Advanced Mathematical and Computational Tools in Metrology V
Edited by P Ciarlini, M G Cox, E Filipe, F Pavese & D Richter

DISCRETE B-SPLINE APPROXIMATION IN A VARIETY OF NORMS

I. J. ANDERSON, D. A. TURNER
School of Computing and Mathematics,
University of Huddersfield, Huddersfield, HD1 3DH, UK
Email: i.j.anderson@hud.ac.uk
d.a.turner@hud.ac.uk

B-splines can be used to approximate data in a variety of metrology applications. In general, the approximation is carried out so as to minimize a criterion defined in terms of the least-squares norm. However, often this choice of norm is inappropriate, and a measure such as the ℓ_1 or minimax norm should be adopted. In this paper, we introduce the topic of B-splines, and we discuss the approximation of data in the ℓ_1, least-squares and minimax norms. In particular, we consider how the specific properties of B-splines can be exploited to allow an approximant to be developed in an efficient manner, and we describe a new algorithm for ℓ_1 B-spline approximation.

1 Introduction

In many metrology applications, it is desirable to replace a prescribed set of data by an approximating curve or surface. Broadly speaking, the approximant can take one of two forms. First, if the data are obtained from a process for which a set of physical relationships is known, then the data can be approximated by a *physical* model. More frequently, the model equations governing the behaviour of the model are unknown, and so an *empirical* model of the data is developed. There are a number of different types of empirical model, although polynomials, splines and radial basis functions (RBFs) are used most frequently in metrology. These forms can in fact be represented in the same general linear form, which for curves is written as

$$f(x) = \sum_j c_j \phi_j(x).$$

The set of functions $\{\phi_j(x)\}$ are termed the basis functions, and their particular form defines the characteristics of the curve. The set of coefficients $\{c_j\}$ define the actual shape of the curve.

Each of the empirical models mentioned above have properties that make them useful in metrology applications. For example, polynomial models are simple to develop, while RBFs are ideally suited for multivariate approximation. However, the flexibility of spline curves means that they are most useful in metrology, especially if we wish to approximate two or three-dimensional data. Spline curves are also widely used in metrology in the solution of differential equations.

In this paper, we discuss spline approximation and consider the role that it can play in metrology. The paper is organized as follows. Section 2 introduces the topic of splines, in particular B-splines, and in Section 3 we consider some applications of B-splines in metrology. In Section 4, we discuss in detail algorithms for developing B-spline approximations to data. In particular, we emphasise that by exploiting the specific structure of the approximation problem, it is possible to make substantial computational savings. These ideas are extended in Section 5, where we show that the B-spline approximation can be obtained by using a piecewise approach. In this method, the global approximation problem is reduced to a set of smaller, local approximation problems. Consequently, the storage and operational requirements are reduced significantly. In particular, we present a new algorithm for ℓ_1 B-spline approximation. Finally, we give our conclusions in Section 6.

2 Splines

The simplicity and effectiveness of polynomials have made them popular for representing functions or data. However, polynomials have limitations, such as the large oscillatory behaviour that may be exhibited by high degree polynomials. This effect is known as the Runge phenomenon.

An alternative to the family of polynomials is needed, and in many respects this need is satisfied by spline functions. A spline function is a *piecewise polynomial* that has continuity conditions imposed between neighbouring sections. The point at which two sections meet is called a *knot*. Thus, a spline has the same continuity conditions as a polynomial except at the knots, where the continuity is based on the degree of the polynomial pieces used to construct the spline.

For polynomials, more shape variation is generally achieved by using a higher degree. With splines, a similar affect can be obtained by having more sections, and so it is not necessary to use high degree polynomial pieces. In fact, most spline applications use cubic polynomials throughout. This avoids any unnecessary oscillatory behaviour that might otherwise occur. Indeed, a pleasing property of splines is their smooth behaviour. The cubic spline $s(x)$ is the twice continuously differentiable function that interpolates a set of data (x_i, f_i) in such a way as to minimize uniquely the curvature norm [1]

$$\int_{x_{\min}}^{x_{\max}} \left(f''(x)\right)^2 dx.$$

2.1 B-spline Representation of a Spline

In order to represent a spline, we need to define (a) a form for the representation, and (b) the parameters or coefficients which are used to distinguish one spline from another.

As with polynomials, there are several different forms available. In the case of splines, the best form is the B-spline basis function representation, which was introduced in 1946 by Schoenberg[2] for uniformly spaced knots. Later, they were generalized to arbitrary knot spacings. B-splines were originally defined to be the nth divided difference of a truncated power function of order n (degree $n-1$), $(x - \cdot)_+^{n-1}$. Thus,

$$M_{n,j}(x) = [\lambda_{j-n}, \ldots, \lambda_j](x - \cdot)_+^{n-1},$$

where $\{\lambda_j\}$ is a set of non-decreasing knots,

$$\lambda_{1-n} = \cdots = \lambda_0 < \lambda_1 \le \lambda_2 \le \cdots \lambda_{q-n} < \lambda_{q-n+1} = \cdots = \lambda_q,$$

with $q+n$ being the number of knots, and $x_{\min} = \lambda_0$ and $x_{\max} = \lambda_{q-n+1}$ being the endpoints of the interval over which the spline is defined.

Figure 1 shows a single cubic B-spline (order $n = 4$), and also all of the other cubic B-splines that are used in the representation. The vertical lines represent the positions of the knots. Clearly, for $x < \lambda_{j-n}$, the B-spline is zero and it is well-known that nth divided differences annihilate polynomials of order n or less. Hence, it follows that for $x > \lambda_j$, the B-spline is also zero and so the B-spline has *compact support.*

Unfortunately, the use of divided differences is numerically ill-conditioned and so this method of evaluation is no longer used. Instead we use the stable recurrence relation for *normalized* B-splines which is defined to be

$$N_{n,j}(x) = \left(\frac{x - \lambda_{j-n}}{\lambda_{j-1} - \lambda_{j-n}}\right) N_{n-1,j-1}(x) + \left(\frac{\lambda_j - x}{\lambda_j - \lambda_{j-n+1}}\right) N_{n-1,j}(x), \quad (1)$$

with

$$N_{1,j}(x) = \begin{cases} 1, & \text{for } x \in (\lambda_{j-1}, \lambda_j], \\ 0, & \text{otherwise.} \end{cases}$$

This stable three term recurrence relation (or a variation of it) was introduced independently by Cox, de Boor and Mansfield.[3,4] The standard B-spline and the normalized B-spline are related by the scaling factor

$$N_{n,j}(x) = (\lambda_j - \lambda_{j-n}) M_{n,j}(x).$$

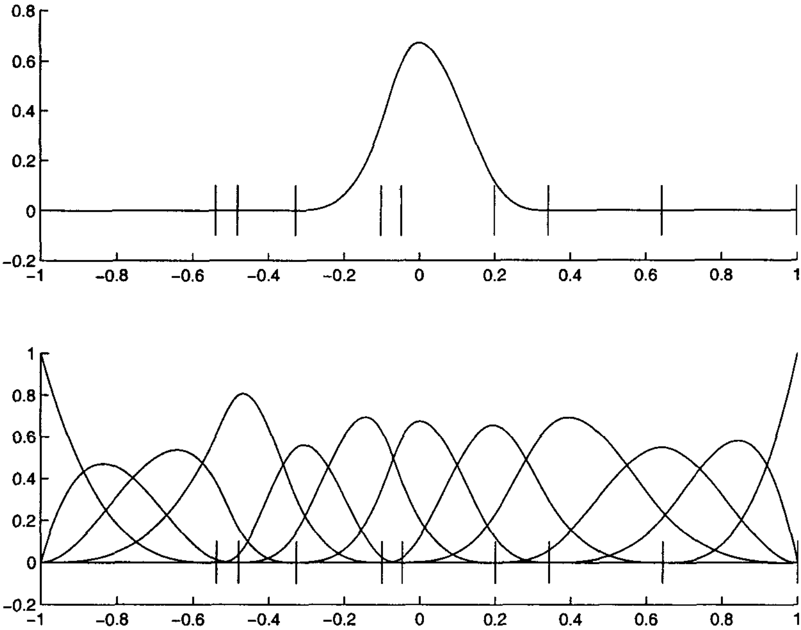

Figure 1: B-splines of order 4.

Now that we have a set of basis functions it is possible to express the spline function as a linear combination of them, namely

$$s(x) = \sum_{j=1}^{q} c_j N_{n,j}(x), \qquad x \in [x_{\min}, x_{\max}],$$

where $\mathbf{c} = \{c_j\}$ are the B-spline parameters or coefficients. In the literature, the coefficients are sometimes referred to as the control polygon. This phrase is particularly popular in computer aided design (CAD) and computer aided geometric design (CAGD) applications.

Because of the compact support property of the B-spline, all but n of these basis functions will be zero for a given value of x. Hence, we have

$$s(x) = \sum_{j=\jmath+1}^{\jmath+n} c_j N_{n,j}(x), \quad \text{for } x \in (\lambda_{\jmath-1}, \lambda_\jmath].$$

In order to evaluate the spline at a given abscissa value, we begin by considering that this abscissa lies between two knots. We find these two knots and then use the recurrence relation to construct the n non-zero B-spline basis functions. Finally, we combine these values with the corresponding coefficients.

If we have several data points, we calculate the vector of spline values $\mathbf{s}$ by repeating this process for each point in turn. Note that we can express $\mathbf{s}$ mathematically as the matrix-vector product of an observation matrix A and a vector of the spline coefficients $\mathbf{c}$. The matrix A is an $m \times q$ matrix whose elements correspond to the appropriate basis function value, so that

$$A = \{N_{n,j}(x_i)\}.$$

Due to the compact support of B-splines, at most n elements of each row of A are non-zero. Furthermore, if the abscissae values are arranged such that

$$x_1 \leq x_2 \leq \cdots \leq x_m,$$

then A is a banded matrix of bandwidth n. An example of the typical structure of the matrix A is given in Figure 2.

3 Metrology Applications

A key question that arises is the type of metrology applications for which B-spline curves and surfaces are appropriate, and we address this issue in this section. In general, B-spline curves and surfaces are developed for metrology applications in one of two ways. First, a user chooses the knots and the spline coefficients manually in order to generate a desired shape, such as in CAD modelling. Second, we use a data approximation technique in which the knots and the coefficients are calculated so as to provide a good representation of the data set. A typical application of this second approach occurs when we wish to develop a CAD model which can be used as the design basis for the manufacture of an artifact. B-spline curves and surfaces are particularly useful in this application, since they offer the user a large amount of flexibility. Furthermore, the compact support of B-splines allows the user to alter a local region of the CAD model without altering its global shape. It should be noted that often parametric rather than explicit B-splines are used in this situation. See Dierckx[5] for a brief introduction to parametric B-splines.

The second situation occurs frequently in metrology. As stated in Section 1, B-splines can be used to represent a set of data by an empirical model. There are a wide variety of metrology applications for which B-spline approximations are regularly used. For example, we may wish to model the voltage response of a fast-time oscilloscope to a step input,[6] or approximate a set of data measured by a coordinate measuring machine (CMM) in the surface of

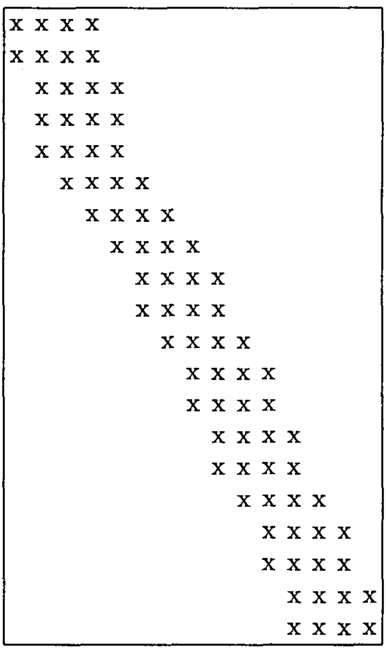

Figure 2: An example of the typical structure of an observation matrix for cubic B-spline approximation.

a clay model of a tooth.[7] In particular, B-splines have proven to be successful in modelling phenomena which exhibit complicated responses. For example, B-spline curves and surfaces are able to cope with oscillatory behaviour, non-uniform behaviour, and derivative discontinuities.

4 Approximation

There are many different ways of approximating data in metrology applications. Here we consider explicit univariate spline fitting which is the most straightforward type of approach. The various other methods of approximation are generalizations of this basic type. See Cox[6] and Dierckx[5] for example.

In the univariate spline approximation problem, we are given a set of m data, $(x_i, f_i)_{i=1}^{m}$. We then wish to calculate the spline curve $s(x)$ which minimizes in some norm the residual differences between the data ordinates $\mathbf{f} = \{f_i\}$ and the value of the spline evaluated at the data abscissae, namely $\{s(x_i)\}$. Thus, we wish to compute the spline coefficients $\mathbf{c}$ which minimize the expression

$$||\mathbf{r}|| = ||\mathbf{f} - A\mathbf{c}||,$$

where the ith component of $\mathbf{r}$ is defined to be $r_i = f_i - s(x_i)$, and the matrix A is defined in Section 2.1.

4.1 Least-squares Approximation

The choice of norm should be based on the nature of the data. For example, if the residuals are known to be drawn from a Gaussian distribution, then the ℓ_2 or least-squares norm should be used. The approximation problem then involves minimizing

$$||\mathbf{r}|| = \sum_{i=1}^{m} r_i^2. \tag{2}$$

Gaussian residuals occur frequently in metrology applications due to natural statistical phenomena. Consequently, least-squares approximation is extremely popular in metrology and several efficient algorithms have been developed for solving this problem. Many of these algorithms involve constructing and solving the *normal equations*, specifically

$$A^{\mathrm{T}}A\mathbf{c} = A^{\mathrm{T}}\mathbf{f}.$$

This system of equations is appealing because it is a square system, and so it can be solved by using elementary techniques such as Gaussian elimination. However, this approach cannot be recommended, since the condition number of the matrix $A^{\mathrm{T}}A$ is dependent upon the square of the condition number of A.[8] As a result, any ill-conditioning in A is exacerbated if an approach involving the normal equations is used to minimize expression (2).

A more numerically stable approach to minimizing expression (2) is to employ the QR factorization.[9] In particular, it is possible to exploit the banded structure of the matrix A by using the techniques described by Cox.[10] In this approach, orthogonal factorizations are applied in a structured manner in order to calculate the coefficients $\mathbf{c}$. These orthogonal methods are both efficient and numerically stable. In fact, by exploiting the banded structure of the matrix A, it is possible to compute the least-squares approximation in $\mathcal{O}(mn^2)$ floating-point operations (flops). However, if QR factorization is used to minimize expression (2) without exploiting any structure, then $\mathcal{O}(mq^2)$ flops are required.

4.2 Approximation in Other Norms

There are many situations that arise in metrology for which it is inappropriate to approximate a set of data in the least-squares norm. Two such examples are quality control and robust approximation. In quality control, we are interested in tolerance assessment, i.e., we wish to assess whether or not an object such as a manufactured artifact is fit for purpose. In this case, it is widely accepted within metrology that the residuals should be minimized in the minimax norm,[11,12] i.e.,

$$\|\mathbf{r}\| = \max_i |r_i|.$$

In robust approximation, we are interested in finding approximations that are able to cope well with presence of outliers. These outliers are often produced from non-Gaussian error sources. One such example of this type of noise might be a speck of dust on the surface of an artifact which is being measured by a high precision instrument. The ℓ_1 norm is appropriate for such data sets,[13] and so we minimize

$$\|\mathbf{r}\| = \sum_{i=1}^{m} |r_i|.$$

Figure 3 shows the linear B-spline fits to a set of data in the ℓ_1, ℓ_2 and ℓ_∞ norms. The data represent a cross-section of a surface with grooves in it and are represented as a piecewise linear interpolant for clarity. In this figure, the ℓ_1 fit is shown as the solid line, the ℓ_2 fit is shown as the dashed line, and the ℓ_∞ fit is shown as the dashed-dotted line. It is apparent from Figure 3 that the ℓ_1 solution is least affected by the grooves in the data. This is because the ℓ_1 norm treats the data corresponding to the grooves as outliers, and consequently these data have only a small effect on the ℓ_1 approximation.

Both the ℓ_1 and the ℓ_∞ spline approximation problems can be represented as linear programming problems, and hence they can be solved by using simplex-based linear algebra techniques. Two particular popular simplex methods are the Barrodale-Roberts[14] and Barrodale-Philips[15] algorithms for ℓ_1 and ℓ_∞ approximation respectively.

Unfortunately, simplex-based methods of optimization tend to be numerically unstable. Small perturbations to the matrix A can result in large differences in the computed coefficients $\mathbf{c}$. The reason for these numerical instabilities is that simplex methods use pivoting strategies that can lead to large cancellation errors. In some cases, the resulting spline can bear little or no relation to the original data due to the accumulation of rounding errors, and this appears to be particularly common when using a B-spline basis and the ℓ_∞ norm.

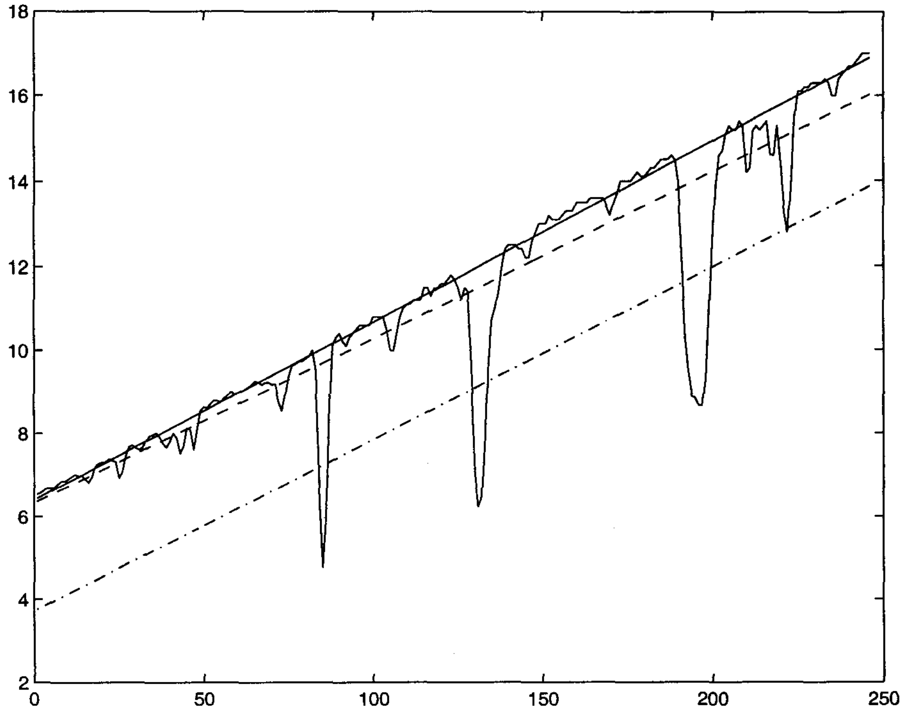

Figure 3: Linear B-spline fits in the ℓ_1, ℓ_2 and ℓ_∞ norms to a set of data.

A further disadvantage of standard simplex methods is that they are unable to exploit the banded structure and general sparseness of the observation matrix A. This can be a major disadvantage in metrology, since often approximation problems arise for which the number of data is very large. Furthermore, in these problems there is frequently significant variation of shape over the domain of the data, and so there is a need for many polynomial pieces within the spline. Thus, it is usual to have

$$n \ll q \ll m.$$

For such data sets, obtaining an ℓ_1 or ℓ_∞ approximation to a set of data can be much more computationally expensive than obtaining a least-squares approximation. For example, in the previous section we commented that a least-squares B-spline fit to a set of data can be computed in $\mathcal{O}(mn^2)$ flops. In contrast, the flop count is at least $\mathcal{O}(mq^2)$ for ℓ_1 approximation. Consequently, it is sometimes common for the least-squares method to be used purely because of speed even though the method is inappropriate for the data.

5 Piecewise Methods for Efficient Spline Approximation

In this section, we address the issue of efficient spline approximation. In particular, we discuss how the compact support of B-splines can be exploited to achieve an efficient algorithm for computing a least-squares approximant to data. We then show how this algorithm can be extended to solving approximation problems in the ℓ_1 norm.

The compact support property of the B-spline basis functions suggests that the function values of data points that lie at one end of the domain should have little influence over the shape and quality of the fit at the opposite end of the domain. In fact, for the case of univariate splines it has been shown that a B-spline basis is locally stable.[16] This means that each coefficient of the fit is bounded by a multiple of the maximum absolute value of the spline function in a region contained within the support of the B-spline corresponding to the coefficient. The value of this multiple increases as the order of the spline increases. Furthermore, the multiple is dependent only on the order of the spline being used. Thus, most of the information about the B-spline fit in a subregion of the domain can be obtained from the data in and around the subregion, and therefore we should be able to calculate good local approximations.

The idea of a piecewise method for approximating a set of data is that rather than trying to compute all of the coefficients $\mathbf{c}$ in one step, we split the array of coefficients into several smaller subarrays and compute the coefficients in one subarray independently from the other subarrays. Let $\mathbf{c}'$ be the subarray of coefficients that we wish to calculate by using a local approximation. Since each coefficient corresponds to a column of A, we need only use the columns of A corresponding to the relevant coefficients in order to compute $\mathbf{c}'$. Let this submatrix be A'. The local approximation then involves minimizing $\|\mathbf{r} - A'\mathbf{c}'\|$ rather than $\|\mathbf{f} - A\mathbf{c}\|$, where $\mathbf{r} = \mathbf{f} - A\mathbf{c}$ are the residuals for the current spline approximation. This local approximation method is then applied in an iterative manner to the various coefficient subarrays in turn. The value of $\mathbf{c}'$ represents the update to the current estimate for the spline coefficients $\mathbf{c}$.

The piecewise method can be applied to any locally stable basis functions. However, the B-spline basis has another advantage which is due to the compact support property. For a particular subarray of coefficients, only a limited number of data will lie within the support of the B-splines corresponding to these coefficients, while the remaining data lie outside the support. Consequently, these latter points do not contribute to the local approximation problem and so the corresponding rows can be removed from A'. Figure 4 shows an example of an observation matrix A with $m = 20$ and $q = 15$ together with the corresponding submatrices. The coefficient array has been divided into three subarrays with 5 elements in each. Notice that for each of the three submatrices, the rows that have been ignored correspond to rows of A' which contained only zero elements. As an example, consider the subarray corresponding to the first five coefficients. From the figure we note that the first eight data correspond to points that lie within the support of the first five B-spline basis functions, and that the remaining twelve points lie outside the support. Consequently, the submatrix A' has only eight rows and is therefore a fraction of the size of the original matrix A.

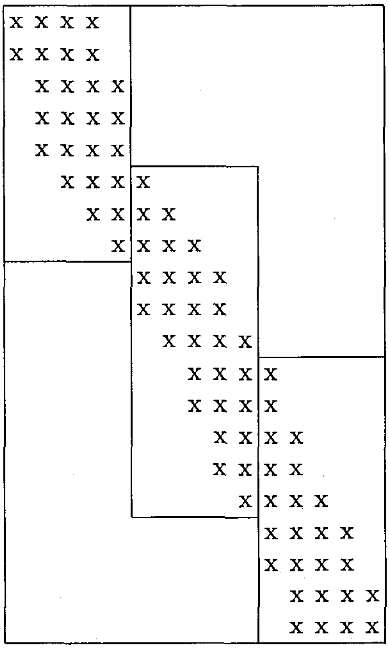

Figure 4: A typical example of an observation matrix for cubic B-spline approximation and its corresponding submatrices.

The local stability of the B-spline basis implies that these local approximations are likely to produce coefficient values which are similar to the values obtained from the global approach. Anderson[17] describes this piecewise algorithm in more detail and proves that for the ℓ_2 approximation problem, the algorithm converges to the global solution of the approximation problem.

5.1 *Piecewise Approximation in the ℓ_1 Norm*

Recent work has concentrated on extending the piecewise approach to the ℓ_1 norm. In principle, there appears to be no difficulty in doing this since each iteration is sufficiently generic and the local approximation to each subset of data merely involves minimizing a norm. We have not specified what the norm must be, nor how the local approximation should be implemented, and so we should be able to use the Barrodale-Roberts algorithm to solve the local approximations.

The main problem with such a technique is that it does not, in general, converge to the best ℓ_1 approximation. In order to appreciate why this problem arises, it is necessary to understand that the best ℓ_1 approximant to the data will interpolate q data points, where q is the number of basis functions. Similarly, given that on a particular subinterval we are working with q' basis functions ($q' < q$), it is reasonable to expect that a local approximation for this subinterval would interpolate only q' points. However, in practice it is generally the case that there are more than q' interpolation points. This is due to the fact that the region corresponding to the support of the basis functions used in one local approximation overlaps slightly with the neighbouring region from another local approximation. Consequently, any points of interpolation that lie within this overlap region will be common interpolation points for both local approximations. This results in a non-unique local ℓ_1 approximation problem, and a tiny perturbation to the residuals produces a slightly different local approximant. The iterative method either gets trapped in a cycle of approximation estimates which have the same error norm, or converges to a single estimate which is not optimal.

Fortunately, the problem of non-uniqueness only arises when there are more than q' interpolation points for a local approximation. This only happens when the current spline estimate successfully interpolates many of the q points that are necessary to characterize the global approximant. Thus, problems with the piecewise technique only arise when we are very near to the required approximant. Indeed, the first few steps of the method perform exceedingly well by reducing the norm of the errors significantly. Moreover, the piecewise technique obtains an approximation in $\mathcal{O}(m(q')^2)$ flops. Although this approximation is not the best ℓ_1 solution, it is still useful in that it can be used as a starting value for a global method. For example, we can employ the standard Barrodale-Roberts algorithm,[14] or alternatively a more stable version such as that described by Lei et al.[18]. In general, these global approaches converge to the ℓ_1 solution from the piecewise solution in substantially fewer iterations than from the ℓ_2 solution, for example.

The piecewise ℓ_1 method for B-spline approximation is illustrated in Figure 5, where we apply one full iteration of the piecewise method to fit a cubic B-spline to a set of response data. This data was kindly supplied to us by National Physical Laboratory (NPL), Teddington, UK. The data set is divided into three subregions, and the figure shows how an approximation is constructed by fitting a local spline curve to each sub-region in turn. We note that we achieve a reasonable approximation after only one full iteration, although further improvement to the quality of the approximation can be made by performing additional iterations.

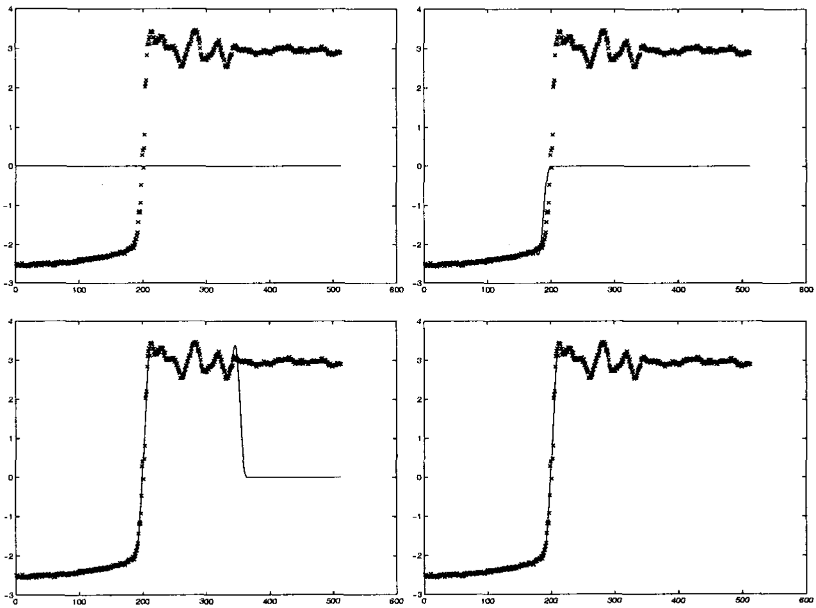

Figure 5: One full iteration of the piecewise ℓ_1 method for approximating a set of response data by a cubic B-spline.

6 Conclusions

B-splines are probably the most useful type of curve and surface for solving the empirical approximation problems that occur in metrology. They are relatively simple to develop, while their flexibility allows a metrologist to model accurately a set of data that represents complex physical behaviour.

In this paper, we have introduced the topic of B-spline approximation, with an emphasis on fitting univariate B-splines to data. This is because all other types of B-spline approximation techniques are generalizations of the univariate approximation methods. In metrology, most approximations to data are computed in the least-squares norm. However, often this choice of norm is inappropriate, and in particular we may be required to approximate the data by using an ℓ_1 or an ℓ_∞ fit. Therefore, we have considered standard algorithms for B-spline approximation in each of these choices of norm.

A key property of B-splines is that of compact support, and in B-spline approximation this results in a banded observation matrix. In least-squares approximation techniques, the banded structure of the observation matrix can be exploited to allow substantial computational savings. However, in both the ℓ_1 and ℓ_∞ norms, it is not possible to exploit directly the banded structure. Therefore, in this paper we have considered a different type of method for exploiting the compact support of B-splines, namely that of piecewise methods. In this approach, the data are divided into a number of groups, and local approximations are computed to each of these groups. The global approximation can then be constructed from these local approximations. We initially discussed a piecewise method in the least-squares norm, before presenting a new piecewise method for ℓ_1 approximation. Future research will involve the refinement of this algorithm so that it reliably obtains the best ℓ_1 approximation to a set of data, and the extension of the piecewise approach to ℓ_∞ approximation.

References

1. C. de Boor. *A Practical Guide to Splines.* Springer-Verlag, New York, NY, US, 1978.
2. I. J. Schoenberg. Contributions to the problem of approximation of equidistant data by analytic functions. *Qu. Appl. Math.*, 4:45–99 and 112–141, 1946.
3. M. G. Cox. The numerical evaluation of B-splines. *Journal of the Institute of Mathematics and its Applications*, 10:134–149, 1972.
4. C. de Boor. On calculating with B-splines. *Journal of Approximation Theory*, 6:50–62, 1972.
5. P. Dierckx. *Curve and Surface Fitting with Splines.* Oxford University Press, Oxford, UK, 1993.
6. M. G. Cox. Survey of numerical methods and metrology applications: Discrete processes. In P. Ciarlini, M. G. Cox, F. Pavese, and D. Richter,

editors, *Advanced Mathematical Tools in Metrology*, pages 1–21. World Scientific, Singapore, 1993.

7. V. Jovanovski, L. Zou, W. M. Tay, E. Lynch, and M. G. Cox. Assessment of three-dimensional structures in clinical dentistry. In P. Ciarlini, M. G. Cox, F. Pavese, and D. Richter, editors, *Advanced Mathematical Tools in Metrology III*, pages 263–265. World Scientific, Singapore, 1993.
8. Å. Björck. *Numerical Methods for Least Squares Problems.* SIAM, Philadelphia, PA, US, 1996.
9. N. J. Higham. *Accuracy and Stability of Numerical Algorithms.* SIAM, Philadelphia, PA, USA, 1996.
10. M. G. Cox. The least squares solution of overdetermined linear equations having band or augmented band structure. *IMA Journal of Numerical Analysis*, 1:3–22, 1981.
11. B. P. Butler, A. B. Forbes, and P. M. Harris. Algorithms for geometric tolerance assessment. Technical Report DITC 228/94, National Physical Laboratory, Teddington, UK, 1994.
12. B. P. Butler, A. B. Forbes, and P. M. Harris. Geometric tolerance assessment problems. In P. Ciarlini, M. G. Cox, F. Pavese, and D. Richter, editors, *Advanced Mathematical Tools in Metrology*, pages 95–104. World Scientific, Singapore, 1994.
13. G. A. Watson. Some robust methods for fitting parametrically defined curves and surfaces to measured data. In P. Ciarlini, A. B. Forbes, F. Pavese, and D. Richter, editors, *Advanced Mathematical Tools in Metrology IV*, pages 256–273. World Scientific, Singapore, 1997.
14. I. Barrodale and F. D. K. Roberts. An improved algorithm for discrete ℓ_1 linear approximation. *SIAM Journal of Numerical Analysis*, 10(5):839–848, 1973.
15. I. Barrodale and C. Phillips. An improved algorithm for discrete Chebyshev linear approximation. In *Proceedings of the Fourth Manitoba Conference on Numerical Mathematics*, pages 177–190. University of Manitoba, Winnipeg, Canada, 1974.
16. C. de Boor. On local linear functionals which vanish at all B-splines but one. In A. G. Law and B. N. Sahney, editors, *Theory of Approximation with Applications*, pages 120–145. Academic Press, New York, NY, US, 1976.
17. I. J. Anderson. A piecewise approach to piecewise approximation. *Numerical Algorithms*, 15:139–152, 1997.
18. D. Lei, I. J. Anderson, and M. G. Cox. These proceedings.