

Smooth supersaturated models

Ron A. Bates^a, Hugo Maruri-Aguilar^{a,*}, Henry P. Wynn^a

^a*Department of Statistics, London School of Economics, London WC2A 2AE, UK*

Abstract

In areas such as kernel smoothing and non-parametric regression there is emphasis on smooth interpolation and smooth statistical models. Splines are known to have optimal smoothness properties in one and higher dimensions. It is shown, with special attention to polynomial models, that smooth interpolators can be constructed by first extending the monomial basis and then minimising a measure of smoothness with respect to the free parameters in the extended basis. Algebraic methods are a help in choosing the extended basis which can also be found as a saturated basis for an extended experimental design with dummy design points. One can get arbitrarily close to optimal smoothing for any dimension and over any region, giving a simple alternative models of spline type. The relationship to splines is shown in one and two dimensions. A case study is given which includes benchmarking against kriging methods.

Key words: regression, splines, kernel smoothing, non-parametric, computer experiments.

* Corresponding author.

Email address: H.Maruri-Aguilar@lse.ac.uk (Hugo Maruri-Aguilar).

1 Introduction

There is a considerable literature on smooth interpolation and its statistical counterparts. The area of non-parametric regression is an example. The optimal smoothness properties of splines have a substantial literature. The optimality result for one dimensions is attributed to Holladay [1957] and for two dimensions, where thin-plate splines are optimal, to Duchon [1976]; see Micula [2002] for a nice review spline optimality and Kimeldorf and Wahba [1970] for an overview. In computer experiments Bayesian kriging using Gaussian kernel stochastic process models have been preferred to splines, Sacks et al. [1989], Kennedy and O'Hagan [2001], and have also become popular in machine learning: Rasmussen and Williams [2005]. Of course, the connection between kriging and spline is thoroughly researched and, for example, splines can arise as kriging (conditional expectation) interpolators for special Gaussian stochastic processes: Kimeldorf and Wahba [1970].

Raw polynomial interpolation is known in general not to have optimal rates of interpolation unless special sampling (design) points are used such as in Tchebychev approximation. On the other hand the *existence* of polynomial interpolators over an arbitrary design is at the core of the newer theory of algebraic statistics: for any arbitrary design in d dimensions there is always a monomial basis out of which we can build a polynomial interpolator. This was introduced into statistics by Pistone and Wynn [1996], covered at length in the monograph Pistone et al. [2001] and was also the basis for Bates et al. [2003] which can be seen as the forerunner of the present paper.

The aim of the present paper is to try to have the best of both worlds: to

draw a little on the algebraic theory but principally to show, in a rather elementary way, how to construct smooth polynomial interpolators or statistical models. This is achieved by extending the model basis and using this freedom to optimise a measure of smoothness. It should be pointed out that the use of polynomials to build kernels with pre-specified properties is familiar in signal processing, see Lin et al. [2004]. By extending the model basis we can show that our interpolators get arbitrarily close to optimal interpolators, which are typically in the spline family.

1.1 Monomial bases and extended bases

Recent work in the area of “algebraic statistics” shows how to construct estimable (identifiable) monomial bases for polynomial regression and we start with a very short description. Having said this, it is not necessary to use these methods, nor indeed to use polynomials. For example a Fourier (trigonometric) basis may be used. The point is that we shall need an extended basis with certain conditions and the algebra is one way of achieving this.

We start with a set of factors $x = (x_1, \dots, x_d)$. For a set of nonnegative integers $\alpha = (\alpha_1, \dots, \alpha_d)$, a monomial, such as $x_1^2 x_2$, is written $x^\alpha = x_1^{\alpha_1} \dots x_d^{\alpha_d}$, and a polynomial is a linear combination of monomials. A design D_n is a set of n distinct points in d dimensions, $D_n = \{x^{(1)}, \dots, x^{(n)}\}$, $x^{(i)} \in \mathbb{R}^d, i = 1, \dots, n$.

The algebraic methods give us the following: *given an experimental design, D_n , it is always possible to find a saturated non-singular monomial basis $B_L = \{x_\alpha, \alpha \in L\}$. Thus, the size of the basis is equal to the size of the design $|L| = |D_n| = n$ and the $n \times n$ X -matrix, $X = \{x^\alpha\}_{x \in D_n, \alpha \in L}$ is non-singular.*

We call such a basis a *good saturated basis* for the design. The intuition behind algebraic methods is simple: terms are included in the good saturated basis according to a term ordering and a rank inclusion criterion. For details on term orderings see Cox et al. [1997], and for description of the algebraic technology see Pistone et al. [2001].

Example 1 Let D_{24} to be the first 24 points of a bidimensional Sobol's space filling sequence. An implementation of the description of Sobol' sequence by Bratley and Fox [1988] is available in the language R, see Ihaka and Gentleman [1996]. Then by selecting terms with a degree lexicographic term order $x_1 \succ x_2$, a good saturated basis with 24 monomials is identified for D_{24} . This model includes the monomials $x_2^6, x_1x_2^5, x_1^2x_2^4$ plus all the terms of a model of total degree five. This basis will be extended in the example of Section 3.2.

It will be critical in our development that we may extend a basis. By this we mean we keep the design D_n fixed but take a larger set of $N > n$ monomials, hence the term "supersaturated" in the title of the paper. But we require a condition contained in the following definition.

Definition 1 *Given a design D_n , with sample size n , a good supersaturated basis is a basis $B_M = \{x_\alpha, \alpha \in M\}$ with $|B| = N > n$ such that there is a non-singular sub-basis of size n .*

Here is an example to show that we have to be a little careful. Let us start with a rather poor design in two dimensions: $D_4 = \{(0, 0), (1, 1), (2, 2), (3, 3)\}$. Then, and this is obvious without any algebra, there are only two good saturated model bases $\{1, x_1, x_1^2, x_1^3\}$ or $\{1, x_2, x_2^2, x_2^3\}$. From this we can see that the extended basis $\{1, x_1, x_1^2, x_2, x_2^2\}$ with five terms is not good as there is no good sub-basis of size four.

If we start with a non-singular basis for a design D_n and extend it, in any way, then we always obtain a good supersaturated basis. But there is a revealing way of generating a good supersaturated basis and that is by extending the design D_n to a design D_N with N points and finding a good saturated basis for larger design, which contains the good basis for D_n . The algebra shows that this is always possible. This leads to a second, and equivalent, way of producing the smooth models which will be called the “dummy design” method, covered in sub-section 2.2.

2 Smooth interpolators

The basic idea of this paper may seem at first to be somewhat contradictory. We start with given polynomial interpolator and by extending the basis make the interpolator smoother. Although one may naturally associates higher order polynomial terms with lack of smoothness, we can, in fact, extend the basis and use the freedom this gives to *increase* smoothness.

Let the experimental design be D_n and y_1, \dots, y_n be real values (observations) at the design points $x^{(i)} \in D_n, i = 1, \dots, n$, respectively. Let B_M be a good supersaturated basis for the design D_n and let

$$y(x) = \sum_{\alpha \in M} \theta_\alpha x^\alpha \tag{1}$$

be a polynomial in that basis. A good supersaturated model will be sought for using a measure of smoothness.

In one dimension ($d = 1$) we shall adopt the following measure of smoothness

based on the second derivative

$$\Psi_2 = \int_{\mathcal{X}} |y''(x)|^2 dx, \quad (2)$$

where the integration is carried out in a desired region $\mathcal{X} \subset \mathbb{R}$. For higher dimensions the Hessian is

$$H(y(x)) = \left\{ \frac{\partial^2 y(x)}{\partial x_i \partial x_j} \right\},$$

and we have

$$\sum_{ij} \left(\frac{\partial^2 y(x)}{\partial x_i \partial x_j} \right)^2 = \|H(y(x))\|^2 = \text{trace} \left(H(y(x))^2 \right). \quad (3)$$

Then define

$$\Psi_2 = \int_{\mathcal{X}} \|H(y(x))\|^2 dx, \quad (4)$$

for some desired region $\mathcal{X} \subset \mathbb{R}^d$.

The smooth interpolator is $\hat{y}(x) = \sum_{\alpha \in M} \hat{\theta}_\alpha x^\alpha$, where the coefficients $\hat{\theta}_\alpha$ are selected to minimise smoothness subject to the interpolation condition, i.e. solving the constrained optimisation problem

$$\min_{\theta} \Psi_2(y(x)) \quad \text{subject to} \quad y_i = \hat{y}(x^{(i)}), \quad i = 1, \dots, n \quad (5)$$

In the next subsection we give the solution of this constrained problem and the in the second subsection give the dummy design method, which is equivalent.

2.1 The constrained problem

The only technical difficulty arises from the fact that linear parts of the model make no difference to the criterion Ψ_2 but do affect the interpolation. It is necessary to partition the X -matrix to take account of this.

Let $f(x)$ and θ respectively be the vectors which hold the good supersaturated basis and the parameters so that we can write (1) as $y(x) = \theta^T f(x)$. Denote $f^{(ij)} = \frac{\partial^2 f(x)}{\partial x_i \partial x_j}$ and define

$$K = \int_{\mathcal{X}} \left(\sum_{i,j=1}^k f^{(ij)} f^{(ij)T} \right) dx, \quad (6)$$

Then we see that

$$\Psi_2(y(x)) = \theta^T K \theta \quad (7)$$

The technical difficulty mention above means that K may not be full rank. In particular any linear term in the models basis will give zero entries. Call this entries *structural zeros*. Permute the rows and columns of K so that the structural zeros are adjacent:

$$K = \begin{bmatrix} 0 & 0 \\ 0 & \tilde{K} \end{bmatrix} \quad (8)$$

Let $X = [X_0, X_1]$, $f = (f_0^T : f_1^T)^T$ and $\theta = (\theta_0^T : \theta_1^T)^T$ be the corresponding rearranged and partitioned versions of X_n , f and θ , respectively. The matrix X has n rows and as many columns as terms in f . Let y , be the column vector with n observations and note that $\Psi_2 = \theta_1^T \tilde{K} \theta_1$.

With this the constrained quadratic problem (5) is:

$$\min_{\theta} \theta_1^T \tilde{K} \theta_1 \quad \text{subject to} \quad X_0 \theta_0 + X_1 \theta_1 = y \quad (9)$$

Let 2λ be an $n \times 1$ vector of Lagrange multipliers (2 is for convenience) so that the Lagrangian is

$$\theta_1^T \tilde{K} \theta_1 - 2\lambda(X_0 \theta_0 + X_1 \theta_1).$$

After differentiation the full set of equations for θ_0, θ_1 and λ can be written in block form

$$\begin{bmatrix} X_0 & X_1 & 0 \\ 0 & \tilde{K} & -X_1^T \\ 0 & 0 & X_0^T \end{bmatrix} \begin{bmatrix} \theta_0 \\ \theta_1 \\ \lambda \end{bmatrix} = \begin{bmatrix} y \\ 0 \\ 0 \end{bmatrix} \quad (10)$$

If the matrix on the left hand side is nonsingular we obtain a unique solution $\hat{\theta}_0, \hat{\theta}_1, \hat{\lambda}$. The following three conditions guarantee this.

- (i) The full basis is a good supersaturated basis for D_n so that X is full rank.
- (ii) X_0 is full rank.
- (iii) \tilde{K} is full rank and thus invertible.

The full matrix inverse with solutions $\hat{\theta}_0, \hat{\theta}_1, \hat{\lambda}$ are given in Appendix 1. Finally, using these results, we express the smooth estimator as

$$\hat{y}(x) = \hat{\theta}_0 f_0 + \hat{\theta}_1 f_1 = \hat{\theta} f(x)$$

and the optimal Ψ_2 as

$$\Psi_2^* = \hat{\theta}_1^T \tilde{K} \hat{\theta}_1.$$

In applications, as is common with quadratic programme, we simply invert the matrix on the right hand side of (9) using a fast numerical method. Thus, given the design D_n , the good supersaturated basis and \tilde{K} , the method is fairly straightforward to implement.

It is revealing to consider the case where K is nonsingular. Then we do not

need the partition of Equation (8) and instead can write Equation (10) as

$$\begin{bmatrix} X & 0 \\ \tilde{K} & -X \end{bmatrix} \begin{bmatrix} \theta \\ \lambda \end{bmatrix} = \begin{bmatrix} y \\ 0 \end{bmatrix}$$

Which has the solution:

$$\hat{\theta} = (X^T X + K(I - P)K)^{-1} X^T y$$

where $P = X^T(XX^T)^{-1}X$ is the projector onto the row space of X . Thus, although $X^T X$ is not invertible, because we have a supersaturated model, the second term $K(I - P)K$ on the left hand side can be seen as a smoothness induced regularisation of the problem which compensates for this singularity.

2.2 The dummy design method

For simplicity of development we assume that K is non-singular in the present case. Let D_N be a large design, with $N > n$ distinct points, which contains the original design D_n and write

$$D_N = D_n \cup D_q,$$

where $q = N - n$. Let $h(x)$ be a good saturated basis for D_n , and let $f(x)$ be an (extended) good saturated basis for D_N , $f(x) = (h(x)^T, g(x)^T)^T$. Also extend the observation vector to $z = (y^T, z^T)^T$ where, as before y holds the “true” observations taken at points in D_n , and z can be thought of as dummy observations on the design D_q . The extended model we write

$$y(x) = f(x)^T \theta = h^T(x)\beta + g^T(x)\gamma \tag{11}$$

and assume, as in the last section, that $y(x)$ interpolates the observations y over D_n .

We now minimize Ψ_2 over the choice of dummy observations z which is now an unconstrained optimization problem, but with a reduced set of free parameters, namely z . The constrained optimization (8) and this unconstrained optimization are equivalent in the case that the full basis is a good for the full design, D_N . This is because of the one to one correspondence between observations and parameters and the fact that the interpolation constraint is the same in both cases.

The unconstrained problem is:

$$\min_z (y^T : z^T) X_N^{-1T} K X_N^{-1} \begin{pmatrix} y \\ z \end{pmatrix}. \quad (12)$$

Where X_N is the X -matrix for the full large model $f(x)$. First, let the following matrix be partitioned according to the model bases $f(x) = (h(x)^T, g(x)^T)^T$:

$$A = X_N^{-1T} K X_N^{-1} = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}.$$

Then after expanding (11) and differentiating, the optimal z is

$$\hat{z} = -A_{22}^{-1} A_{21} y$$

and the minimum value of the smoothness is

$$\Psi_2^* = y^T Q y,$$

where $Q = A_{11} - A_{22}^{-1}A_{21}$. The smooth interpolator is

$$\hat{y}(x) = f^T(x)X_N^{-1}\begin{pmatrix} y \\ \hat{z} \end{pmatrix} = f^T(x)X_N^{-1}\begin{pmatrix} I \\ -A_{22}^{-1}A_{21} \end{pmatrix}y = f^T(x)K^{-1}(X_{11} : X_{12})Qy \quad (13)$$

where

$$X_N = \begin{pmatrix} X_{11} & X_{12} \\ X_{21} & X_{22} \end{pmatrix}$$

is the appropriate partition of X_N , i.e. the rows of X_N are indexed by D_n and D_q , while the columns are indexed by $h(x)$ and $g(x)$.

The last equality and the equivalence to the solution in the last subsection is shown for the case that K is non-singular. The equivalence in general holds under conditions (i), (ii) and (iii) in that section. We note that the solution does not depend on the dummy design D_q , except in so far as it is involved in guaranteeing that we have a good supersaturated basis.

3 One and two dimensions

3.1 A one dimensional example: spline-like behavior

In this example, smooth saturated models are used for interpolating a known univariate function. The function considered is the sine cardinal

$$m(x) = \text{sinc}(ax + b) = \sin(ax + b)/(ax + b)$$

with $a = 15\pi/2$ and $b = -10\pi/2$. The region over which the interpolators will be smoothed is $\mathcal{X} = [0, 1]$.

Suppose that the design D_6 is a uniform design in $[0, 1]$, and that the response vector y contains the values of $m(x)$ at points in D_6 . The choice of a good saturated and supersaturated models can be driven by algebraic methods. For the present case, an obvious candidate is $h(x) = (1, x, \dots, x^5)^T$. Call \hat{y}_0 to the interpolator fitted solely with $h(x)$. Now a process of smoothing is carried out by adding dummy points, one at a time. While adding dummy points $h(x)$ remains unchanged. With only one dummy point, a clear candidate for $g(x)$ is $g(x) = (x^6)$, while for q dummy points, $g(x) = (x^6, \dots, x^{6+q-1})$ could be used. Call \hat{y}_q to the smooth interpolator obtained by adding q dummy points, $q = 1, \dots, 5$. The value of smoothness for \hat{y}_q quickly drops down so that a similar smoothness to that of a spline is achieved with \hat{y}_4 (only four extra terms), see Table 1. The progressive smoothing achieved with extra terms can be appreciated graphically as well. Figure 1 shows the interpolator and smooth saturated models.

Model	\hat{y}_0	\hat{y}_1	\hat{y}_2	\hat{y}_3	\hat{y}_4	\hat{y}_5	Spline
Ψ_2^*	76.543	74.698	33.153	33.020	27.767	27.745	26.744

Table 1

Value of Ψ_2^* for supersaturated models interpolating $m(x)$ over D_6 of Section 3.1.

A comparison between the smooth supersaturated method and cubic splines, which are optimally smooth, is carried out as follows. First, for a uniform design D_n on $[0, 1]$, a saturated model \hat{y}_0 is fitted to the values of $m(x)$ at the design points. Call $\Psi_2^*(0)$ the value of smoothness for \hat{y}_0 . Then, using extra q basis terms, a smooth supersaturated model \hat{y}_q is fitted. Call $\Psi_2^*(q)$ the corresponding value of smoothness. Additionally, a spline is fitted to the same data and call $\Psi_2^*(\text{sp})$ its smoothness value. The important feature is that the $\Psi_2^*(0), \Psi_2^*(1), \dots$ form a decreasing sequence which con-




Fig. 1. Sequence of smooth saturated models: \hat{y}_0 is a polynomial of fifth degree (- -), $\hat{y}_1, \dots, \hat{y}_4$ (—) are supersaturated models. True model $m(x)$ (...) and design points are also shown.

verges surprisingly quick to $\Psi_2^*(\text{sp})$. This behavior can be quantified by plotting the ratio $\sqrt{\Psi_2^*(q)/\Psi_2^*(\text{sp})}$ against the number of terms added to smooth the model. Figure 2 shows such comparison when D_n are uniform designs of size $n = 5, 10, 15, 20$.

3.2 A two dimensional example: alternative to thin-plate splines?

The objective of this example is to compare the performance of smooth supersaturated interpolators against thin plate splines, but there is also interest




Fig. 2. Logarithm of smoothness ratio $R(q) = \sqrt{\Psi_2^*(q)/\Psi_2^*(sp)}$ against number of smoothing terms added q : sample sizes $n = 5, 10, 15$ (- · -, ···, —). The line for $n = 20$ is indistinguishable from $R(q) = 1$.

to make comparisons against a kriging interpolator. Initially, interpolators of the three kinds above are constructed for a known function at given design points and then predictions over new design points are used to compare the performance of the interpolators. The known function is $m(x_1, x_2)$, which is constructed as $m(x_1, x_2) = p(4x_1 - 2, 4x_2 - 2)$, where $p(x_1, x_2)$ is the **peaks** function from MATLAB®. The objective of scaling and shifting $p(x_1, x_2)$ is to include interesting features into the smoothing region $\mathcal{X} = [0, 1]^2$.

In order to allow a good covering of the design region \mathcal{X} without an excessive number of points, we use Sobol's space filling design D_{24} and $h(x)$ to be the good saturated model of Example 1. The response vector y contains the values of $p(x_1, x_2)$ at points in D_{24} .

A smooth supersaturated model was then fitted to this data using the 91 terms of a good supersaturated complete model of degree twelve in x_1, x_2 . Call this model \hat{y} . A thin plate spline interpolator model was also fitted to the same data, which we refer to as \hat{y}_{sp} . A kriging interpolator, \hat{y}_{kr} , was also fitted using

the model

$$Y(x) = \beta + Z(x), \quad (14)$$

where $Z(x)$ is a stochastic process with exponential covariance structure, i.e. $\text{Cov}(Z(r), Z(s)) = \exp(\sum_{i=1}^2 \theta_i |r_i - s_i|^{p_i})$.




Fig. 3. Smooth supersaturated predictions \hat{y} against spline \hat{y}_{sp} and kriging predictions \hat{y}_{kr} for the extra design points in Section 3.2.

For comparison, a set of predictions were generated for each model at new design points. The new design points were the next 500 points from the Sobol' sequence used for the first step. The predictions obtained with the smooth supersaturated model \hat{y} are closely correlated with those of the spline \hat{y}_{sp} and the kriging \hat{y}_{kr} models, see Figure 3 (a) and (b), only showing bias for low predicted values, especially when comparing with the kriging model. Additionally, the root mean square error (RMSE) was computed using the true values $g(x_1, x_2)$ and the predictions for each of the three models at the extra design points. The values of RMSE for the smooth supersaturated, spline and kriging models are 1.117, 1.009, 0.640, respectively. This figures represent the 7.7%, 7.0% and 4.4% of the response range, respectively. The results show that the smooth supersaturated models are a good alternative to splines for inter-

polation, which can also be seen in Figure 4 against the simulated response.




Fig. 4. Smooth supersaturated \hat{y} , spline \hat{y}_{sp} and kriging predictions \hat{y}_{kr} against true simulated values y for the extra design points in Section 3.2.

4 From interpolators to statistical models

4.1 Designs points versus knots

The bulk of the development in this paper concerns the use of the smooth function as interpolators. However they can be used as statistical models in a

straightforward way. Recall that the solution are of the form

$$\hat{y}(x) = \hat{\theta}^T f(x) = y^T Bf(x)$$

for the matrix B , in one of the equivalent forms in the development. We see that $\hat{y}(x)$ is linear in the observations y . The idea is to make y a free parameter, that is to change the role of y . Indeed we could relabel y as ϕ and write the model as

$$\hat{y} = \phi^T Bf(x)$$

The design point in D_n become *knots* and we are parameterizing the model by the values at the knots. This is somewhat familiar in splines. With this change we are free to fit the models using any regression, stepwise regression, penalised method etc we choose. There is no requirement to observe at the knots. But when we have carried out the fitting and write \hat{y} instead of we have the level of smoothness achieved by replacing y by \hat{y} in our formula for Ψ_2 . Moreover we are free to choose the location of the knots and the “real” experimental design at which to observe. In terms of the dummy design method, this amounts to a double-dummying: once for the knots and once for the smoothness; even before we actually take observations.

The function $k(x) = Bf(x)$ can be considered as special kernels each with a value unity at a design point and zero at other design points and we can write the model as $\sum_i k_i(x)y_i$ when the y_i are observations or, in the parametric case just described, as $\sum_i k_i(x)\phi_i$.

We restrict the discussion to the case that K is non-singular, again for simplicity. Then

$$\Psi_2^* = y^T Q y = y^T (X K^{-1} X^T)^{-1} y$$

We first note that the design D_n , via the design model matrix X , affects the value of the smoothness in the interpolation case, even without any statistical considerations. Given that we have to choose the design *before* we observe y one may consider that some measure of the size of $Q = (X K^{-1} X^T)^{-1}$ may be important. We may borrow criteria from the optimal design of experiments and seek to minimize some function of Q . In the case that K is non-singular $\det(Q)$ may be used, but as pointed out, since K is not typically full rank, nor is Q .

We consider a small example. Let $n = 3, N = 5$ and $d = 1$ and take the saturated basis as $1, x, x^2, x^3, x^4$ and let both the design interval and the integration interval be \mathcal{X} be $[-1, 1]$. We need to minimize $\Psi_2 = y^T Q y$ with respect to the choice of four design points in $[-1, 1]$. After some analysis it can be shown that the optimal design take the form $\{-1, -a, a, 1\}$ for some positive a . As expected, because of the two linear terms, the matrix Q has rank two. The largest eigenvalue of Q takes the value

$$\frac{12(1 + a^2)}{a^2(1 - 2a^2 + a^4)}$$

Minimisation of the largest eigenvalue of Q leads to an optimal value of $a = 1/2\sqrt{-3 + \sqrt{17}} \approx 0.52988$. Minimising the product of the eigenvalues of Q gives $a \approx 0.40570$.

In the case that the design $D_{(n)}$ becomes a set of knots we are free to choose

the actual design points separately. If we fit using smooth supersaturated models this gives an optimal design problem with the kernels $\{k_j\}$ given above. Continuing with the above example and guessing that the D -optimal on $[0, 1]$ for the optimally smooth kernels obtained by the first solution takes the form $\{-1, -b, b, 1\}$ we find that D -optimal solution as

$$b = \frac{1}{35} \sqrt{1925 + 175\sqrt{17} - 35\sqrt{2785 + 480\sqrt{17}}} \approx 0.43402,$$

which can, indeed, be confirmed to be the D -optimum design by checking against the Kiefer-Wolfowitz General Equivalence Theorem. One sees that these are not the same as the optimal knots.

But now an attractive possibility arises. Optimal design experimental design for splines has received some attention in the literature, but it has been considered a somewhat intractable problem. Now, given that splines can be found as the limit of polynomial models it may be considered that optimal design for splines can be found approximately by taking smooth supersaturated models with large bases, and using one of a number of optimum design algorithms to find the (approximate) solution. One exchanges a problem of handling real splines analytically with that of high dimensional linear algebra. This will be the subject of further research.

In the case that we are free to choose the knots and the design points separately, a conceptually simple approach, then, to carry out two separate separate optimal “design” problems one for knot placement for smoothness, as above, and a second for, say, D -optimality of the design points.

It becomes conceptually harder if we wish to take into account smoothness and statistical precision in a joint analysis. One might seek to minimize some

portmanteau criterion with respect to a simultaneous optimizations over design points and knots. If, moreover, Ψ_0 is a statistical criterion such as from D -optimality, we might take as a criterion some weighted combination:

$$(1 - \lambda)\Psi_0 + \lambda\Psi_2$$

As the y -values at the knots are now unknown parameters ϕ_i , in a linear model we have that the true smoothness is $\Psi_2 = \phi^T Q \phi$ is non-linear in ϕ .

5 A case study: Engine Emissions Data

The performance of a smooth supersaturated model is evaluated against a kriging model using the engine emissions data set analysed in Bates et al. [2003]. This data set comes from a computer experiment and comprises 48 observations in five factors N, C, A, B and M . An extra set of 49 observations is available for validation purposes. The smooth supersaturated model \hat{y} is constructed with 100 terms fitted to the set of 48 observations. For this model, 48 terms correspond to the good saturated basis proposed in [Bates et al., 2003, Section 6.3], and this forms $h(x)$. A set of 22 terms are added to complement missing terms of total degree three and then a set of extra 30 terms of total degree four were added. All the extra 52 terms described form $g(x)$ and were added using a degree lexicographic order. Call \hat{y}_{sp} and \hat{y}_{kr} to the spline and kriging models constructed with the first data set. The kriging model \hat{y}_{kr} was built with a five dimensional extension of the covariance structure used in Equation (14).

In the validation stage, predictions at the extra 49 design points were built using the three models \hat{y} , \hat{y}_{sp} and \hat{y}_{kr} . The values of RMSE for \hat{y} , \hat{y}_{sp} and \hat{y}_{kr} are

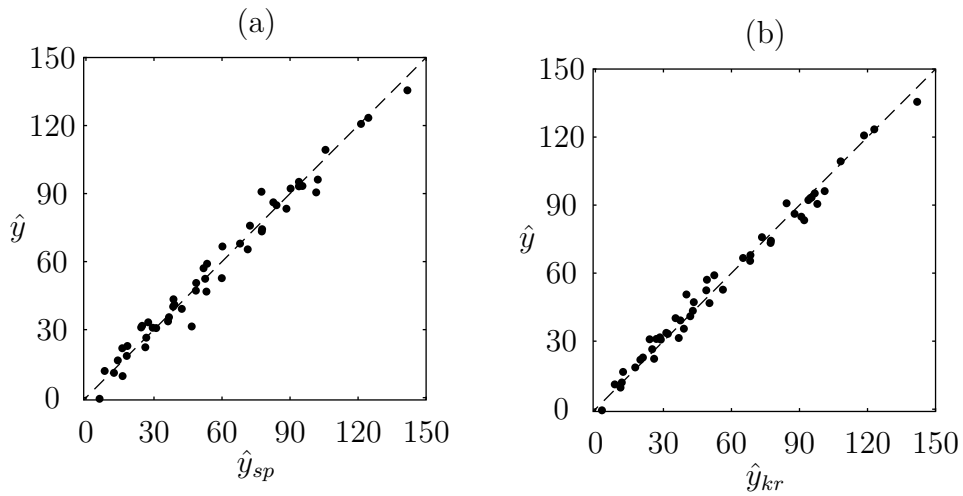


Fig. 5. Smooth supersaturated predictions (\hat{y}) against spline (\hat{y}_{sp}) and kriging predictions (\hat{y}_{kr}) for the validation data set of Section 5.

5.844, 5.896 and 4.450 respectively, which respectively represent the 4.4%, 4.5% and 3.4% of the range of the response values. The smooth supersaturated model \hat{y} compares well with both spline and kriging. Figure 5 shows that the predictions with the smooth supersaturated model are also closely correlated to those obtained with spline and kriging models. Figure 6 also shows the smooth supersaturated model to be a good predictor of the true response.

6 Discussion and further research

We have tried to show in this paper that the simple idea of extending a basis in regression and using the free parameters which that gives to increase smoothness give interpolators which have the same order of magnitude error as the two main alternative: splines and kriging. For smaller dimensions not too many additional additional basis terms are need to give a large decrease in accuracy. Although there is still work to be done on the theory it seems clear that one can get arbitrarily close to the theoretically smoothest func-

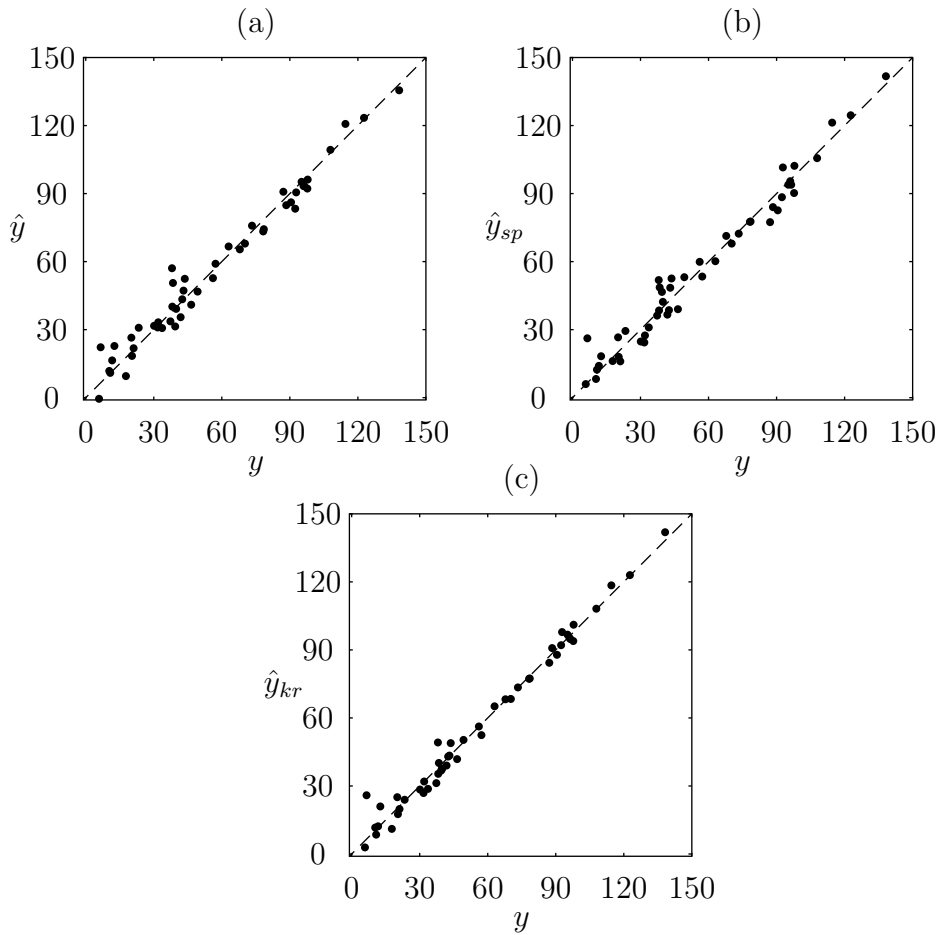


Fig. 6. True values (y) against smooth supersaturated predictions (\hat{y}), spline (\hat{y}_{sp}) and kriging predictions (\hat{y}_{kr}) for the validation data set of Section 5.

tions, namely splines. Moreover this can be achieved for complex regions of integration and sets of observation points (designs), limited only by a rank condition.

There a number of ways in which one can generalise or adapt these methods, which we discuss briefly.

- (1) The same analysis will go through for weighted criteria:

$$\Psi_2 = \int_{\mathcal{X}} ||H(y(x))||^2 w(x), dx,$$

where $w(x)$ is a non-negative weight function. This simply changes the

definition of K and \tilde{K} .

- (2) The smoothness criteria we adopted is one of a number in a wider quadratic class such as

$$\Psi_1 = \int_{\mathcal{X}} \|\Delta(y(x))\|^2 dx,$$

where $\Delta(y(x))$ is the gradient vector. Another is the deviation from a target

$$\Psi_{0,t} = \int_{\mathcal{X}} |y(x) - t(x)|^2 dx,$$

and one could have weighted versions of them or even weighted combinations of different criteria.

- (3) We have ignored analysis based on building in additional, more statistical criteria, such as cross-validation to have a trade off between smoothness and statistical variation. A simple way of taking this forward would be to consider smooth supersaturated as adding to the catalogue of kernels which are now studied in many fields such computer experiments, non-parametric regression, imaging, machine learning and signal processing. They would be candidates for analysis using stepwise methods, AIC, BIC, LASSO and so on.
- (4) A possible advantage of the kernels we have developed is that their polynomial nature makes them more tractable than, say, splines in some circumstances; for example for differentiation in sensitivity analysis, error propagation or integration.
- (5) We summarize that given detailed attention to computational issues, it is possible to develop optimal experimental designs for the high degree, but smooth, kernel models which arise from the present methods. As mentioned, this may be a way of tackling optimal design for complex regions.

(6) The same methods can be applied for other bases, for example Fourier bases in one and higher dimensions. Again as the basis order gets larger one will tend to the optimal spline-like kernels. For Fourier bases one can gain smoothness by using higher frequencies, in seeming, but not actual, contradiction to the Nyquist sample theorem.

7 Appendix

7.1 Appendix 1: solution for $\hat{\theta}_0$ and $\hat{\theta}_1$

It is possible to use block matrix inverse methods, but they are a little cumbersome. We first find $\hat{\theta}_0$. Writing the equations out we have

$$X_0\theta_0 + X_1\theta_1 = y$$

$$K\theta_1 - X_1^T\lambda = 0$$

$$X_0\lambda = 0$$

Solving for λ from the second two equations we have

$$\lambda = (X_1K^{-1}X_1^T + X_0X_0^T)^{-1}X_1\theta_1$$

Using this to eliminate θ_1 from the first equation we have

$$X_0^T(X_1K^{-1}X_1^T + X_0X_0^T)^{-1}X_0\theta_0 = X_0^T(X_1K^{-1}X_1^T + X_0X_0^T)^{-1}y,$$

giving

$$\hat{\theta}_0 = (X_0^T(X_1K^{-1}X_1^T + X_0X_0^T)^{-1}X_0)^{-1}X_0^T(X_1K^{-1}X_1^T + X_0X_0^T)^{-1}y,$$

Writing $y^* = y - X_0\hat{\theta}_0$ we obtain reduced matrix equation:

$$\begin{bmatrix} X_1 & 0 \\ \tilde{K} & -X_1^T \\ 0 & X_0^T \end{bmatrix} \begin{bmatrix} \theta_1 \\ \lambda \end{bmatrix} = \begin{bmatrix} y^* \\ 0 \\ 0 \end{bmatrix}$$

Left multiplying by the transpose of the matrix on the left and inverting we have

$$\hat{\theta}_1 = (X_1^T X_1 + \tilde{K}(I - X_1^T (X X^T)^{-1} X_1) \tilde{K})^{-1} X_1 y^* \quad (15)$$

Note that in the case that X_0 and X_1 have orthogonal columns we reduce to the standard form

$$\hat{\theta}_0 = (X_0^T X_0)^{-1} X_0^T y$$

This can be achieved by rewriting the supersaturated basis so that the terms with degree higher than linear (degree one) are orthogonal to the linear terms with respect to the design. Of course, the definition of \tilde{K} should be changed accordingly.

7.2 Equivalence of forms in the case K nonsingular

The three forms for $\hat{\theta} = B y$ where B is one of the following:

(i) $B_1 = (X_1^T X_1 + K(I - P)K)^{-1} X_1^T y$

(i) $B_2 = K^{-1}(X_{11}, X_{12})^T Q y$

(ii) $B_3 = X^{-1} \begin{pmatrix} I \\ -A_{22}^{-1} A_{21} \end{pmatrix}$

To show that $B_1 = B_2$ multiply both by $X_1^T X_1 + K(I - P)K$ and note that

$PX^T = 0$ to obtain respectively X^T and $X^T XK^{-1}X^T Q$. But from the definition of Q and using block the partition inverse formula we see that that $XK^{-1}X^T = Q^{-1}$ and we are done (reversing the steps).

To show that $B_2 = B_3$ we multiply both by $X^{-1T}K$. Then B_2 gives

$$X^{-1T}KK^{-1}(X_{11}, X_{12})^T QQ^{-1} = X^{-1T}(X_{11}, X_{12})^T = \begin{pmatrix} I \\ 0 \end{pmatrix},$$

and B_3 gives

$$\begin{aligned} X^{-1T}KX^{-1}\begin{pmatrix} I \\ -A_{22}^{-1}A_{21} \end{pmatrix}Q^{-1} &= A\begin{pmatrix} I \\ -A_{22}^{-1}A_{21} \end{pmatrix}Q^{-1} = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \begin{pmatrix} I \\ -A_{22}^{-1}A_{21} \end{pmatrix}Q^{-1} \\ &= \begin{pmatrix} A_{11}-A_{12}A_{22}^{-1}A_{21} \\ A_{21}-A_{22}A_{22}^{-1}A_{21} \end{pmatrix}Q^{-1} = \begin{pmatrix} A_{11}-A_{12}A_{22}^{-1}A_{21} \\ 0 \end{pmatrix}Q^{-1} = \begin{pmatrix} I \\ 0 \end{pmatrix}. \end{aligned}$$

Again, reversing the steps we obtain our result.

Acknowledgments

The first and third authors acknowledge the EPSRC grant GR/S63502/01, while the second and third authors acknowledge the EPSRC grant EP/D048893/1 (MUCM project).

References

- Bates, R., Giglio, B., and Wynn, H. (2003). A global selection procedure for polynomial interpolators. *Techno.*, 45(3):246–255.
- Bratley, P. and Fox, B. L. (1988). ALGORITHM 659 Implementing Sobol’s quasirandom sequence generator. *ACM Trans. Math. Soft.*, 14(1):88–100.

- Cox, D., Little, J., and O’Shea, D. (1997). *Ideals, Varieties, and Algorithms*. Springer-Verlag, New York. Second Edition.
- Duchon, J. (1976). Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *R.A.I.R. Analyses Numrique*, 10(3):5–12.
- Holladay, J. (1957). A smoothest curve approximation. *Maths. Tables Aids Compute.*, 11(3):233–243.
- Ihaka, R. and Gentleman, R. (1996). R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics*, 5(3):299–314.
- Kennedy, M. and O’Hagan, A. (2001). Bayesian calibration of computer models. *J. Roy. Statist. Soc. B.*, 63(3):425–2001.
- Kimeldorf, G. and Wahba, G. (1970). A correspondance between bayesian estimation of stochastic processes and smoothing by splines. *Ann. Statist.*, 41:495–502.
- Lin, Z., Xu, L., and Wu, W. (2004). Applications of Gröbner bases to signal processing: a survey. *Lin. Alg. Appl.*, 391(3):169–202.
- Micula, G. (2002). A variational approach to spline functions theory. *General Mathematics*, 10(1-2):21–50.
- Pistone, G., Riccomagno, E., and Wynn, H. P. (2001). *Algebraic Statistics*, volume 89 of *Monographs on Statistics and Applied Probability*. Chapman & Hall/CRC, Boca Raton.
- Pistone, G. and Wynn, H. (1996). Generalised confounding with Gröbner bases. *Biometrika*, 83(3):653–666.
- Rasmussen, C. and Williams, C. (2005). *Gaussian processes for machine learning*. MIT Press, Cambridge, Mass.
- Sacks, J., Welch, W., Mitchell, T., and Wynn, H. (1989). The design and analysis of computer experiments. *Statistical Science*, 4(4):409–439.