. In the next section we will consider piecewise polynomial models, which are better suited for flexible regression. These are based on the idea of splitting the input domain and fitting low-order polynomials in each interval. As we can see from figure 3.5(a) fitting polynomials independently of each other does not yield satisfactory results. We will thus introduce additional constraints which make the function continuous and (potentially) differentiable (cf. panel (b)).

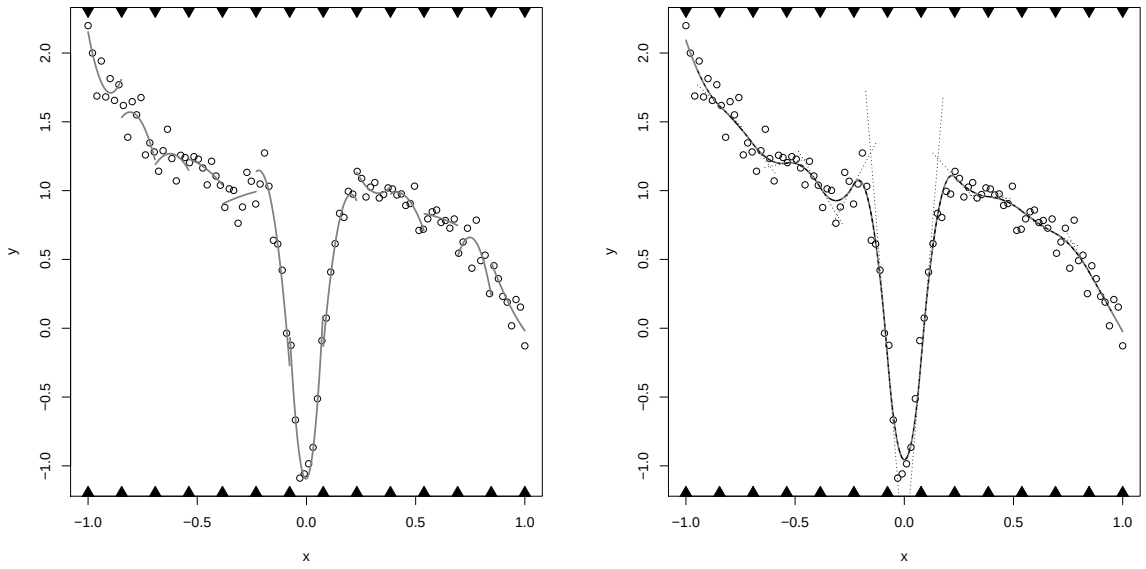

(a) Discontinuous piecewise polynomials

(b) Piecewise polynomials which form a continuously differentiable function (derivatives at knots shown as dashed lines)

Figure 3.5. Piece-wise polynomials fitted to the data from example 3.3 with an without smoothness constraints. The back triangles show the positions of the knots.

3.2.2 Polynomial splines

In this section we will introduce polynomial splines which are piecewise polynomials, which "glued together" at the knots so that the resulting function is r-times continuously differentiable.

[1] The condition number of a matrix is defined as the ratio of the largest singular value divided by the smallest singular value. For a symmetric positive-definite matrix this is the same as the ratio of the largest over the smallest eigenvalue. The condition number is of measure of how numerically unstable matrix operations like taking the inverse will be.

Definition 3.1 (Polynomial spline). *Given a set of knots $a = \kappa_1 < \kappa_2 < \ldots < \kappa_l = b$, a function $m : [a, b] \to \mathbb{R}$ is called a (polynomial) spline of degree r if*

– *$m(\cdot)$ is a polynomial of degree r on each interval (κ_j, κ_{j+1}) ($j = 1, \ldots, l-1$).*
– *$m(\cdot)$ is $r-1$ times continuously differentiable.*[2]

Historically, a spline was an elastic ruler used to draw technical designs, notably in shipbuilding and the early days of aircraft engineering. Figure 3.6 shows such a spline.[3]

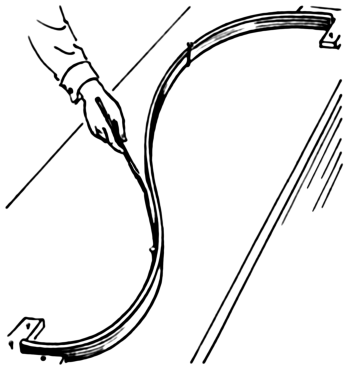

Figure 3.6. A spline.

Choice of degree r. The degree r of the spline controls the smoothness in the sense of controlling its differentiability. For $r = 0$ the spline is a discontinuous step function. For $r = 1$ the spline is a polygonal line. For larger values of r the spline is increasingly smooth, but also behaves more and more like one global polynomial. It is worth noting that assuming too smooth a function can have significant detrimental effects on the fitted regression function (*e.g.* oscillations, ballooning). In practice it is rarely necessary to go beyond $r = 3$.

Example 3.4 (Radiocarbon dating). In a scientific experiment high-precision measurements of radiocarbon were performed on Irish oak. To construct a calibration curve we need to learn the relationship between the radiocarbon age and the calendar age. Figure 3.7 shows spline fits to the data using splines of different degrees. ◁

Choice of the number of knots l. In an (unpenalised) spline the number of knots acts as a smoothing parameter. The more knots are used, the more flexible the regression function can become.

[2] For a spline of degree 0 the function $m(\cdot)$ does not need to be continuous. For a spline of degree 1 the function $m(\cdot)$ needs to be continuous, but does not need to be differentiable.

[3] See `http://pages.cs.wisc.edu/~deboor/draftspline.html` for a picture (probably from the 1960's) of a Boeing engineer using a spline.

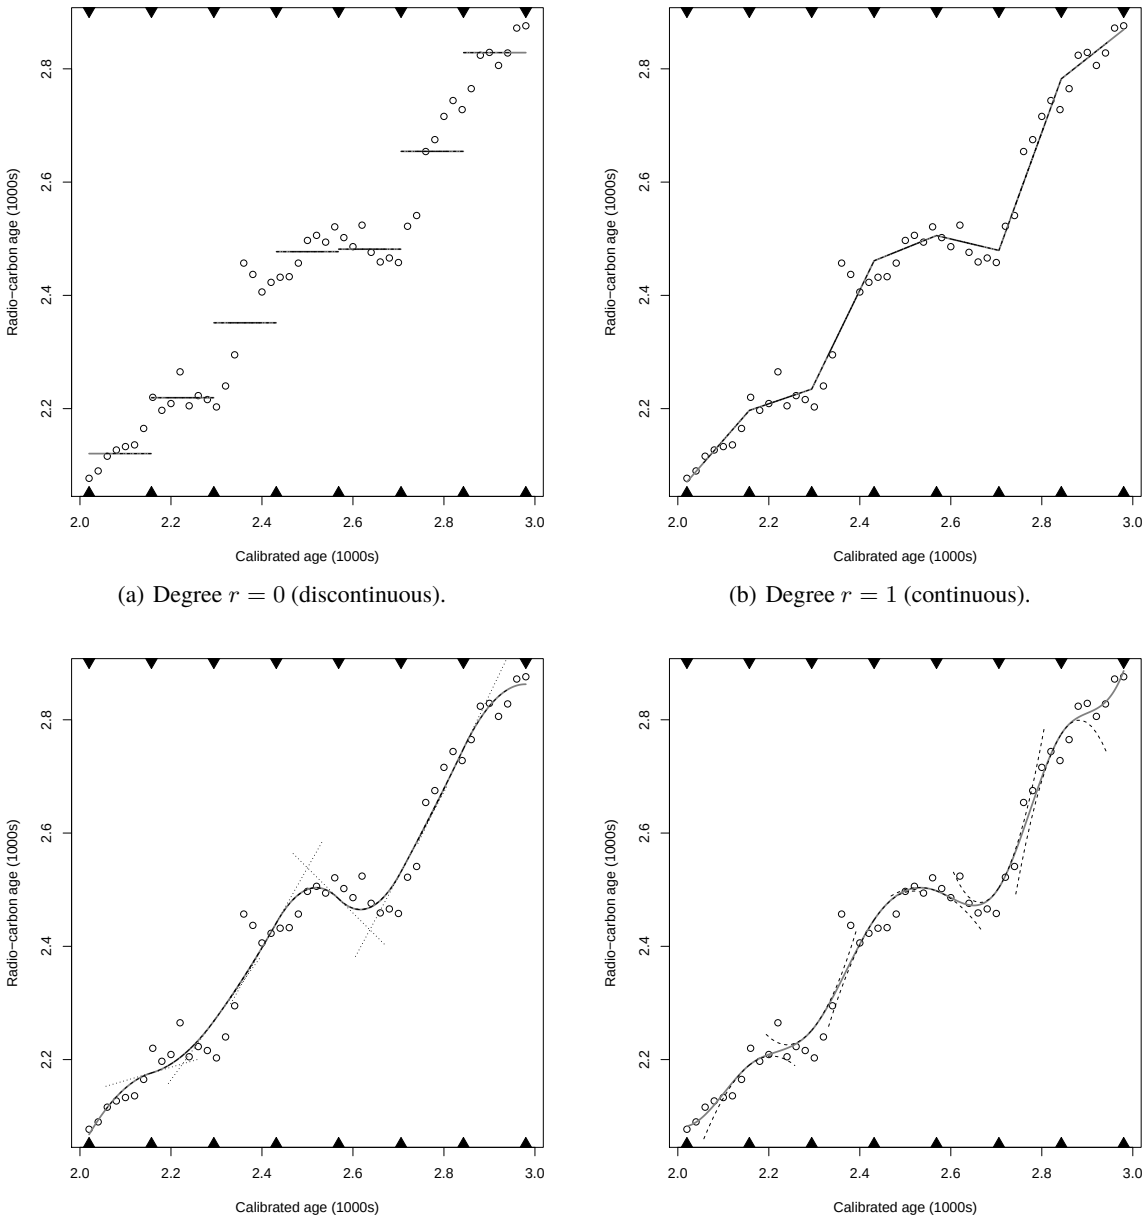

(a) Degree $r = 0$ (discontinuous).

(b) Degree $r = 1$ (continuous).

(c) Degree $r = 2$ (continuous first derivative).

(d) Degree $r = 3$ (continuous second derivative).

Figure 3.7. Splines of degree $r \in \{0, 1, 2, 3\}$ fitted to the radiocarbon data.

As we have seen in chapter **??** a more flexible regression function has a lower bias, but a higher variance.

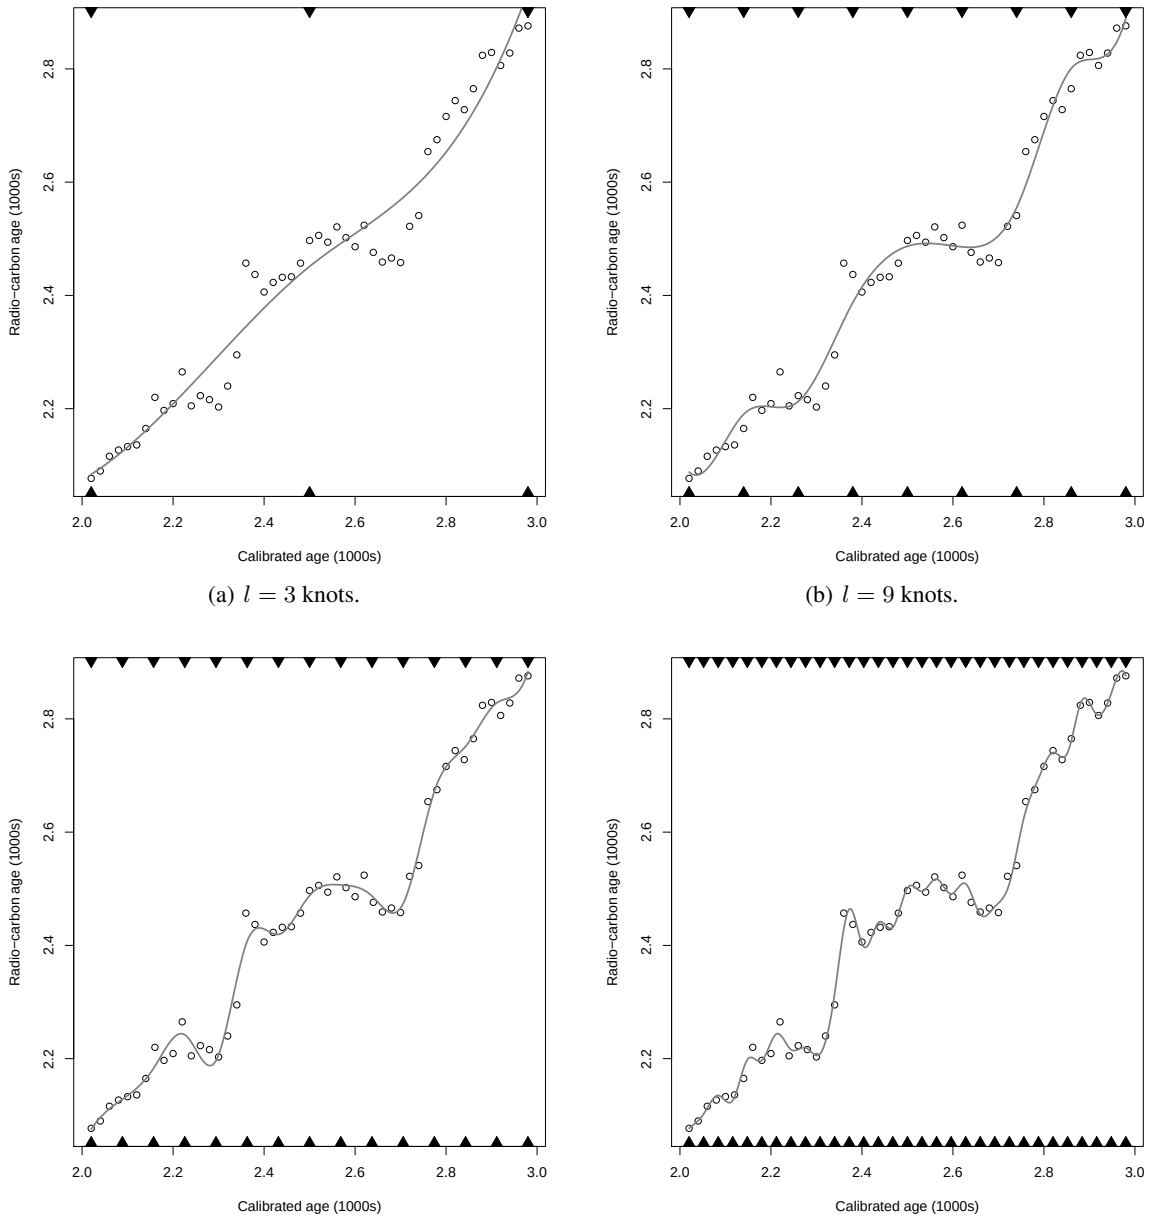

(a) $l = 3$ knots.

(b) $l = 9$ knots.

(c) $l = 15$ knots.

(d) $l = 31$ knots.

Figure 3.8. Cubic spline with different number of knots $l \in \{3, 9, 15, 31\}$ fitted to the radiocarbon data.

Example 3.5 (Radiocarbon dating (continued)). Figure 3.8 shows a cubic spline fitted to the radiocarbon data using an increasing number of knots. Too few knots lead to an underfit to the data: the fitted function does not fully represent the relationship between radiocarbon age and calendar age. Too many knots on the other hand lead to an overfit: the spline does not only pick up the signal, but also adapts to artefacts in the noise. ◁

Especially when the number of knots is small, the positioning of the knots can be important. The simplest strategy consist of using a set of equally spaced knots; this is computationally the simplest. Alternatively, we can place the knots according to the quantiles of the covariate. This makes the spline more flexible in regions with more data (and thus potentially more information) and less flexible in areas with less data (and potentially less information). A third strategy consists of trying to find an optimal placement of the knots. This usually is computationally very demanding.

Yet another approach consists of using "too many" knots — one knot per observation in the most extreme case — and use a penalty term to control for the smoothness. This avoid the need to select the number of knots altogether. We will study two such approaches in sections 3.2.3 and 3.3.

Splines as a vector space. For a given set of l knots and given degree r, the space of polynomial splines is a vector space, i.e. the sum of two splines as well as a scalar multiples of each spline are again splines. To find the dimension of the vector space have to find the number of "free parameters".

- Each polynomial has $r+1$ parameters and there are $l-1$ polynomials. Thus the spline model has $(r+1)\cdot(l-1)$ parameters. However we cannot choose all these parameters freely, as the resulting function needs to be $r-1$ times continuously differentiable.
- At the $l-2$ inside knots we have to guarantee that $m(\cdot)$ is $r-1$ times continuously differentiable. This corresponds to r constraints ($r-1$ constraints for each derivative and one for $m(\cdot)$ to be continuous). Thus there are $r\cdot(l-2)$ constraints (which are all linearly independent).

Thus there are $(r+1)\cdot(l-1)-r\cdot(l-2)=r+l-1$ free parameters. Thus the vector space of polynomial splines of degree r with l knots is $r+l-1$.

In section 3.2.4 we will explore different ways of constructing a basis for this space. The dimension will come in handy when proving that a given set of basis functions is indeed a basis of this space, as we only need to show that the basis functions are independent and that we use the correct number of basis functions.

Natural cubic splines. Finally, we will introduce the concept of a natural cubic spline. It is based on the idea that it is "safer" (or more "natural") to assume that the curvature of the spline at the first and last knot is zero. If we were to extrapolate, we would then extrapolate linearly.

Definition 3.2 (Natural cubic spline). *A polynomial spline $m:[a,b]\to\mathbb{R}$ of degree 3 is called a natural cubic spline if $m''(a)=m''(b)=0$.*

Given a set of l knots the vector space of all cubic splines has dimension $l + 2$. Natural cubic splines introduce two additional constraints, thus they form a vector space of dimension l. This makes natural cubic splines perfectly suited for interpolation.

Proposition 3.3. *A set of l points (x_i, y_i) can be exactly interpolated using a natural cubic spline with the $x_1 < \ldots < x_l$ as knots. The interpolating natural cubic spline is unique.*

Proof. The space of natural cubic splines with knots at $x_1, \ldots, x_l$ is vector space of dimension l. Introducing l additional constraints ($y_i = m(x_i)$ for $i = 1, \ldots, l$) yields a system of l equations and l free parameters, which yields a unique solution.[4] □

In the next section we will show that natural cubic spline have an important optimality property.

3.2.3 Optimality of splines

This section provides a theoretical justification for the choice of splines for flexible regression.

In this section we will ask a rather general question. Given a data set (x_i, y_i) with $a \leq x_i \leq b$ we try to find, amongst all twice continuously differentiable functions, the function which "best" models the relationship between response y_i and covariate x_i.

First of all, we need to specify what we mean by "best". We could look for the function $m(\cdot)$ which yields the smallest least-squares criterion

$$\sum_{i=1}^{n} (y_i - m(x_i))^2$$

This is however not a good idea. Any function which interpolates all the observations (x_i, y_i) would be optimal in this sense, yet such a function would typically not describe the relationship between x_i and y_i but rather model the artefacts of the random noise. Thus we will consider a so-called *penalised* (or *regularised*) criterion which tries to balance out two aspects which are important to us:

Fit to the data. We want $m(\cdot)$ to follow the data closely.

Simplicity/smoothness. We want the function $m(\cdot)$ not to be too complicated so that it generalises well to future unseen data.

We will thus the following penalised fitting criterion

[4] Strictly speaking, we would need to show that the system of equations cannot be rank-deficient, which could cause the solution to be either non-unique or non-existing.

$$\underbrace{\sum_{i=1}^{n}(y_i - m(x_i))^2}_{\text{Fit to the data}} + \lambda \underbrace{\int_a^b m''(x)^2\, dx}_{\text{Roughness penalty}}, \tag{3.1}$$

where $\lambda > 0$ is a tuning parameter which controls the trade off between following the data and preventing $m(\cdot)$ from being too rough.

We will now show that the minimiser of (3.2) over all twice continuously differentiable functions has to be a natural cubic spline, i.e. natural cubic splines with knots at each of the unique x_i are in this sense the optimal class functions.

We will start by showing that natural cubic splines are optimal interpolators, in the sense that they minimise the roughness penalty $\int_a^b m''(x)^2\, dx$.

Lemma 3.4. *Amongst all functions on $[a, b]$ which are twice continuously differentiable and which interpolate the set of points (x_i, y_i), a natural cubic spline with knots at the x_i yields the smallest roughness penalty*

$$\int_a^b m''(x)^2\, dx.$$

Proof. Let $m(\cdot)$ be the natural cubic spline with knots at the x_i, interpolating the data. Suppose there is another function $g(\cdot)$, which is twice continuously differentiable and which also interpolates the data. Denote by $h(x) = g(x) - m(x)$ the difference between the two functions.

1. We will first of all show that we can decompose

$$\int_a^b g''(x)^2\, dx = \int_a^b m''(x)^2\, dx + \int_a^b h''(x)^2\, dx$$

   i. As both $m(\cdot)$ and $g(\cdot)$ interpolate the (x_i, y_i) we have that $m(x_i) = g(x_i) = y_i$, thus $h(x_i) = g(x_i) - m(x_i) = 0$.

   ii. Using integration by parts we obtain that

$$\begin{aligned}
\int_a^b m''(x)h''(x)\, dx &= \underbrace{[m''(x)h'(x)]_{x=a}^{b}}_{=0 \text{ (as } m''(a) = m''(b) = 0)} - \int_a^b m'''(x)h'(x)\, dx \\
&= -\sum_{i=1}^{n-1} \int_{x_i}^{x_{i+1}} m'''(x)h'(x)\, dx \\
&= -\sum_{i=1}^{n-1} m'''\left(\frac{x_i + x_{i+1}}{2}\right) \cdot \underbrace{\int_{x_i}^{x_{i+1}} h'(x)\, dx}_{=\underbrace{h(x_{i+1})}_{=0} - \underbrace{h(x_i)}_{=0} = 0} \\
&= 0
\end{aligned}$$

   In the second line we have used that the natural cubic spline is piece-wise cubic polynomial, i.e. between two knots x_i and x_{i+1} the third derivative $m'''(x)$ is constant.

iii. Thus

$$\begin{aligned}\int_a^b g''(x)^2\,dx &= \int_a^b (g''(x) - m''(x) + m''(x))^2\,dx = \int_a^b (h''(x) + m''(x))^2\,dx \\ &= \int_a^b h''(x)^2\,dx + 2\underbrace{\int_a^b h''(x)m''(x)\,dx}_{=0} + \int_a^b m''(x)^2\,dx \\ &= \int_a^b h''(x)^2\,dx + \int_a^b m''(x)^2\,dx\end{aligned}$$

2. Because of $\int_a^b h''(x)^2\,dx \geq 0$ we have that

$$\int_a^b g''(x)^2\,dx = \int_a^b m''(x)^2\,dx + \underbrace{\int_a^b h''(x)^2\,dx}_{\geq 0} \geq \int_a^b m''(x)^2,$$

i.e. the natural cubic spline cannot have a larger roughness penalty.

3. In the above inequality equality holds if and only if $\int_a^b h''(x)^2\,dx = 0$, which, given that $h(x_i) = 0$, can only be the case if $g(x) = m(x)$ for all $x \in [a, b]$. □

Spline-based interpolation is implemented in the Rfunctions `spline` and `splinefun`.

We will now generalise the result about interpolation to the case of smoothing.

$$\underbrace{\sum_{i=1}^n (y_i - m(x_i))^2}_{\text{Model fit}} + \lambda \underbrace{\int_a^b m''(x)^2\,dx}_{\text{Roughness penalty}}, \tag{3.2}$$

where λ is a tuning parameter which controls the trade off between following the data and preventing $m(\cdot)$ from being too rough.

Theorem 3.5. *The minimiser of*

$$\sum_{i=1}^n (y_i - m(x_i))^2 + \lambda \cdot \int_a^b m''(x)^2\,dx$$

amongst all twice continuously differentiable functions on $[a, b]$ *is given by a a natural cubic spline with knots in the unique* x_i*.*

This is an extremely powerful theorem. Even though we consider the entire infinite-dimensional vector space of all twice continuously differentiable functions, we only need to consider the finite-dimensional vector space of natural cubic splines. We have thus reduced the complexity of the optimisation problem to the comparatively simple problem of finding the optimal coefficients of the natural cubic spline. This can be done using least-squares.

Proof. Let $g(\cdot)$ be a twice continuously differentiable function. We will now create a competitor to $g(\cdot)$, which is a natural cubic spline with knots in the x_i. We will now show that, unless $g(\cdot)$ is already a natural cubic spline, $m(\cdot)$ leads to a smaller value of the objective function. We choose the natural cubic spline $m(\cdot)$ such that it interpolates the fitted values $g(\cdot)$ generates, i.e. $m(x_i) = g(x_i)$. Thus $\sum_{i=1}^{n}(y_i - m(x_i))^2 = \sum_{i=1}^{n}(y_i - g(x_i))^2$, i.e. both functions model the data equally well, however as we have shown in Lemma 3.4 the natural cubic spline $m(\cdot)$ has the smaller roughness penalty. □

Note that the proof did not make use of the fact that we have used the least-squares loss function. In fact, the theorem holds for any pointwise loss function.

The technique of *smoothing splines* is based on this theoretical result and finds the natural cubic spline minimising (3.2), and, due to the theorem, the optimal function amongst all twice continuously differentiable functions. This approach is implemented in the Rfunction `smooth.spline`. We will study this method in more detail in section 3.3.

3.2.4 Constructing splines

In this section we will studies two ways of constructing a basis for the vector space of polynomial splines: the truncated power basis and the B-spline basis. We will only cover the case of generic polynomial splines. However one can modify these bases to only span the space of natural cubic splines.

Truncated power basis. The simplest basis for polynomial splines is the truncated power basis.

Definition 3.6 (Truncated power basis). *Given a set of knots $a = \kappa_1 < \ldots < \kappa_l = b$ the truncated power basis of degree r is given by*

$$\left(1, x, \ldots, x^{r-1}, (x-\kappa_1)_+^r, (x-\kappa_2)_+^r, \ldots, (x-\kappa_{l-1})_+^r\right),$$

where $(z)_+^r = \begin{cases} z^r & \text{for } z > 0 \\ 0 & \text{otherwise.} \end{cases}$

The truncated power basis has $r+l-1$ basis functions. It is easy to see that they are linearly independent. Thus the truncated power basis is indeed a basis of the vector space of polynomial splines. Figure 3.9 shows the truncated power series basis of degree 3 for six equally spaced knots.

To fit a polynomial spline to data we can exploit the fact the truncated power basis is a basis of the vector space of polynomial splines of the given degree and with the given set of knots. Thus we can write any spline $m(\cdot)$ as a linear combination of the basis functions, i.e.

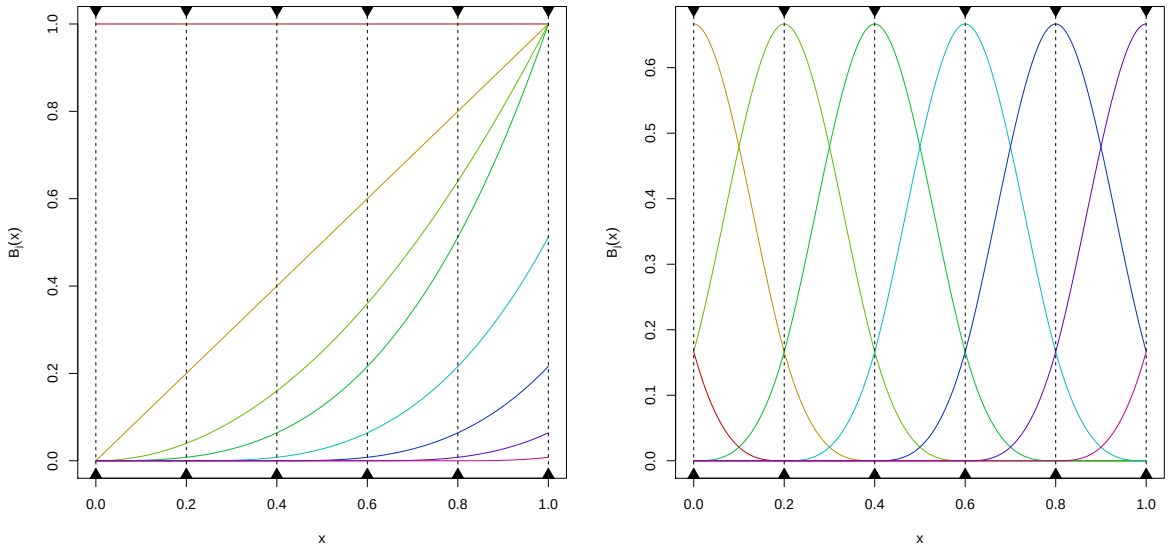

(a) Truncated power series (b) B-splines

Figure 3.9. Basis functions $B_j(x)$ of the cubic truncated power series basis (left panel) and B-splines (right panel). The vertical lines indicate the location of the knots.

$$m(x) = \beta_0 + \beta_1 x + \ldots + \beta_{r-1} x^{r-1} + \beta_r (x - \kappa_1)_+^r + \ldots + \beta_{r+l-2} (x - \kappa_{l-1})_+^r$$

We can thus find the optimal spline $m(\cdot)$ by just finding the optimal set of coefficients β_j, which is nothing other than a linear regression problem with design matrix

$$\mathbf{B} = \begin{pmatrix} 1 & x_1 & \ldots & x_1^{r-1} & (x_1 - \kappa_1)_+^r & \ldots & (x_1 - \kappa_{l-1})_+^r \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_n & \ldots & x_n^{r-1} & (x_n - \kappa_1)_+^r & \ldots & (x_n - \kappa_{l-1})_+^r \end{pmatrix}$$

We can use the design matrix $\mathbf{B}$ in exactly the same way as we would use the design matrix of a classical linear model.

We can interpret the truncated power series as a regression model in which the leading coefficient changes at each knot. At each knot, the remaining coefficients change as well. However they are fully constrained by the condition that the spline has to be $r - 1$ times continuously differentiable at each knot.

Example 3.6 (Radiocarbon data (continued)). Figure 3.10 illustrates the use of a truncated power series basis for fitting a spline-based flexible regression model for the radiocarbon data.

As one can see from the middle panel of figure 3.10 and from figure 3.11, some of the estimated coefficients are very large: some of the basis functions are scaled up by a factor of more than 1000, with "neighbouring" basis functions having opposite signs. The reason for this is the high correlation between the columns of the design matrix of the truncated power series. The largest correlation between columns is 0.99921, which is very close to 1.

Figure 3.12 shows a scree plot of the singular values of the design matrix $\mathbf{B}$. The condition number of the matrix $\mathbf{B}$ is 225333.0, with the condition number of $\mathbf{B}^\top\mathbf{B}$ being $5,857,413,839$, i.e. $\mathbf{B}^\top\mathbf{B}$ is close to being numerically singular. This suggests that finding the least-squares estimate of the coefficients is close to being numerically unstable. ◁

As we have seen in the above example the truncated power basis can easily lead to numerical instability. Thus we will turn to an alternative basis, the so-called B-spline basis.

B-splines. B-splines form a numerically more stable basis. They also make the definition of meaningful penalty matrices easier, which we will exploit in section 3.3.

The key idea of B-splines is to use basis functions which are local, i.e. only non-zero for a "small" proportion of the range of the covariate and which are bounded above. We can think of B-splines as a sequence of "bumps". Figure 3.13 shows a B-spline basis function for degrees $r \in \{0, 1, 2, 3\}$. We will define B splines recursively.

Definition 3.7 (B-spline basis). *(a) Given a set of l knots the B-spline basis of degree 0 is given by the functions $(B_1^0(x), \ldots, B_{l-1}^0)$ with*

$$B_j^0(x) = \begin{cases} 1 & \text{for } \kappa_j \leq x < \kappa_{j+1} \\ 0 & \text{otherwise.} \end{cases}$$

(b) Given a set of l knots the B-spline basis of degree $r > 0$ is given by the functions $(B_1^r(x), \ldots, B_{l+r-1}^r)$ with

$$B_j^r(x) = \frac{x - \kappa_{j-r}}{\kappa_j - \kappa_{j-r}} B_{j-1}^{r-1}(x) + \frac{\kappa_{j+1} - x}{\kappa_{j+1} - \kappa_{j+1-r}} B_j^{r-1}(x).$$

In order to be able to construct the splines recursively we have to introduce additional outside knots to the left of κ_1 and to the right of κ_l. In order to be able to construct a basis of degree r we need r additional outside knots on each side. Figure 3.15 illustrates this idea. These outside knots are just used to construct the basis.

From their recursive definition one can derive that B-splines have the following properties. These can also be seen in figure 3.15.

- A B-spline basis function of degree r is made up of $r + 1$ polynomials of degree r. Outside these $r + 1$ intervals, the basis function is zero. This makes the basis functions local.
- At every $x \in (a, b)$ only $r + 1$ basis functions are non-zero.
- The basis functions sum to 1 for all $x \in [a, b]$. This implies that we do not need to include an intercept in the design matrix.
- One can show (homework exercise) that the derivative of a B-spline of degree r is a B-spline of degree $r - 1$.

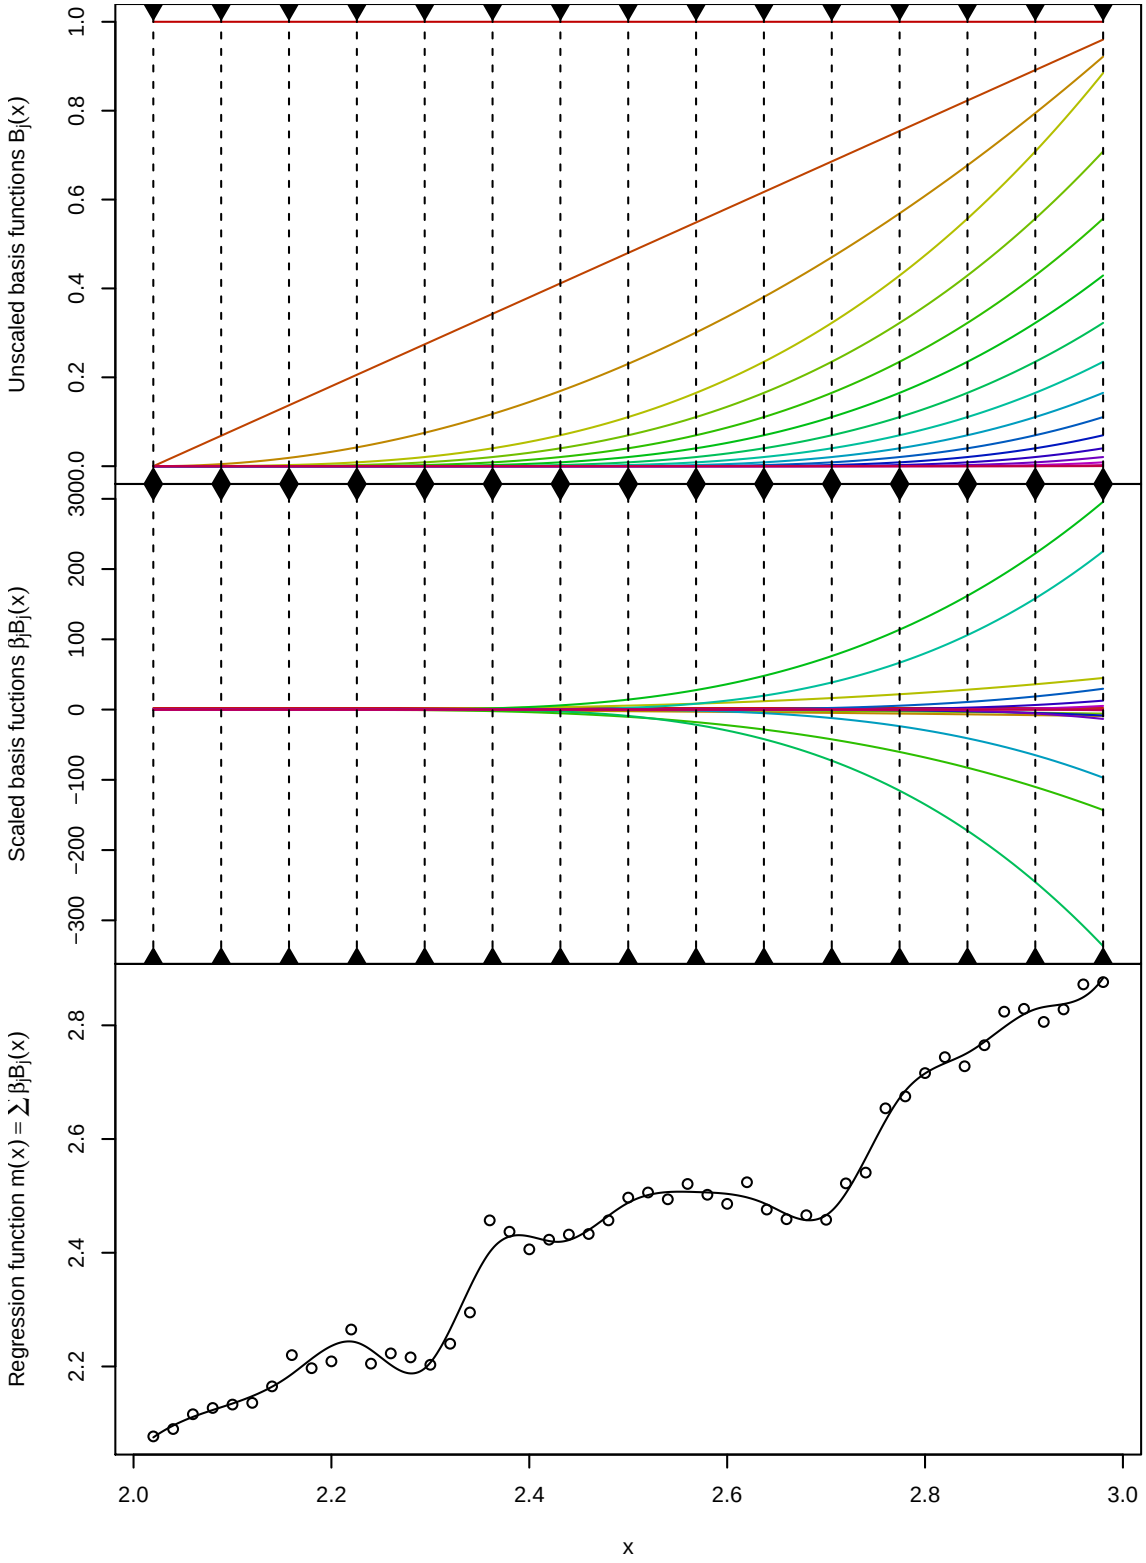

Figure 3.10. Illustration of flexible regression using the truncated power series basis of degree 3 applied to the radiocarbon data. The top panel shows the unscaled basis functions $B_j(x)$. The middle panel shows the scaled basis functions $\hat{\beta}_j B_j(x)$. The bottom panel shows a scatter plot of the data together with the fitted function $\hat{m}(x) = \sum_j \hat{\beta}_j B_j(x)$.

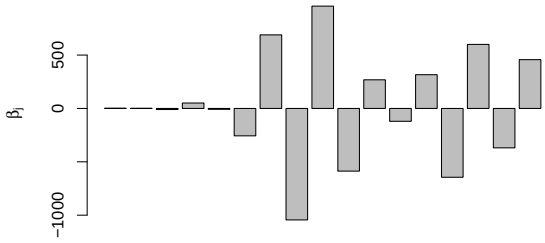

Figure 3.11. Bar plot of the coefficients $\hat{\boldsymbol{\beta}}$ estimated using the truncated power series regression model shown in figure 3.10.

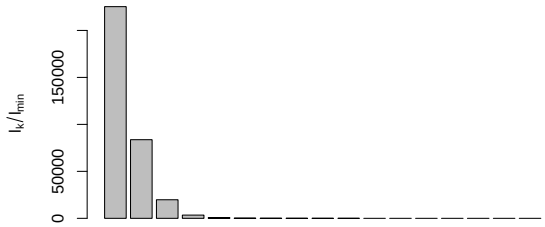

Figure 3.12. Scree plot of the singular values of the design matrix $\mathbf{B}$ (square root of the eigenvalues of the cross-product matrix $\mathbf{B}'\mathbf{B}$) for the truncated power series regression model shown in figure 3.10.

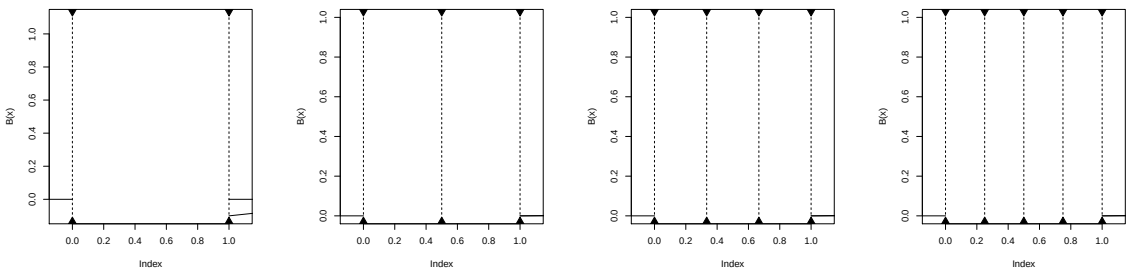

(a) One basis function of degree $r = 0$ (b) One basis function of degree $r = 1$ (c) One basis function of degree $r = 2$ (d) One basis function of degree $r = 3$

Figure 3.13. One basis function of a B-spline basis with degree $r \in \{0, 1, 2, 3\}$ using $r + 1$ knots.

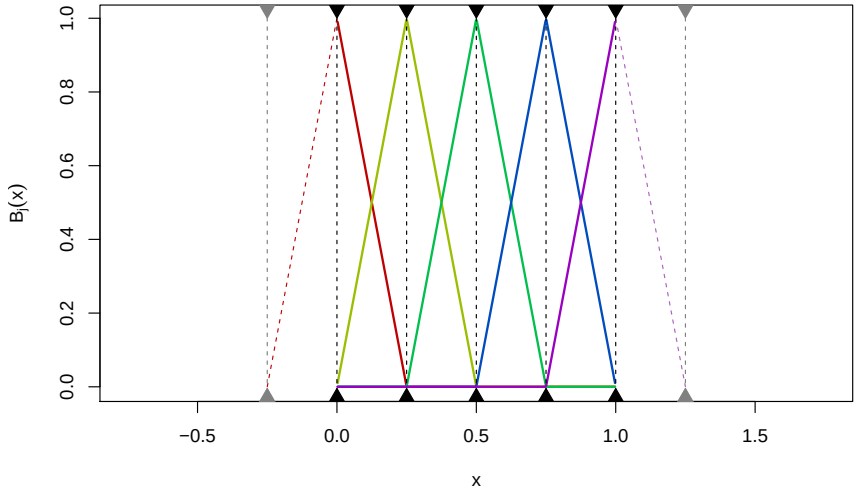

(a) Degree $r = 1$

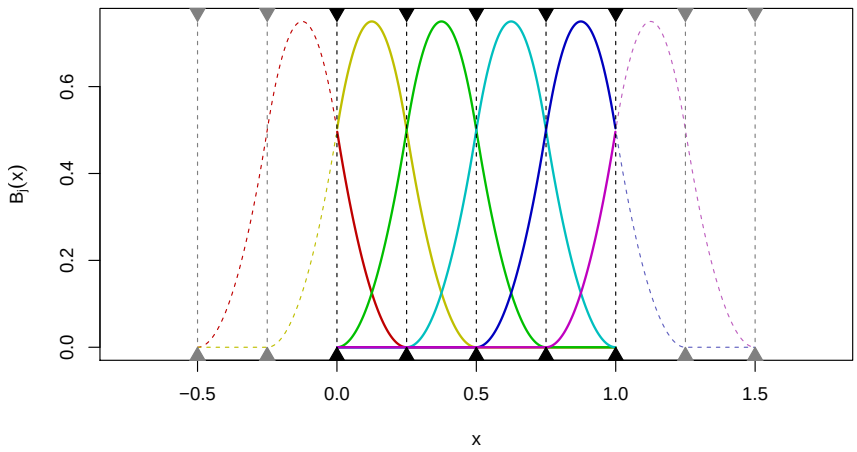

(b) Degree $r = 2$

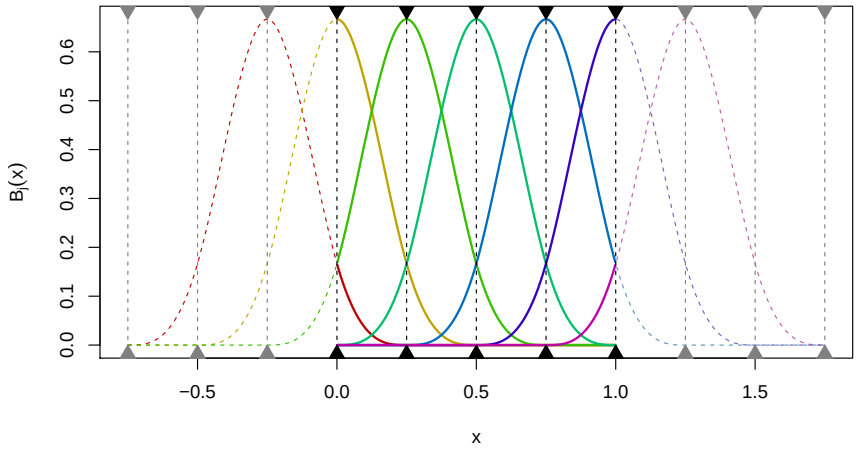

(c) Degree $r = 3$

Figure 3.14. B spline bases for degrees $r \in \{1, 2, 3\}$.

We can fit a B-spline model to data by using the design matrix

$$\mathbf{B} = \begin{pmatrix} B_1^r(x_1) & \dots & B_{l+r-1}^r(x_1) \\ \vdots & \ddots & \vdots \\ B_1^r(x_n) & \dots & B_{l+r-1}^r(x_n) \end{pmatrix}.$$

Example 3.7 (Radiocarbon data (continued)). Figure 3.15 illustrates the use of a B-spline basis for fitting a spline-based flexible regression model for the radiocarbon data.

The B-spline basis is numerically much better behaved. The coefficient values (cf. figure 3.16) are not too large and the columns of the design matrix $\mathbf{B}$ are much less correlated than the columns of the truncated power basis; the maximum correlation is 0.8309. The condition number of $\mathbf{B}$ is 25.664 (cf. figure 3.17) and the condition number of $\mathbf{B}^\top\mathbf{B}$ is 358.263. ◁

3.3 Penalised splines (P-splines)

A reminder of ridge regression

Ridge regression solves the penalised (or regularised) least-squares criterion

$$\|\mathbf{y} - \mathbf{B}\boldsymbol{\beta}\|^2 + \lambda\|\boldsymbol{\beta}\|^2,$$

where $\mathbf{B}$ is the matrix of covariates. The solution of this problem is given by

$$\hat{\boldsymbol{\beta}}_{\text{ridge}} = (\mathbf{B}^\top\mathbf{B} + \lambda\mathbf{I}_p)^{-1}\mathbf{B}^\top\mathbf{y}$$

To compute $\hat{\boldsymbol{\beta}}_{\text{ridge}}$ it is numerically more stable to use a QR decomposition to minimise the augmented system

$$\left\| \begin{pmatrix} \mathbf{y} \\ \mathbf{0} \end{pmatrix} - \begin{pmatrix} \mathbf{B} \\ \sqrt{\lambda}\mathbf{I} \end{pmatrix} \boldsymbol{\beta} \right\|^2$$

When using splines the positioning of the knots can have a large influence on the fitted function, especially if a comparatively small number of basis functions is used. One way of avoiding this problem is to use *penalised splines*. They are based on the idea of *not* using the number of basis functions to control the smoothness of the estimate, but to use a roughness penalty to this end. This is similar in spirit to the approach discussed in section 3.2.3, though in most cases it is not necessary to use one basis function per observation. Around 20 to 30 basis functions should be sufficient. Without including a penalty in the fitting criterion this would

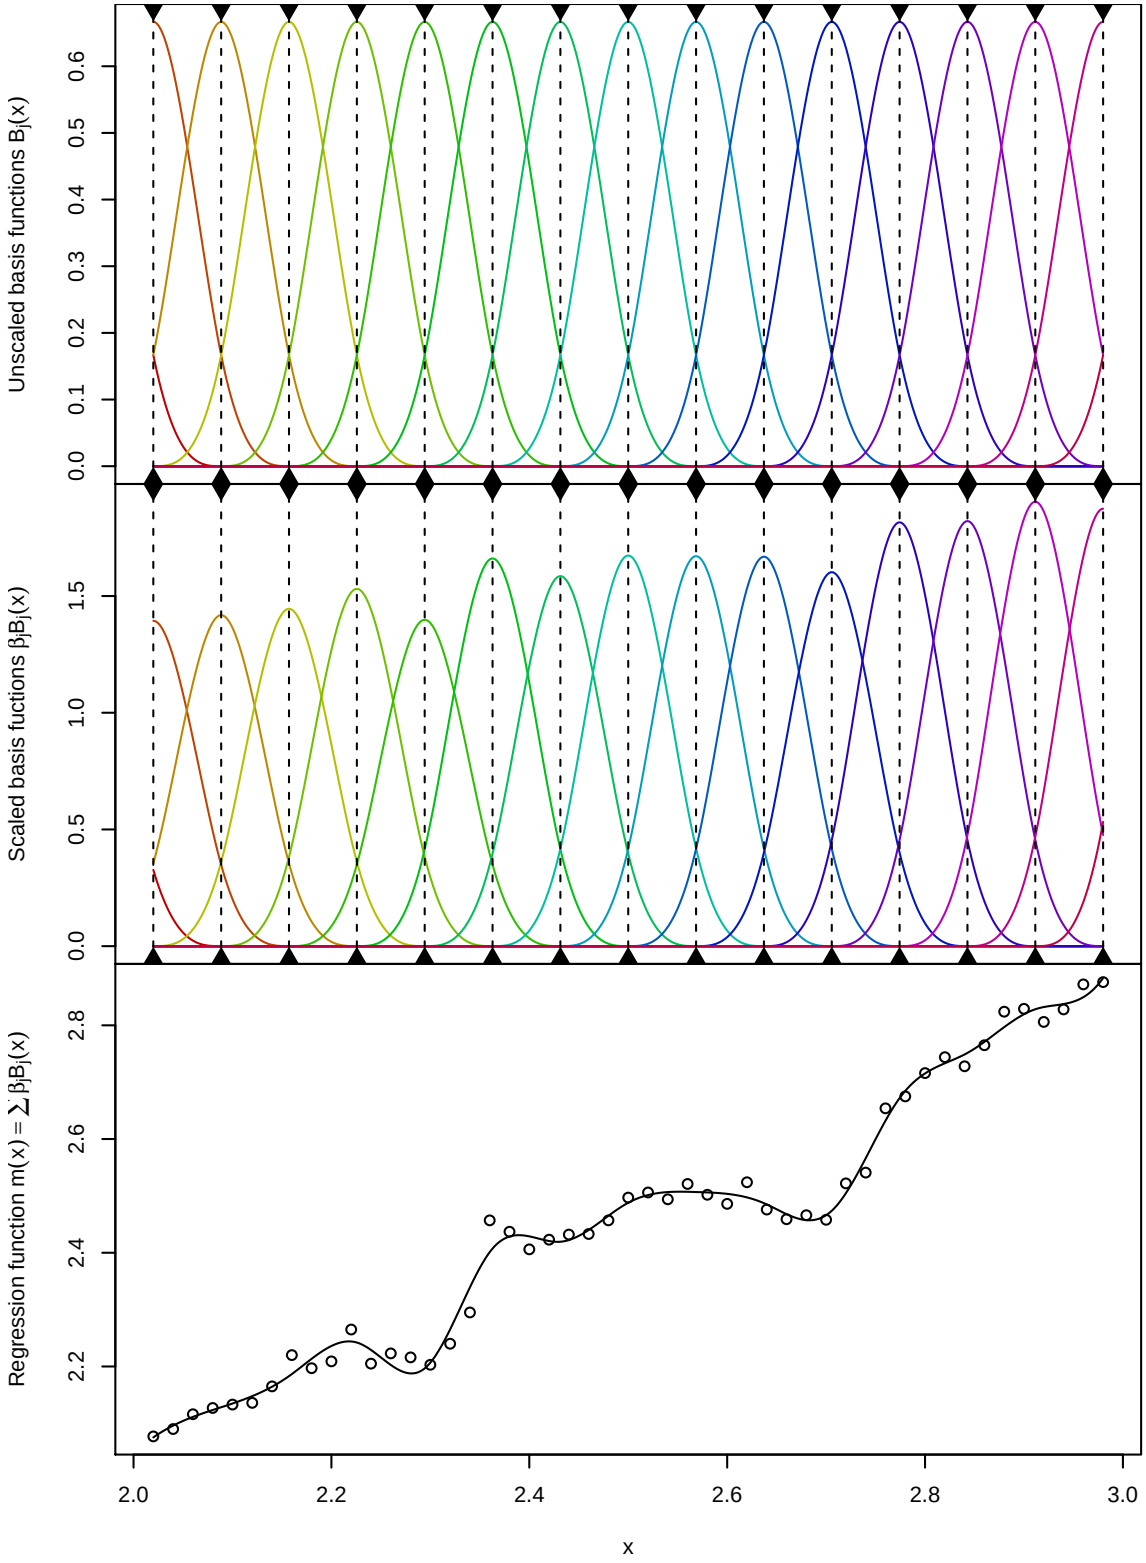

Figure 3.15. Illustration of flexible regression using the B-spline basis applied to the radiocarbon data. The top panel shows the unscaled basis functions $B_j(x)$. The middle panel shows the scaled basis functions $\hat{\beta}_j B_j(x)$. The bottom panel shows a scatter plot of the data together with the fitted function $\hat{f}(x) = \sum_j \hat{\beta}_j B_j(x)$.

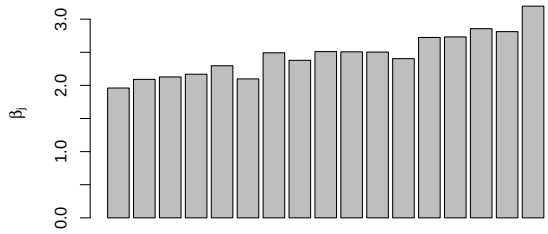

Figure 3.16. Bar plot of the coefficients $\hat{\boldsymbol{\beta}}$ estimated using the B-spline regression model shown in figure 3.15.

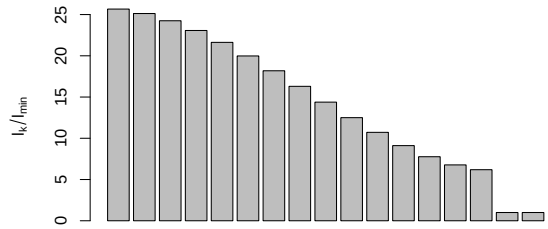

Figure 3.17. Scree plot of the singular values of the design matrix $\mathbf{B}$ (square root of the eigenvalues of the cross-product matrix $\mathbf{B}'\mathbf{B}$) for the B-spline regression model shown in figure 3.15. The condition number of $\mathbf{B}'\mathbf{B}$ is 395.661.

most likely lead to an overfit to the data. Thus we need to consider a penalised criterion which, just like in section 3.2.3, contains a roughness penalty. In this section we will use $\|\mathbf{D}\boldsymbol{\beta}\|^2$ as roughness penalty, i.e. we choose the regression coefficients $\boldsymbol{\beta}$ by minimising

$$\sum_{i=1}^{n}(y_i - m(x_i))^2 + \lambda\|\mathbf{D}\boldsymbol{\beta}\|^2. \tag{3.3}$$

This objective function is, with the exception of the inclusion of the matrix $\mathbf{D}$, the objective function of ridge regression. As before, λ controls the trade-off between following the data (small λ) and obtaining a strongly regularised curve (large λ). In analogy with ridge regression one can show that the optimal $\boldsymbol{\beta}$ is given by

$$\boldsymbol{\beta} = (\mathbf{B}^\top\mathbf{B} + \lambda\mathbf{D}^\top\mathbf{D})^{-1}\mathbf{B}^\top\mathbf{y},$$

where $\mathbf{B}$ is the design matrix corresponding to the B-spline basis used for $m(\cdot)$. Numerically, it is more advantageous to represent the penalty term $\lambda\|\mathbf{D}\boldsymbol{\beta}\|^2$ by including it into an expanded design matrix, i.e. to solve

$$\left\| \begin{pmatrix} \mathbf{y} \\ \mathbf{0} \end{pmatrix} - \begin{pmatrix} \mathbf{B} \\ \sqrt{\lambda}\mathbf{D} \end{pmatrix} \boldsymbol{\beta} \right\|^2$$

using a QR decomposition.

There are (at least) two possible approaches for choosing $\mathbf{D}$. We can choose $\mathbf{D}$ to be a difference matrix, or we can choose $\mathbf{D}$ such that $\|\mathbf{D}\boldsymbol{\beta}\|^2 = \int_a^b m''(x)^2\,dx$. The former is both conceptually and computationally simpler; the latter is closer to what the theory suggests as optimal.

3.3.1 Difference penalties

The simplest choice of $\mathbf{D}$ is to use a difference penalty. Using the identity matrix for $\mathbf{D}$, as we would in ridge regression, is usually not appropriate: it shrinks all coefficients to zero, i.e. it shrinks the regression function $m(\cdot)$ to zero as well, which rarely desirable (cf. figure 3.18(a)). As we can see from the middle panel of figure 3.15, we obtain a smooth function when neighbouring β_j's are similar.

This can be achieved by using one of the following choices. We assume that we are using equally-spaced knots.

First-order differences. We can set

$$\mathbf{D}_1 = \begin{pmatrix} 1 & -1 & \dots & 0 \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 1 & -1 \end{pmatrix}.$$

This calculates the roughness penalty as the sum of the squared first-order differences between the neighbouring β_j, i.e.

$$\|\mathbf{D}_1\boldsymbol{\beta}\|^2 = \sum_{j=1}^{l+r-2} (\beta_{j+1} - \beta_j)^2$$

This penalty shrinks the coefficients towards a common constant (cf. figure 3.18(b)) and thus shrinks the regression function $m(\cdot)$ towards a constant function. Adding a constant to $m(\cdot)$ does thus not change the penalty.

This penalty is the natural choice if B-splines of order 2 are used.

Second-order differences. We can set

$$\mathbf{D}_2 = \begin{pmatrix} 1 & -2 & 1 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 1 & -2 & 1 \end{pmatrix}.$$

This calculates the roughness penalty as the sum of the squared second-order differences between the neighbouring β_j, i.e.

$$\|\mathbf{D}_2\boldsymbol{\beta}\|^2 = \sum_{j=1}^{l+r-3} (\beta_{j+2} - 2\beta_{j+1} + \beta_j)^2$$

This penalty shrinks the coefficients towards a linear sequence (cf. figure 3.18(c)) and thus shrinks the regression function $m(\cdot)$ towards a linear function. Adding a linear function to $m(\cdot)$ does thus not change the penalty.

This penalty is the natural choice if B-splines of order 3 are used.

Higher-order differences. Higher-order difference matrices can be constructed using the recursive formula $\mathbf{D}_r = \mathbf{D}_1\mathbf{D}_{r-1}$ where $\mathbf{D}_r$ denotes the penalty matrix of order r.

Penalty interpretation: Only an all-zero coefficient vector incurs no penalty.
Bayesian interpretation: Independent zero-mean Gaussian prior.

Penalty interpretation: Only an all constant coefficient vector incurs no penalty.
Bayesian interpretation: Conditional distribution of β_2 given β_1 is Gaussian with mean β_1.
(First-order random walk)

Penalty interpretation: Only a coefficient vector which forms a linear sequence incurs no penalty.
Bayesian interpretation: Conditional distribution of β_3 given β_1 and β_2 is Gaussian with mean $2 \cdot \beta_2 - \beta_1$.
(Second-order random walk)

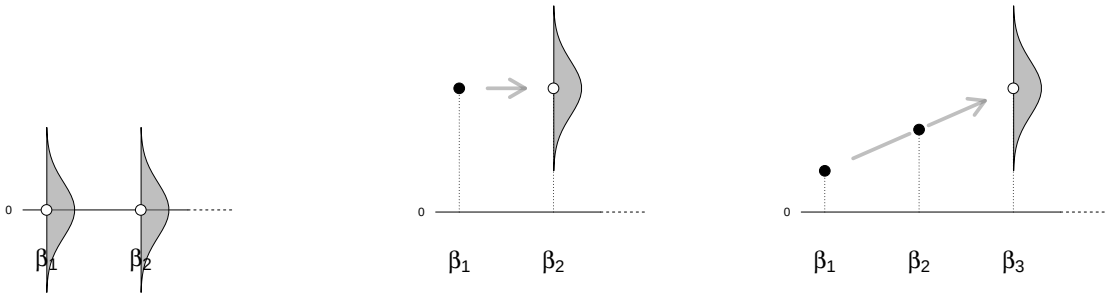

(a) Illustration of a 0-th order penalty (ridge regression).

(b) Illustration of a first-order penalty.

(c) Illustration of a second-order penalty.

Figure 3.18. Illustration of difference penalties of order 0 to 2.

Example 3.8 (Radiocarbon dating (continued)). Figure 3.19 shows the model fit obtained when fitting a P-spline model with different values of the smoothing parameter λ. The smaller λ the closer the fitted function $\hat{m}(\cdot)$ is to the data, which leads for very small values of λ to an overfit to the data. ◁

3.3.2 Other penalties

Difference penalties are not the only choice of penalty matrix. An alternative choice consists of choosing $\mathbf{D}$ such that $\|\mathbf{D}\boldsymbol{\beta}\|^2 = \int_a^b m''(x)^2 \, dx$, which is the roughness penalty we have used in section 3.2.3.

Using that $m''(x) = \sum_{j=1}^{l+r-1} \beta_j B_j''(x)$ we have that

$$\begin{aligned}\int_a^b m''(x)^2 \, dx &= \sum_{j=1}^{l+r-1} \sum_{k=1}^{l+r-1} \beta_j \beta_k \int_a^b B_j''(x) B_k''(x) \, dx \\ &= \boldsymbol{\beta}^\top \begin{pmatrix} \int_a^b B_1''(x) B_1''(x) \, dx & \dots & \int_a^b B_1''(x) B_{l+r-1}''(x) \, dx \\ \vdots & \ddots & \vdots \\ \int_a^b B_1''(x) B_{l+r-1}''(x) \, dx & \dots & \int_a^b B_{l+r-1}''(x) B_{l+r-1}''(x) \, dx \end{pmatrix} \boldsymbol{\beta}\end{aligned}$$

Thus we just need to choose $\mathbf{D}$ such that

$$\mathbf{D}^\top \mathbf{D} = \begin{pmatrix} \int_a^b B_1''(x) B_1''(x) \, dx & \dots & \int_a^b B_1''(x) B_{l+r-1}''(x) \, dx \\ \vdots & \ddots & \vdots \\ \int_a^b B_1''(x) B_{l+r-1}''(x) \, dx & \dots & \int_a^b B_{l+r-1}''(x) B_{l+r-1}''(x) \, dx \end{pmatrix}.$$

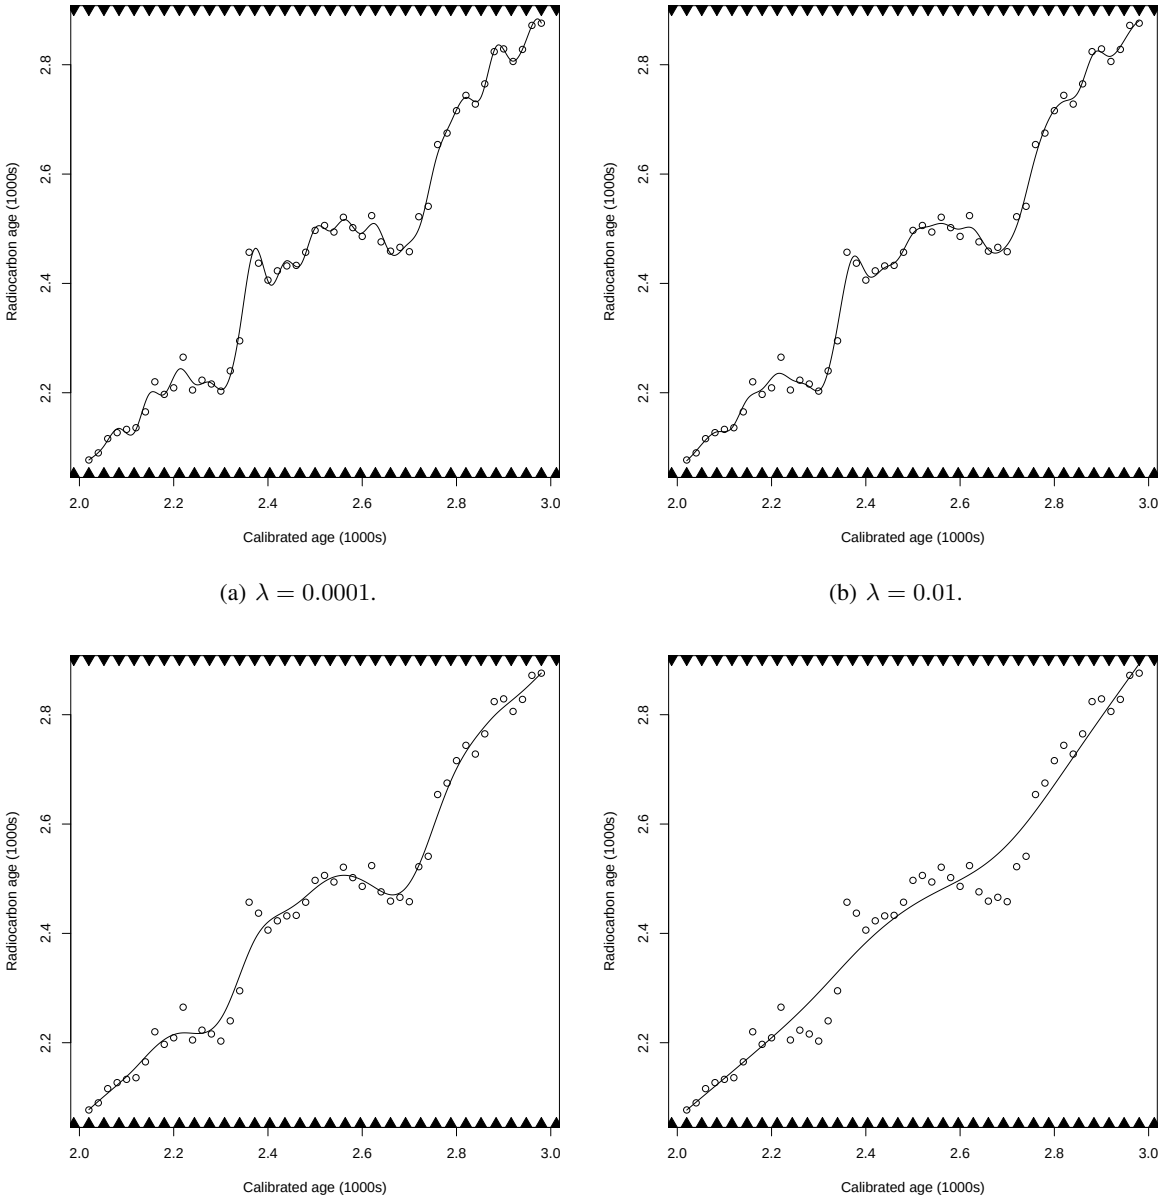

(a) $\lambda = 0.0001$. (b) $\lambda = 0.01$.

(c) $\lambda = 1$. (d) $\lambda = 10$.

Figure 3.19. P-spline with different values of the smoothing parameter λ fitted to the radiocarbon data.

3.3.3 Effective degrees of freedom

Finally we introduce the notion of effective degrees of freedom, also sometimes called the effective number of parameters. In an un-penalised regression problem, the number of parameters provides us with information about the complexity of the model. More complex models have more parameters than simpler models. For penalised regression problems counting the parameters is however not meaningful. Due to the roughness penalty not all parameters are "free". Recall that in linear regression the hat matrix $\mathbf{S} = \mathbf{B}(\mathbf{B}^\top\mathbf{B})^{-1}\mathbf{B}^\top$ is a projection matrix and thus the trace $\text{tr}(\mathbf{S})$ equals the number of parameters. We can generalise this to penalised models and define the effective degrees of freedom as

$$\text{edf}(\lambda) = \text{tr}(\mathbf{S}_\lambda),$$

where $\mathbf{S}_\lambda = \mathbf{B}(\mathbf{B}^\top\mathbf{B} + \lambda\mathbf{D}^\top\mathbf{D})^{-1}\mathbf{B}^\top$.

3.3.4 Random effects interpretation

Random effect models – Likelihood

In the random effects model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\alpha} + \mathbf{Z}\boldsymbol{\gamma} + \boldsymbol{\varepsilon}$$

with error term $\boldsymbol{\varepsilon} \sim \mathsf{N}(\mathbf{0}, \sigma^2\mathbf{I})$ and random effect $\boldsymbol{\gamma} \sim \mathsf{N}(\mathbf{0}, \tau^2\mathbf{I})$ twice the loglikelihood is (ignoring the variance parameters) given by

$$-\frac{1}{\sigma^2}\sum_{i=1}^{n}(y_i - \mathbf{x}_i^\top\boldsymbol{\alpha} - \mathbf{z}_i^\top\boldsymbol{\gamma})^2 - \frac{1}{\tau^2}\sum_{j=1}^{q}\gamma_j^2$$

Comparing the penalised least squares criterion (3.3) to the loglikelihood suggests that we can interpret the penalised regression model as a random effects model with no fixed effect and random effect $\boldsymbol{\beta}$. However the problem is that, at least for difference matrices, $\mathbf{D}^\top\mathbf{D}$ is not of full rank, thus we cannot take its inverse matrix square root. In order to obtain a proper random-effects representation we need to "split" $\boldsymbol{\beta}$ into an (unpenalised) fixed effect and a (penalised) random effect.

In the following we will only consider the case of a difference penalty of order 1 or 2. In the case of a first-order difference penalty we define $\mathbf{G} = (1, \ldots, 1)$. For a second-order difference penalty we define $\mathbf{G} = \begin{pmatrix} 1 & 1 & \ldots & 1 \\ 1 & 2 & \ldots & l+r-1 \end{pmatrix}$. The rows in $\mathbf{G}$ are parameter sequences

which do not incur a penalty, i.e. $\mathbf{GD} = \mathbf{0}$. We also define $\mathbf{H} = \mathbf{D}^\top(\mathbf{DD}^\top)^{-1}$. We can now write

$$\boldsymbol{\beta} = \mathbf{G}\boldsymbol{\alpha} + \mathbf{H}\boldsymbol{\gamma}$$

Because $\mathbf{DD}^\top$ is of full rank we have that $\mathbf{DH} = \mathbf{DD}^\top(\mathbf{DD}^\top)^{-1} = \mathbf{I}$. Plugging this into the objective function (3.3) gives

$$\begin{aligned}
&\|\mathbf{y} - \mathbf{BG}\boldsymbol{\alpha} - \mathbf{BH}\boldsymbol{\gamma}\|^2 + \lambda\Big(\underbrace{\boldsymbol{\alpha}\mathbf{G}^\top\mathbf{D}^\top\mathbf{DG}\boldsymbol{\alpha}}_{=0} + 2\underbrace{\boldsymbol{\alpha}\mathbf{G}^\top\mathbf{D}^\top\mathbf{DH}\boldsymbol{\gamma}}_{=0} + \boldsymbol{\gamma}^\top\underbrace{\mathbf{H}^\top\mathbf{D}^\top\mathbf{DH}}_{=\mathbf{I}}\boldsymbol{\gamma}\Big) \\
&= \|\mathbf{y} - \mathbf{BG}\boldsymbol{\alpha} - \mathbf{BH}\boldsymbol{\gamma}\|^2 + \lambda\|\boldsymbol{\gamma}\|^2
\end{aligned}$$

Defining $\mathbf{X} = \mathbf{BG}$ and $\mathbf{Z} = \mathbf{BH}$ and denoting rows of $\mathbf{X}$ and $\mathbf{Z}$ by $\mathbf{x}_i$ and $\mathbf{z}_i$ respectively, this is equivalent to

$$\sum_{i=1}^{n}(y_i - \mathbf{x}_i^\top\boldsymbol{\alpha} - \mathbf{z}_i^\top\boldsymbol{\gamma})^2 + \lambda\sum_{j=1}^{q}\gamma_j^2,$$

which is $-\sigma^2$ times the loglikelihood of a random-effects model, which we have stated above. Hereby we have used $\lambda = \sigma^2/\tau^2$.

Thus the penalised regression model is nothing other than a random effects effect and we can use standard mixed model software to fit these models. Most importantly we can estimate the variances σ^2 and τ^2 is a mixed model (using (restricted) maximum likelihood, which gives us a way of estimating the otherwise rather elusive smoothing parameter $\hat{\lambda} = \hat{\sigma}^2/\hat{\tau}^2$.

3.3.5 Bayesian interpretation

Rather than interpreting the penalised fitting criterion as a random effects model we can treat the penalised regression model as a fully Bayesian model with the following prior and data model.

$$\begin{aligned}
\mathbf{D}\boldsymbol{\beta}|\tau^2 &\sim \mathsf{N}(\mathbf{0}, \tau^2\mathbf{I}) \\
\mathbf{y}|\boldsymbol{\beta}, \sigma^2 &\sim \mathsf{N}(\mathbf{B}\boldsymbol{\beta}, \sigma^2\mathbf{I})
\end{aligned}$$

The prior distribution of $\boldsymbol{\beta}$ is improper if $\mathbf{D}$ is not of full rank, which is the case for all difference penalties. However in the case of difference penalties the prior distribution of $\boldsymbol{\beta}$ can be expressed in terms of random walks (cf. figure 3.18).

First-order random walk The first-order penalty corresponds to an improper flat prior on β_1 and $\beta_j|\beta_{j-1} \sim \mathsf{N}(\beta_{j-1}|\tau^2)$ (for $j \geq 2$).

Second-order random walk The second-order penalty corresponds to an improper flat prior on β_1 and β_2 and $\beta_j|\beta_{j-1}, \beta_{j-2} \sim \mathsf{N}(2\beta_{j-1} - \beta_{j-2}|\tau^2)$ (for $j \geq 3$).

It seems natural to complement the model with priors for σ^2 and τ^2

$$\sigma^2 \sim \mathsf{IG}(a_{\sigma^2}, b_{\sigma^2})$$

$$\tau^2 \sim \mathsf{IG}(a_{\tau^2}, b_{\tau^2})$$

Inference can then be carried out efficiently using a Gibbs sampler. This model and many other Bayesian smoothing models are implemented in the software `BayesX`.

Rather than placing a independent inverse-gamma prior on τ^2 we can set $\tau^2 = \sigma^2/\lambda$ and place a prior of our choice on λ. In this model the posterior distribution distribution of λ does not follow a known distribution, but can be evaluated efficiently, as all the other parameters can be integrated out in closed form. The drawback is that the integration over λ would need to be carried out numerically, which suggests that this approach is better suited for an empirical Bayes strategy for estimating λ.

4

Flexible regression in more than one dimension

In this chapter, methods of extending flexible regression to more than one covariate will be explored. For one covariate, spline methods have been discussed in some detail in earlier chapters while a local fitting approach was outlined in the *preliminary material*. We will revisit the local linear approach and see how that can be extended, before returning to splines. A more general approach known as *additive modeling* will then be described. First of all there is a reminder of the example which was used earlier.

Example 4.1 (Great Barrier Reef data). A survey of the fauna on the sea bed lying between the coast of northern Queensland and the Great Barrier Reef was carried out. The sampling region covered a zone which was closed to commercial fishing, as well as neighbouring zones where fishing was permitted. The variables are:

Zone	an indicator for the closed (1) and open (0) zones
Year	an indicator of 1992 (0) or 1993 (1)
Latitude	latitude of the sampling position
Longitude	longitude of the sampling position
Depth	bottom depth
Score1	catch score 1
Score2	catch score 2

The details of the survey and an analysis of the data are provided by Poiner et al. (1997), *The effects of prawn trawling in the far northern section of the Great Barrier Reef*, CSIRO Division of Marine Research, Queensland Dept. of Primary Industries. ◁

4.1 The local fitting approach with a single covariate

When we have regression data on a response variable y and a single covariate x, simple nonparametric model has the form

$$y_i = m(x_i) + \varepsilon_i,$$

where the data (x_i, y_i) are described by a smooth curve m plus independent errors ε_i. One approach to estimating m is to use a model we know and fit it locally. For example, we can construct a *local linear regression*. This involves solving the least squares problem

$$\min_{\alpha,\beta} \sum_{i=1}^{n} \{y_i - \alpha - \beta(x_i - x)\}^2 \, w(x_i - x \, ; h)$$

and taking as the estimate at x the value of $\hat{\alpha}$, as this defines the position of the local regression line at the point x. An even simpler approach is to fit a local mean. Specifically, at any point of interest x, we choose our estimator of the curve there as the value of μ which minimises

$$\sum_{i=1}^{n} \{y_i - \mu\}^2 w(x_i - x; h)$$

and this is easily shown to produce the 'running mean'

$$\hat{m}(x) = \frac{\sum_{i=1}^{n} w(x_i - x; h) \, y_i}{\sum_{i=1}^{n} w(x_i - x; h)}.$$

If we do the algebra to minimise the sum-of-squares in the local linear approach, then an explicit formula for the local estimator can be derived as

$$\hat{m}(x) = \frac{1}{n} \sum_{i=1}^{n} \frac{\{s_2(x; h) - s_1(x; h)(x_i - x)\} w(x_i - x; h) y_i}{s_2(x; h) s_0(x; h) - s_1(x; h)^2},$$

where $s_r(x; h) = \{\sum (x_i - x)^r w(x_i - x; h)\}/n$.

In both the local mean and the local linear cases, the estimator is seen to be of the form $\sum_i \kappa_i y_i$, where the weights κ_i sum to 1. There is a broad sense then in which even the local linear method is 'locally averaging' the data. In fact, many other forms of nonparametric regression can also be formulated in a similar way.

The *preliminary material* explored the properties of these estimators and showed that the local linear approach has some advantages. Some useful descriptions of the behaviour of the local linear approach is that

$$\mathbb{E}\{\hat{m}(x)\} \approx m(x) + \frac{h^2}{2} m''(x),$$
$$\mathsf{var}\{\hat{m}(x)\} \approx \frac{1}{nh} \left\{ \int w(u)^2 du \right\} \sigma^2 \frac{1}{f(x)},$$

where σ^2 denotes the variance of the error terms ε_i.

It is helpful to express the fitted values of the nonparametric regression as

$$\hat{m} = Sy,$$

where $\hat{m}$ denotes the vector of fitted values, S denotes a *smoothing matrix* whose rows consist of the weights appropriate to estimation at each evaluation point, and y denotes the observed responses in vector form. This linear structure applies with both local fitting and spline approaches and it is very helpful.

For example, it gives us a route to defining *degrees of freedom* by analogy with what happens with the usual linear model, where the number of parameters is the trace of the projection matrix. An approximate version of these can be constructed for nonparametric models as

$$\text{df} = \text{tr}\left\{S\right\}.$$

Similarly, we can construct an estimate of the error variance σ^2 through the residual sum-of-squares, which in a nonparametric setting is simply $\text{RSS} = \sum\{y_i - \hat{m}(x_i)\}^2$. This leads to the estimator of the error variance $\hat{\sigma}^2 = \text{RSS}/\text{df}$. The linear structure of the fitted values also makes it very easy to produce standard errors which quantify the variability of the estimate at any value of x. Unfortunately, we can't easily produce confidence intervals for the curve because of the bias mentioned above. However, by adding and subtracting two standard errors at each point on the curve we can produce *variability bands* which express the variation in the curve estimate. In fact, we don't need to rely on the asymptotic formula for variance. If $\hat{m}$ denotes the estimated values of m at a set of evaluation points then

$$\mathsf{var}\{\hat{m}\} = \mathsf{var}\{Sy\} = SS^T\sigma^2$$

and so, by plugging in $\hat{\sigma}^2$ and taking the square root of the diagonal elements, the standard errors at each evaluation point are easily constructed. The plot below illustrates this on the Reef data.

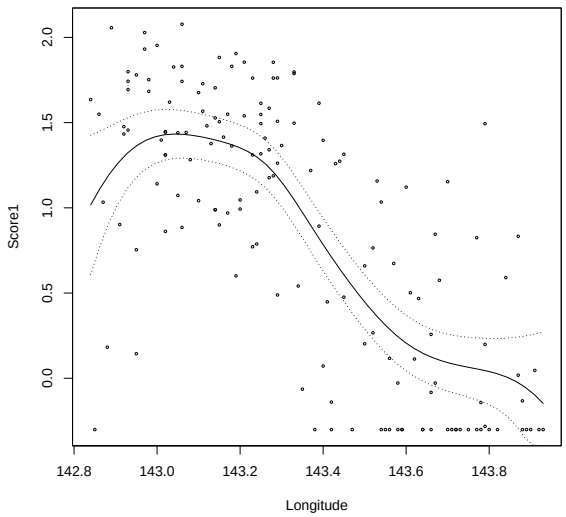

4.2 Local fitting with more than one covariate

It is rare to have problems which involve only a single covariate. For the Reef data a natural extension is to look at the relationship between the catch score and both latitude (x_1) and longitude (x_2), in a model

$$y_i = m(x_{1i}, x_{2i}) + \varepsilon_i.$$

The local linear approach is particularly easy to extend to this setting. If the observed data are denoted by $\{x_{1i}, x_{2i}, y_i; i = 1, \ldots, n\}$, then for estimation at the point (x_1, x_2) the weighted least squares formulation is

$$\min_{\alpha,\beta,\gamma} \sum_{i=1}^{n} \{y_i - \alpha - \beta(x_{1i} - x_1) - \gamma(x_{2i} - x_2)\}^2 \, w(x_{1i} - x_1; h_1) \, w(x_{2i} - x_2; h_2).$$

The value of the fitted surface at (x_1, x_2) is simply $\hat{\alpha}$. With careful thought, the computation can be performed efficiently.

This is illustrated below with one year of Reef data from the closed zone.

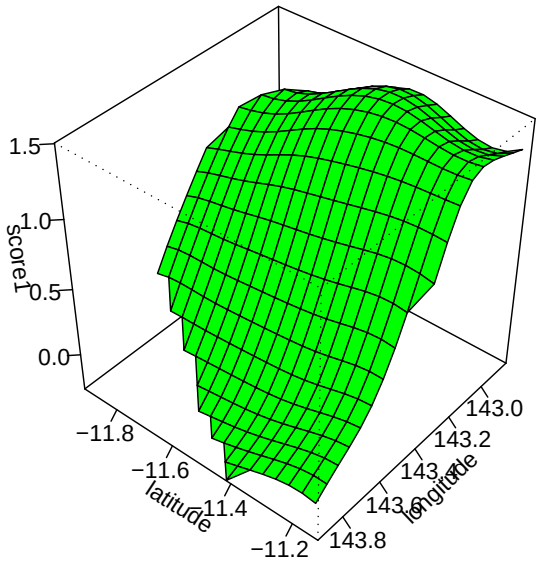

The effect of longitude dominates, as we see from the earlier nonparametric regression. However, a small effect of latitude is also suggested. The methods by which surfaces of this type can be constructed will be discussed in the next two chapters.

Methods for choosing the smoothing parameters (now two of them) such as generalised cross-validation and AIC extend to the two-dimensional case without difficulty in principle.

4.3 Splines with more than one covariate

As we have seen earlier, a b-spline basis provides a set of building blocks for a flexible regression function. The basic idea is to use the combination of a set of overlapping functions and a set of weights for each component to control the shape of the surface created. If we have two covariates then these basis functions need to be functions of both covariates.

$$y = \sum_{i=1}^{p}\sum_{j=1}^{p} \beta_{ij} b_{ij}(x_1, x_2) + \varepsilon.$$

A very simple way of achieving this is to create the basis functions from pairwise products of functions of a single covariate. Specifically,

$$b_{ij}(x_1, x_2) = b_i(x_1) b_j(x_2).$$

The graphics below try to illustrate what is going on in terms of the basis functions.

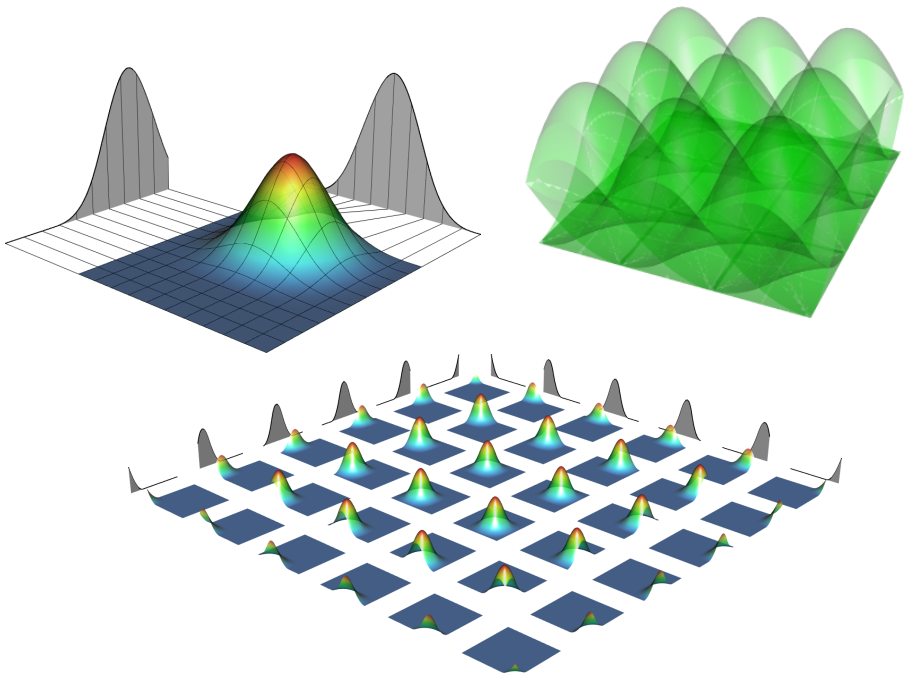

In terms of a matrix formulation, a b-spline representation for a flexible function of x_1 is $y = B_1\alpha_1 + \varepsilon$, where y now represents the vector of observed responses, x_1 represents a vector of observed covariate values, B_1 is a matrix with columns which evaluate each basis function at all the values of x_1, and ε represents a vector of errors. Similarly $y = B_2\alpha_2 + \varepsilon$ is a model for a flexible regression based on x_2. A simple way of combining both covariates is as $y = B_{12}\alpha_{12} + \varepsilon$, where the columns of the matrix B_{12} are constructed from all the pairwise products of the columns of B_1 and B_2 and α_{12} is the corresponding set of parameters. In fact, the parameters in α_{12} could be annotated as α_{ij}, where i and j correspond to the rows and columns of the layout of basis functions. That notation makes it easier to consider suitable penalty functions, if we wish to employ them, for example

$$\lambda_1 \sum_{i=2}^{p} \sum_{j=1}^{p} (\alpha_{ij} - \alpha_{i-1,j})^2 + \lambda_2 \sum_{i=1}^{p} \sum_{j=2}^{p} (\alpha_{ij} - \alpha_{i,j-1})^2 +$$

4.4 A simple additive model

It would be unrealistic to generalise this much further, by modelling additional covariates through functions of ever-increasing dimension. However, now that we have tools available to estimate smooth curves and surfaces, linear regression models can be extended to *additive models* as

$$y_i = \beta_0 + m_1(x_{1i}) + \ldots + m_p(x_{pi}) + \varepsilon_i, \qquad i = 1, \ldots, n,$$

where the m_i are functions whose shapes are unrestricted, apart from an assumption of smoothness. This gives a very flexible set of modelling tools. To see how these models can be fitted, consider the case of only two covariates,

$$y_i = \beta_0 + m_1(x_{1i}) + m_2(x_{2i}) + \varepsilon_i, \qquad i = 1, \ldots, n,$$

A rearrangement of this as $y_i - \beta_0 - m_2(x_{2i}) = m_1(x_{1i}) + \varepsilon_i$ suggests that an estimate of component m_1 can then be obtained by smoothing the residuals of the data after fitting $\hat{m}_2$,

$$\hat{m}_1 = S_1(y - \bar{y} - \hat{m}_2)$$

and that, similarly, subsequent estimates of m_2 can be obtained as

$$\hat{m}_2 = S_2(y - \bar{y} - \hat{m}_1).$$

Repetition of these steps gives a simple form of the *backfitting* algorithm. The same idea applies when we have more than two components on the model. At each step we smooth over a particular variable using as response the y variable with the current estimates of the other components subtracted.

A simple example of an additive model for the Reef data is shown below.

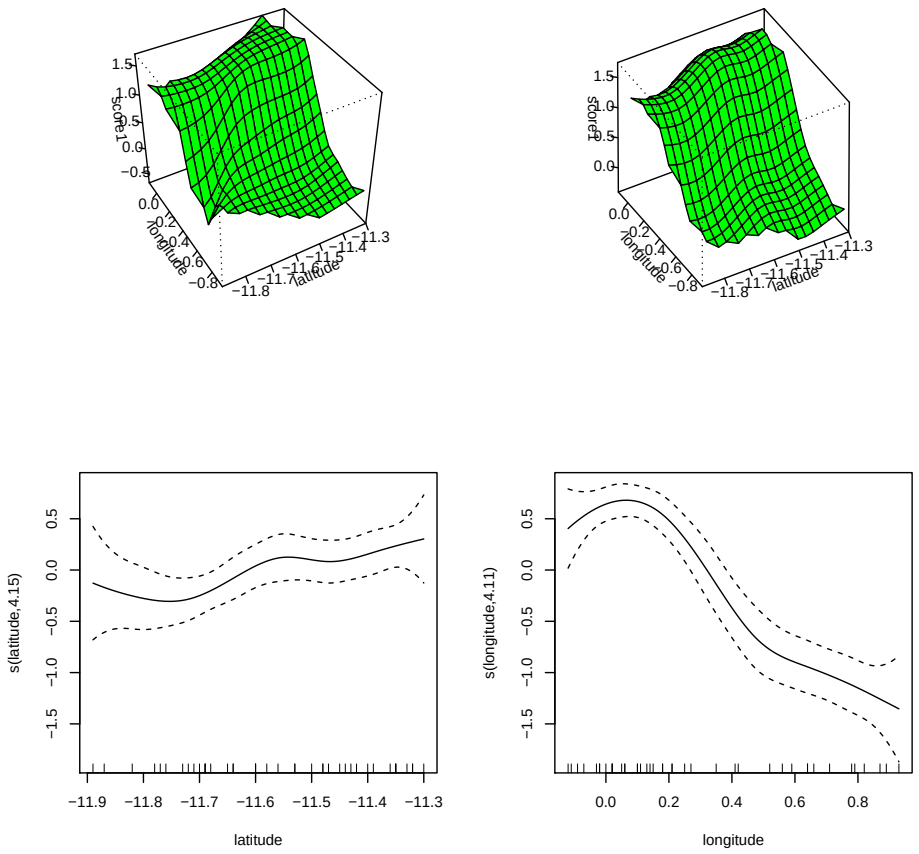

4.5 More general additive models

For the more general model

$$y_i = \alpha + m_1(x_{1i}) + \ldots + m_p(x_{pi}) + \varepsilon_i.$$

a simple extension of the steps outlined for two covariates gives a form of the *backfitting* algorithm. In order to ensure identifiability, we assume that $\sum_i m_j(x_{ji}) = 0$, for each j. At each step we smooth over a particular variable using as response the y variable with the current estimates of the other components subtracted.

The backfitting algorithm can be expressed as:

$$\hat{m}_j^{(r+1)} = S_j \left(y - \hat{\alpha}\mathbf{1} - \sum_{k<j} \hat{m}_k^{(r+1)} - \sum_{k>j} \hat{m}_k^{(r)} \right).$$

We can also express these in terms of projection matrices.

$$P_j^{(l)} = (I_n - P_0) S_j (I_n - \sum_{k<j} P_k^{(l)} - \sum_{k>j} P_k^{(l-1)}),$$

$$\hat{y} = Py = (P_0 + \sum_{j=1}^{p} P_j)y$$

If a regression splines or p-splines model is adopted, the each of the functions $m_i(x)$ is represented by a linear expression and so the model itself remains linear. It can then be fitted be standard linear regression, incorporating a set of penalties in the p-splines case. This has the advantage of direct, rather than iterative, fitting but it has the potential disadvantage of needing to invert very large matrices if the model has many terms.

The plots below show data from a survey of dissolved oxygen (DO) in the River Clyde at a single sampling station, related to potential explanatory variables of interest. The additive terms usefully capture the underlying trends.

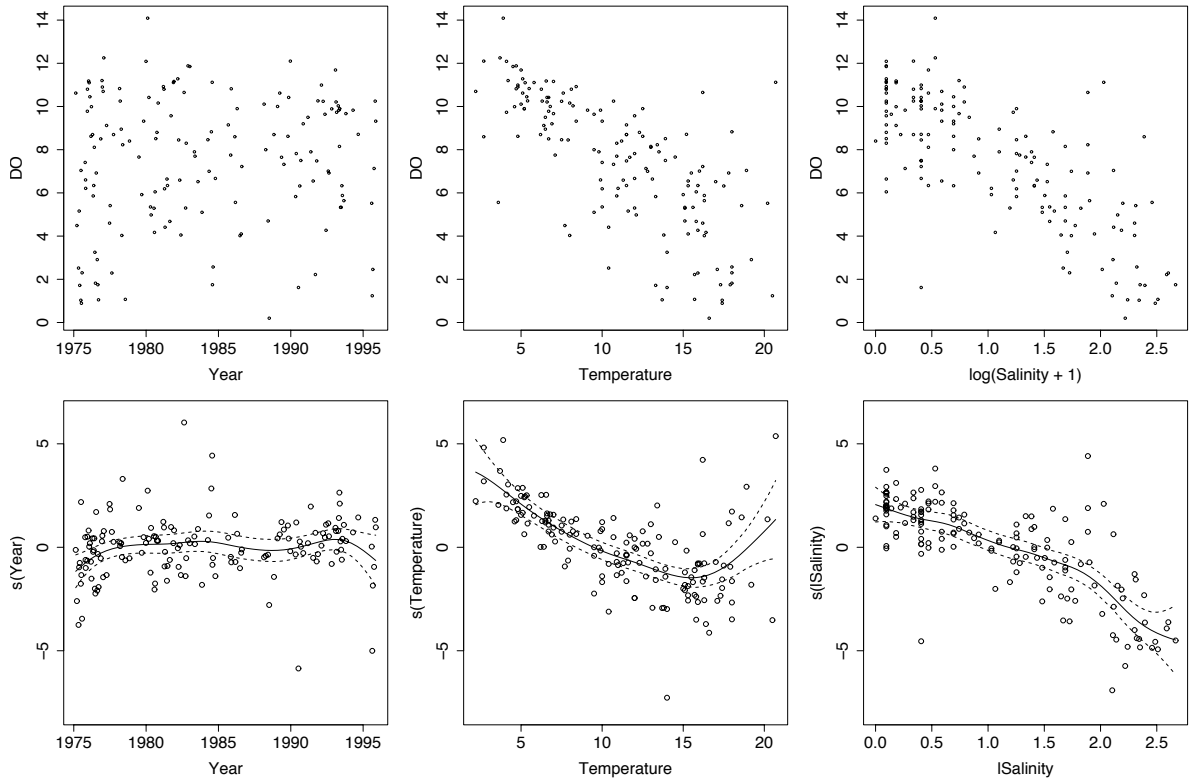

4.6 Comparing additive models

While models of this type provide very flexible and visually informative descriptions of the data, it is also necessary to consider how models can be compared and inferences drawn. Hastie & Tibshirani (1990) recommend the use of residual sums-of-squares and their associated approximate degrees of freedom to provide guidance for model comparisons.

For an additive model, the residual sum-of-squares can easily be defined as

$$\text{RSS} = \sum_{i=1}^{n} (y_i - \hat{y}_i)^2,$$

where $\hat{y}_i$ denotes the fitted value, produced by evaluating the additive model at the observation x_i. We can write the residual sum-of-squares as

$$\text{RSS} = \sum_{i=1}^{n}(y_i - \hat{y}_i)^2 = y^\top (I - P)^\top (I - P)y,$$

where P denotes the projection matrix discussed earlier. The approximate degrees of freedom for error can be defined as

$$\text{df} = \text{tr}\left\{(I - P)^\top (I - P)\right\}.$$

In an obvious notation, comparisons of two models can expressed quantitatively in

$$F = \frac{(\text{RSS}_2 - \text{RSS}_1)/(\text{df}_2 - \text{df}_1)}{\text{RSS}_1/\text{df}_1},$$

by analogy with the F-statistic used to compare linear models. Unfortunately, this analogy does not extend to distributional calculations and no general expression for the distribution of this test statistic is available. However, Hastie and Tibshirani (1990, sections 3.9 and 6.8) suggest that at least some approximate guidance can be given by referring the observed nonparametric F-statistic to an F distribution with $(\text{df}_2 - \text{df}_1)$ and df_1 degrees of freedom. Wood (2006) gives a formulation in terms of testing whether relevant spline coefficients might be 0.

The reef data provide a simple illustration of how model comparisons may be made. There are three obvious models of interest.

Model	RSS	df
1: $\beta_0 + m_1(\text{Latitude}) + m_2(\text{Longitude})$	4.541	33.99
2: $\beta_0 + m_2(\text{Longitude})$	6.128	37.46
3: $\beta_0 + m_1(\text{Latitude})$	27.306	37.54

The observed F-statistic for the latitude component is

$$\frac{(6.128 - 4.541)/(37.46 - 33.99)}{4.541/33.99} = 3.42.$$

Referring this to an $F_{3.47,33.99}$ distribution produces an approximate p-value of 0.023. This therefore suggests that there is some evidence that the underlying regression surface changes with latitude.

The observed F-statistic for the longitude component is 48.08 on 3.54 and 33.99 degrees of freedom, which is highly significant, and so confirms the importance of the effect of longitude.

4.7 Further examples of additive models

4.7.1 Mackerel eggs in the Eastern Atlantic

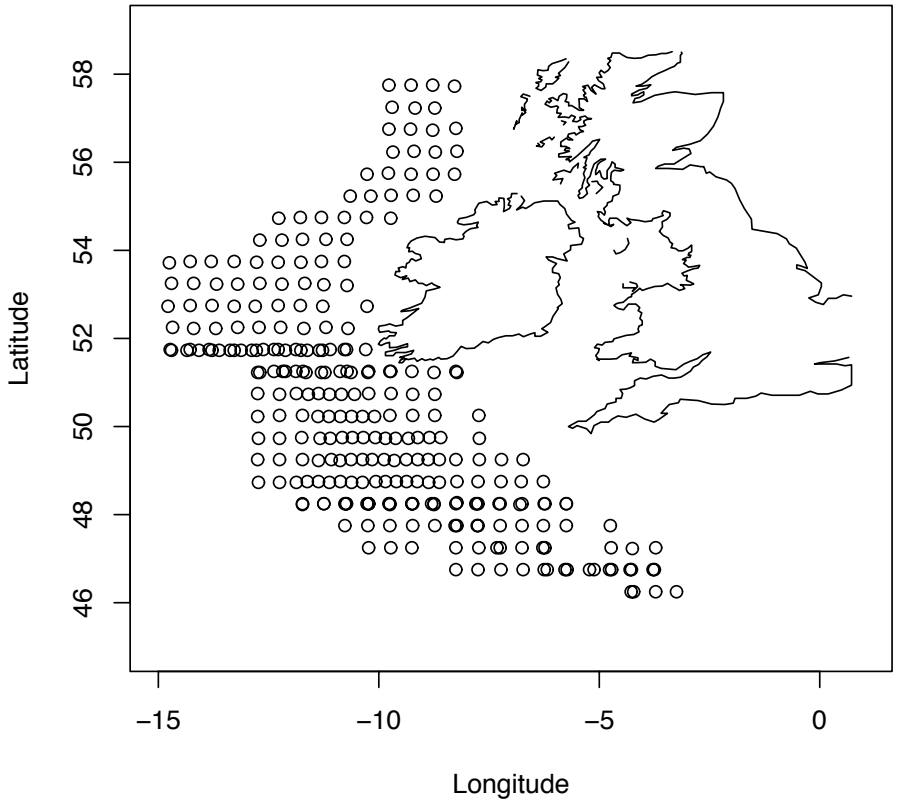

A further example uses data from a multi-country survey of mackerel eggs in the Eastern Atlantic. An additive model for egg density might reasonably contain terms for depth and temperature, plus a joint term for latitude and longitude, to reflect spatial position. This leads to the model

$$y = \beta_0 + m_{12}(x_1, x_2) + m_3(x_3) + m_4(x_4) + \varepsilon,$$

where m_{12} represents a smooth two-dimensional function of latitude (x_1) and longitude (x_2), and m_3 and m_4 represent additive terms of the usual type for depth (x_3) and temperature (x_4). Two-dimensional terms require restrictions to define the functions uniquely, as in the one-dimensional case. A simple convention is $\sum_{i=1}^{n} m_{12}(x_{1i}, x_{2i}) = 0$.

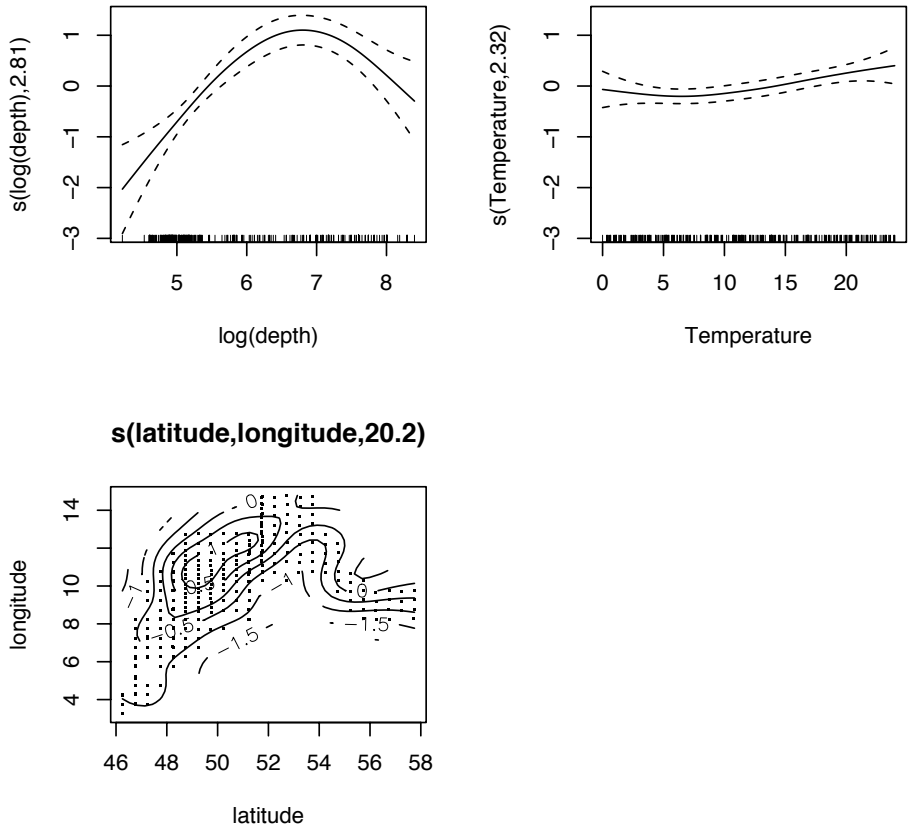

Model	RSS	df
1: $\beta_0 + m_{12}(\text{Lat, Long}) + m_3(\log(\text{Depth})) + m_4(\text{Temp})$	261.13	262.80
2: $\beta_0 + m_{12}(\text{Lat, Long}) + m_4(\text{Temp})$	360.24	266.51
3: $\beta_0 + m_{12}(\text{Latitude, Longitude}) + m_3(\log(\text{Depth}))$	272.08	266.10
4: $m_3(\text{Depth}) + m_4(\text{Temp})$	335.53	270.99

The large change in residual sum-of-squares between models 1 and 2 confirms that depth is an important variable. Similarly, the change between models 1 and 4 shows that there are additional geographical effects which should be accounted for in the model by the presence of the term involving latitude and longitude. However, the F-value for the temperature effect, namely

$$\frac{(272.08 - 261.13)/(266.10 - 262.80)}{262.13/262.80} = 3.33$$

when compared to an $F_{3.30,262.80}$ distribution suggests that the effect of temperature may also be significant.

4.7.2 Dissolved oxygen in the River Clyde

The following plots build a model for the whole river, using data at many sampling stations. Some care has to be taken here because of the repeated measures nature of the data. Also, we are likely to need interaction terms and these are shown in the surface plots.

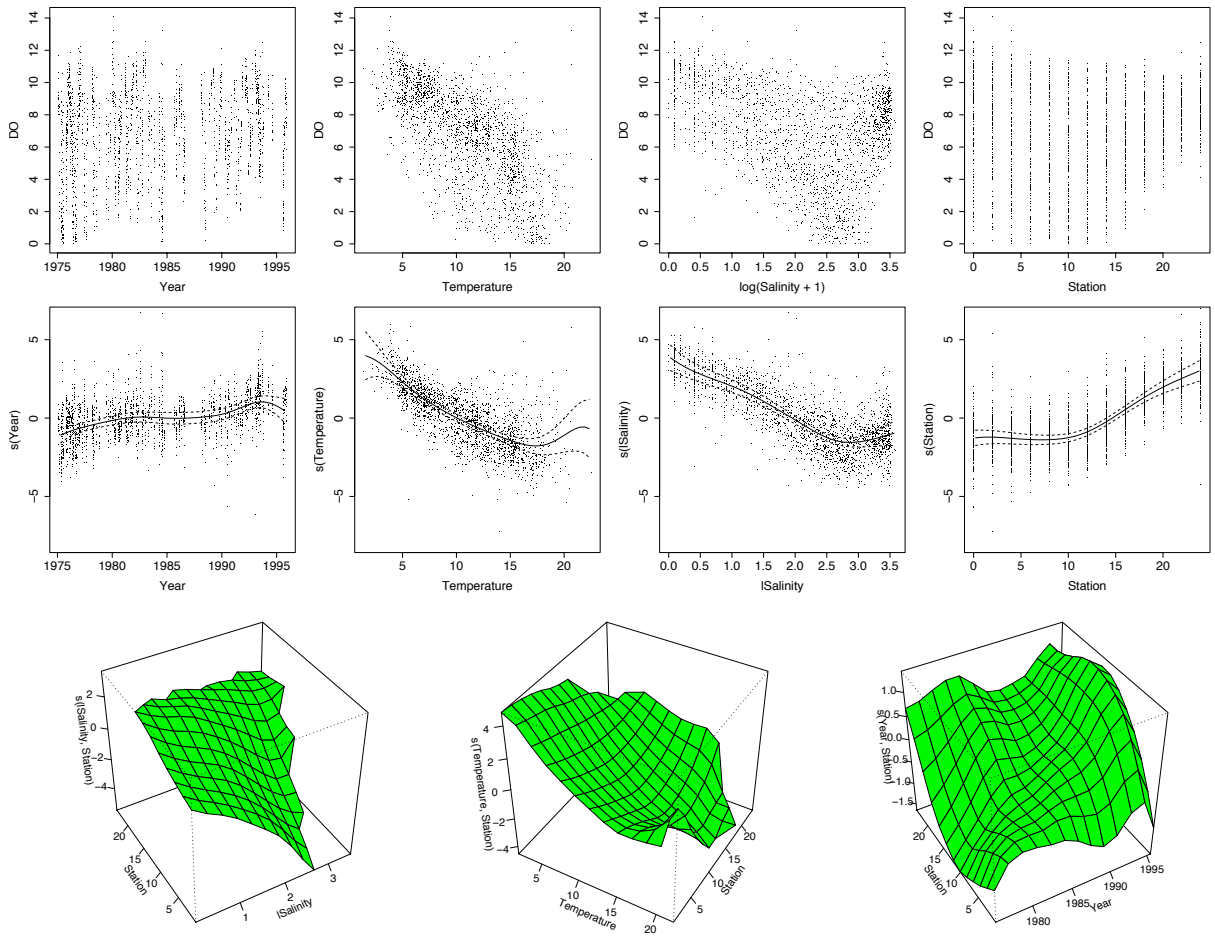
DO
Year
Temperature
log(Salinity + 1)
Station
s(Year)
s(Temperature)
s(lSalinity)
s(Station)
lSalinity
s(lSalinity, Station)
s(Temperature, Station)
s(Year, Station)

5

Bayesian nonparametrics

In this section we will introduce two Bayesian nonparametric methods, Gaussian processes and Dirichlet processes. Gaussian processes are a Bayesian model for function estimation and Dirichlet processes are a Bayesian model for density estimation. They have in common that the object for which we want to perform inference is an infinite-dimensional object rather than a finite-dimensional vector of parameters.

5.1 Gaussian Processes

5.1.1 Bayesian Linear Model

In this section we will quickly revise the Bayesian Linear Model, i.e. we assume the linear regression model

$$y_i|\boldsymbol{\beta} \sim \mathsf{N}(\mathbf{x}_i^\top \boldsymbol{\beta}, \sigma^2) \text{ i.i.d.}$$

or equivalently,

$$\mathbf{y}|\boldsymbol{\beta} \sim \mathsf{N}(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \cdot \mathbf{I}) \tag{5.1}$$

Rather than using the usual normal-inverse-gamma prior jointly placed on $(\boldsymbol{\beta}, \Sigma^2)$, we will for the moment assume that σ^2 is known and just place a Gaussian prior on $\boldsymbol{\beta}$, i.e.

$$\boldsymbol{\beta} \sim \mathsf{N}(0, \tau^2 \cdot \mathbf{I}). \tag{5.2}$$

We can write the p.d.f. of the posterior distribution of $\boldsymbol{\beta}$ as

$$\begin{aligned} f(\boldsymbol{\beta}|y_1, \ldots, y_n) &\propto \underbrace{\left(\prod_{i=1}^{n} f(y_i|\boldsymbol{\beta})\right)}_{\text{Likelihood}} \cdot \underbrace{f(\boldsymbol{\beta})}_{\text{prior}} \\ &= \left(\prod_{i=1}^{n} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(y_i - \mathbf{x}_i^\top \boldsymbol{\beta})^2}{2\sigma^2}\right)\right) \cdot \left(\frac{1}{\sqrt{2\pi\tau^2}}\right)^p \exp\left(-\frac{\sum_{j=1}^{p} \beta_j^2}{2\tau^2}\right) \end{aligned}$$

Collecting terms, taking logs and keeping only terms involving $\boldsymbol{\beta}$ yields the log-posterior density

$$\log f(\boldsymbol{\beta}|y_1,\ldots,y_n) = \text{const} - \frac{1}{2\sigma^2}\sum_{i=1}^{n}(y_i - \mathbf{x}_i^\top\boldsymbol{\beta})^2 - \frac{1}{2\tau^2}\sum_{j=1}^{p}\beta_j^2,$$

which is, up to a multiplicative constant, the objective function used in ridge regression with $\lambda = \frac{\sigma^2}{\tau^2}$.

One can show (by completing the square) that the posterior distribution of $\boldsymbol{\beta}$ is

$$\boldsymbol{\beta}|y_1,\ldots,y_n \sim \mathsf{N}\left(\left(\mathbf{X}^\top\mathbf{X} + \frac{\sigma^2}{\tau^2}\mathbf{I}\right)^{-1}\mathbf{X}^\top\mathbf{y}, \left(\mathbf{X}^\top\mathbf{X} + \frac{\sigma^2}{\tau^2}\mathbf{I}\right)^{-1}\right)$$

Thus the ridge regression estimate $\hat{\boldsymbol{\beta}}^{\text{ridge}} = (\mathbf{X}^\top\mathbf{X} + \lambda\mathbf{I})^{-1}\mathbf{X}^\top\mathbf{y}$ is the Bayesian maximum-a-posteriori (MAP) estimate of $\boldsymbol{\beta}$.

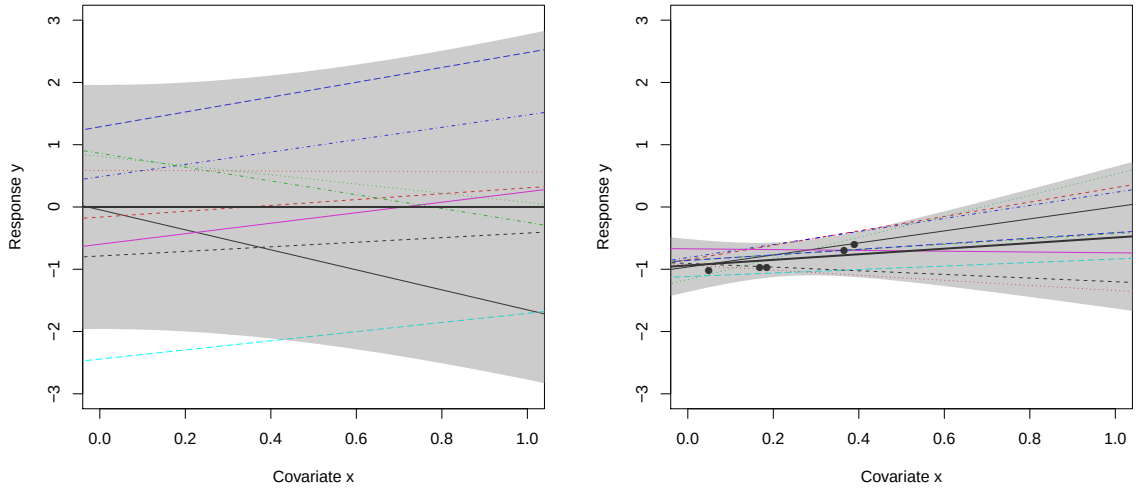

(a) Samples from the prior distribution. (b) Data and samples from the posterior distribution.

Figure 5.1. Draws from the prior distribution and the posterior distribution of a Bayesian linear model. The bold line corresponds to the mean, the shaded area corresponds to pointwise 95% credible intervals.

Figure 5.1 illustrates this idea of Bayesian inference for a linear model with design matrix $\mathbf{X} = \begin{pmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_n \end{pmatrix}$. Panel (a) shows ten draws from the prior distribution, whereas panel (b) shows draws from the posterior distribution given the data.

So far we have stated the model structure of the Bayesian linear model in terms of the mean, i.e. $\mathbb{E}(\mathbf{y}|\boldsymbol{\beta}) = \mathbf{X}\boldsymbol{\beta}$. We will now explain that we can rewrite this, such that exactly the same model structure can be expressed in terms of covariances.

To achieve this, we try to re-express the model without reference to $\boldsymbol{\beta}$, i.e. we suppose that we are not interested in the regression coefficients $\boldsymbol{\beta}$, but only in predictions for future observations. Essentially we have to combine (5.1) and (5.2) to find the marginal distribution of $\mathbf{y}$. The theory of the normal distribution tells us that the marginal distribution of $\mathbf{y}$ is also a

normal distribution, so we only need to find its expected values and its variance.

$$\begin{aligned}\mathbb{E}(\mathbf{y}) &= \mathbb{E}_{\boldsymbol{\beta}}\left(\mathbb{E}_{\mathbf{y}|\boldsymbol{\beta}}(\mathbf{y})\right) = \mathbb{E}_{\boldsymbol{\beta}}(\mathbf{X}\boldsymbol{\beta}) = \mathbf{X}\mathbb{E}_{\boldsymbol{\beta}}(\boldsymbol{\beta}) = \mathbf{0}\\ \mathbf{Var}(\mathbf{y}) &= \mathbf{Var}_{\boldsymbol{\beta}}\left(\mathbb{E}_{\mathbf{y}|\boldsymbol{\beta}}(\mathbf{y})\right) + \mathbb{E}_{\boldsymbol{\beta}}\left(\mathbf{Var}_{\mathbf{y}|\boldsymbol{\beta}}(\mathbf{y})\right)\\ &= \mathbf{Var}_{\boldsymbol{\beta}}(\mathbf{X}\boldsymbol{\beta}) + \mathbb{E}_{\boldsymbol{\beta}}\left(\sigma^2\mathbf{I}\right) = \mathbf{X}\mathbf{Var}_{\boldsymbol{\beta}}(\boldsymbol{\beta})\mathbf{X}^\top + \sigma^2\mathbf{I}\\ &= \tau^2\mathbf{X}\mathbf{X}^\top + \sigma^2\mathbf{I}\end{aligned}$$

thus the Bayesian linear model is equivalent to assuming that

$$\mathbf{y} \sim \mathsf{N}\left(\mathbf{0}, \tau^2\mathbf{X}\mathbf{X}^\top + \sigma^2\mathbf{I}\right). \tag{5.3}$$

What we have achieved by eliminating $\boldsymbol{\beta}$[1] is to move the structural assumption of the Bayesian linear model regression from the mean into the covariance of the Gaussian distribution. This is the key idea which allows us to generalise the Bayesian linear model to Gaussian processes is that we can replace $\mathbf{X}\mathbf{X}^\top$ by a more general matrix.

5.1.2 Definition of a Gaussian process

We start by defining what a Gaussian process actually is. We define a Gaussian process to be a collection of random variables $y_i = y(\mathbf{x}_i)$ $(i = 1, 2, 3, \ldots)$ depending on covariates $\mathbf{x}_i$ such that any *finite* subset of random variables $\mathbf{y} = (y_1, \ldots, y_n) = (y(\mathbf{x}_1), \ldots, y(\mathbf{x}_n))$ has a multivariate Normal distribution. In geostatistics, this model is known as a kriging model[2].

We assume that we can only make observations subject to noise and assume (without any loss of generality) a mean of 0, i.e.

$$\mathbf{y} \sim \mathsf{N}\left(\mathbf{0}, \mathbf{K} + \sigma^2\mathbf{I}\right). \tag{5.4}$$

with

$$\mathbf{K} = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \ldots & k(\mathbf{x}_1, \mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}_n, \mathbf{x}_1) & \ldots & k(\mathbf{x}_n, \mathbf{x}_n) \end{bmatrix}. \tag{5.5}$$

If we write $y_i = m_i + \varepsilon_i$ with $m_i = m(\mathbf{x}_i)$ and $\varepsilon_i \sim \mathsf{N}(0, \sigma^2)$ then (5.4) is equivalent to

$$\mathbf{y}|\mathbf{m} \sim \mathsf{N}(\mathbf{m}, \sigma^2\mathbf{I}) \qquad\qquad \mathbf{m} \sim \mathsf{N}(\mathbf{0}, \mathbf{K}).$$

i.e. $\mathrm{Cov}(m_i, m_j) = k(\mathbf{x}_i, \mathbf{x}_j)$

The function $k(\cdot, \cdot)$ is called *covariance function* or *kernel function*. We are free to choose any $k(\cdot, \cdot)$ as long as it is symmetric in its arguments and the matrix $\mathbf{K}$ from (5.5) is positive semi-definite.

[1] To be mathematically more precise, we have integrated out $\boldsymbol{\beta}$.

[2] named after Daniel Gerhardus Krige, a South African mining engineer and professor at the University of the Witwatersrand, who first suggested kriging to model mineral deposits.

Note that the Bayesian linear model (5.3) is a special case of a Gaussian process with covariance function $k(\mathbf{x}_i, \mathbf{x}_j) = \tau^2 \cdot \mathbf{x}_i^\top \mathbf{x}_j$.

We often make the assumption that the Gaussian process is *stationary*, which is the case if and only if

$$k(\mathbf{x}_i, \mathbf{x}_j) = k(\mathbf{x}_i - \mathbf{x}_j).$$

Figure 5.2 shows a draw from a non-stationary Gaussian process.

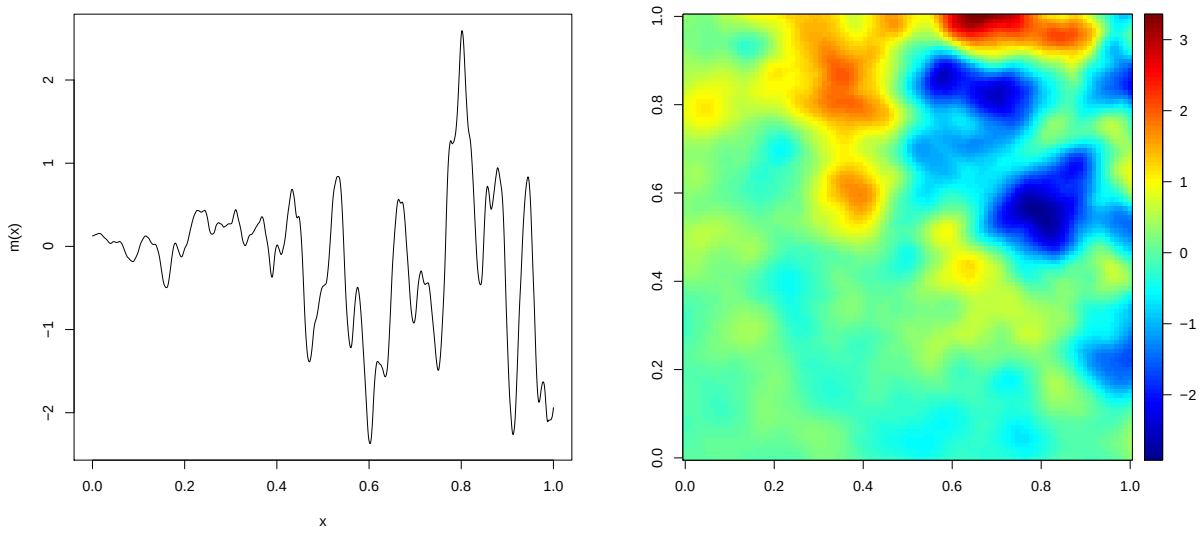

(a) The variance of $f(\cdot)$ is smaller to the left than to the right.

(b) The variance of $f(\cdot)$ is smaller at the bottom-left.

Figure 5.2. Draw from a one-dimensional and two-dimensional Gaussian process, which are not stationary.

An additional simplifying assumption is that the process is *isotropic*, which is the case if and only if

$$k(\mathbf{x}_i, \mathbf{x}_j) = k(\|\mathbf{x}_i - \mathbf{x}_j\|),$$

i.e. only distance, but not direction matters. For the remainder we will assume that the Gaussian process is stationary and isotropic. Figure 5.3 shows a draw from a non-isotropic Gaussian process.

We will discuss different choices of $k(\cdot, \cdot)$ later on in section 5.1.4. In geostatistics it is quite common to use a different parametrisation and work with the so-called *(semi-)variogram*

$$\gamma(\mathbf{x}_i, \mathbf{x}_j) = \frac{1}{2}\text{Var}\,(m_i - m_j) = \frac{1}{2}(k(\mathbf{x}_i, \mathbf{x}_i) + k(\mathbf{x}_j, \mathbf{x}_j)) - k(\mathbf{x}_i, \mathbf{x}_j)$$

instead of the covariance function. There is a one-to-one mapping between the two, so you can either work with the (semi-)variogram or the covariance function.

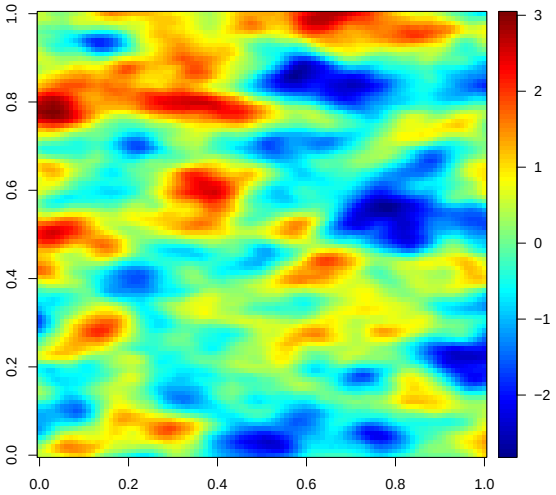

Figure 5.3. Draw from a non-isotropic two-dimensional Gaussian process. The variability in the horizontal direction is less than the one in the vertical direction.

5.1.3 Predictions for Gaussian processes

Conditionals of Gaussian distributions

Assume that

$$\begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} \sim \mathsf{N}\left(\begin{pmatrix} \boldsymbol{\mu}_1 \\ \boldsymbol{\mu}_2 \end{pmatrix}, \begin{pmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{pmatrix} \right)$$

Then the conditional distribution of $\mathbf{y}_2$ given $\mathbf{y}_1$ is

$$\mathbf{y}_2|\mathbf{y}_1 \sim \mathsf{N}\left(\boldsymbol{\mu}_2 + \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}(\mathbf{y}_1 - \boldsymbol{\mu}_1), \boldsymbol{\Sigma}_{22} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\boldsymbol{\Sigma}_{12}\right)$$

We can compute predictions for a new observation with covariates $\mathbf{x}_0$ by looking at the joint distribution

$$\begin{pmatrix} \mathbf{y} \\ y_0 \end{pmatrix} \sim \mathsf{N}\left(\begin{pmatrix} \mathbf{0} \\ 0 \end{pmatrix}, \begin{pmatrix} \mathbf{K} + \sigma^2\mathbf{I} & \mathbf{k}_0 \\ \mathbf{k}_0^\top & k_{00} + \sigma^2 \end{pmatrix} \right),$$

where $\mathbf{K}$ is as defined is the preceding section, $\mathbf{k}_0 = (k(\mathbf{x}_0, \mathbf{x}_1), \ldots, k(\mathbf{x}_0, \mathbf{x}_n))$ is the covariance between the training data and the test case and $k_{00} = k(\mathbf{x}_0, \mathbf{x}_0)$. Then using the formula for the conditional distribution of a Gaussian we obtain

$$y_0|\mathbf{y} \sim \mathsf{N}\left(\mathbf{k}_0^\top \left(\mathbf{K} + \sigma^2\mathbf{I}\right)^{-1}\mathbf{y}, \left(k_{00} - \mathbf{k}_0^\top \left(\mathbf{K} + \sigma^2\mathbf{I}\right)^{-1}\mathbf{k}_0\right) + \sigma^2\right)$$

The mean of the posterior distribution of y_0 can be shown to the best linear unbiased predictor (BLUP). The formula above gives the variance to be used for a prediction interval for a new observation. If we want to get the variance for a confidence interval for its mean we have to

omit the "+σ^2" accounting for the error on the unseen data, i.e. the variance of the predicted mean is $\left(k_{00} - \mathbf{k}_0^\top \left(\mathbf{K} + \sigma^2 \mathbf{I}\right)^{-1} \mathbf{k}_0\right)$.

Figure 5.4 shows five draws each from the prior distribution (panel (a)) and the posterior distribution (panel (b)) from a simple Gaussian process fitted to data.

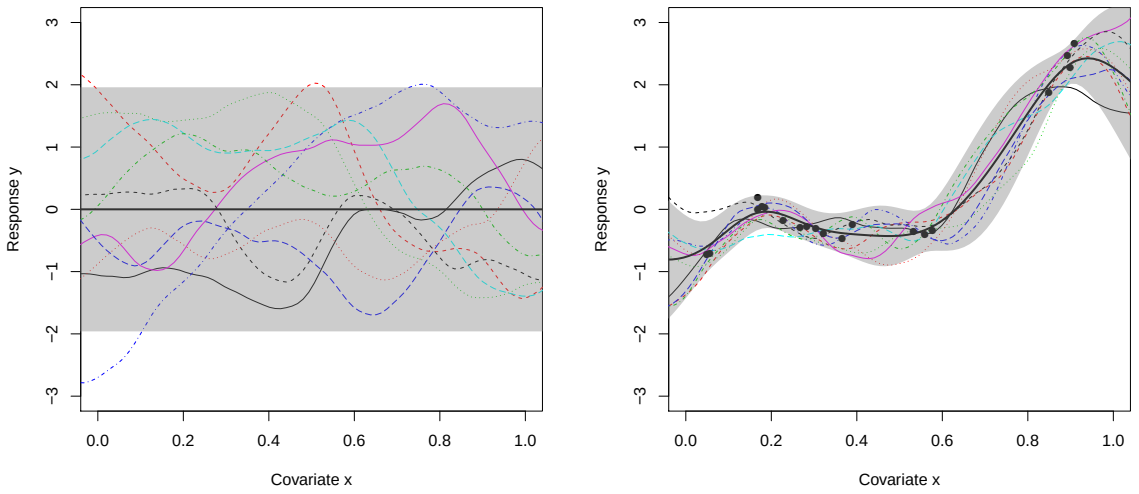

(a) Samples from the prior distribution.

(b) Samples from the posterior distribution.

Figure 5.4. Draws from the prior distribution and the posterior distribution of a simple Gaussian process (Matérn covariance with $\kappa = 2.5$). The bold line corresponds to the mean, the shaded area corresponds to pointwise 95% credible intervals.

5.1.4 Covariance functions (kernel functions)

Squared exponential (SE) The squared exponential (or, Gaussian) covariance function is defined as

$$k(\mathbf{x}_i, \mathbf{x}_j) = \tau^2 \cdot \exp(-\rho \|\mathbf{x}_i - \mathbf{x}_j\|^2)$$

The squared exponential covariance function generates very smooth processes: their paths are infinitely differentiable, which is often unrealistically smooth.

Exponential covariance function – Ornstein-Uhlenbeck (OU) process The exponential covariance function is defined as

$$k(\mathbf{x}_i, \mathbf{x}_j) = \tau^2 \cdot \exp(-\rho \|\mathbf{x}_i - \mathbf{x}_j\|)$$

It leads to a continuous, but not a differentiable process, which is often unrealistically rough. The OU process is the continuous equivalent of an $AR(1)$ process.

γ-exponential One can generalise the above two covariance functions by considering

$$k(\mathbf{x}_i, \mathbf{x}_j) = \tau^2 \cdot \exp(-\rho \|\mathbf{x}_i - \mathbf{x}_j\|^\gamma)$$

with $0 < \gamma \leq 2$, which allows choosing any model between the rough OU process and the squared exponential. However it is less flexible than the Mateérn class.

Matérn class The Matérn covariance function[3] is more flexible than the γ-exponential covariance function, however also much more complex.

$$k(\mathbf{x}_i, \mathbf{x}_j) = \tau^2 \cdot \frac{1}{\Gamma(\kappa)2^{\kappa-1}}(2\sqrt{\kappa}\rho\|\mathbf{x}_i - \mathbf{x}_j\|)^\kappa K_\kappa \left(2\sqrt{\kappa}\rho\|\mathbf{x}_i - \mathbf{x}_j\|\right),$$

where $K_\kappa(\cdot)$ is the modified Bessel function of the second kind. Special cases of the Matérn covariance function are the OU process ($\kappa = \frac{1}{2}$) and the squared exponential ($\kappa \to +\infty$). Figure 5.5 shows functions drawn the from Matérn class for different values of κ.

For all of the above covariance functions, the parameter ρ controls how fast the correlation decays. The larger ρ the quicker the decay of the correlation. The parameter τ^2 controls the prior variance of the signal.

5.1.5 Estimation of hyperparameters

We have so far assumed that the hyperparameters (σ^2 and parameters of the kernel function) are known. In practice however, these need to be estimated from the data. This is best done using the marginal log-density of $\mathbf{y}$,

$$\log f(\mathbf{y}) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log\det(\mathbf{K} + \sigma\mathbf{I}) - \frac{1}{2}\mathbf{y}^\top(\mathbf{K} + \sigma^2\mathbf{I})^{-1}\mathbf{y}$$

We can use an empirical Bayes strategy (sometime also referred to as maximum-likelihood) and maximise the density with respect to the hyperparameters.

However, a Gaussian process can use many hyperparameters and there is often little information in the data about the hyperparameters. This is especially true for the parameter κ of the Matérn covariance function. Full Bayesian models thus typically fare better as they take into account the uncertainty about the values of the hyperparameters. However, with the possible exception of σ^2, none of the hyperparameters can be integrated out in closed form, thus one has to resort to either using a discrete grid or sampling techniques such as Markov Chain Monte Carlo (MCMC).

5.1.6 Classification using Gaussian processes and other extensions

So far we have only studied Gaussian processes for regression. Suppose the response variable of interest is z_i, which does not have a Gaussian distributions, i.e. z_i could be a -1/1 indicator, which would correspond to binary classification. We can handle this problem using Gaussian processes by assuming that there is an unobserved Gaussian process $(y_1, \ldots, y_n)$ and that we can only observe a process $z_i = z_i(y_i)$, which is defined such that z_i only depends on $\mathbf{y}_i$[4]. In the case of binary classification $z_i \in \{-1, 1\}$ and we could assume a probit model, i.e. $\mathbb{P}\{z_i = 1|y_i\} = \Phi(y_i)$, where $\Phi(\cdot)$ is the Gaussian c.d.f.

This idea of latent Gaussian processes can be applied to many others settings.

[3] named after Bertil Matén, a Swedish statistician.

[4] More precisely, this corresponds to assuming that the z_i are conditionally independent given the y_i.

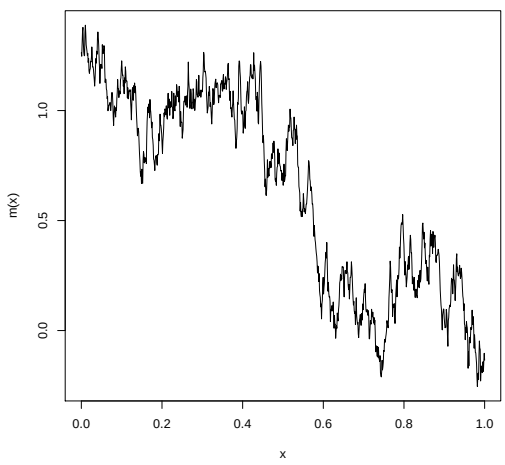

(a) $\kappa = 0.5$ (OU process)

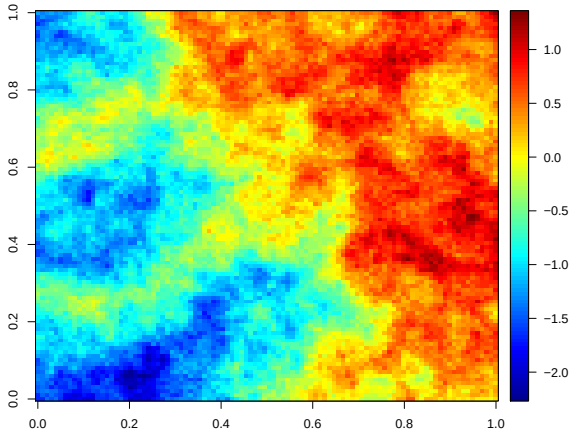

(b) $\kappa = 0.5$ (OU process)

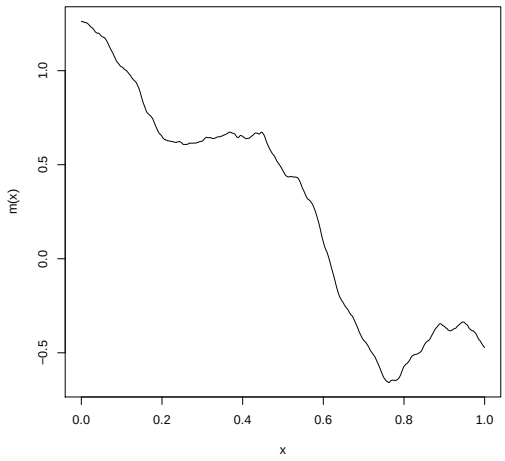

(c) $\kappa = 1.5$

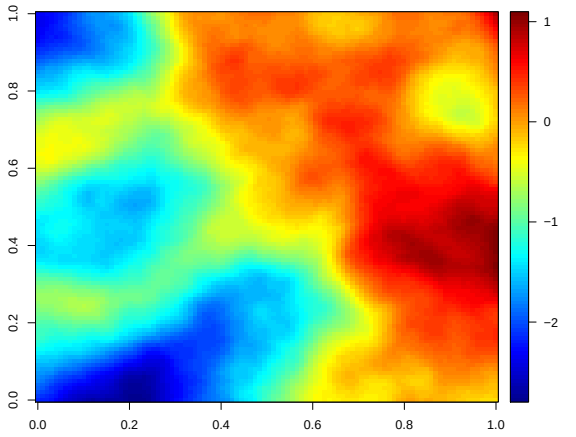

(d) $\kappa = 1.5$

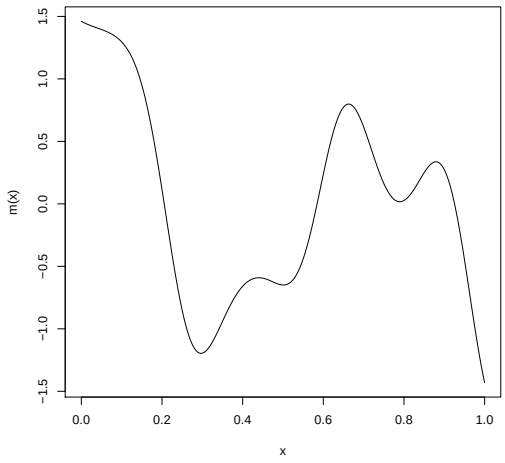

(e) $\kappa = +\infty$ (SE)

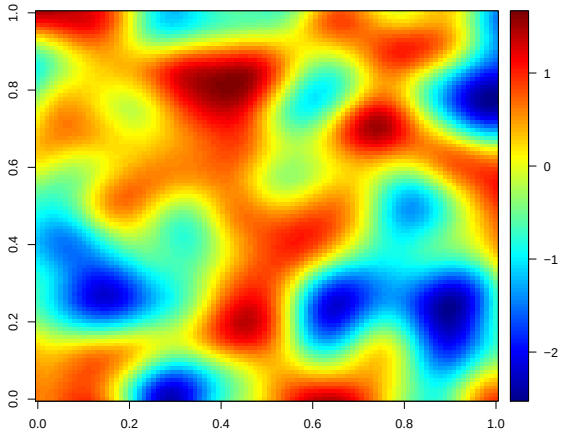

(f) $\kappa = +\infty$ (SE)

Figure 5.5. Samples drawn from the prior distribution of a Gaussian process with a Matérn covariance function for of $\kappa \in \{0.5, 1.5, +\infty\}$. The parameter ρ was chosen so that the covariance at lag $\frac{1}{2}$ is the same for all plots.

5.1.7 Excursion: Kernel methods in more general

An alternative derivation of Gaussian processes In this section we will look again at the derivation of Gaussian processes, more specifically at the transition from equation (5.3) to equation (5.4). We start by the Bayesian linear model, which we showed to be equivalent to

$$\mathbf{y} \sim \mathsf{N}\left(\mathbf{0}, \tau^2 \mathbf{X}\mathbf{X}^\top + \sigma^2 \mathbf{I}\right).$$

The matrix $\mathbf{X}\mathbf{X}^\top$ is nothing other than the matrix of inner products

$$\begin{aligned}\mathbf{X}\mathbf{X}^\top &= \begin{bmatrix} \sum_{j=1}^p x_{1j}^2 & \dots & \sum_{j=1}^p x_{1j}x_{nj} \\ \vdots & \ddots & \vdots \\ \sum_{j=1}^p x_{nj}x_{1j} & \dots & \sum_{j=1}^p x_{nj}^2 \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{x}_1^\top \mathbf{x}_1 & \dots & \mathbf{x}_1^\top \mathbf{x}_n \\ \vdots & \ddots & \vdots \\ \mathbf{x}_n^\top \mathbf{x}_1 & \dots & \mathbf{x}_n^\top \mathbf{x}_n \end{bmatrix} = \begin{bmatrix} \langle \mathbf{x}_1, \mathbf{x}_1 \rangle & \dots & \langle \mathbf{x}_1, \mathbf{x}_n \rangle \\ \vdots & \ddots & \vdots \\ \langle \mathbf{x}_n, \mathbf{x}_1 \rangle & \dots & \langle \mathbf{x}_n, \mathbf{x}_n \rangle \end{bmatrix}\end{aligned}$$

What we have achieved so far is that we have shown that the model (5.3) from above only depends on the covariates through inner products. Note that this is only true for the Bayesian linear model or ridge regression, but not for classical frequentist linear models. We will now exploit this to turn the Bayesian linear model into a generic non-parametric methods, namely Gaussian processes. The advantage of this derivation is that it is in no way tied to the Bayesian linear model and can be applied to a variety of other methods.

Inner products

An inner product space is a vector space with a function $\langle \cdot, \cdot \rangle$, called inner product, which satisfies the following three properties.

(Conjugate) symmetry $\langle \mathbf{x}, \mathbf{y} \rangle = \overline{\langle \mathbf{y}, \mathbf{x} \rangle}$ (in case of a real-valued inner products $\langle \mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{y}, \mathbf{x} \rangle$)

Linearity $\langle \alpha \mathbf{x}_1 + \beta \mathbf{x}_2, \mathbf{y} \rangle = \alpha \langle \mathbf{x}_1, \mathbf{y} \rangle + \beta \langle \mathbf{x}_2, \mathbf{y} \rangle$.

Together with the (conjugate) symmetry, linearity in the first argument implies (conjugate) symmetry in the second argument.

Positive-definiteness $\langle \mathbf{x}, \mathbf{x} \rangle \geq 0$ with equality if and only if $\mathbf{x} = \mathbf{0}$.

One can show that then $\|\mathbf{x}\| = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$ is a norm and thus $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\| = \sqrt{\langle \mathbf{x} - \mathbf{y}, \mathbf{x} - \mathbf{y} \rangle}$ is a distance. Examples of inner products are:

- In the vector space of $\mathbb{R}^p$ the classical inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^p x_i y_i$ satisfies the above definition.
- In the vector space of $\mathbb{C}^p$ the classical inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^p x_i \bar{y}_i$ satisfies the above definition.

– In the vector space of random variables with finite variance, the covariance $\mathrm{Cov}(X, Y)$ satisfies the properties of an inner product.

Rather than working with the data $\mathbf{x}_i$ itself, we now consider a basis expansion $\mathbf{b}(\mathbf{x}_i)$, i.e. an extended design matrix

$$\mathbf{B} = \begin{bmatrix} \mathbf{b}(\mathbf{x}_1)^\top \\ \vdots \\ \mathbf{b}(\mathbf{x}_n)^\top \end{bmatrix} = \begin{bmatrix} b_1(\mathbf{x}_1) & \dots & b_q(\mathbf{x}_1) \\ \vdots & \ddots & \vdots \\ b_1(\mathbf{x}_n) & \dots & b_q(\mathbf{x}_n) \end{bmatrix}.$$

Now

$$\mathbf{B}\mathbf{B}^\top = \begin{bmatrix} \sum_{j=1}^q b_j(\mathbf{x}_1)^2 & \dots & \sum_{j=1}^q b_j(\mathbf{x}_1) b_j(\mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ \sum_{j=1}^q b_j(\mathbf{x}_1) b_j(\mathbf{x}_n) & \dots & \sum_{j=1}^q b_j(\mathbf{x}_n)^2 \end{bmatrix}$$

$$= \begin{bmatrix} \mathbf{b}(\mathbf{x}_1)^\top \mathbf{b}(\mathbf{x}_1) & \dots & \mathbf{b}(\mathbf{x}_1)^\top \mathbf{b}(\mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ \mathbf{b}(\mathbf{x}_1)^\top \mathbf{b}(\mathbf{x}_n) & \dots & \mathbf{b}(\mathbf{x}_n)^\top \mathbf{b}(\mathbf{x}_n) \end{bmatrix} = \begin{bmatrix} \langle \mathbf{b}(\mathbf{x}_1), \mathbf{b}(\mathbf{x}_1) \rangle & \dots & \langle \mathbf{b}(\mathbf{x}_1), \mathbf{b}(\mathbf{x}_n) \rangle \\ \vdots & \ddots & \vdots \\ \langle \mathbf{b}(\mathbf{x}_1), \mathbf{b}(\mathbf{x}_n) \rangle & \dots & \langle \mathbf{b}(\mathbf{x}_n), \mathbf{b}(\mathbf{x}_n) \rangle \end{bmatrix}$$

If we now define $k(\mathbf{x}_i, \mathbf{x}_j) = \tau^2 \langle \mathbf{b}(\mathbf{x}_i), \mathbf{b}(\mathbf{x}_j) \rangle$, then

$$\tau^2 \mathbf{B}\mathbf{B}^\top = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \dots & k(\mathbf{x}_1, \mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}_1, \mathbf{x}_n) & \dots & k(\mathbf{x}_n, \mathbf{x}_n) \end{bmatrix}$$

Plugging this into (5.3) yields equation (5.4), which is the definition of Gaussian processes.

Note that we do not need to know much about the basis expansion $\mathbf{b}(\cdot)$, all we need to be able to do is to compute inner products $k(\mathbf{x}_i, \mathbf{x}_j) = \langle \mathbf{b}(\mathbf{x}_i), \mathbf{b}(\mathbf{x}_j) \rangle$. Actually, we don't even need to specify $\mathbf{b}(\cdot)$, we can simply write down the function $k(\cdot, \cdot)$. An important result from functional analysis, called Mercer's theorem[5], guarantees that if

i. $k(\cdot, \cdot)$ is symmetric in its arguments, i.e. $k(\mathbf{x}_i, \mathbf{x}_j) = k(\mathbf{x}_j, \mathbf{x}_i)$, and
ii. $k(\cdot, \cdot)$ is positive semi-definite, i.e. for any choice of $\mathbf{x}_1, \dots, \mathbf{x}_n$ the matrix

$$\mathbf{K} = \begin{pmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \dots & k(\mathbf{x}_1, \mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}_n, \mathbf{x}_1) & \dots & k(\mathbf{x}_n, \mathbf{x}_n) \end{pmatrix}$$

is positive semi-definite,

then there exists a unique basis expansion ("feature map") $\mathbf{b}(\cdot)$ such that $k(\mathbf{x}_i, \mathbf{x}_j) = \langle \boldsymbol{\phi}(\mathbf{x}_i), \boldsymbol{\phi}(\mathbf{x}_j) \rangle$. Note that the conditions of Mercer's theorem just say that $k(\cdot, \cdot)$ needs to be a valid covariance function, so we can use the same functions as discussed in section 5.1.4.

[5] named after James Mercer FRS (1883 –1932), a British mathematician.

This way of turning a linear method (such as the Bayesian linear model) into a nonlinear method is known as the kernel trick in the Machine Learning community. This trick can be applied to any linear technique, as long as it only depends on the data through inner products.

Support Vector Machines In this section we will introduce Support Vector Machines, a popular Machine Learning methods, which also make use of the kernel trick.

Introduction The history of support vector machines reaches back to the 1960s. The "Generalised Portrait" algorithm, which constructs a separating hyperplane with maximal margin, was originally proposed by the Soviet mathematicians Vapnik and Chernovenkis. Over last decade, support vector machines have become an increasingly popular learning algorithm.

Though support vector machines are mostly used for classification, we will only cover support vector machines for regression.

Robust Statistics You probably remember from your introductory undergraduate Statistics course that the median is more robust than the mean. In other words, the median is less affected by outliers than the mean. In this section we will relate this to loss functions and use these to propose robust methods for regression.

One can show that the mean $\bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i$ of a sample minimises the sum of squares, i.e.

$$\sum_{i=1}^{n}(y_i - a)^2$$

is minimal for $a = \bar{y}$. Similarly one can show than the median $\tilde{y}$ minimises the sum of absolute differences, i.e.

$$\sum_{i=1}^{n}|y_i - a|$$

is minimal for $a = \tilde{y}$.

Remember that in standard linear regression we choose the regression coefficients $\boldsymbol{\beta}$ such that

$$\sum_{i=1}^{n}(y_i - \mathbf{x}_i^\top \boldsymbol{\beta})^2,$$

is minimal, i.e. (standard) linear regression is using a quadratic loss function (just like the mean). One can obtain a robust version of linear regression by choosing $\boldsymbol{\beta}$ such that

$$\sum_{i=1}^{n}|y_i - \mathbf{x}_i^\top \boldsymbol{\beta}|,$$

is minimal. This robust approach to regression yields an algorithm that can cope much better with outliers. However computing the regression coefficients is computationally more demanding and there is no "nice" theory for tests and confidence / prediction intervals.

A compromise between two loss functions is Huber's loss function which is defined as

$$\sum_{i=1}^{n} L_\delta^H(y_i - \mathbf{x}_i^\top \boldsymbol{\beta}) \qquad \text{where} \qquad L_\delta^H(z) = \begin{cases} \frac{z^2}{2} & \text{for } -\delta \leq z \leq \delta \\ \delta(|z| - \delta/2) & \text{otherwise,} \end{cases}$$

where $\delta > 0$ is a suitably chosen constant. Huber's loss function is implemented in the function `rlm` in `MASS`. Figure 5.6 (a) to (c) compares the three loss functions.

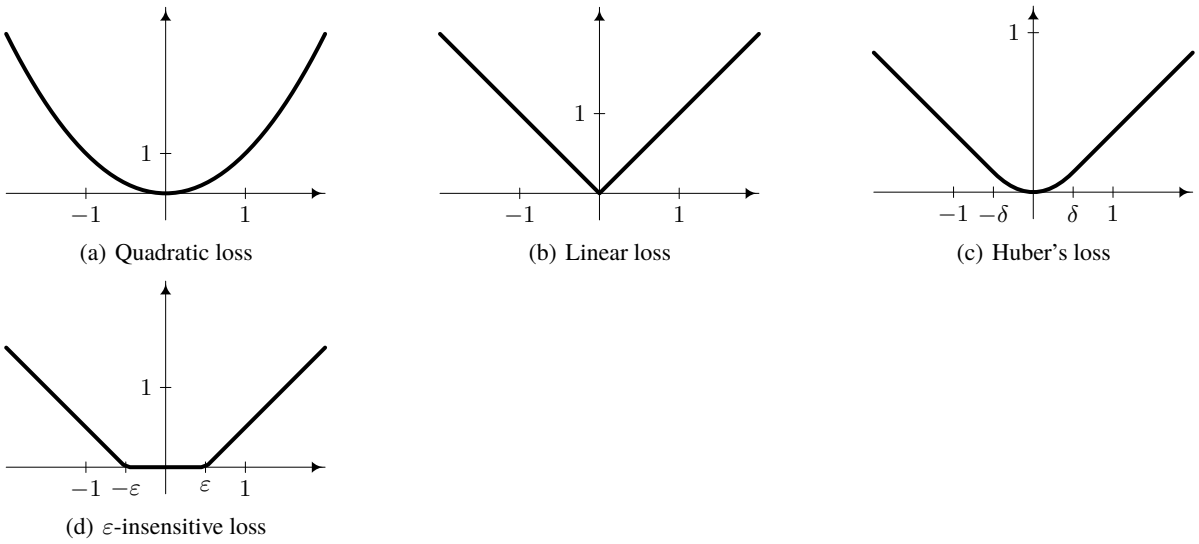

(a) Quadratic loss (b) Linear loss (c) Huber's loss

(d) ε-insensitive loss

Figure 5.6. Different loss functions for regression.

Support Vector Regression Support vector regression is yet another way of performing robust regression. All methods described in the previous section yield estimates of the regression coefficients which depend on all observations. In order to obtain a sparse solution which depends only on a small subset of the observations a modification of the above loss functions is used. This new "ε-insensitive" loss function is defined as

$$L_\varepsilon(z) = (|z| - \varepsilon)_+ = \begin{cases} 0 & \text{for } -\varepsilon \leq z \leq \varepsilon \\ |z - \varepsilon| & \text{otherwise,} \end{cases}$$

where ε is a suitably chosen constant. Small errors (i.e. errors less than ε) will not incur any loss with this new loss function, thus ε is typically chosen rather small (often as small as 10^{-3}).

As common in the support vector literature we will denote the regression coefficients by $\mathbf{w}$, rather than $\boldsymbol{\beta}$. In linear support vector regression we fit a linear function $b + \langle \mathbf{x}_i, \mathbf{w} \rangle$ to a response y_i by minimising the criterion

$$\underbrace{\frac{1}{2}\|\mathbf{w}\|^2}_{\text{regularisation}} + \underbrace{C \sum_{i=1}^{n} L_\varepsilon \left(y_i - b - \langle \mathbf{x}_i, \mathbf{w} \rangle \right)}_{\text{training loss}}$$

Note that the objective function is almost the same as in ridge regression. The only difference is that we use the ε-insensitive loss function for the training loss rather than the quadratic loss used in ridge regression.

Before we go into the details of solving the optimisation problem we will first generalise the problem to the nonlinear case. Suppose we want to use a feature map $\mathbf{b}(\cdot)$, i.e. fit the function $b + \langle \mathbf{b}(\mathbf{x}_i), \mathbf{w} \rangle$. In this case the objective function becomes

$$\frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{i=1}^{n} L_\varepsilon\left(y_i - b - \langle \mathbf{b}(\mathbf{x}_i), \mathbf{w}\rangle\right)$$

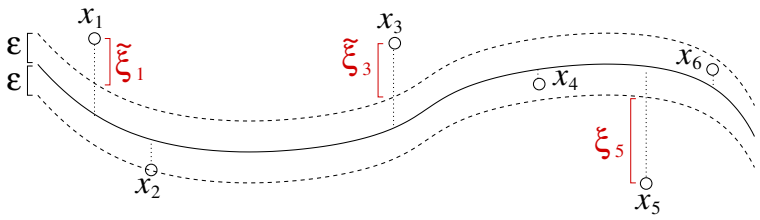

Figure 5.7. Slack variables $\tilde{\xi}_i$ and ξ_i used in support vector regression

The above optimisation can be written as an optimisation problem using slack variables $\tilde{\xi}_i$ and ξ_i. We want to minimise

$$\frac{1}{2}\|\mathbf{w}\|^2 + C\sum_{i=1}^{n}(\tilde{\xi}_i + \xi_i)$$

subject to $y_i - (\langle \mathbf{b}(\mathbf{x}_i), \mathbf{w}\rangle + b) \leq \varepsilon + \tilde{\xi}_i$ and $(\langle \mathbf{b}(\mathbf{x}_i), \mathbf{w}\rangle + b) - y_i \leq \varepsilon + \xi_i$ with slack variables $\tilde{\xi}_i, \xi_i \geq 0$. This is illustrated in figure 5.7. The corresponding dual is

$$D(\tilde{\boldsymbol{\alpha}}, \boldsymbol{\alpha}) = -\frac{1}{2}\sum_{i,j}(\tilde{\alpha}_i - \alpha_i)(\tilde{\alpha}_j - \alpha_j)\left\langle \mathbf{b}(\mathbf{x}_i), \mathbf{b}(\mathbf{x}_j)\right\rangle - \varepsilon\sum_{i=1}^{n}(\tilde{\alpha}_i + \alpha_i) + \sum_{i=1}^{n} y_i(\tilde{\alpha}_i - \alpha_i),$$

which is to be maximised over $\tilde{\alpha}_i, \alpha_i \in [0, C]$ with $\sum_{i=1}^{n}(\tilde{\alpha}_i - \alpha_i) = 0$ and $\tilde{\alpha}_i\alpha_i = 0$. The estimated regression curve can be expressed as a function of $(\tilde{\alpha}_i - \alpha_i)$ and b:

$$\hat{m}(\mathbf{x}_0) = \sum_{i=1}^{r}(\tilde{\alpha}_i - \alpha_i)\left\langle \mathbf{b}(\mathbf{x}_0), \mathbf{b}(\mathbf{x}_i)\right\rangle + b$$

Once again the optimisation problem and its solution only depend on a subset of the learning dataset (those vectors having $\tilde{\alpha}_i > 0$ or $\alpha_i > 0$) and only through inner products, i.e. we can use a kernel $k(\mathbf{x}_i, \mathbf{x}_j) = \langle \mathbf{b}(\mathbf{x}_i), \mathbf{b}(\mathbf{x}_j))\rangle$. In this case the estimated regression curve becomes

$$\hat{m}(\mathbf{x}_0) = \sum_{i=1}^{r}(\tilde{\alpha}_i - \alpha_i)k(\mathbf{x}_0, \mathbf{x}_i) + b.$$

5.2 Dirichlet processes

5.2.1 Bayesian mixture models

We have studied mixture models in section 2.3. One challenge was estimating the number of components. This problem can be overcome by considering a Bayesian approach. For this we need to place priors on the model parameters:

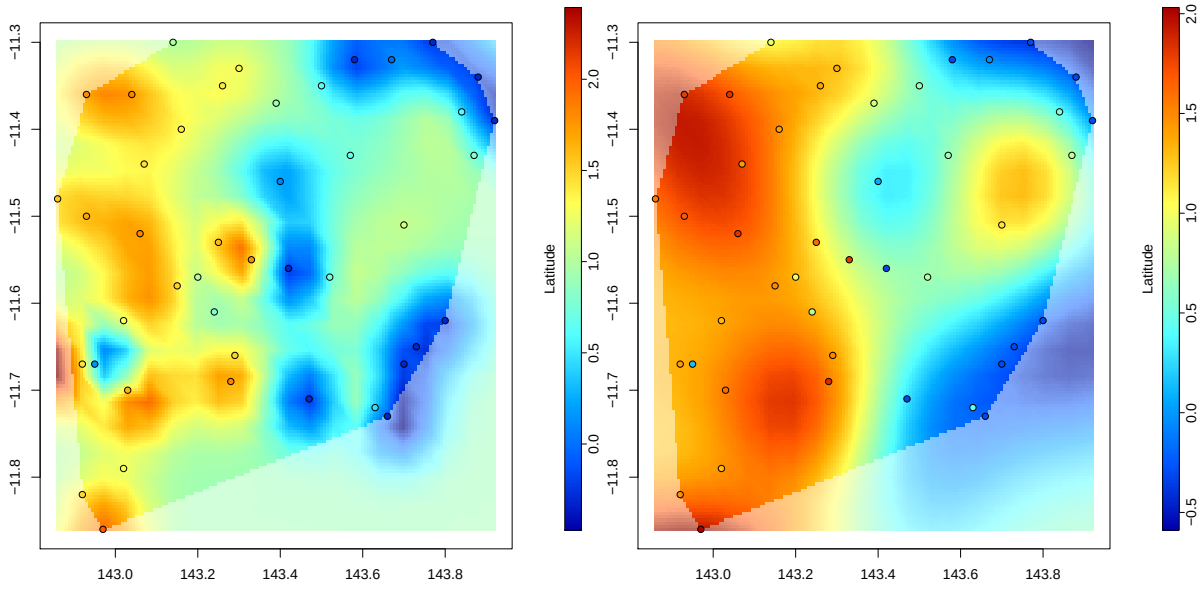

(a) $C = 0.1, \rho = 0.1$ (b) $C = 1, \rho = 1$

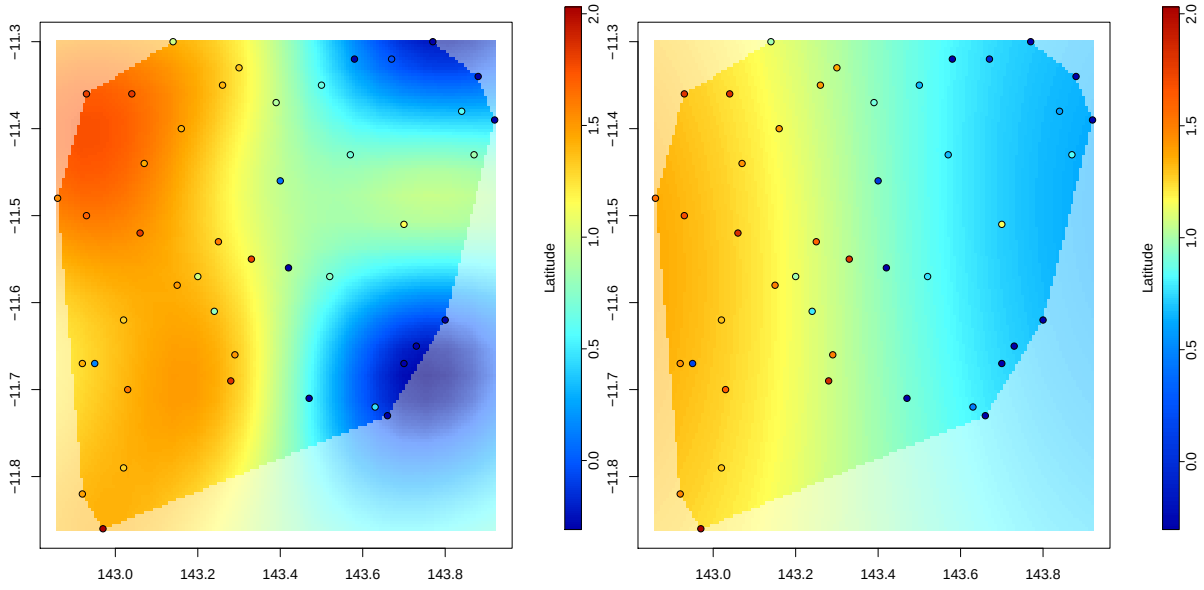

(c) $C = 10, \gamma = 1$ (d) $C = 10, \gamma = 10$

Figure 5.8. Suppport vector machine fits to the Great Barrier Reef data for different values of the hyperparameters using a squared exponential kernel.

– We place a Dirichlet distribution on $\boldsymbol{\phi} = (\pi_1, \ldots, \pi_K)$. The Dirichlet distribution is a generalisation of the Beta distribution and generates realisations such that $\sum_{k=1}^{K} \pi_k = 1$ with $\pi_k > 0$. Its probability density function is, for $pi_k > 0$ with $\sum_{k=1}^{K} \pi_k = 1$, given by

$$f(\pi_1, \ldots, \pi_K) = \frac{\Gamma(\delta_1 + \ldots + \delta_K)}{\Gamma(\delta_1) \cdots \Gamma(\delta_K)} \pi_1^{\delta_1 - 1} \cdots \pi_K^{\delta_K - 1}$$

In this context we typically choose all δ_k to be equal, i.e. $\delta_1 = \ldots = \delta_K = \delta$. Figure 5.9 shows the distribution for a trivariate Dirichlet distribution for $\delta \in \{0.75, 1, 1.5\}$. The figure shows the density only as a function of π_1 and π_2 as $\pi_3 = 1 - \pi_1 - \pi_2$.

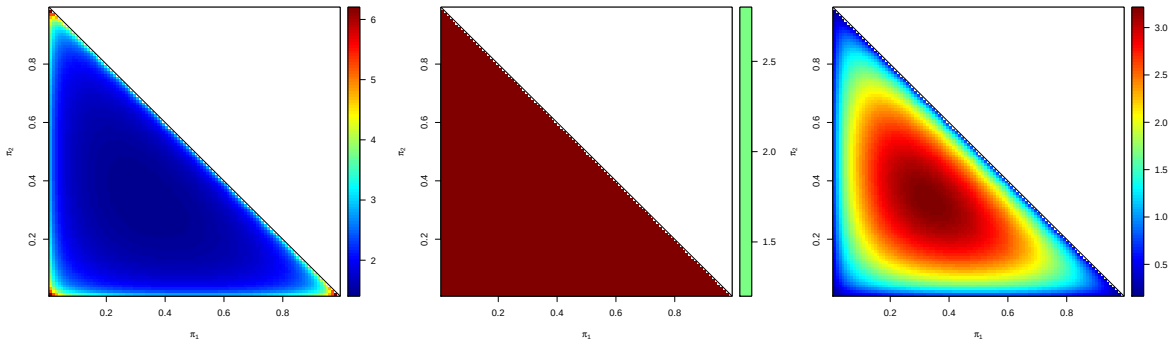

(a) $\delta = 0.75$ (b) $\delta = 1$ (c) $\delta = 1.5$

Figure 5.9. Probability density function of the trivariate Dirichlet distribution for $\delta \in \{0.75, 1, 1.5\}$.

– We place a conjugate Normal-Wishart distribution on $\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k$.

In order to perform posterior inference using a Gibbs sampler we introduce the same latent variables S_i (giving the cluster allocation) as have for the EM algorithm. The Gibbs sampler iterates, just like the EM algorithm, between updating the cluster allocations and then updating the means and variances in each cluster.

Note that our model currently assumes a fixed K. If we want to infer K from the data, we have to include K as a parameter in our model. This however introduces the problem that the dimension of the parameter space changes as K is chganged (we need to introduce a new pair $(\boldsymbol{\mu}_K, \boldsymbol{\Sigma}_K)$ or remove one). This requires the use of transdimensional MCMC techniques, such as reversible jump Markov Chain Monte Carlo (RJMCMC), which is rather complex. In the next section we will consider a model which gets around this.

5.2.2 Dirichlet processes

An alternative model for estimating densities (and clustering) in a Bayesian context is given by the Dirichlet process. The key advantage of the the Dirichlet process does not require the specification of a number of clusters or transimensional MCMC if one wants to lern the number of clusters from data.

Dirichlet process prior The Dirichlet process prior is a Bayesian model for a random measure, i.e. it is a probability distribution on distributions.

Thse Dirichlet process (DP) can be characterised using a number of equivalent definitions. The first definition was given by Ferguson (1973). The Dirichlet process has two parameters, a so-called base measure (a probability distribution) G_0 around which the Dirichlet process is centred and a concentration parameter $\alpha > 0$. Note that G_0 is itself a probability distribution, i.e. it assigns probabilities to sets A, which are subsets of the domain of interest Ω. for example our base measure G_0 could be the $\mathsf{N}(0, 1)$ distribution in which case $\Omega = \mathbb{R}$ and $G_0((a, b)) = \int_a^b \phi(z)\, dz = \Phi(b) - \Phi(a)$.

Ferguson's definition The original definition given by Ferguson Dirichlet processes is based on considering finite partitions $(A_1, \ldots, A_q)$ of Ω.[6] G is a Dirichlet processes with parameters G_0 (base measure) and $\alpha > 0$ (concentration) if and only if

$$(G(A_1), \ldots, G(A_q)) \sim \mathsf{Dirichlet}(\alpha G_0(A_1), \ldots, \alpha G_0(A_q))$$

for every partition $(A_1, \ldots, A_q)$ of Ω.

Using the formulae for the mean and the variance of the Dirichlet distribution we can show that

$$\begin{aligned}
\mathbb{E}(G(A_j)) &= \frac{\alpha G_0(A_j)}{\alpha G_0(A_1) + \ldots + \alpha G_0(A_q)} = G_0(A_j) \\
\text{Var}\,(G(A_j)) &= \frac{G_0(A_j)(1 - G_0(A_j))}{\alpha + 1},
\end{aligned}$$

i.e. on average the random measure G assigns probbability $G_0(A)$ to a set A, so G is centered around G_0. α controls the spread of G: the smaller α the higher the variance of the random probabilities of a set A.

Though this is the original definition of the Dirichlet process, many of the key properties of Dirichlet processes are not obvious from this definition.

One, at first sight very surprising, feature of the Dirichlet process is that draws from the Dirichlet process exhibit ties, i.e. if we were to draw a sample from G, then we would draw some values more than once with non-zero probability. This not the case the continuous distributions.[7] So the distribution G is discrete with probability 1, even if G_0 is continuous.

Stick-breaking construction (Sethuraman, 1994) An equivalent definition of the Dirichlet process is based on a "stick-breaking" construction. Instead of characterising G, we write down the stick-breaking construction in terms of drawing samples from G.

We can draw a sample $\theta_1, \theta_2, \ldots \sim G$ as follows.

1. Draw a sequence of realisations θ_i from the base measure, i.e. $\theta_1, \theta_2, \ldots \sim G_0$.

[6] A finite partition of Ω is a finite set of subsets $A_j \subset \Omega$, which are disjoint ($A_j \cap A_k = \varnothing$ for $j \neq k$) and whose union spans Ω, i.e. $A_1 \cup \ldots \cup A_q = \Omega$. This way every $y \in \Omega$ is in exaclty one of the A_j's.

[7] If we draw two realisations from say the $\mathsf{N}(0, 1)$ distribution, these would be different with probability 1.

2. Draw a sequence of weights $V_1, V_2, \ldots \sim \mathsf{Beta}(1, \alpha)$.
3. Construct a discrete probability distribution G using a stick breaking construction

$$G(x) = \begin{cases} V_1 & \text{for } x = \theta_1 \\ (1 - V_1)V_2 & \text{for } x = \theta_2 \\ (1 - V_1)(1 - V_2)V_3 & \text{for } x = \theta_3 \\ \ldots & \ldots \\ 0 & \text{otherwise} \end{cases}$$

4. We can draw samples $\phi_1, \phi_2, \ldots$ from G by drawing from the discrete distribution G from step 3.

One key advantage of the stick-breaking definition is that it makes it clear that the distribution G is discrete (albeit with random range and random probabilities), thus a sample $\phi_1, \ldots, \phi_m$ from G is likely to exhibit ties. How many ties we would expect depends on the concentration parameter α.

– If $\alpha \to 0$ all the ϕ_i are (almost surely) identical with $\phi_1 \sim G_0$.
– If $\alpha \to +\infty$ the ϕ_i are i.i.d. draws from G_0, i.e. we would observe no ties with probability 1.

Chinese Restaurant process (CRP) Yet another way of characterising the Dirichlet process is given by the Chinese restaurant process. One can show that samples can be drawn recursively from G as follows:

1. Draw $\phi_1 \sim G_0$.
2. For $i = 2, 3, \ldots$
   a) Construct a frequency table of the values $\phi_1, \ldots, \phi_{i-1}$ drawn so far.

Value	θ_1	θ_2	θ_3	$\ldots$
Absolute frequency	n_1	n_2	n_3	$\ldots$

   b) Draw $\theta_0 \sim G_0$.
   c) Draw ϕ_i from the discrete distribution given by

$$p(\phi) = \begin{cases} \frac{\alpha}{i-1+\alpha} & \text{for } \phi = \theta_0 \\ \frac{n_1}{i-1+\alpha} & \text{for } \phi = \theta_1 \\ \frac{n_2}{i-1+\alpha} & \text{for } \phi = \theta_2 \\ \ldots & \ldots \\ 0 & \text{otherwise.} \end{cases}$$

We can imagine this as a sequence of customers coming into a Chinese restaurant. Assume that in the restaurant all customers (representing random draws from G) sitting at a table will be served the same dish (draw from G). We will denote by θ_j the dish on the j-th table and by

ϕ_i the dish ordered by the i-th customer. If customer i sits at table j then $\phi_i = \theta_j$. Figure 5.10 illustrates this. Due to the exchangeability of the Dirichlet process[8], this can view forms the basis of an efficient Gibbs sampler, referred to later on in this section.

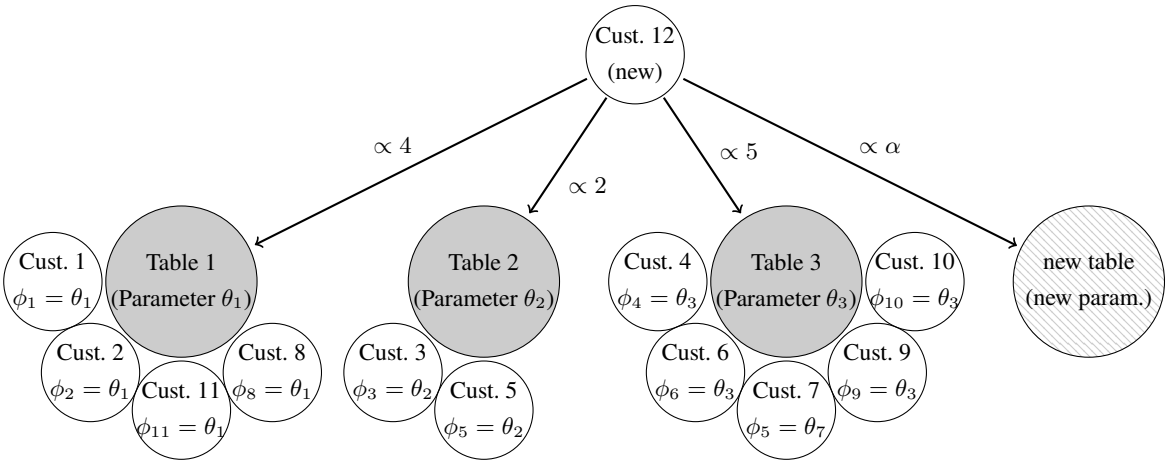

Figure 5.10. Illustration of the Chinese restaurant view of the Dirichlet process, assuming 11 customers have already chosen their table.

One equivalent view is to assume that the i-th customer entering the restaurant chooses an existing customer next to whom he wants to sit (each with probability $\frac{1}{i-1+\alpha}$). With probability $\frac{\alpha}{i-1+\alpha}$ the customer chooses to sit at a new table.

One key conclusion from this way of viewing the Dirichlet process is that it leads to a varying number of clusters (tables), depending on how the customers are seated.

The Dirichlet process has a "rich gets richer"property. Tables with many customers are more likely to attact future customers.

The CRP view of the Dirichlet process makes it very easy come up with generalisations (just think of other ways of seating customers). It is also (due to the exchangeability of the Dirichlet process) the basis for constructing Gibbs samplers.

Limit of finite mixture models The finite mixture model from section 5.2.1 converges to a Dirichlet process if we let the number of components $K \to +\infty$ and $\delta \to 0$ such that $K\delta \to \alpha$.

Posterior distribution So far we have just stated different (equivalent) ways of defining the Dirichlet process. We will now use it to perform Bayesian inference. In this model we want to infer the probability distribution of a sample $Y_1, \ldots, Y_n$.

We start by placing a Dirichlet process prior with base measure G_0 and concentration α on the distribution of distributions. One can now show that the posterior distribution is again given by a Dirichlet process, this time with concentration $\alpha + n$ and base measure

$$\frac{\alpha}{\alpha + n} G_0 + \frac{1}{\alpha + n} \sum_{i=1}^{n} \delta_{y_i},$$

[8] A sequence of random variable is called exchangeable if the joint distribution of any permutation of them has teh same distribution. Any i.i.d. sample is exchangeable, but the converse is not necessarily the case.

where δ_{y_i} denotes a distribution which takes the value y_i with probability 1. In other words, the base measure of the posterior distribution is a mixture of the base distribution of the prior (with weight $\frac{\alpha}{\alpha+n}$) and the empirical distribution of the sample $Y_1, \ldots, Y_n$ (with weight $\frac{n}{\alpha+n}$). Figure 5.11 gives an example of this. This is however rather unsatisfying: even if we assume that our data comes from a continuous distribution the distribution resulting from posterior inference has discrete components. This is why Dirichlet processes are usually not used to model the distribution of the $Y_1, \ldots, Y_n$ directly. Instead one typically assumes that Y_i comes from a distribution with parameter ϕ_i, where ϕ_i is modelled using a Dirichlet process. This setup effectively adds "jitter" such that the distributions resulting from posterior inference are continuous.

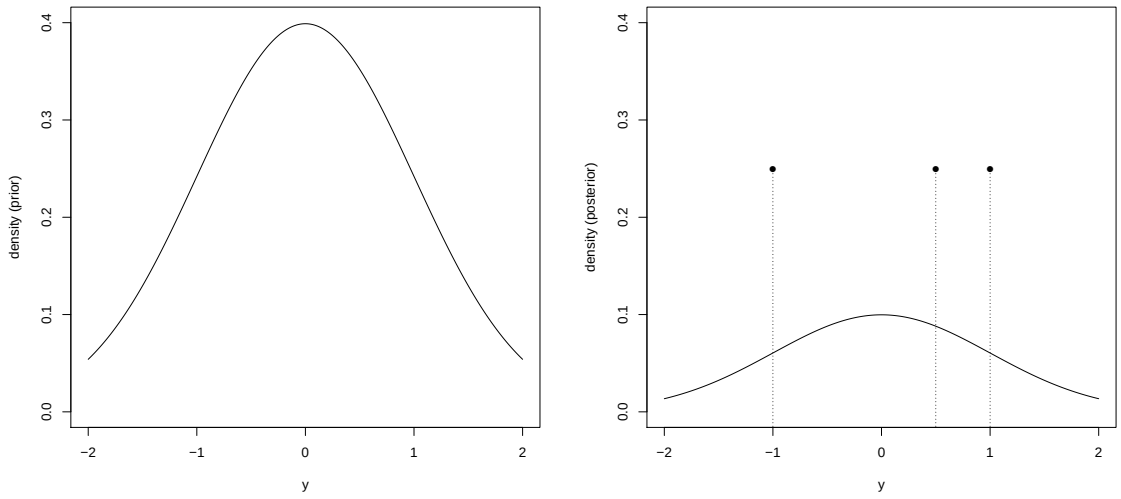

(a) Base measure of the prior distribution

(b) Base measure of the posterior distribution

Figure 5.11. Base measure of the prior and posterior distribution in a Dirichlet process model with the $\mathsf{N}(0, 1)$ distribution as prior base measure and a prior concentration parameter of 1. The observed sample was $y_1 = -1$, $y_2 = 1$ and $y_3 = 0.5$.

Dirichlet process mixture models In this section we will explain how a Dirichlet process can be used to model a mixture of normals.

We have seen in the preceding section that we do not want to use a Dirichlet process to directly model the Y_i. Instead we assume that the mean vector and covariance matrix for each observation come from a Dirichlet process, i.e.

$$Y_i|\boldsymbol{\phi}_i \sim \mathsf{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i) \qquad \text{using } \boldsymbol{\phi}_i = (\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$$

We now place a Dirichlet process prior on the distribution of the $\boldsymbol{\phi}_i$. This is likely to result in ties, i.e. more than one observation will have the same $\boldsymbol{\mu}_i$ and $\boldsymbol{\Sigma}_i$, which corresponds to nothing other than clusters.

Inference in this model can be performed efficiently using a Gibbs sampler (see e.g. Neal (2000)).

Figures 5.12 and 5.13 show the results when applying the method to the aircraft data from chapter 2.1.

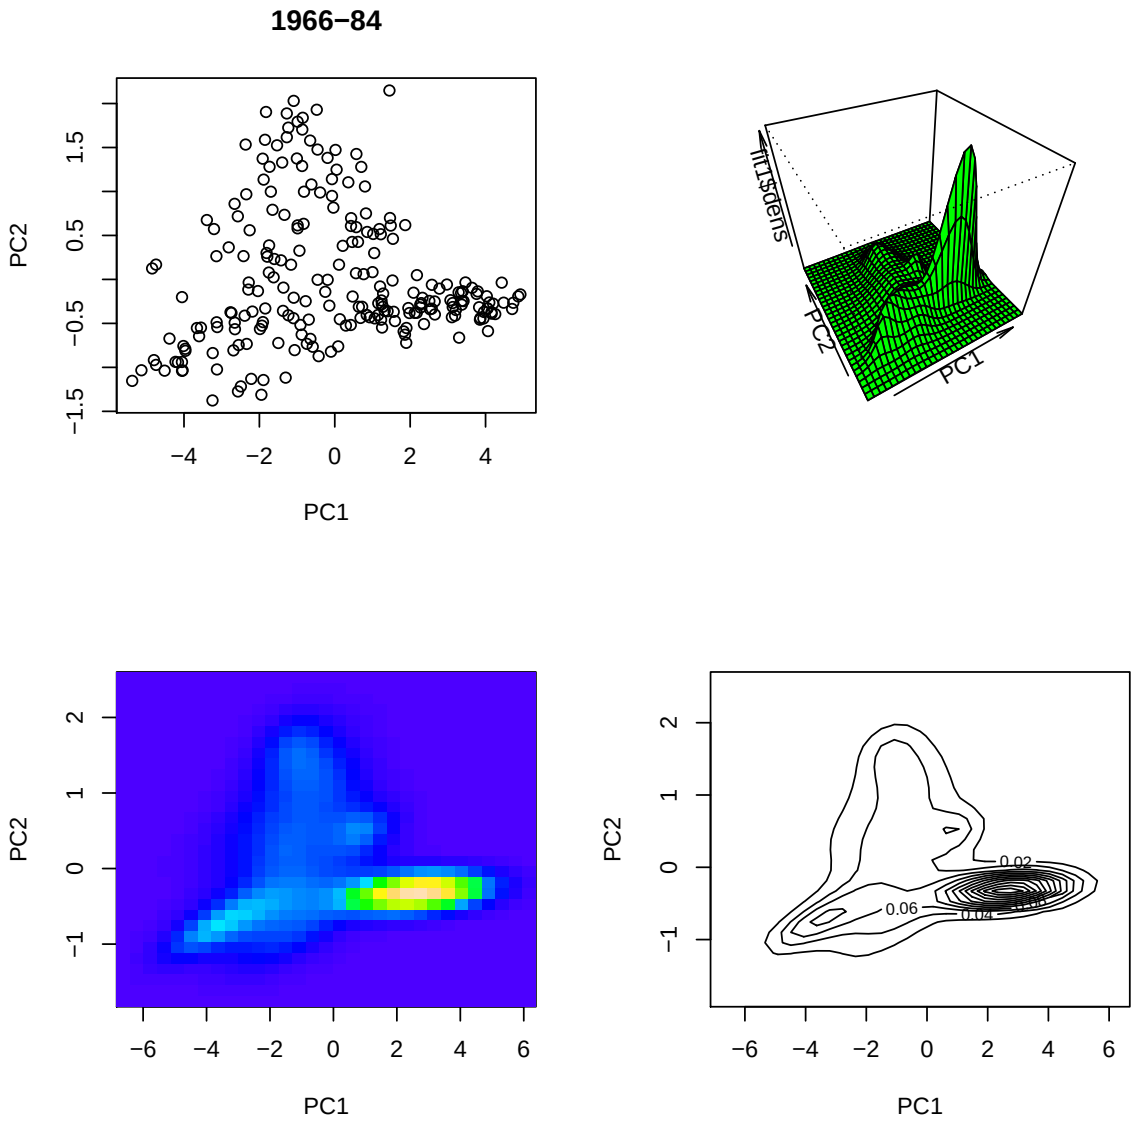

Figure 5.12. Estimated density using a Dirichlet process mixtures of normals models for the first two principal components of aircraft data (using the R package `DPpackage`).

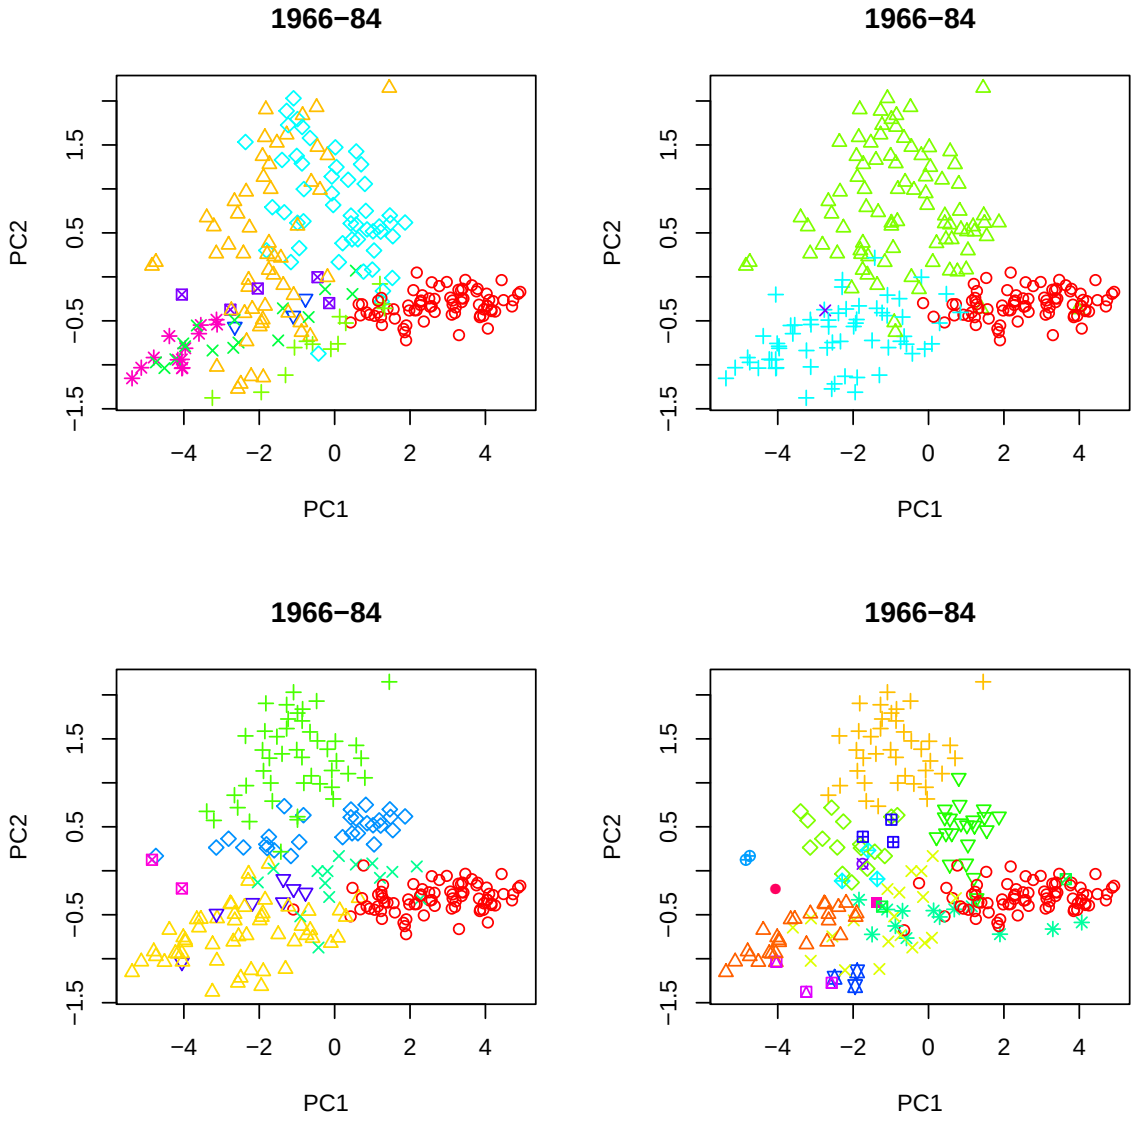

Figure 5.13. Clustering at four different iterations of the MCMC algorithm obtained for the principal components of aircraft data.

6

Spatiotemporal models

In the previous chapter, Gaussian processes were introduced as a powerful description for random variation over space. One use of these techniques is to define prior distributions over function spaces, in the context of building Bayesian models in spatial and other settings where flexible functional descriptions are required. This has been an area of very considerable research activity in recent years. In particular, spatiotemporal data, where observations are indexed over both space and time, have become very common. An excellent general introduction to this topic is given by Cressie and Wikle (2011). In this chapter, some approaches to modelling spatiotemporal data will be outlined, with an emphasis on the use of flexible regression, as described in earlier chapters.

6.1 Spatiotemporal processes

A spatiotemporal process can be defined in an analogous way to a spatial process, with values now indexed by space and time, rather than just space. If we denote the process by $Z(s,t)$, where s denotes spatial position, t denotes the time-point and Z denotes the random variable which described the value of the process, than a spatiotemporal Gaussian process can be defined by the fact that for any finite set of space-time locations $\{(s_1,t_1),\ldots,(s_n,t_n)\}$ the corresponding set of random variables $\{Z(s_1,t_1),\ldots,Z(s_n,t_n)\}$ has a multivariate normal distribution, with some mean μ and covariance matrix Σ. This definition doesn't give much intuition about what a spatiotemporal Gaussian process looks like so the pictures below illustrate the values of the process at three adjacent time points. The key idea is that the spatial shape evolves with a degree of smoothness.

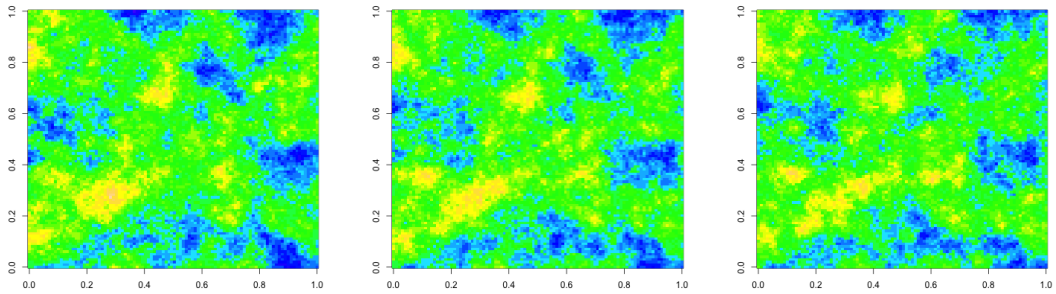

The covariance function

$$\text{cov}\{Z(s_1, t_1), Z(s_2, t_2)\}$$

is the key to this, as we saw earlier in the purely spatial case. To keep things simple, let's assume that the process is stationary and isotropic so that we can write

$$\text{cov}\{Z(s_1, t_1), Z(s_2, t_2)\} = C(|s_1 - s_2|, |t_1 - t_2|)$$

for some suitable function C. A further simplification arises if the process is *separable* in space and time so that we can write

$$C(|s_1 - s_2|, |t_1 - t_2|) = C_s(|s_1 - s_2|)C_t(|t_1 - t_2|)$$

for suitable functions C_s and C_t.

Considerable research continues to be devoted to models based on spatiotemporal processes, recognising that non-stationary, non-separable models will sometimes be needed. A discussion of that area is beyond the scope of the present course, so we will deal with the topic by placing the emphasis on flexible *mean* functions of s and t.

6.2 Additive models for spatiotemporal data

Example 6.1 (SO_2 air pollution over Europe). In the 1970's and 1980's there was considerable concern about SO_2 air pollution. This was emitted by power stations and other installations and the material rises high in the atmosphere and can travel long distances, causing pollution problems in neighbouring countries. The `SO2` dataset documents values of SO_2, on a log scale, from monitoring stations across Europe from 1990 to 2001. The aim of the monitoring stations was to assess whether increasing European regulatory control of SO_2 emissions was effective.

The data were collected through the *European monitoring and evaluation programme* (EMEP) and they are available at `www.emep.int`. The data recorded here have been organised into a convenient form for analysis. The data file consists of six variables:

`site`	a site code for the monitoring station
`longitude`	longitude of the monitoring station
`latitude`	latitude of the monitoring station
`year`	year of measurement
`month`	month of measurement
`logSO2`	SO2 measurement on a log scale

Analysis of these data is reported in *Spatiotemporal smoothing and sulphur dioxide trends over Europe*, A. W. Bowman, M. Giannitrapani and E. M. Scott; *Applied Statistics*, 58 (2009), 737-752. ◁

Our earlier discussion of additive models gives us a possible starting point, namely

$$y = \mu + m_s(s) + m_t(t) + m_z(z) + \varepsilon,$$

where y denotes log(SO_2), s denotes spatial position indexed by latitude and longitude (or equivalent parameters), t denotes time on a scale of years, z denotes month of the year, and the functions m_s, m_t, m_z denote flexible descriptions of the effects on the mean of these three variables. These could be constructed by local linear smoothing but, as we have seen, a b-spline approach have strong advantages in terms of low-dimensional representation of the functions involved.

However, it seems very unlikely that a model of this simplicity will do. It implies that the spatial effect is unchanged with time or season and this seems very implausible. Using standard model formula notation, we need to include interactions to create a model of the form

$$y = \mu + m_s(s) + m_t(t) + m_z(z) + m_s(s) : m_t(t) + m_s(s) : m_z(z) + m_t(t) : m_z(z) + \varepsilon.$$

The interaction terms can be constructed through column products of the b-spline design matrices for the constituent terms.

A model of this type needs to be viewed carefully to understand the nature of the interactions. In the lecture we will use some animation to explore the effects involved.

If there is concern that spatial and temporal correlation remains in $varepsilon$ despite the flexible nature of the mean function which has been constructed, then the residuals of the model can be examined, for example through autocovariance functions. Measures of uncertainty can be adjusted to incorporate a fitted covariance function, for example based on a separable model with suitable spatial and temporal covariance functions.

6.3 A Bayesian approach

A p-spline model is fitted by minimising an expression of the form

$$(y - B\alpha)^T (y - B\alpha) + \lambda \alpha^T D^T D \alpha.$$

This leads naturally to a Bayesian formulation, with conjugate distributions, as

$$\begin{aligned} Y|\alpha, \sigma^2, M_\lambda &\sim \mathcal{N}\left(B\alpha, \sigma^2 I\right) \\ \alpha, \sigma^2|M_\lambda &\sim \mathcal{NIG}_m\left(\mu = 0,\, V^{-1} = \lambda D^T D, a, b\right) \end{aligned}$$

where M_λ denotes the model resulting from a particular value of λ.

The posterior density for λ can be shown to be

$$f_{M_\lambda|Y} \propto \lambda^{\frac{\text{rank}(D^T D)}{2}} \frac{|B^T B + \lambda D^T D|^{-\frac{1}{2}}}{\left\{2b + y^T \left[I - B(B^T B + \lambda D^T D)^{-1} B^T\right] y\right\}^{a+\frac{n}{2}}} f_{M_\lambda}$$

where f_{M_λ} denotes the prior for λ. The values $a = b = 0.0001$ produce a flat prior for σ^2 while an improper uniform distribution gives an uninformative prior for λ.

This formulation uses a single value of the penalty parameter λ. Often different penalty parameters for space and time will be required but there are high computational efficiencies possible when a single penalty parameter is used and, where this is needed, some differentiation between space and time might be achieved by the use of different numbers of basis functions in different dimensions.

The plots below show an example of data on groundwater pollution around an industrial installation. The analysis of the data will be discussed in the lecture.

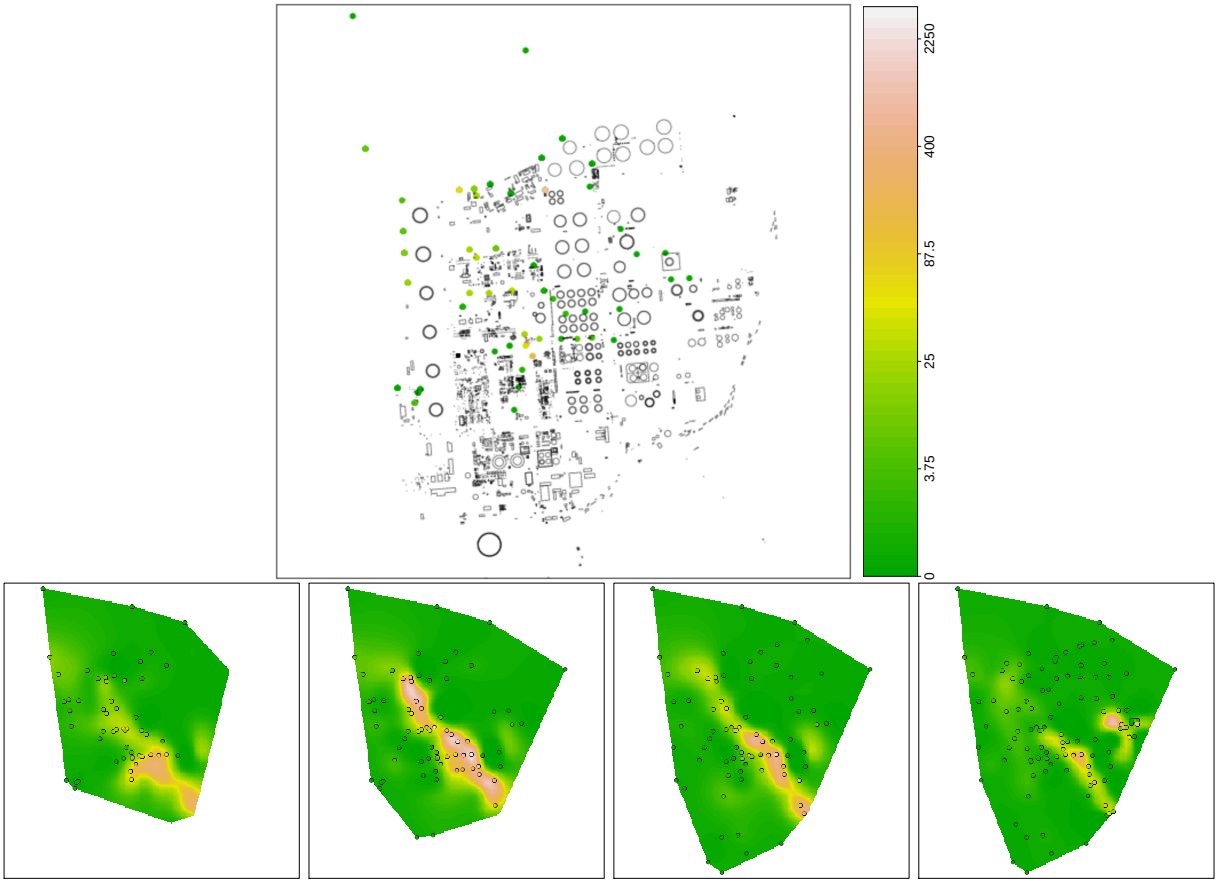

6.4 A case study in brain imaging

A final case study of using flexible models for spatiotemporal data will be discussed in the context of brain imaging. Here a key feature of interest is the periodic nature of the signal but the period changes over time. Flexible models again provide tools to track this kind of behaviour.

7

Extending the concepts

Earlier chapters have discussed a variety of concepts associated with the construction of flexible estimators of regression functions. This provides a rich set of tools which we can use to build models for a wide variety of other data structures. We will briefly look at how this works for several different non-Gaussian and non-Euclidean models, and for cases where the observations are themselves functions.

7.1 Smoothing binary data

Example 7.1 (Ascaris lumbricoides). A survey of the occurence of the human parasitic worm infection *Ascaris lumbricoides* was carried out in residents of a rural community in China. The variables are:

Age	age of the resident
Infection	presence (1) or absence (0) of infection
Sex	male (1) or female (2)

The background to the data, and an analysis, are described by Weidong et al. (1996), *Ascaris, people and pigs in a rural community of Jiangxi province, China*, Parasitology 113, 545-57. ◁

A natural extension to non-normal data is to apply weights to the log-likelihood rather than the sum-of-squares as

$$\sum_i \ell_i(\alpha, \beta)\, w(x_i - x; h),$$

where $\ell_i(\alpha, \beta)$ is the contribution to the usual log-likelihood from the ith observation. For example, in the case of logistic regression for binary responses, we have the model

$$\log\left(\frac{p_i}{1-p_i}\right) = \alpha + \beta x_i \qquad (i = 1, \ldots, n)$$

and then the likelihood contributions

$$\ell_i(\alpha, \beta) = y_i \log\left(\frac{p_i}{1-p_i}\right) + \log(1-p_i).$$

Here p_i denotes the probability of a 1 at design point x_i, and as usual the logit link function is assumed to relate p_i to a linear predictor in x_i. Maximisation of the weighted likelihood with respect to (α, β) provides local estimates $(\hat{\alpha}, \hat{\beta})$ and by solving the model definition for p we obtain a fitted value

$$\hat{m}(x) = \frac{\exp(\hat{\alpha} + \hat{\beta}x)}{1 + \exp(\hat{\alpha} + \hat{\beta}x)}.$$

This sounds a bit complicated but it is really just fitting a glm locally by using weights to pay attention only to data points which are near the point of interest x. The procedure will be illustrated below on the worm data.

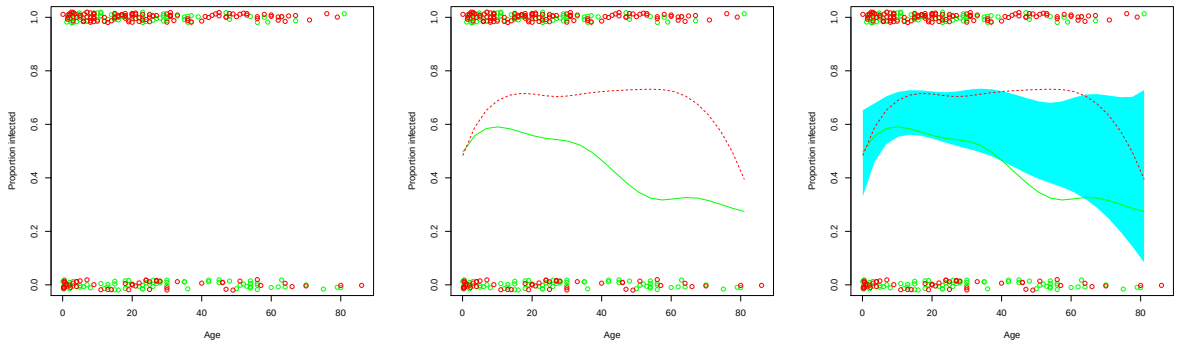

The form of smoothing described above corresponds to the 'local fitting' approach to nonparametric modelling. The penalty approach is also available. Recall from earlier APTS work that a generalised linear model involves a linear predictor plus an error distribution, usually from the exponential family, and that the model can be fitted by maximum likelihood through an iteratively reweighted least squares algorithm. In the nonparametric setting, we can simply replace the linear predictor by a linear term in spline bases and fit the model by a similar algorithm, but this time adding to the log-likelihood a penalty term based on differencing of the spline coefficients. Wood (2006) gives the details of this.

7.2 Generalised additive models

In the mackerel survey above, the data collected by Spanish vessels in the Bay of Biscay exhibit rather different features from the remainder of the survey. One of these features is that no eggs were detected at all at a substantial number of the sampling points. The sampling positions, together with an indication of the presence or absence of eggs, are displayed below. Depth and sea surface temperature are again available as potential covariates.

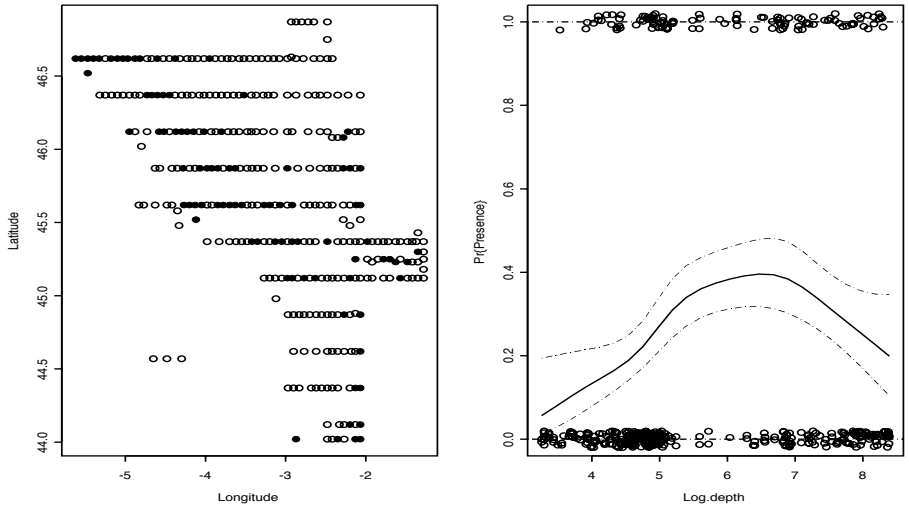

The right panel shows a local logistic regression curve fitted to the relationship between presence and depth, on a log scale. As in the earlier investigation of the density of eggs, log(depth) appears to have an approximately quadratic effect on presence, with an optimal depth around $\exp(6) \approx 400$m.

The framework of generalized linear models is extremely useful in providing a unified description of a very wide class of parametric models, along with methods of fitting and analyses. In the nonparametric setting a corresponding framework of generalised additive models provides a very flexible form of extension.

Excellent overviews of this area are provided in the texts by Hastie and Tibshirani (1990), Green and Silverman (1994) and Wood (2006). A brief introduction to the main ideas is provided in this section, based on an illustration. Another very useful text is Ruppert *et al.* (2003) which discusses *semiparametric regression*.

In the case of logistic regression, a linear model with four covariates takes the form

$$\log\left(\frac{p}{1-p}\right) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_4 x_4,$$

A logistic form of a generalised additive model therefore extends this by replacing each linear component with a nonparametric one.

For the Spanish mackerel data, where the covariates represent latitude, longitude, log(depth) and temperature respectively, a natural model is

$$\log\left(\frac{p}{1-p}\right) = \beta_0 + m_{12}(x_1, x_2) + m_3(x_3) + m_4(x_4),$$

since latitude and longitude merely define a convenient two-dimensional co-ordinate system.

In order to fit a generalised additive model the corresponding fitting procedure for generalised linear models again provides a helpful guide. The likelihood function is the natural starting point and a Newton-Raphson, or Fisher scoring, procedure allows the parameter estimates to be located by an iterative algorithm. Each step of these algorithms can be formulated

as a weighted least squares linear regression. This immediately provides provides a natural analogue in the generalised additive model setting, by employing a weighted nonparametric regression at each step.

Hastie & Tibshirani (1990), Green & Silverman (1994) and Wood (2006) provide detailed discussion of this approach.

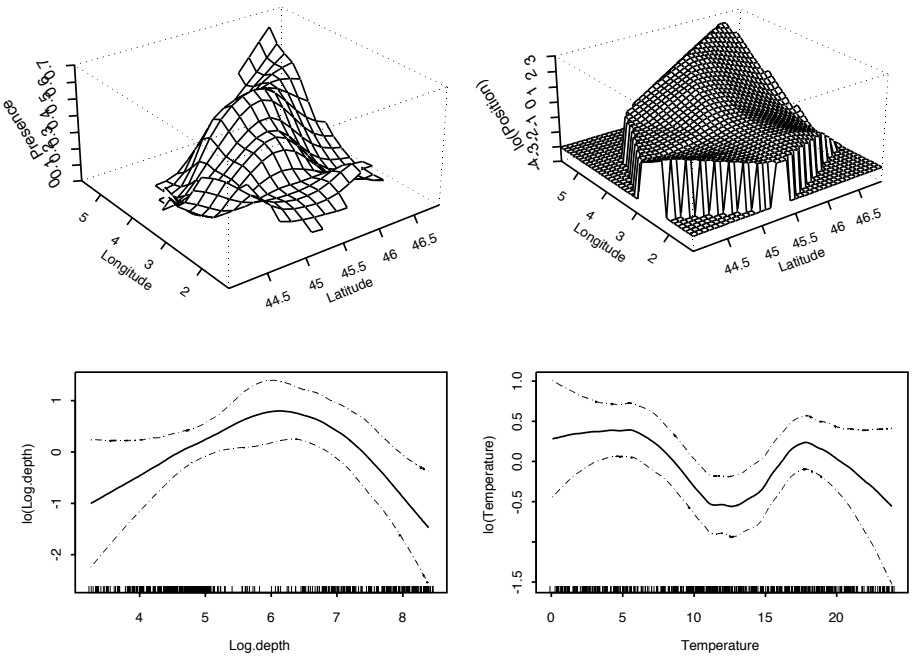

The expected concave effect of log(depth) is apparent. A more complex non-linear curve describes the effect of temperature and a two-dimensional function in latitude and longitude captures the additional spatial variation.

In order to explore which of these terms can be regarded as evidence of an underlying systematic effect, rather than random variation, different models for the data can be compared. A deviance can also be defined in an analogous manner for a generalised additive model. As with additive models for data with normal responses, general distribution theory for model comparisons is not available, even asymptotically. However, by applying suitable quadratic approximations, degrees of freedom can be associated with each model, and so some guidance on model comparisons can be taken by comparing differences in deviance to χ^2 distributions indexed by the difference in the approximate degrees of freedom.

For the Spanish mackerel data, the deviances for a number of models of interest are shown below.

Model	Deviance	df
1: $\beta_0 + m_{12}(\text{Lat, Long}) + m_3(\log(\text{Depth})) + m_4(\text{Temp})$	384.22	401.30
2: $\beta_0 + m_{12}(\text{Lat, Long}) + m_4(\text{Temp})$	394.32	404.75
3: $\beta_0 + m_{12}(\text{Lat, Long}) + m_3(\log(\text{Depth}))$	398.43	404.68
4: $m_3(\text{Depth}) + m_4(\text{Temp})$	431.95	409.20

An analysis of this type suggests that all of the three components in model 1 contribute significantly. The temperature effect looks rather implausible from a scientific point of view and should perhaps be treated with some caution. Hastie and Tibshirani (1986, section 6.10) provide some discussion on the dangers of over-interpreting additive fits in models of this kind.

7.3 Functional data analysis

The type of data which are now routinely collected can have quite complex structures, rather than simply having a single measurements of a response variable. For example, a response might be in the form of a function collected by a monitoring device which effectively collects data continuously over time. Although in practice the data may be discretised on a grid of time points, it can be helpful to think of this as representing a function. This leads to the concept of *functional data analysis* which has attracted considerable interest over the last couple of decades. There are strong links here with the techniques we have been discussing, as methods of flexible regression provide curve descriptions which reduce noise or have compact representations through basis functions.

The left hand plot below gives an example of daily temperature measurements from 35 different locations in Canada, averaged over 1960 to 1994. These data are provided in the `fda` packaged associated with the book on FDA by Ramsay and Silverman (1997). (This example will therefore give the opportunity to experiment further with the code provided in the package.) Each curve can be regarded as an observation of the seasonal pattern of temperature from a particular location.

When we deal with observations of multivariate data Y, a vector of length p, principal components allow us to explore the structure by creating a set of new low-dimensional variables which aim to capture as much of the variation as possible. These new variables are constructed as linear combinations of the original variables. Specifically, we create a new variable $z_1 = \sum_{i=1}^{p} \beta_j Y_j$ which maximises the variance, Var $(Z_1) = \beta^T \Sigma \beta$, where Σ is the covariance matrix of Y. To make the problem identifiable, we require that $||\beta|| = 1$.

The functional analogue is to consider Y and β as functions $Y(s)$ and $\beta(s)$, where s is a continuous index. An elegant body of theory based on orthonormal eigenfunctions, and associated eigenvalues, allow principal components to be implemented in functional form. As a simple indication of what this produces, the right hand panel of the plot below shows the mean temperature function with dotted lines to indicate the deviations associated with one standard deviation, corresponding to the first functional principal component. This expresses the fact that variation is most pronounced in the winter and less pronounced in the summer.

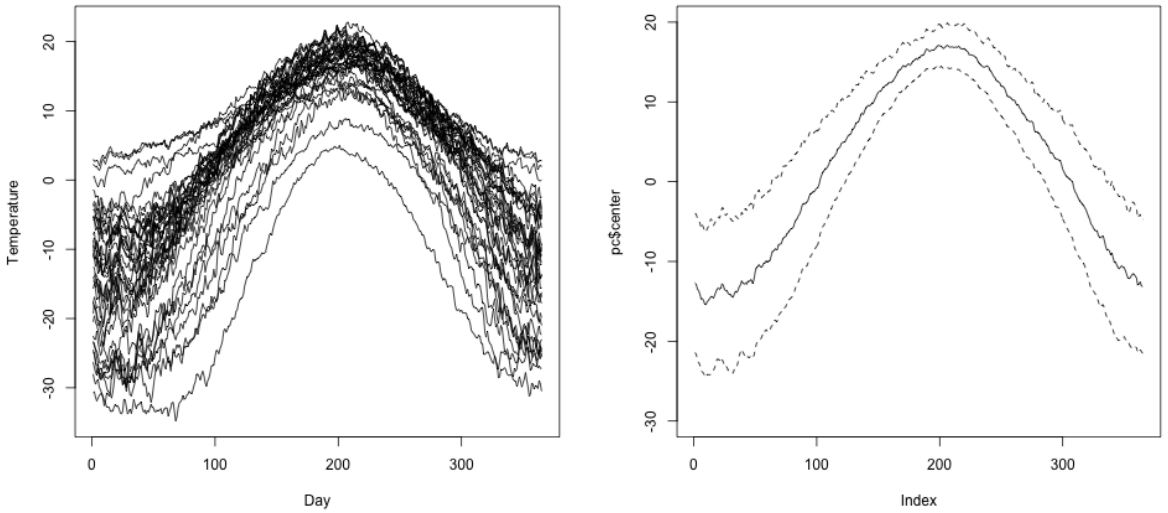

Functional versions of regression, with the response and/or the explanatory variable may be functions, and of various other types of statistical analysis can be constructed in this way. The book by Ramsay and Silverman (1997) and the associated `fda` package for `R` are good starting points to explore this.

7.4 Quantile regression

Regression models usually target the mean function. However, there can be interest in other characteristics of the distribution of the data. In particular, quantiles of the conditional distribution of the response variable given a covariate value $(Y|x)$ can be useful, particularly in settings such as growth curves, where interest can lie in quantiles which characterise the behaviour of more extreme cases.

The idea of quantile regression was introduced by Koenker and Bassett (1978). Suppose we have data $\{(y_1, x_1), \ldots, (y_n, x_n)\}$ and a predictor function $\nu(x, \alpha)$ for a set of parameters α. For example, ν may simply be a linear function $\alpha_1 + \alpha_2 x$. If we choose α to minimise the criterion

$$\sum_{i=1}^{n} \{y_i - \nu(x_i, \alpha)\}^2$$

then this provides an estimate of the mean function. However, if we choose α to minimise the criterion

$$\sum_{i=1}^{n} \rho_\theta(y_i - \nu(x_i, \alpha))$$

where ρ is the *check function* defined as

$$\rho_\theta(z) = \begin{cases} \theta z, & \text{if } z \geq 0 \\ (\theta - 1)z, & \text{if } z < 0 \end{cases}$$

then this provides an estimate of the θth quantile of the distribution of y as a function of x. In the case of the median, when $\theta = 0.5$, the check function becomes the absolute value function (or L_1 norm). For other values of θ, the check function is asymmetric.

If, instead of using a simple parametric function, we use a b-spline basis, $\nu(x, \alpha) = \sum_{j=1}^{p} \alpha_j \phi_j(x)$ then we immediately create a flexible regression function. If required, a difference penalty on α can be used to control the smoothness. This approach is described by Bollaerts *et al.* (2006). There are several other ways in which flexible estimators could be produced, using the other principles outlined earlier in the course.

A simple example is provided by the `quantregGrowth` and `quantreg` packages from `R`, used with some simulated data provided in the first of these packages. The two graphs below display the deciles estimated by linear and flexible models.

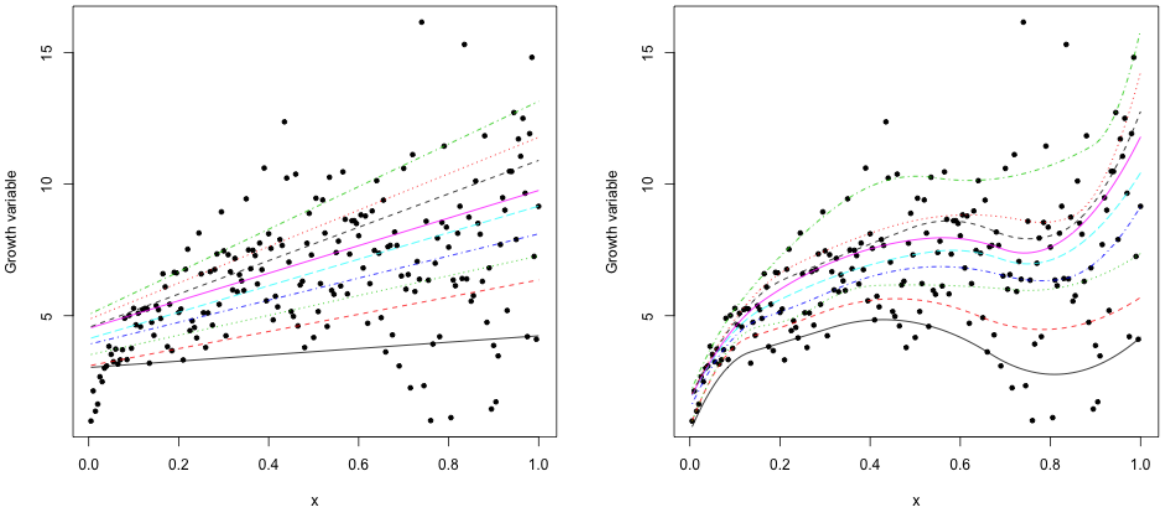

7.5 River networks

A final case study of flexible regression on data collected over a river network will be provided. The picture below shows measurements of nitrate pollution in the River Tweed. The challenge is to construct a model which allows the pollution over the whole river system to be produced, while respecting the network nature of the data.

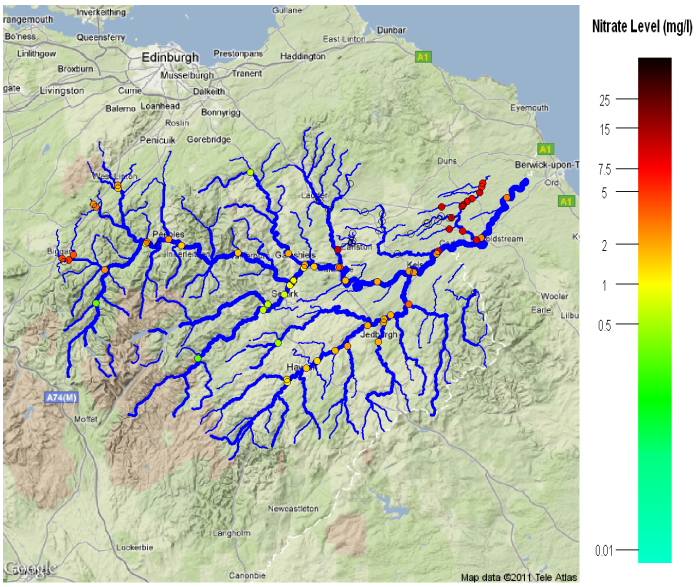

References

Bollaerts, K., Eilers, P. H., and Aerts, M. (2006). Quantile regression with monotonicity restrictions using p-splines and the l1-norm. *Statistical Modelling 6*(3), 189–207.

Cressie, N. and Wikle, C. K. (2011). *Statistics for spatio-temporal data*. John Wiley & Sons.

Green, P. and Silverman, B. (1994). *Nonparametric Regression and Generalized Linear Models*. Chapman and Hall.

Hastie, T. and Tibshirani, R. (1986). Generalized additive models (with discussion). *Statistical Science 1*, 297–318.

Hastie, T. and Tibshirani, R. (1990). *Generalized Additive Models*. London: Chapman and Hall.

Koenker, R. and Bassett, G. (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, 33–50.

Ramsay, J. and Silverman, B. (1997). *Functional Data Analysis. 1997*. Springer, New York.

Ruppert, D., Wand, M. P., and Carroll, R. (2003). *Semiparametric regression*. London: Cambridge University Press.

Scott, D. (1992). *Multivariate Density Estimation: Theory, Practice, and Visualization*. New York: John Wiley.

Silverman, B. (1986). *Density estimation for statistics and data analysis*. London: Chapman and Halland Hall.

Simonoff, J. S. (1996). *Smoothing methods in statistics*. New York: Springer.

Wand, M. P. and Jones, M. C. (1995). *Kernel smoothing*. London: Chapman and Hall.

Wood, S. (2006). *Generalized Additive Models: an introduction with R*. London: Chapman and Hall/CRC.