Adaptive Regularization of B-Spline Models for Scientific Data^{*}

David Lenz¹[0000-0002-2587-2783]</sup>, Raine Yeh²[0000-0002-1175-6212]</sup>, Vijay Mahadevan¹[0000-0002-3337-2607]</sup>, Iulian Grindeanu¹[0000-0002-0264-8236]</sup>, and Tom Peterka¹[0000-0002-0525-3205]

¹ Argonne National Laboratory, Lemont, IL, USA {dlenz, mahadevan, iulian, tpeterka}@anl.gov
² Google Inc., New York, NY, USA raineyeh@google.com

Abstract. B-spline models are a powerful way to represent scientific data sets with a functional approximation. However, these models can suffer from spurious oscillations when the data to be approximated are not uniformly distributed. Model regularization (i.e., smoothing) has traditionally been used to minimize these oscillations; unfortunately, it is sometimes impossible to sufficiently remove unwanted artifacts without smoothing away key features of the data set. In this article, we present a method of model regularization that preserves significant features of a data set while minimizing artificial oscillations. Our method varies the strength of a smoothing parameter throughout the domain automatically, removing artifacts in poorly-constrained regions while leaving other regions unchanged. The behavior of our method is validated on a collection of two- and three-dimensional data sets produced by scientific simulations.

Keywords: B-Spline · Regularization · Functional Approximation

1 Introduction

Data sets assembled from scientific simulations or experimental readings are often defined as a list of position-value pairs, where each data point consists of a measurement and the corresponding location of that measurement. These point locations can form structured grids, unstructured meshes, or unconnected point clouds, depending on the application. Methods for analyzing these data often apply only to particular layouts; usually, numerical analysis techniques become more complex as the geometry of the point locations becomes more general (e.g. point clouds). Even seemingly straightforward tasks such as interpolation can be computationally burdensome on unstructured point clouds and numerically inaccurate on highly nonuniform meshes. One way to avoid these challenges is

^{*} This work is supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research under Contract DE-AC02-06CH11357, Program Manager Margaret Lentz.

by representing a data set with a mathematical function and then analyzing the function instead of the original data. This can substantially streamline the process of interpolation and differentiation away from data points, simplify visualization tasks, and make resampling the data almost trivial.

The focus of this article is the approximation of scientific data sets by (tensorproduct) B-splines. B-splines are a family of smooth functions used ubiquitously throughout geometric modeling [14] and form the underpinnings of isogeometric analysis (IGA) [11]. Recent study has shown that large, complex data sets produced by scientific simulations at extreme scale can be effectively modeled by B-splines [16]. Similar results have also been obtained for nonuniform rational B-splines, or NURBS, which are a generalization of B-splines [15].

B-splines have a number of properties that make them useful as a functional representation of data. B-splines are high-order approximants, and evaluating, differentiating, and integrating a B-spline model is fast and numerically stable [3]. Crucially, differentiation and integration can be computed in closed-form and incur no additional loss of accuracy, unlike finite differences or Riemann sums. Thus, once a sufficiently accurate spline has been computed to represent the data, it is often more productive to analyze the functional model than the original data.

However, computing a best-fit B-spline requires solving a linear system which may be ill-conditioned or rank-deficient. A common cause of this ill-conditioning is an input data set that contains both sparse and dense patches of points in proximity to each other. Without additional effort, solving this system can produce a function that oscillates strongly between data points or even diverges in regions where input data are very sparse. This problem can be hard to detect automatically, since error metrics are usually defined in terms of the pointwise error between the original data and the model. Spurious oscillations occurring away from the input data will not be captured by these metrics.

To address these challenges, we developed a new method for fitting B-splines to unstructured data that reduces or eliminates oscillations while leaving critical features of the data set unchanged. Our method regularizes the solution to the B-spline fitting problem by adding a variable-strength smoothing parameter that automatically adapts based on characteristics of the input data set. This additional term smooths out spike artifacts in regions where the data set is very sparse but does not do any smoothing where data points are densely packed, thereby preserving accuracy in these regions. In addition, our method creates well-defined spline models even for data sets with irregular boundaries. No knowledge of the boundary is required; the method automatically handles areas outside the boundary that contain no data points.

The remainder of this paper is organized as follows. A review of related ideas and methods is given in Section 2. In Section 3, we provide a primer on the mathematical details used to describe B-splines throughout the paper. Our main result, a method for adaptive regularization of B-spline models, is described in Section 4. We then exhibit the performance of this method in Section 5 with a series of numerical examples. We summarize directions for further research in Section 6 and present conclusions in Section 7.

2 Related Work

Creating B-spline models to represent unstructured data sets is a particular example of scattered data approximation (SDA), a broad area of study concerned with defining continuous functions that interpolate or approximate spatially scattered inputs. SDA is often applied to image reconstruction problems, where an experimental or physical constraint prohibits the collection of uniformly-spaced samples, such as medical [1], seismic [5], or astronomical [18] imaging. An introductory comparison of SDA methods was compiled by Francis et al. [8].

Ill-conditioned numerical methods are a persistent challenge throughout the SDA literature, and a number of techniques have been proposed to increase numerical stability. Our approach is most similar to the variational methods of SDA, in which the magnitude of the approximating function's derivative (or "roughness") is minimized. The early work of Duchon [4] is a canonical example. Historically, roughness minimization has been achieved through the use of smoothing splines; a thorough exposition of smoothing splines can be found in the book by Gu [10]. The application of smoothing splines requires a trade-off between accuracy and roughness minimization, since aggressively penalizing roughness tends to degrade accuracy. Therefore, much work has been devoted to parametrizing this trade-off appropriately. Craven and Wahba [2] developed the influential "cross-validation" approach, which is expanded upon by Gu [9].

The functional approximation used in this article is based on global tensorproduct B-splines, but a number of other spline-based regression methods have been proposed. Truncated thin-plate splines were used by Wood [19] to improve the efficiency of thin-plate regression splines while maintaining their characteristic stability. Lee et al. [13] utilized hierarchies of B-splines to fit unstructured data points, but the instability arising from sparse point distributions was not treated explicitly. Francis et al. [8] consider a two-step process for resampling unstructured point clouds with variable point density onto unstructured grids. While this method does not construct a functional approximation, it does show good performance as a resampling methodology.

Our novel adaptive regularization procedure was first explored in the dissertation of the second author [20]. The method is also directly inspired by the work of El-Rushaidat et al. [6], in which a two-level regularization process was introduced in the context of resampling unstructured data onto structured meshes. However, their method requires an ad-hoc selection of the criteria to switch between high and low regularization strengths, as well as an application-dependent overall level of smoothing. A notable contribution in our work is a continuously varying regularization strength (not two-level) that is adapted automatically.

3 Background on B-Splines

In this section, we provide a brief overview of the basic definitions and constructions necessary to describe B-spline models for scientific data. A thorough presentation on the fundamental theory of B-splines can be found in the books by de Boor [3] and by Piegl and Tiller [17].

3.1 B-Spline Curves

A one-dimensional B-spline curve of degree p in \mathbb{R}^D is a parameterized curve

$$C(u) = \sum_{j=0}^{n-1} N_{j,p}(u) P_j,$$
(1)

where each $N_{j,p}$ is a piecewise-polynomial function of degree p, and each $P_j \in \mathbb{R}^D$ is a "control point" in D-dimensional space.

The B-spline basis functions, denoted $N_{j,p}$, are defined on the parameter space $[0,1] \subset \mathbb{R}$, which is divided by a nondecreasing sequence of "knots" $t_0 \leq t_1 \leq \ldots \leq t_{n+p} \in [0,1]$. Each basis function $N_{j,p}$ is a bump function in $[t_j, t_{j+p+1}]$ and zero elsewhere.³ In this paper, we will assume that the degree of the B-spline is fixed and drop the *p* subscript, instead denoting the *j*th function as N_j .

In order to simplify notation when describing high-dimensional tensor product splines, we use multi-indices to index quantities in multiple dimensions simultaneously. A multi-index $\alpha = (\alpha^1, \ldots, \alpha^d)$ is a *d*-tuple of nonnegative indices, where the sum of components of α is denoted $|\alpha| = \sum_k \alpha_k$.

We will often consider index sets for our multi-indices in the form of

$$A = \{ \text{all } \alpha \in \mathbb{N}^d \text{ such that } 0 \le \alpha^k < n_k \text{ for } 1 \le k \le d \},$$
(2)

where n_k are previously defined positive numbers. We impose a lexicographic ordering on these sets, and denote by $[\alpha]_A$ the index of α in the lexicographic ordering of A. In the following sections, we consider matrices in which each column corresponds to a multi-index. In this scenario, we list multi-indices in lexicographic order; thus, the multi-index α corresponding to the j^{th} column satisfies $[\alpha]_A = j$.

3.2 Tensor Product B-Splines

Tensor product B-splines are a natural extension of B-spline curves to higherdimensional manifolds, such as surfaces, volumes, and hypervolumes. Here, we denote by d the dimension of the tensor product volume and D the dimension of the ambient space (for instance, a 2D surface in 3D space would correspond to d = 2, D = 3). The parameter space for a d-dimensional tensor product B-spline is $[0,1]^d$, which is divided by d different knot vectors $\mathbf{t}_k = \{t_k^j\}_{j=0}^{n_k+p}, k = 1, \ldots, d$.

Given a tuple $u = (u^1, \ldots, u^d) \in [0, 1]^d$, the tensor product basis functions are defined as

$$N_{\alpha}(u) = \prod_{k=1}^{d} N_{\alpha^{k}}^{k}(u^{k}).$$
 (3)

where α is a multi-index as described above and $N_{\alpha^k}^k$ is the $(\alpha^k)^{th}$ basis function with respect to the knot vector \mathbf{t}_k . With $n_k + p + 1$ total knots in each dimension,

³ We consider only "clamped" knot sequences in this paper; thus, the first p+1 knots are always 0 and the last p+1 knots are always 1.

there are n_k basis functions in each dimension.⁴ Therefore, the total number of tensor product basis functions is $n_{tot} = \prod_{k=1}^{d} n_k$, which is also the total number of control points for the tensor product spline.

A d-dimensional tensor product B-spline in \mathbb{R}^D is a function of the form

$$C(u) = \sum_{\alpha \in A} N_{\alpha}(u) P_{\alpha}, \tag{4}$$

where A is the set of all basis functions and $P_{\alpha} \in \mathbb{R}^{D}$ for each α .

3.3 Optimal Control Points

Given a collection of knot vectors and polynomial degree, the best-fit B-spline to a given data set is determined by a linear least-squares minimization problem. Let $\{Q_i\}_{i=0}^{m-1}$ be the list of points in \mathbb{R}^D to be approximated with a *d*-dimensional tensor product spline. For each $0 \leq i < m$, let $v_i \in [0, 1]^d$ be the parameter tuple corresponding to the point Q_i . The optimal control points are determined by the least-squares minimization problem:

$$\{\hat{P}_j\} = \underset{P_j}{\operatorname{argmin}} \sum_{i=0}^{m-1} \|Q_i - C(v_i)\|^2.$$
(5)

This minimization problem can be rewritten in normal form by differentiating the objective function in Equation (5) with respect to each of the control points. The normal system reduces to the matrix equation $\mathbf{N}^T \mathbf{N} \mathbf{P} = \mathbf{N}^T \mathbf{Q}$, where

$$\mathbf{N}_{ij} = N_{\alpha}(v_i) \quad \text{where } [\alpha]_A = j, \qquad \mathbf{P}_{ij} = P_{\alpha}^j, \quad \text{where } [\alpha]_A = i, \qquad \mathbf{Q}_{ij} = Q_i^j.$$
(6)

The superscripts in the above equations index the components of the vectors P_{α} and Q_i . **N** is a $m \times n_{tot}$ matrix, often called the "B-spline collocation matrix," **P** is an $n_{tot} \times D$ matrix with each row containing a control point, and **Q** is a $m \times D$ matrix with each row containing an input point.

Typically, the matrix **N** is very sparse, and this system may be solved with an iterative method or sparse direct solver. However, as we show in the following section, this system is ill-conditioned when the sample density of the input points P_i varies from region to region.

4 Adaptive Regularization

A significant challenge when modeling unstructured data with tensor product B-splines is the (ill-)conditioning of the fitting procedure. Generally speaking, tensor product B-spline models can oscillate strongly due to overfitting in regions where input data is sparse (see Figure 1). Our method of adaptive regularization

⁴ Here we assume for simplicity that the degree of the B-spline is the same in each dimension, but the degree can vary in practice if desired.

produces a unique solution to systems which would otherwise be rank-deficient and improves the overall conditioning of the system. In practice, the adaptively regularized models possess fewer oscillatory artifacts and do not exhibit any divergent behavior in our testing. In contrast to standard regularization techniques, our method does not smooth out the model indiscriminately – instead, it regularizes only those regions that require smoothing.



Fig. 1. Left: A data set with nonuniform point density. Center: Best-fit B-spline model without regularization. Right: B-spline model with adaptive regularization. The center image is cropped; spike artifacts in this model extend well outside the frame.

This technique employs a spatially-varying regularization strength that is computed automatically as a function of the relative positioning of input data points to the B-spline knots. In general, the regularization strength increases in regions with little to no input data and decreases (potentially to zero) in regions "saturated" with input points. When the regularization strength is zero throughout a region of the domain, no smoothing is performed in that region; therefore, any sharp features present in densely sampled regions of the domain will be preserved by the adaptive regularization procedure.

Standard roughness minimization can be formulated as a penalized leastsquares minimization problem, similar to Equation (5). The penalty term weights the size of the second derivative at each point with a new parameter, which we denote as $\lambda > 0$. The control points of the regularized spline are defined by:

$$\{\hat{P}_j\} = \operatorname*{argmin}_{P_j} \left(\sum_{i=0}^{m-1} \|Q_i - C(v_i)\|^2 + \lambda^2 S(C) \right),$$
(7)

where S(C) approximates the size of the second derivatives of C.

Let $w_{\alpha} \in [0,1]^d$ be the parameter that maximizes the value of N_{α} . Let ∂^{δ} denote the partial derivative where the order of derivative in each dimension is given by the components of multi-index δ^{5} . We define S(C) to be

$$S(C) = \sum_{\alpha \in A} \sum_{|\delta|=2} \left\| \partial^{\delta} C(w_{\alpha}) \right\|^{2}.$$
(8)

Note that the summation above is a sum over all derivatives of order 2, including mixed partial derivatives.

⁵ For example,
$$\partial^{(2,0)}f = \partial^2 f/\partial x_1^2$$
, and $\partial^{(0,2)}f = \partial^2 f/\partial x_2^2$, while $\partial^{(1,1)}f = \partial^2 f/(\partial x_1 \partial x_2)$.

Equation (7) can be converted into a system of equations in the same way as Equation (5). The only additional step is computing the derivative of S(C)with respect to the control points. In matrix form, the system is

$$\left(\mathbf{N}^T \ \lambda \mathbf{M}^T\right) \begin{pmatrix} \mathbf{N} \\ \lambda \mathbf{M} \end{pmatrix} \mathbf{P} = \mathbf{N}^T \mathbf{Q}, \qquad \text{where } \mathbf{M} = \begin{pmatrix} \mathbf{M}_{\delta_1} \\ \vdots \\ \mathbf{M}_{\delta_n} \end{pmatrix}$$
(9)

and $(\mathbf{M}_{\delta})_{i,j} = \partial^{\delta} N_{\beta}(w_{\alpha})$, $[\alpha]_{A} = i$, $[\beta]_{A} = j$. Intuitively, each column of \mathbf{M}_{δ} describes the ∂^{δ} partial derivative of an individual B-spline basis function. The matrix \mathbf{M} is the concatenation of all the individual \mathbf{M}_{δ} matrices, where $|\delta| = 2$. **N**, **P**, and **Q** are defined as in Section 3.

The novel improvement of our adaptive regularization scheme is to modify the above system of equations by varying the size of λ for each column of **M**. Since each column of this matrix corresponds to a B-spline basis function and control point, variation in the size of λ provides a mechanism to modify the smoothing conditions on each control point of the spline individually. Due to the local support property of B-splines, setting $\lambda_j = 0$ for control points in a given region "disables" the regularization in that region, while still allowing for smoothing to be applied elsewhere. Algebraically, we replace the scalar parameter λ by a diagonal matrix $\mathbf{\Lambda} = \text{diag}(\lambda_1, \ldots, \lambda_{n_{tot}})$, where each $\lambda_j \geq 0$, and consider the new linear system

$$\left(\mathbf{N}^T \ (\mathbf{M}\mathbf{\Lambda})^T\right) \begin{pmatrix} \mathbf{N} \\ \mathbf{M}\mathbf{\Lambda} \end{pmatrix} \mathbf{P} = \mathbf{N}^T \mathbf{Q}.$$
 (10)

The value of each λ_i is computed automatically as a function of the relative positioning between input data points and B-spline knots.

To better control this function, we introduce a user-specified parameter called the "regularization threshold," denoted s^* . Changing the regularization threshold adjusts the criterion by which some regions of the domain are smoothed and others are not. As s^* increases, smoothing constraints will be applied to larger and larger regions in the domain.

Let s_j denote the j^{th} column sum of **N** and \tilde{s}_j the j^{th} column sum of **M**. Given $s^* \ge 0$, we define

$$\lambda_j = \frac{\max(s^* - s_j, 0)}{\widetilde{s}_j}.$$
(11)

Thus, Λ is defined such that every column sum of $\binom{\mathbf{N}}{\mathbf{M}\Lambda}$ is no less than s^* .

Adapting the regularization strengths λ_j this way has a number of important results. If $s^* = 0$, then $\mathbf{\Lambda} = 0$ and the minimization becomes the usual least-squares problem. When s^* is small, λ_j will be zero unless the j^{th} column sum of \mathbf{N} is small, which is indicative of an ill-conditioned system. Here, adaptive regularization smooths out only those control points which are poorly constrained.

This formulation also explains why the adaptive regularization method preserves sharp, densely sampled features while smoothing out oscillatory artifacts.

In regions of the domain that are densely sampled, control points will be constrained by many data points and thus the corresponding column sum in **N** will be relatively large. By choosing s^* to be sufficiently small, all control points in this region will have a regularization strength of zero; i.e. $\lambda_j = 0$. Therefore, the best-fit spline in this region will not be artificially smoothed.

Finally, we remark that the adaptive regularization framework, while defined above in terms of second derivatives, can easily be extended to other derivatives. We find that constraining first and second derivatives simultaneously is particularly helpful when modeling data sets with no points at all in certain regions. This typically happens when the data represent an object with an interior hole or irregular boundary. To consider both first and second derivatives, the only change is to matrix **M** in Equation (10). Originally, **M** is the concatenation of all matrices \mathbf{M}_{δ} , where δ describes a second derivative. To minimize first derivatives as well, we change this to the concatenation of all \mathbf{M}_{δ} such that δ describes a first or second derivative. Both versions of the method are described in Section 5.

5 Results

We demonstrate the effectiveness of our method with a series of numerical experiments. First, we compare adaptive regularization to uniform regularization where the smoothing parameter has been chosen manually. Next, we study the reconstruction of an analytical signal from sparse samples with varying levels of sparsity. For each sparsity level, we report the error and condition number for and unregularized and adaptively regularized model. We then test the performance of adaptive regularization on data sets with no data in certain regions. In these problems, we construct a B-spline model that extrapolates into regions with no pointwise constraints, and check that the adaptive regularization method produces a reasonable result.

5.1 Data Sets

The performance of the adaptive regularization method was studied on a collection of two- and three-dimensional point clouds with different characteristics. Some data sets were sampled from analytical functions so that we could compute pointwise errors relative to a ground truth, while other data sets were generated by scientific experiments and simulations.

2D Polysinc. The polysinc data set is a two-dimensional point cloud sampling the function $f(x, y) = \operatorname{sinc} (x^2 + y^2) \operatorname{sinc} (2(x-2)^2 + (y+2)^2).^6$ 360,000 point locations are uniformly sampled from the box domain $[-4\pi, 4\pi] \times [-4\pi, 4\pi]$, except at four disk-shaped regions where the sample rate is 50× lower.

XGC Fusion. The XGC fusion data set represents a normalized derivative of electrostatic potential in a single poloidal plane of a Tokamak fusion simulation. The data set contains 56,980 points with an irregular boundary and was produced by the XGC code [12] in a simulation of the gyrokinetic equations.

⁶ We consider the unnormalized sinc function: $\operatorname{sin}(x) = \frac{\sin(x)}{x}$, with $\operatorname{sinc}(0) = 1$.

CMIP6 Climate. The CMIP6 climate data set represents ocean surface temperature in a projected box region around Antarctica. The data set contains 585,765 points with a large hole (representing Antarctica) in the center and was produced by a Coupled Model Intercomparison Project (CMIP6) [7] simulation.

sahex Nuclear. The sahex nuclear data set is derived from a simulation of a single nuclear reactor component, produced with the SHARP toolkit [21]. The three-dimensional data are bounded by a hexagonal prism and point density is coarser in the z dimension than x and y. The data set contains 63,048 points.

5.2 Comparison of Adaptive versus Uniform Regularization

Applying a uniform regularization strength to an entire model can produce unsatisfactory results, because sufficiently smoothing oscillatory artifacts can also smooth out sharp features. Figure 2 compares our adaptive regularization scheme (with $s^* = 6$) against three strengths of uniform regularization. The data in Figure 2 is the XGC fusion data set, which contains sharp peaks in a ring but is flat inside the ring. Data are sparse or nonexistent outside the ring. The two images at right show a model with uniform regularization that is too weak, causing artifacts (top), or too strong, dampening the features (bottom). The best uniform regularization strength we could find is given at bottom-left, but even in this example the characteristic peaks in the data are smoothed down.



Fig. 2. Comparison of uniform vs adaptive regularization. Clockwise from top-left: Adaptive regularization, uniform regularization with $\lambda = 10^{-6}$, uniform regularization with $\lambda = 10^{-5}$.

5.3 Accuracy on Analytical Signals

To quantify the accuracy of B-spline models with adaptive regularization, we consider the oscillatory polysinc function with a highly nonuniform input data set (Figure 3). We illustrate two B-spline models, one fit without regularization and one with our adaptive regularization $(s^* = 1)$. Both models are degree

four with a 300 \times 300 grid of control points. Without regularization, the model diverges in the regions of low sample density; with adaptive regularization, the model produces an accurate representation even where sample density is low. A top-down view of the error profiles is given in the second row of Figure 3. A close comparison of the ground-truth (top left) and adaptively-regularized spline (top right) shows that the spline model is not artificially smoothed in dense regions, even though the sparse regions are smoothed. In particular, our regularization procedure preserves the distinctive oscillations and peaks in the signal.



Fig. 3. Top row: Synthetic polysinc signal (left), model with no regularization (center), model with adaptive regularization (right). Bottom row: Top down view of input distribution (left), error profile with no regularization (center), error profile with adaptive regularization (right). Area of interest for error calculation is in red at bottom-left.

The degree of sparsity in the input data strongly influences the accuracy of a B-spline model. Table 1 lists the errors in each model for varying levels of sparsity in the voids. The errors are measured in a box around two voids (see Figure 3) in order to pinpoint the behavior of the models in this region. When the point density is equal inside and outside of the voids (sparsity = 1.0), error for both models is low. As the voids become more sparse, the error in the unregularized model increases by four orders of magnitude while error in the adaptively regularized model stays essentially flat.

Data sparsity also affects the condition number of the least-squares minimization. Table 1 gives the condition number of the matrices N ('noreg') and $\binom{N}{M\Lambda}$ ('reg') for each sparsity level. As sparsity is increased, the condition number of $\binom{N}{M\Lambda}$ remains lower and steady but the condition number of N starts higher and eventually becomes infinite. Condition numbers were computed with the Matlab routine svds.

Table 1. Model errors as a function of sparseness. Maximum and L^2 (average) errors are computed for both adaptively regularized and unregularized models in the vicinity of two voids (see Figure 3). Condition numbers for both minimization problems are reported at bottom.

Sparsity	0.02	0.08	0.16	0.32	0.64	1.00
Max Error (reg)	3.25e-2	2.89e-2	2.71e-2	2.39e-2	1.20e-2	1.16e-2
Max Error (no reg)	1.29e2	6.27e2	2.36e-1	3.65e-2	1.20e-2	1.16e-2
L^2 Error (reg)	1.93e-3	1.53e-3	1.13e-3	7.04e-4	5.56e-4	5.44e-4
L^2 Error (no reg)	2.17e0	4.68e0	3.42e-3	7.35e-4	5.56e-4	5.44e-4
Condition $\#$ (reg)	177	980	289	198	121	189
Condition $\#$ (noreg)	inf	\inf	6.97 e4	2.58e3	1.57e3	3.74e3

5.4 Extrapolation into Unconstrained Regions

When a large region of the domain does not contain any data points to constrain the best-fit B-spline problem, the least-squares minimization will be illposed and the resulting model can exhibit extreme oscillations. However, data sets with empty regions or "holes" are very common in scientific and industrial applications. For example, some climate models measure ocean temperatures or land temperatures, but not both simultaneously. Industrial simulations often model objects with irregular boundaries, and data from physics simulations are shaped by the locations of detectors.

Although empty regions are usually omitted in subsequent analysis, it is still important to understand and control the behavior of a B-spline model in empty regions. Extreme oscillations near the boundary of a hole can distort the derivative of the model away from this boundary. In addition, attempting to compute simple statistics about the model (minimum, maximum, mean) can be biased if the model exhibits unpredictable behavior in empty regions.

Figure 4 shows a scenario from a simulation of ocean temperatures. The data set is centered around the continent of Antarctica, for which no temperature values are given. Exhibited in the figure are two models, both of degree 2 with a 400 × 400 grid of control points. The unregularized model at center oscillates between $\pm 10^8$ along the coast (while the input data range from -2 to 10). In contrast, the regularized model (with $s^* = 5$) smoothly transitions at the coast to a near-constant value over the landmass. Regularization was performed with constraints on first and second derivatives, as described at the end of Section 4.

We next consider a three-dimensional data set representing power produced in a component of a nuclear reactor (Figure 5). The data are contained inside a hexagonal prism, but the B-spline model is defined on the bounding box of this prism. Hence the corners of this box are devoid of any data. Without regularization, the least-squares minimization does not converge, so we exhibit only our regularized model (with $s^* = 10$) in Figure 5. The adaptive regularization method (Figure 5, center and right) produces an accurate model of the six interior "pins" and is well defined in the corner regions. Some artifacts are observed in the corners, but they are not significant enough to affect the interior of the



Fig. 4. Ocean temperature simulation around Antarctica. Top row: Input data on a hexagonal mesh (left), B-spline model without regularization (center), B-spline model with adaptive regularization (right). Bottom row: Detail view from the top row.

model. Without regularization, the condition number of the system is infinite; with adaptive regularization, the condition number is 2.18×10^5 . The model was produced by constraining first and second derivatives simultaneously.



Fig. 5. Power production in a nuclear reactor simulation. Input data (left), spline model (center), spline model top view (right).

6 Future Work

In the construction of our method, we imposed constraints on the first and second derivatives of the B-spline in regions where the density of input data was low or vanishing. However, these additional constraints could have been based on different orders of derivative or been unrelated to derivatives altogether. For instance, a different type of constraint would be one that penalizes deviation from a known baseline value. Another option would be to penalize B-spline values

that exceed a given range (for example, the original bounds of the input data). We remark that our procedure for adaptive regularization of tensor product B-splines is separate from the type of artificial constraint imposed. Depending on the application, different constraints may be more useful, and our method allows for those to be used instead.

Another direction for future research is an investigation of the parameter s^* . While our method automatically varies regularization strength throughout the domain, these strengths are all relative to the parameter s^* . Choosing a value of s^* too large can lead to overly-smoothed models. Further research into heuristics or iterative schemes to select s^* automatically would allow this method to be applied with no user interaction at all.

7 Conclusions

Modeling unstructured data sets with tensor product B-splines can be difficult due to the ill-conditioning of the fitting problem. In general, data sets with large variations in point density or regions without data exacerbate this problem to the point that artificial smoothing is necessary. However, smoothing an entire model can wash out sharp features in the data.

We introduced a regularization procedure for B-spline models that preserves features by adapting the regularization strength throughout the domain. Our method automatically varies the smoothing intensity as a function of input point density and relies on a single user-specified parameter, which we call the regularization threshold. We observe that adaptive regularization performs better than typical uniform regularization schemes that may over-smooth some regions while under-smoothing others. We also showed that our method can fit B-spline models to data sets with regions of extremely sparse point density and remain well-defined even in areas without data points. Overall, adaptive regularization of B-spline models produces smooth and accurate models for data sets which would otherwise be difficult to fit.

References

- Arigovindan, M., Suhling, M., Jansen, C., Hunziker, P., Unser, M.: Full motion and flow field recovery from echo doppler data. IEEE Transactions on Medical Imaging 26(1), 31–45 (2006). https://doi.org/10.1109/TMI.2006.884201
- Craven, P., Wahba, G.: Smoothing noisy data with spline functions. Numerische mathematik 31(4), 377–403 (1978). https://doi.org/10.1007/BF01404567
- 3. De Boor, C.: A practical guide to splines. Springer-Verlag New York (2001)
- Duchon, J.: Splines minimizing rotation-invariant semi-norms in sobolev spaces. In: Schempp, W., Zeller, K. (eds.) Constructive Theory of Functions of Several Variables. pp. 85–100. Springer Berlin Heidelberg (1977). https://doi.org/10.1007/ BFb0086566
- Duijndam, A., Schonewille, M., Hindriks, C.: Reconstruction of band-limited signals, irregularly sampled along one spatial direction. Geophysics 64(2), 524–538 (1999). https://doi.org/10.1190/1.1444559

- 14 D. Lenz et al.
- El-Rushaidat, D., Yeh, R., Tricoche, X.M.: Accurate parallel reconstruction of unstructured datasets on rectilinear grids. Journal of Visualization pp. 1–20 (2021). https://doi.org/10.1007/s12650-020-00740-0
- Eyring, V., Bony, S., Meehl, G.A., Senior, C.A., Stevens, B., Stouffer, R.J., Taylor, K.E.: Overview of the coupled model intercomparison project phase 6 (cmip6) experimental design and organization. Geoscientific Model Development 9(5), 1937– 1958 (2016). https://doi.org/10.5194/gmd-9-1937-2016
- Francis, B., Viswanath, S., Arigovindan, M.: Scattered data approximation by regular grid weighted smoothing. Sādhanā 43(1), 1–16 (2018). https://doi.org/10. 1007/s12046-017-0765-y
- 9. Gu, C.: Cross-validating non-gaussian data. Journal of Computational and Graphical Statistics 1(2), 169–179 (1992). https://doi.org/10.1080/10618600.1992. 10477012
- Gu, C.: Smoothing spline ANOVA models, vol. 297. Springer (2013). https://doi. org/10.1007/978-1-4757-3683-0
- Hughes, T., Cottrell, J., Bazilevs, Y.: Isogeometric analysis: Cad, finite elements, nurbs, exact geometry and mesh refinement. Computer Methods in Applied Mechanics and Engineering 194(39), 4135 – 4195 (2005). https://doi.org/10.1016/j. cma.2004.10.008
- Ku, S., Hager, R., Chang, C.S., Kwon, J., Parker, S.E.: A new hybrid-lagrangian numerical scheme for gyrokinetic simulation of tokamak edge plasma. Journal of Computational Physics **315**, 467–475 (2016). https://doi.org/10.1016/j.jcp.2016. 03.062
- Lee, S., Wolberg, G., Shin, S.: Scattered data interpolation with multilevel bsplines. IEEE Transactions on Visualization and Computer Graphics 3(3), 228–244 (1997). https://doi.org/10.1109/2945.620490
- Lin, H., Maekawa, T., Deng, C.: Survey on geometric iterative methods and their applications. Computer-Aided Design 95, 40 – 51 (2018). https://doi.org/10.1016/ j.cad.2017.10.002
- Nashed, Y.S., Peterka, T., Mahadevan, V., Grindeanu, I.: Rational approximation of scientific data. In: Rodrigues, J.M.F., Cardoso, P.J.S., Monteiro, J., Lam, R., Krzhizhanovskaya, V.V., Lees, M.H., Dongarra, J.J., Sloot, P.M. (eds.) Computational Science – ICCS 2019. pp. 18–31. Springer (2019). https://doi.org/10.1007/ 978-3-030-22734-0_2
- Peterka, T., Nashed, Y., Grindeanu, I., Mahadevan, V., Yeh, R., Trixoche, X.: Foundations of multivariate functional approximation for scientific data. In: Proceedings of 2018 IEEE Symposium on Large Data Analysis and Visualization (2018). https://doi.org/10.1109/LDAV.2018.8739195
- Piegl, L., Tiller, W.: The NURBS book. Springer-Verlag, 2 edn. (1997). https: //doi.org/10.1007/978-3-642-59223-2
- Vio, R., Strohmer, T., Wamsteker, W.: On the reconstruction of irregularly sampled time series. Publications of the Astronomical Society of the Pacific 112(767), 74 (2000). https://doi.org/10.1086/316495
- Wood, S.N.: Thin plate regression splines. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 65(1), 95–114 (2003). https://doi.org/10.1111/ 1467-9868.00374
- 20. Yeh, R.: Efficient Knot Optimization for Accurate B-spline-based Data Approximation. Ph.D. thesis, Purdue University Graduate School (2020)
- Yu, Y., Shemon, E., Mahadevan, V.S., Rahaman, R.O.: Sharp multiphysics tutorials. Tech. Rep. ANL/NE-16/1, Argonne National Lab.(ANL), Lemont, IL (United States) (2016). https://doi.org/10.2172/1250465