

Additive Models with Trend Filtering

Veeranjaneyulu Sadhanala

Ryan J. Tibshirani

Abstract

We study additive models built with trend filtering, i.e., additive models whose components are each regularized by the (discrete) total variation of their k th (discrete) derivative, for a chosen integer $k \geq 0$. This results in k th degree piecewise polynomial components, (e.g., $k = 0$ gives piecewise constant components, $k = 1$ gives piecewise linear, $k = 2$ gives piecewise quadratic, etc.). Analogous to its advantages in the univariate case, additive trend filtering has favorable theoretical and computational properties, thanks in large part to the localized nature of the (discrete) total variation regularizer that it uses. On the theory side, we derive fast error rates for additive trend filtering estimates, and show these rates are minimax optimal when the underlying function is additive and has component functions whose derivatives are of bounded variation. We also show that these rates are unattainable by additive smoothing splines (and by additive models built from linear smoothers, in general). On the computational side, we use backfitting, to leverage fast univariate trend filtering solvers; we also describe a new backfitting algorithm whose iterations can be run in parallel, which (as far as we can tell) is the first of its kind. Lastly, we present a number of experiments to examine the empirical performance of trend filtering.

1 Introduction

As the dimension of the input space grows large, nonparametric regression turns into a notoriously difficult problem. In this work, we adopt the stance taken by many others, and consider an *additive model* for responses $Y^i \in \mathbb{R}$, $i = 1, \dots, n$ and corresponding input points $X^i = (X_1^i, \dots, X_d^i) \in \mathbb{R}^d$, $i = 1, \dots, n$, of the form

$$Y^i = \mu + \sum_{j=1}^d f_{0j}(X_j^i) + \epsilon^i, \quad i = 1, \dots, n,$$

where $\mu \in \mathbb{R}$ is an overall mean parameter, each f_{0j} is a univariate function with $\sum_{i=1}^n f_{0j}(X_j^i) = 0$ for identifiability, $j = 1, \dots, d$, and the errors ϵ^i , $i = 1, \dots, n$ are i.i.d. with mean zero. A comment on notation: here and throughout, when indexing over the n samples we use superscripts, and when indexing over the d dimensions we use subscripts, so that, e.g., X_j^i denotes the j th component of the i th input point. (Exceptions will occasionally be made, but the role of the index should be clear from the context.)

Additive models are a special case of the more general *projection pursuit regression* model of [Friedman and Stuetzle \[1981\]](#). Additive models for the Cox regression and logistic regression settings were studied in [Tibshirani \[1983\]](#) and [Hastie \[1983\]](#), respectively. Some of the first asymptotic theory for additive models was developed in [Stone \[1985\]](#). Two algorithms closely related to (backfitting for) additive models are the *alternating least squares* and *alternating conditional expectations* methods, from [van der Burg and de Leeuw \[1983\]](#) and [Breiman and Friedman \[1985\]](#), respectively. The work of [Buja et al. \[1989\]](#) advocates for the use of additive models in combination with linear smoothers, a surprisingly simple combination that gives rise to flexible and scalable multidimensional regression tools. The book by [Hastie and Tibshirani \[1990\]](#) is the definitive practical guide for additive models for exponential family data distributions, i.e., generalized additive models.

More recent work on additive models is focused on high-dimensional nonparametric estimation, and here the natural goal is to induce sparsity in the component functions, so that only a few select

dimensions of the input space are used in the fitted additive model. Some nice contributions are given in [Lin and Zhang \[2006\]](#), [Ravikumar et al. \[2009\]](#), [Meier et al. \[2009\]](#), all primarily focused on fitting splines for component functions and achieving sparsity through a group lasso type penalty. In other even more recent and interesting work sparse additive models, [Lou et al. \[2016\]](#) consider a semiparametric (partially linear) additive model, and [Petersen et al. \[2016\]](#) study componentwise fused lasso (i.e., total variation) penalization in an additive model.

The literature on additive models (and by now, sparse additive models) is vast and the above is far from a complete list of references. In this paper, we examine a method for estimating additive models wherein each component is fit in a way that is *locally adaptive* to the underlying smoothness along its associated dimension of the input space. The literature on this line of work, as far as we can tell, is much less extensive. First, we review linear smoothers in additive models, motivate our general goal of local adaptivity, and then describe our specific proposal.

1.1 Review: additive models and linear smoothers

The influential paper by [Buja et al. \[1989\]](#) studies additive minimization problems of the form

$$\begin{aligned} \min_{\theta_1, \dots, \theta_d \in \mathbb{R}^n} \quad & \left\| Y - \bar{Y} \mathbb{1} - \sum_{j=1}^d \theta_j \right\|_2^2 + \lambda \sum_{j=1}^d \theta_j^T Q_j \theta_j \\ \text{subject to} \quad & \mathbb{1}^T \theta_j = 0, \quad j = 1, \dots, d, \end{aligned} \tag{1}$$

where $Y = (Y^1, \dots, Y^n) \in \mathbb{R}^n$ denotes the vector of responses, and $Y - \bar{Y} \mathbb{1}$ is its centered version, with $\bar{Y} = \frac{1}{n} \sum_{i=1}^n Y^i$ denoting the sample mean of Y , and $\mathbb{1} = (1, \dots, 1) \in \mathbb{R}^n$ the vector of all 1s. Each vector $\theta_j = (\theta_j^1, \dots, \theta_j^n) \in \mathbb{R}^n$ represents the evaluations of the j th component function f_j in our model, i.e., tied together by the relationship

$$\theta_j^i = f_j(X_j^i), \quad i = 1, \dots, n, \quad j = 1, \dots, d.$$

In the problem (1), $\lambda \geq 0$ is a regularization parameter and Q_j , $j = 1, \dots, d$ are penalty matrices. As a typical example, we might consider Q_j to be the Reinsch penalty matrix for smoothing splines along the j th dimension of the input space, for $j = 1, \dots, d$. Under this choice, a backfitting (block coordinate descent) routine for (1) would repeatedly cycle through the updates

$$\theta_j = (I + \lambda Q_j)^{-1} \left(Y - \bar{Y} \mathbb{1} - \sum_{\ell \neq j} \theta_\ell \right), \quad j = 1, \dots, d, \tag{2}$$

where the j th update fits a smoothing spline to the j th partial residual, over the j th dimension of the input points, denoted by $X_j = (X_j^1, X_j^2, \dots, X_j^n) \in \mathbb{R}^n$. At convergence, we arrive at an additive smoothing spline estimate, which solves (1).

Modeling the component functions as smoothing splines is arguably the most common formulation for additive models, and it is the standard in several statistical software packages like the R package `gam`. However, as [Buja et al. \[1989\]](#) explain, the backfitting perspective suggests a more algorithmic approach to additive modeling: one can replace the operator $(I + \lambda Q_j)^{-1}$ in (2) by S_j , a particular (user-chosen) *linear smoother*, meaning, a linear map that performs univariate smoothing across the j th dimension of inputs X_j . The linear smoothers S_j , $j = 1, \dots, d$ could correspond to smoothing splines, regression splines (regression using a spline basis with given knots), kernel smoothing, local polynomial smoothing, or a combination of these, across the input dimensions.

The class of linear smoothers is broad enough to offer fairly flexible, interesting mechanisms for smoothing, and simple enough to understand precisely. [Buja et al. \[1989\]](#) provide a unified analysis of additive models with linear smoothers, in which they derive the effective degrees of freedom of these estimators and a generalized cross-validation routine for tuning; they also study fundamental

properties such as uniqueness of the component fits, and convergence of the backfitting steps. Most of the work following [Buja et al. \[1989\]](#) remains in keeping with the idea of using linear smoothers in combination with additive models.

1.2 The limitations of linear smoothers

The beauty of linear smoothers lies in their simplicity. However, with this simplicity comes serious limitations, in terms of their ability to adapt to varying local levels of smoothness. In the univariate setting, the seminal theoretical work by [Donoho and Johnstone \[1998\]](#) makes this idea precise. With $d = 1$, suppose that underlying regression function f_0 lies in the univariate function class

$$\mathcal{F}_k(C) = \{f : \text{TV}(f^{(k)}) \leq C\}, \quad (3)$$

for a constant $C > 0$, where $\text{TV}(\cdot)$ is the total variation operator, and $f^{(k)}$ the k th weak derivative of f . The class in (3) allows for greater fluctuation in the local level of smoothness of f_0 than, say, more typical function classes like Holder and Sobolev spaces. The results of [Donoho and Johnstone \[1998\]](#) (see also Section 5.1 of [Tibshirani \[2014\]](#)) imply that the minimax error rate for estimation over $\mathcal{F}_k(C)$ is $n^{-(2k+2)/(2k+3)}$, but the minimax error rate when we consider only linear smoothers (linear transformations of Y) is $n^{-(2k+1)/(2k+2)}$. This difference is highly nontrivial, e.g., for $k = 0$ this is a difference of $n^{-2/3}$ (optimal) versus $n^{-1/2}$ (optimal among linear smoothers) for estimating a function f_0 of bounded variation.

It is important to emphasize that this shortcoming is not just a theoretical one; it is also clearly noticeable in basic practical examples. Just as linear smoothers will struggle in the univariate setting, an additive estimate based on linear smoothers will not be able to efficiently track local changes in smoothness, across any of the input dimensions. This could lead to a loss in accuracy even if only some of the components f_{0j} , $j = 1, \dots, d$ possesses heterogeneous smoothness across its domain.

Two well-studied univariate estimators that are locally adaptive, i.e., that attain the minimax error rate over the k th order total variation class in (3), are wavelet smoothing and locally adaptive regression splines, as developed by [Donoho and Johnstone \[1998\]](#) and [Mammen and van de Geer \[1997\]](#), respectively. There is a substantial literature on these methods in the univariate case (especially for wavelets), but fewer authors have considered them in the additive models context. Some notable exceptions are [Zhang and Wong \[2003\]](#), [Sardy and Tseng \[2004\]](#), [Petersen et al. \[2016\]](#), with the latter work especially related to our focus in this paper.

1.3 Additive trend filtering

We consider additive models that are constructed using *trend filtering* (instead of linear smoothers, wavelets, or locally adaptive regression splines) as their componentwise smoother. Proposed independently by [Steidl et al. \[2006\]](#) and [Kim et al. \[2009\]](#), trend filtering is a relatively new approach to univariate nonparametric regression. As explained in [Tibshirani \[2014\]](#), it can be seen as a discrete-time analog of the locally adaptive regression spline estimator. Denoting by $X = (X^1, \dots, X^n) \in \mathbb{R}^n$ the vector of univariate input points, where we assume $X^1 < \dots < X^n$, the trend filtering estimate of order $k \geq 0$ is defined as the solution of the optimization problem

$$\min_{\theta \in \mathbb{R}^n} \frac{1}{2} \|Y - \theta\|_2^2 + \lambda \|D^{(X,k+1)}\theta\|_1, \quad (4)$$

where $\lambda \geq 0$ is a tuning parameter, and $D^{(X,k+1)} \in \mathbb{R}^{(n-k-1) \times n}$ is a k th order difference operator, constructed based on X . These difference operators can be defined recursively, as in

$$D^{(X,1)} = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 & 0 \\ 0 & -1 & 1 & \dots & 0 & 0 \\ \vdots & & & & & \\ 0 & 0 & 0 & \dots & -1 & 1 \end{bmatrix} \in \mathbb{R}^{(n-1) \times n}, \quad (5)$$

and for $k = 1, 2, 3, \dots$,

$$D^{(X,k+1)} = D^{(X,1)} \text{diag} \left(\frac{k}{X^k - X^1}, \dots, \frac{k}{X^n - X^{n-k+1}} \right) D^{(X,k)} \in \mathbb{R}^{(n-k-1) \times n}. \quad (6)$$

(The leading matrix $D^{(X,1)}$ in (6) is the $(n-k-1) \times (n-k)$ version of the difference operator in (5).) Intuitively, the interpretation is that the problem (4) penalizes the sum of absolute $(k+1)$ st order discrete derivatives of $\theta^1, \dots, \theta^n$ across the input points X^1, \dots, X^n . Thus, at optimality, the coordinates of the trend filtering solution $\hat{\theta}^1, \dots, \hat{\theta}^n$ obey a k th order piecewise polynomial form.

This intuition is formalized in Tibshirani [2014] and Wang et al. [2014], where it is shown that the components of the k th order trend filtering estimate $\hat{\theta}$ are precisely the evaluations of a fitted k th order piecewise polynomial function across the inputs, and that the trend filtering and locally adaptive regression spline estimates of the same order k are asymptotically equivalent. When $k = 0$ or $k = 1$, in fact, there is no need for asymptotics, and the equivalence between trend filtering and locally adaptive regression spline estimates is exact in finite samples. It is also worth pointing out that when $k = 0$, the trend filtering estimate reduces to the 1d fused lasso estimate [Tibshirani et al., 2005], which is known as 1d total variation denoising in signal processing [Rudin et al., 1992].

Over the k th order total variation function class defined in (3), Tibshirani [2014], Wang et al. [2014] prove that k th order trend filtering achieves the minimax optimal $n^{-(2k+2)/(2k+3)}$ error rate, just like k th order locally adaptive regression splines. Another important property, as developed by Kim et al. [2009], Tibshirani [2014], Ramdas and Tibshirani [2016], is that trend filtering estimates are relatively cheap to compute—much cheaper than locally adaptive regression spline estimates—owing to the bandedness of the difference operators in (5), (6), which means that specially implemented convex programming routines can solve (4) in an efficient manner.

It is this computational efficiency, along with its capacity for local adaptivity, that makes trend filtering a particularly desirable candidate to extend to the additive model setting. Specifically, we consider the *additive trend filtering* estimate of order $k \geq 0$, defined as a solution in the problem

$$\begin{aligned} \min_{\theta_1, \dots, \theta_d \in \mathbb{R}^n} \quad & \frac{1}{2} \left\| Y - \bar{Y} \mathbb{1} - \sum_{j=1}^d \theta_j \right\|_2^2 + \lambda \sum_{j=1}^d \|D^{(X_j, k+1)} S_j \theta_j\|_1 \\ \text{subject to} \quad & \mathbb{1}^T \theta_j = 0, \quad j = 1, \dots, d. \end{aligned} \quad (7)$$

As before, $Y - \bar{Y} \mathbb{1}$ is the centered response vector, $\lambda \geq 0$ is a regularization parameter, and now $S_j \in \mathbb{R}^{n \times n}$ in (7) is a permutation matrix that sorts the j th component of inputs $X_j = (X_j^1, X_j^2, \dots, X_j^n)$ into increasing order, i.e.,

$$S_j X_j = (X_j^{(1)}, X_j^{(2)}, \dots, X_j^{(n)}), \quad j = 1, \dots, d.$$

Also, $D^{(X_j, k+1)}$ in (7) is the $(k+1)$ st order difference operator, as in (5), (6), but defined over the sorted j th dimension of inputs $S_j X_j$, for $j = 1, \dots, d$. With backfitting (block coordinate descent), computation of a solution in (7) is still quite efficient, since we can leverage the efficient routines for univariate trend filtering.

1.4 A motivating example

Figure 1 shows a simulated example that compares the additive trend filtering estimates in (7) (of quadratic order, $k = 2$), to the additive smoothing spline estimates in (1) (of cubic order). In the simulation, we used $n = 3000$ and $d = 3$. We drew input points $X^i \stackrel{\text{i.i.d.}}{\sim} \text{Unif}[0, 1]^3$, $i = 1, \dots, 3000$, and drew responses $Y^i \stackrel{\text{i.i.d.}}{\sim} N(\sum_{j=1}^3 f_{0j}(X_j^i), \sigma^2)$, $i = 1, \dots, 3000$, where $\sigma = 1.72$ was set to give a signal-to-noise ratio of about 1. The underlying component functions were defined as

$$f_{01}(t) = \min(t, 1-t)^{0.2} \sin \left(\frac{2.85\pi}{0.3 + \min(t, 1-t)} \right), \quad f_{02}(t) = e^{3t} \sin(4\pi t), \quad f_{03}(t) = -(t-1/2)^2,$$

so that f_{01}, f_{02}, f_{03} possess different levels of smoothness (f_{03} being the smoothest, f_{02} less smooth, and f_{01} the least smooth), and so that f_{01} itself has heterogeneous smoothness across its domain.

The first row of Figure 1 shows the estimated component functions from additive trend filtering, at a value of λ that minimizes the mean squared error (MSE), computed over 20 repetitions. The second row shows the estimates from additive smoothing splines, also at a value of λ that minimizes the MSE. We see that the trend filtering fits adapt well to the varying levels of smoothness, but the smoothing spline fits are undersmoothed, for the most part. In terms of effective degrees of freedom (df), the additive smoothing spline estimate is much more complex, having about 85 df (computed via Monte Carlo over the 20 repetitions); the additive trend filtering has only about 42 df. The third row of the figure shows the estimates from additive smoothing splines, when λ is chosen so that the resulting df roughly matches that of additive trend filtering in the first row. Now we see that the first component fit is oversmoothed, yet the third is still undersmoothed.

Figure 2 displays the MSE curves from additive trend filtering, as a function of df. We see that trend filtering achieves a lower MSE, and moreover, its MSE curve is optimized at a lower df (i.e., less complex model) than that for smoothing splines. This is analogous to what is typically seen in the univariate setting [Tibshirani, 2014].

We note that this motivating example is intended to elucidate the differences in what additive smoothing splines and additive trend filtering can do with a single tuning parameter each; a serious applied statistician, in just $d = 3$ dimensions, would likely use REML or some related technique to fit a multiple tuning parameter smoothing spline model, which would bring it closer to additive trend filtering in performance here; see our later discussion on this topic in Section 5.2.)

1.5 Summary of contributions

A summary of our contributions, and an outline for the rest of this paper, are given below.

- In Section 2, we investigate basic properties of the additive trend filtering model: an equivalent continuous-time formulation, a condition for uniqueness of component function estimates, and a simple formula for the effective degrees of freedom of the additive fit.
- In Section 3, we derive error bounds for additive trend filtering. Assuming that the underlying regression function is additive, denoted by $f_0 = \sum_{j=1}^d f_{0j}$, and that $\text{TV}(f_{0j}^{(k)})$ is bounded, for $j = 1, \dots, d$, we prove that the k th order additive trend filtering estimator converges to f_0 at the rate $n^{-(2k+2)/(2k+3)}$ when the dimension d is fixed (under weak assumptions), and at the rate $dn^{-(2k+2)/(2k+3)}$ when d is growing (under stronger assumptions). We prove that these rates are optimal in a minimax sense, and also show that additive smoothing splines (generally, additive models built from linear smoothers of any kind) are suboptimal over such a class of functions f_0 .
- In Section 4, we study the backfitting algorithm for additive trend filtering models, and give a connection between backfitting and an alternating projections scheme in the additive trend filtering dual problem. This inspires a new parallelized backfitting algorithm.
- In Section 5, we present empirical experiments and comparisons, and we also investigate the use of multiple tuning parameter models. In Section 6, we give a brief discussion.

2 Basic properties

In this section, we derive a number of basic properties of additive trend filtering estimates, starting with a representation for the estimates as continuous functions over \mathbb{R}^d (rather than simply discrete fitted values at the input points).

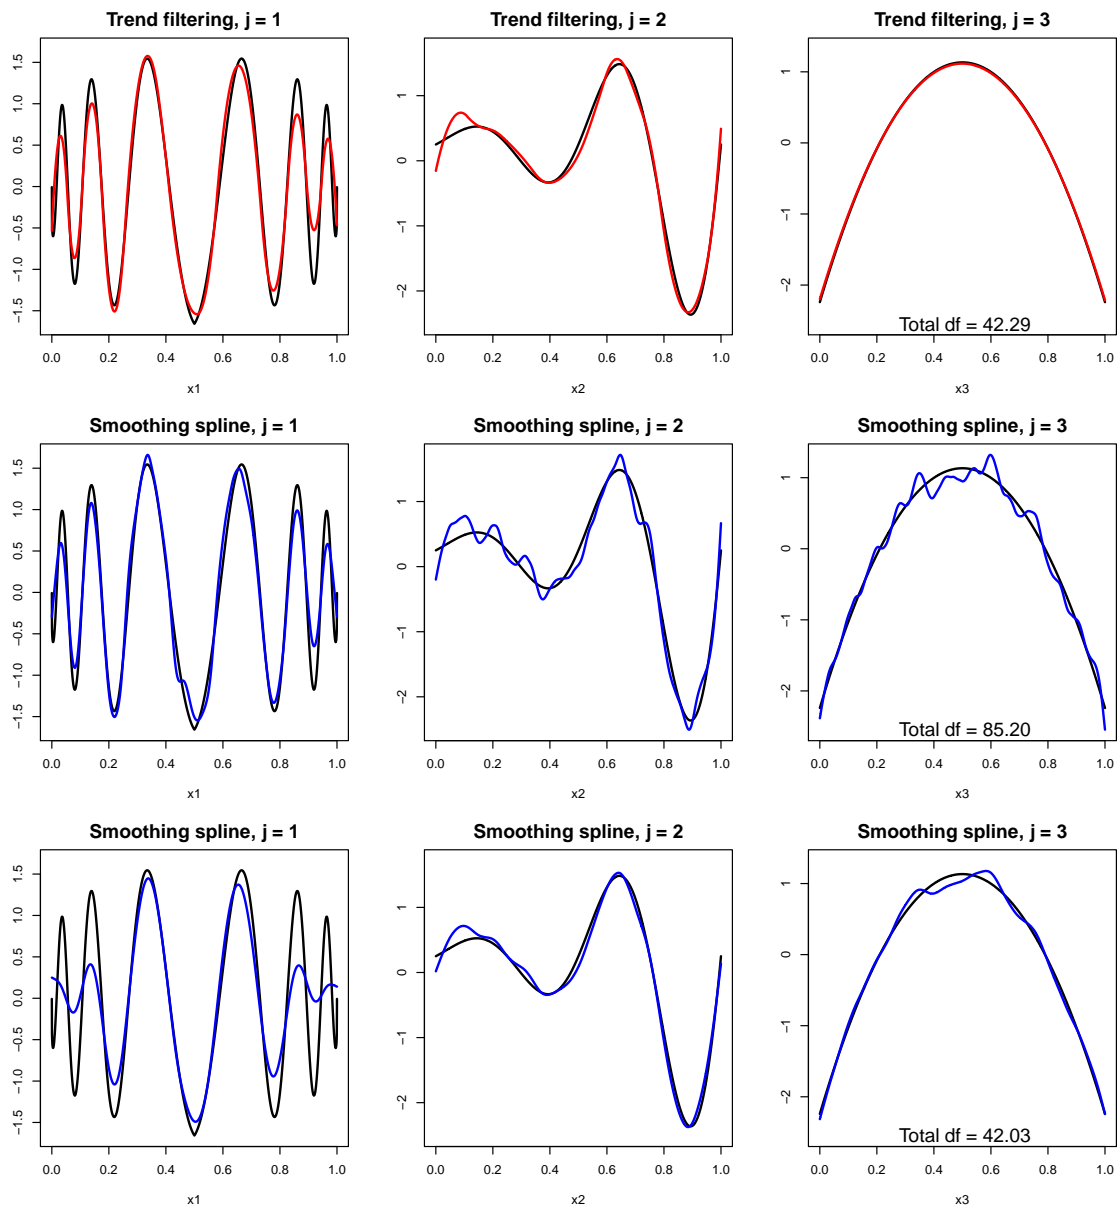


Figure 1: Comparing estimates from additive trend filtering (7) (of quadratic order) and additive smoothing splines (1) (of cubic order), for a simulation with $n = 3000$ and $d = 2$, as described in Section 1.4. In each row, the underlying component functions are plotted in black. The first row shows the estimated component functions using additive trend filtering, in red, at a value of λ chosen to minimize mean squared error (MSE), computed over 20 repetitions. The second row shows the estimates from additive smoothing splines, in blue, again at a value of λ that minimizes MSE. The third row shows the estimates from additive smoothing splines when λ is tuned so that the effective degrees of freedom (df) of the fit roughly matches that of additive trend filtering in the first row.

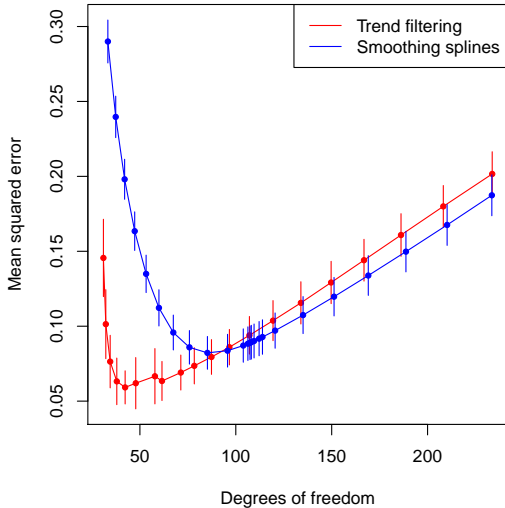


Figure 2: MSE curves for additive trend filtering and additive smoothing splines, computed over 20 repetitions from the same simulation setup as in Figure 1. Vertical segments denote ± 1 standard deviations. The MSE curves are parametrized by degrees of freedom (computed via standard Monte Carlo methods over the 20 repetitions).

2.1 Falling factorial representation

We may describe additive trend filtering in (7) as an estimation problem written in *analysis form*. The components are modeled directly by the parameters θ_j , $j = 1, \dots, d$, and the desired structure is established by regularizing the discrete derivatives of these parameters, through the penalty terms $\|D^{(X_j, k+1)} S_j \theta_j\|_1$, $j = 1, \dots, d$. Here, we present an alternative representation for (7) in *basis form*, where each component is expressed as a linear combination of basis functions, and regularization is applied to the coefficients in this expansion.

Before we derive the basis formulation that underlies additive trend filtering, we first recall the *falling factorial basis* [Tibshirani, 2014, Wang et al., 2014]. Given knot points $t^1 < \dots < t^n \in \mathbb{R}$, the k th order falling factorial basis functions h_1, \dots, h_n are defined by

$$\begin{aligned}
 h_i(t) &= \prod_{\ell=1}^{i-1} (t - t^\ell), \quad i = 1, \dots, k+1, \\
 h_{i+k+1}(t) &= \prod_{\ell=1}^k (t - t^{i+\ell}) \cdot 1\{t > t^{i+k}\}, \quad i = 1, \dots, n-k-1.
 \end{aligned}
 \tag{8}$$

We denote $1\{t > a\} = 1$ when $t > a$, and 0 otherwise. (Also, our convention is to define the empty product to be 1, so that $h_1(t) = 1$.) The functions h_1, \dots, h_n are piecewise polynomial functions of order k , and appear very similar in form to the k th order truncated power basis functions. In fact, when $k = 0$ or $k = 1$, the two bases are equivalent (meaning that they have the same span). Similar to an expansion in the truncated power basis, an expansion in the falling factorial basis,

$$g = \sum_{i=1}^n \alpha^i h_i$$

is a continuous piecewise polynomial function, having a global polynomial structure determined by $\alpha^1, \dots, \alpha^{k+1}$, and exhibiting a knot—i.e., a change in its k th derivative—at the location t^{i+k} when $\alpha^{i+k+1} \neq 0$. But, unlike the truncated power functions, the falling factorial functions in (8) are not splines, and when g (as defined above) has a knot at a particular location, it displays a change not only in its k th derivative at this location, but also in all lower order derivatives (i.e., all derivatives of orders $1, \dots, k-1$).

Tibshirani [2014], Wang et al. [2014] establish a connection between univariate trend filtering and the falling factorial functions, and show that the trend filtering problem can be interpreted as a sparse basis regression problem using these functions. As we show next, the analogous result holds for additive trend filtering.

Lemma 1 (Falling factorial representation). *For $j = 1, \dots, d$, let $h_1^{(X_j)}, \dots, h_n^{(X_j)}$ be the falling factorial basis in (8) with knots $(t^1, \dots, t^n) = S_j X_j$, the j th dimension of the input points, properly sorted. Then the additive trend filtering problem (7) is equivalent to the problem*

$$\begin{aligned} \min_{\alpha_1, \dots, \alpha_d \in \mathbb{R}^n} \quad & \frac{1}{2} \sum_{i=1}^n \left(Y^i - \bar{Y} - \sum_{j=1}^d \sum_{\ell=1}^n \alpha_j^\ell h_\ell^{(X_j)}(X_j^i) \right)^2 + \lambda k! \sum_{j=1}^d \sum_{\ell=k+2}^n |\alpha_j^\ell| \\ \text{subject to} \quad & \sum_{i=1}^n \sum_{\ell=1}^n \alpha_j^\ell h_\ell^{(X_j)}(X_j^i) = 0, \quad j = 1, \dots, d, \end{aligned} \quad (9)$$

in that, at any solutions in (7), (9), we have

$$\hat{\theta}_j^i = \sum_{\ell=1}^n \hat{\alpha}_j^\ell h_\ell^{(X_j)}(X_j^i), \quad i = 1, \dots, n, \quad j = 1, \dots, d.$$

An alternative way of expressing problem (9) is

$$\begin{aligned} \min_{f_j \in \mathcal{H}_j, j=1, \dots, d} \quad & \frac{1}{2} \sum_{i=1}^n \left(Y^i - \bar{Y} - \sum_{j=1}^d f_j(X_j^i) \right)^2 + \lambda \sum_{j=1}^d \text{TV}(f_j^{(k)}) \\ \text{subject to} \quad & \sum_{i=1}^n f_j(X_j^i) = 0, \quad j = 1, \dots, d, \end{aligned} \quad (10)$$

where $\mathcal{H}_j = \text{span}\{h_1^{(X_j)}, \dots, h_n^{(X_j)}\}$ is the span of the falling factorial basis over the j th dimension, and $f_j^{(k)}$ is the k th weak derivative of f_j , $j = 1, \dots, d$. In this form, at any solutions in (7), (10),

$$\hat{\theta}_j^i = \hat{f}_j(X_j^i), \quad i = 1, \dots, n, \quad j = 1, \dots, d.$$

Proof. For $j = 1, \dots, d$, define the k th order falling factorial basis matrix $H^{(X_j, k)} \in \mathbb{R}^{n \times n}$ by

$$H_{i\ell}^{(X_j, k)} = h_\ell^{(X_j)}(X_j^i), \quad i = 1, \dots, n, \quad \ell = 1, \dots, n. \quad (11)$$

Note that the columns of $H^{(X_j, k)}$ follow the order of the sorted inputs $S_j X_j$, but the rows do not; however, for $S_j H^{(X_j, k)}$, both its rows and columns follow the order of $S_j X_j$. From Wang et al. [2014], we know that

$$(S_j H^{(X_j, k)})^{-1} = \begin{bmatrix} C^{(X_j, k+1)} \\ \frac{1}{k!} D^{(X_j, k+1)} \end{bmatrix},$$

for some matrix $C^{(X_j, k+1)} \in \mathbb{R}^{(k+1) \times n}$, i.e.,

$$(H^{(X_j, k)})^{-1} = \begin{bmatrix} C^{(X_j, k+1)} \\ \frac{1}{k!} D^{(X_j, k+1)} \end{bmatrix} S_j. \quad (12)$$

Problem (9) is given by reparameterizing (7) according to $\theta_j = H^{(X_j, k)} \alpha_j$, for $j = 1, \dots, d$. As for (10), the equivalence between this and (9) follows by noting that for $f_j = \sum_{\ell=1}^n \alpha_j^\ell h_\ell^{(X_j)}$, we have

$$f_j^{(k)}(t) = k! + k! \sum_{\ell=k+2}^n \alpha_j^\ell \cdot 1\{t > X_j^{\ell-1}\},$$

and so $\text{TV}(f_j^{(k)}) = k! \sum_{\ell=k+2}^n |\alpha_j^\ell|$, for each $j = 1, \dots, d$. \square

This lemma not only provides an interesting reformulation for additive trend filtering, it is also practically useful in that it allows us to perform interpolation or extrapolation using the additive trend filtering model. That is, from a solution $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_d)$ in (7), we can extend each component fit $\hat{\theta}_j$ to the real line, by forming an appropriate linear combination of falling factorial functions:

$$\hat{f}_j(x_j) = \sum_{\ell=1}^n \hat{\alpha}_j^\ell h_\ell^{(X_j)}(x_j), \quad x_j \in \mathbb{R}. \quad (13)$$

The coefficients above are determined by the relationship $\hat{\alpha}_j = (H^{(X_j, k)})^{-1} \hat{\theta}_j$, and are easily computable given the highly structured form of $(H^{(X_j, k)})^{-1}$, as revealed in (12). Writing the coefficients in block form, as in $\hat{\alpha}_j = (\hat{a}_j, \hat{b}_j) \in \mathbb{R}^{(k+1)} \times \mathbb{R}^{(n-k-1)}$, we have

$$\hat{a}_j = C^{(X_j, k+1)} S_j \hat{\theta}_j, \quad (14)$$

$$\hat{b}_j = \frac{1}{k!} D^{(X_j, k+1)} S_j \hat{\theta}_j. \quad (15)$$

The first $k+1$ coefficients \hat{a}_j index the pure polynomial functions $h_1^{(X_j)}, \dots, h_{k+1}^{(X_j)}$. These coefficients will be generically dense (the form of $C^{(X_j, k+1)}$ is not important here, so we omit it for simplicity, but details are given in Appendix A.1). The last $n-k-1$ coefficients \hat{b}_j index the knot-producing functions $h_{k+2}^{(X_j)}, \dots, h_n^{(X_j)}$, and when $(\hat{b}_j)_\ell = \frac{1}{k!} (D^{(X_j, k+1)} S_j \hat{\theta}_j)_\ell \neq 0$, the fitted function \hat{f}_j exhibits a knot at the $(\ell+k)$ th sorted input point among $S_j X_j$, i.e., at $X_j^{(\ell+k)}$. Figure 3 gives an example.

We note that the coefficients $\hat{\alpha}_j = (\hat{a}_j, \hat{b}_j)$ in (14), (15) can be computed in $O(n)$ operations and $O(1)$ memory. This makes extrapolation of the j th fitted function \hat{f}_j in (13) highly efficient. Details are given in Appendix A.1.

2.2 Uniqueness of component fits

It is easy to see that, for the problem (7), the additive fit $\sum_{j=1}^d \hat{\theta}_j$ is always uniquely determined: denoting $\sum_{j=1}^d \theta_j = T\theta$ for a linear operator T and $\theta = (\theta_1, \dots, \theta_d) \in \mathbb{R}^{nd}$, the loss term $\|y - T\theta\|_2^2$ is strictly convex in the variable $T\theta$, and this, along with the convexity of the problem (7), implies a unique additive fit $T\theta$, no matter the choice of solution $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_d) \in \mathbb{R}^{nd}$.

On the other hand, when $d > 1$, the criterion in (7) is not strictly convex in θ , and hence there need not be a unique solution $\hat{\theta}$, i.e., the individual components fits $\hat{\theta}_j$, $j = 1, \dots, d$ need not be uniquely determined. We show next that uniqueness of the component fits can be guaranteed under some conditions on the input matrix $X = [X_1 \dots X_d] \in \mathbb{R}^{n \times d}$. We will rely on the falling factorial representation for additive trend filtering, introduced in the previous subsection, and on the notion of *general position*: a matrix $A \in \mathbb{R}^{m \times p}$ is said to have columns in general position provided that, for any $\ell < \min\{m, p\}$, subset of $\ell+1$ columns denoted $A_{i_1}, \dots, A_{i_{\ell+1}}$, and signs $s_1, \dots, s_{\ell+1} \in \{-1, 1\}$, the affine span of $\{s_1 A_{i_1}, \dots, s_{\ell+1} A_{i_{\ell+1}}\}$ does not contain any element of $\{\pm A_i : i \neq i_1, \dots, i_{\ell+1}\}$.

Lemma 2 (Uniqueness). *For $j = 1, \dots, d$, let $H^{(X_j, k)} \in \mathbb{R}^{n \times n}$ be the falling factorial basis matrix constructed over the sorted j th dimension of inputs $S_j X_j \in \mathbb{R}^n$, as in (11). Decompose $H^{(X_j, k)}$ into its first $k+1$ columns $P^{(X_j, k)} \in \mathbb{R}^{n \times (k+1)}$, and its last $n-k-1$ columns $K^{(X_j, k)} \in \mathbb{R}^{n \times (n-k-1)}$. The former contains evaluations of the pure polynomials $h_1^{(X_j)}, \dots, h_{k+1}^{(X_j)}$; the latter contains evaluations of the knot-producing functions $h_{k+2}^{(X_j)}, \dots, h_n^{(X_j)}$. Also, let $\tilde{P}^{(X_j, k)}$ denote the matrix $P^{(X_j, k)}$ with its first column removed, for $j = 1, \dots, d$, and $M = I - \mathbf{1}\mathbf{1}^T/n$. Define*

$$\tilde{P} = M [\tilde{P}^{(X_1, k)} \quad \dots \quad \tilde{P}^{(X_d, k)}] \in \mathbb{R}^{n \times dk}, \quad (16)$$

the product of M and the columnwise concatenation of $\tilde{P}^{(X_j, k)}$, $j = 1, \dots, d$. Let UU^T denote the projection operator onto the space orthogonal to the column span of \tilde{P} , where $U \in \mathbb{R}^{n \times (n-kd-1)}$ has orthonormal columns, and define

$$\tilde{K} = U^T M [K^{(X_1, k)} \quad \dots \quad K^{(X_d, k)}] \in \mathbb{R}^{(n-kd-1) \times (n-k-1)d}, \quad (17)$$




Figure 3: An example of extrapolating the fitted additive trend filtering model, where $n = 1000$ and $d = 2$. We generated input points $X^i \stackrel{i.i.d.}{\sim} \text{Unif}[0, 1]^2$, $i = 1, \dots, 1000$, and responses $Y^i \stackrel{i.i.d.}{\sim} N(\sum_{j=1}^2 f_{0j}(X_j^i), \sigma^2)$, $i = 1, \dots, 1000$, where we $f_{01}(x_1) = \sqrt{x_1} \sin(3\pi/(x_1 + 1/2))$ and $f_{02}(x_2) = x_2(x_2 - 1/3)$, and $\sigma = 0.36$. The top row shows three perspectives of the data. The bottom left panel shows the fitted values from additive trend filtering (7) (with $k = 2$ and $\lambda = 0.004$), where points are colored by their depth for visualization purposes. The bottom right panel shows the 2d surface associated with the trend filtering estimate, $\hat{f}_1(x_1) + \hat{f}_2(x_2)$ over $(x_1, x_2) \in [0, 1]^2$, with each component function extrapolated as in (13).

the product of $U^T M$ and the columnwise concatenation of $K^{(X_j, k)}$, $j = 1, \dots, d$. A sufficient condition for uniqueness of the additive trend filtering solution in (7) can now be given in two parts.

1. If \tilde{K} has columns in general position, then the knot-producing parts of all component fits are uniquely determined, i.e., for each $j = 1, \dots, d$, the projection of $\hat{\theta}_j$ onto the column space of $K^{(X_j, k)}$ is unique.
2. If in addition \tilde{P} has full column rank, then the polynomial parts of component fits are uniquely determined, i.e., for each $j = 1, \dots, d$, the projection of $\hat{\theta}_j$ onto the column space of $P^{(X_j, k)}$ is unique, and thus the component fits $\hat{\theta}_j$, $j = 1, \dots, d$ are all unique.

The proof is deferred to Appendix A.2. To rephrase, the above lemma decomposes each component of the additive trend filtering solution according to

$$\hat{\theta}_j = \hat{\theta}_j^{\text{poly}} + \hat{\theta}_j^{\text{knot}}, \quad j = 1, \dots, d,$$

where $\hat{\theta}_j^{\text{poly}}$ exhibits a purely polynomial trend over $S_j X_j$, and $\hat{\theta}_j^{\text{knot}}$ exhibits a piecewise polynomial trend over $S_j X_j$, and hence determines the knot locations, for $j = 1, \dots, d$. The lemma shows that the knot-producing parts $\hat{\theta}_j^{\text{knot}}$, $j = 1, \dots, d$ are uniquely determined when the columns of \tilde{K} are in general position, and the polynomial parts $\hat{\theta}_j^{\text{knot}}$, $j = 1, \dots, d$ are unique when the columns of \tilde{K} are in general position, and the columns of \tilde{P} are linearly independent.

The conditions placed on \tilde{P} , \tilde{K} in Lemma 2 are not strong. When $n > kd$, and the elements of input matrix X are drawn from a density over \mathbb{R}^{nd} , it is not hard to show that \tilde{P} has full column rank with probability 1. We conjecture that, under the same conditions, \tilde{K} will also have columns in general position with probability 1, but do not pursue a proof.

Remark 1 (Relationship to concurvity). *It is interesting to draw a connection to Buja et al. [1989]. In the language used by these authors, when \tilde{P} has linearly dependent columns, we say that the predictor variables display concurvity, i.e., linear dependence after nonlinear (here, polynomial) transformations are applied. Buja et al. [1989] establish that the components in the additive model (1), built with quadratic penalties, are unique provided there is no concurvity between variables. In comparison, Lemma 2 establishes uniqueness of the additive trend filtering components when there is no concurvity between variables, and additionally, the columns of \tilde{K} are in general position. The latter two conditions together can be seen as requiring no generalized concurvity—if \tilde{K} were to fail the general position assumption, then there would be a small subset of the variables that are linearly dependent after nonlinear (here, piecewise polynomial) transformations are applied.*

2.3 Dual problem

Let us abbreviate $D_j = D^{(X_j, k+1)}$, $j = 1, \dots, d$ for the penalty matrices in the additive trend filtering problem (7). Basic arguments in convex analysis, deferred to Appendix A.3, show that the dual of problem (7) can be expressed as:

$$\begin{aligned} \min_{u \in \mathbb{R}^n} \|Y - \bar{Y}\mathbb{1} - u\|_2^2 \quad \text{subject to} \quad u \in U = U_1 \cap \dots \cap U_d, \\ \text{where} \quad U_j = \{S_j D_j^T v_j : \|v_j\|_\infty \leq \lambda\}, \quad j = 1, \dots, d, \end{aligned} \tag{18}$$

and that primal and dual solutions in (7), (18) are related by:

$$\sum_{j=1}^d \hat{\theta}_j = Y - \bar{Y}\mathbb{1} - \hat{u}. \tag{19}$$

From the form of (18), it is clear that we can write the (unique) dual solution as $\hat{u} = \Pi_U(Y - \bar{Y}\mathbb{1})$, where Π_U is the (Euclidean) projection operator onto U . Moreover, using (19), we can express the

additive fit as $\sum_{j=1}^d \hat{\theta}_j = (\text{Id} - \Pi_U)(Y - \bar{Y}\mathbb{1})$, where $\text{Id} - \Pi_U$ is the operator that gives the residual from projecting onto U . These relationships will be revisited in Section 4, where we return to the dual perspective, and argue that the backfitting algorithm for the additive trend filtering problem (7) can be seen as a type of alternating projections algorithm for its dual problem (18).

2.4 Degrees of freedom

In general, given data $Y \in \mathbb{R}^n$ with $\mathbb{E}(Y) = \eta$, $\text{Cov}(Y) = \sigma^2 I$, and an estimator $\hat{\eta}$ of η , recall that we define the *effective degrees of freedom* of $\hat{\eta}$ as [Efron, 1986, Hastie and Tibshirani, 1990]:

$$\text{df}(\hat{\eta}) = \frac{1}{\sigma^2} \sum_{i=1}^n \text{Cov}(\hat{\eta}^i(Y), Y^i),$$

where $\hat{\eta}(Y) = (\hat{\eta}^1(y), \dots, \hat{\eta}^n(Y))$. Roughly speaking, the above definition sums the influence of the i th component Y^i on its corresponding fitted value $\hat{\eta}^i(Y)$, across $i = 1, \dots, n$. A precise understanding of degrees of freedom is useful for model comparisons (recall the x-axis in Figure 2), and other reasons. For linear smoothers, in which $\hat{\eta}(Y) = SY$ for some $S \in \mathbb{R}^{n \times n}$, it is clear that $\text{df}(\hat{\eta}) = \text{tr}(S)$, the trace of S . (This also covers additive models whose components are built from univariate linear smoothers, because in total these are still just linear smoothers: the additive fit is still just a linear function of Y .)

Of course, additive trend filtering is not a linear smoother; however, it is a particular type of generalized lasso estimator, and degrees of freedom for such a class of estimators is well-understood [Tibshirani and Taylor, 2011, 2012]. The next result is a consequence of existing generalized lasso theory, proved in Appendix A.4.

Lemma 3 (Degrees of freedom). *Assume the conditions of Lemma 2, i.e., that the matrix \tilde{P} in (16) has full column rank, and the matrix \tilde{K} in (17) is in general position. Assume also that the response is Gaussian, $Y \sim N(\eta, \sigma^2 I)$, and treat the input points $X^i \in \mathbb{R}^d$, $i = 1, \dots, n$ as fixed and arbitrary, as well as the tuning parameter value $\lambda \geq 0$. Then the additive trend filtering fit from (7) has degrees of freedom*

$$\text{df}\left(\sum_{j=1}^d \hat{\theta}_j\right) = \mathbb{E}\left(\sum_{j=1}^d (\text{number of knots in } \hat{\theta}_j)\right) + kd.$$

Remark 2 (The effect of shrinkage). *Lemma 3 says that for an unbiased estimate of the degrees of freedom of the additive trend filtering fit, we count the number of knots in each component fit $\hat{\theta}_j$ (recall that this is the number of nonzeros in $D^{(X_j, k+1)}\hat{\theta}_j$, i.e., the number of changes in the discrete $(k+1)$ st derivative), add them up over $j = 1, \dots, d$, and add kd . This may seem surprising, as these knot locations are chosen adaptively based on the data Y . But, such adaptivity is counterbalanced by the shrinkage induced by the ℓ_1 penalty in (7) (i.e., for each component fit $\hat{\theta}_j$, there is shrinkage in the differences between the attained k th derivatives on either side of a selected knot). See Tibshirani [2015] for a study of this phenomenon.*

Remark 3 (Easy unbiased degrees of freedom estimation). *It is worth emphasizing that an unbiased estimate from Lemma 3 for the degrees of freedom of the total fit in additive trend filtering is very easy to calculate: we scan the individual component fits and add up the number of knots that appear in each one. The same cannot be said for additive smoothing splines, or additive models built from univariate linear smoothers, in general. Although computing the fit itself is typically cheaper with additive linear smoothers than with additive trend filtering, computing the degrees of freedom is more challenging. For example, for the additive model in (1) built with quadratic penalties, we have*

$$\text{df}\left(\sum_{j=1}^d \hat{\theta}_j\right) = \text{tr}\left(F^T F (F^T F + \lambda Q)^+\right),$$

where $F \in \mathbb{R}^{n \times nd}$ has d copies of the centering matrix $M = I - \mathbb{1}\mathbb{1}^T/n \in \mathbb{R}^{n \times n}$ stacked across its columns, $Q \in \mathbb{R}^{nd \times nd}$ is a block diagonal matrix with blocks MQ_jM , $j = 1, \dots, d$, and A^+ denotes the Moore-Penrose pseudoinverse of a matrix A . The above formula does not obviously decompose into a sum of quantities across components, and is nontrivial to compute post optimization of (1), specifically when a backfitting algorithm as in (2) has been used to compute a solution.

2.5 Two related additive spline estimators

From its equivalent formulation in (10), additive trend filtering is seen to be closely related to two other additive spline estimators, which we introduce here. Consider, for univariate function classes \mathcal{S}_j , $j = 1, \dots, d$, the problem

$$\begin{aligned} \min_{f_j \in \mathcal{S}_j, j=1, \dots, d} \quad & \frac{1}{2} \sum_{i=1}^n \left(Y^i - \bar{Y} - \sum_{j=1}^d f_j(X_j^i) \right)^2 + \lambda \sum_{j=1}^d \text{TV}(f_j^{(k)}) \\ \text{subject to} \quad & \sum_{i=1}^n f_j(X_j^i) = 0, \quad j = 1, \dots, d. \end{aligned} \tag{20}$$

When each \mathcal{S}_j , $j = 1, \dots, d$ is the set of k times weakly differentiable functions, we call the solution in (20) the *additive locally adaptive regression spline* of order $k \geq 0$, as it is the natural extension of the univariate estimator considered in Mammen and van de Geer [1997]. Denote by \hat{f}_j , $j = 1, \dots, d$ this solution; the representation arguments used by these authors apply immediately to the additive setting, and imply that each \hat{f}_j , $j = 1, \dots, d$ is indeed a spline of degree k (justifying the choice of name). The same arguments show that, for $k = 0$ or $k = 1$, the knots of the spline \hat{f}_j lie among the j th dimension of the input points X_j^1, \dots, X_j^n , for $j = 1, \dots, d$, but for $k \geq 2$, this need not be true, and in general the components will be splines with knots at locations other than the inputs.

We can facilitate computation by taking $\mathcal{S}_j = \mathcal{G}_j$, where \mathcal{G}_j is the set of splines of degree k with knots lying among the j th dimension of inputs X_j^1, \dots, X_j^n , for $j = 1, \dots, d$. We call the resulting solution the *restricted additive locally adaptive regression spline* of order $k \geq 0$. More precisely, we require that the splines in \mathcal{G}_j have knots in a set T_j , which, writing $t_j = S_j X_j$ for the sorted inputs along the j th dimension, is defined by

$$T_j = \begin{cases} \{t_j^{k/2+2}, \dots, t_j^{n-k/2}\} & \text{if } k \text{ is even,} \\ \{t_j^{(k+1)/2+1}, \dots, t_j^{n-(k+1)/2}\} & \text{if } k \text{ is odd,} \end{cases} \tag{21}$$

i.e., defined by removing $k + 1$ input points at the boundaries, for $j = 1, \dots, d$. Setting $\mathcal{S}_j = \mathcal{G}_j$, $j = 1, \dots, d$ turns (20) into a finite-dimensional problem, just like (10). When $k = 0$ or $k = 1$, as is evident from their form in (8), the falling factorial functions are simply splines, which implies that $\mathcal{H}_j = \mathcal{G}_j$ for $j = 1, \dots, d$, hence additive trend filtering and restricted additive locally adaptive regression splines are the same estimator. When $k \geq 2$, this is no longer true, and they are not the same. Additive trend filtering will be much easier to compute, since $\text{TV}(g^{(k)})$ does not admit a nice representation in terms of discrete derivatives for a k th order spline (and yet it does for a k th order falling factorial function, as seen in (7)).

To summarize, additive locally adaptive splines, restricted additive locally adaptive splines, and additive trend filtering all solve a problem of the form (20) for different choices of function classes \mathcal{S}_j , $j = 1, \dots, d$. For $k = 0$ or $k = 1$, these three estimators are equivalent. For $k \geq 2$, they will be generically different, though our intuition tells us that their differences should not be too large: the unrestricted problem admits a solution that is a spline in each component; the restricted problem simply forces these splines to have knots at the input points; and the trend filtering problem swaps splines for falling factorial functions, which are highly similar in form. Next, we give theory that confirms this intuition, in large samples.

3 Error bounds

We derive error bounds for additive trend filtering and additive locally adaptive regression splines (both the unrestricted and restricted variants), when the underlying regression function is additive, and has components whose derivatives are of bounded variation. These results are actually special cases of a more general result we prove in this section, on a generic roughness-regularized additive estimator, where we assume a certain decay for the entropy of the unit ball in the roughness operator. We treat separately the settings in which the dimension d of the input space is fixed and growing. We also complement our error rates with minimax lower bounds. We start by introducing helpful notation.

3.1 Notation

Given a distribution Q supported on a set D , and i.i.d. samples X^i , $i = 1, \dots, n$ from Q , denote by Q_n the associated empirical distribution. We define the $L_2(Q)$ and $L_2(Q_n)$ inner products, denoted $\langle \cdot, \cdot \rangle_{L_2(Q)}$ and $\langle \cdot, \cdot \rangle_{L_2(Q_n)}$, respectively, over functions $m, r : D \rightarrow \mathbb{R}$

$$\langle m, r \rangle_{L_2(Q)} = \int_D m(x)r(x) dQ(x), \quad \text{and} \quad \langle m, r \rangle_{L_2(Q_n)} = \frac{1}{n} \sum_{i=1}^n m(X^i)r(X^i).$$

Definitions for the corresponding $L_2(Q)$ and $L_2(Q_n)$ norms, denoted $\|\cdot\|_{L_2(Q)}$ and $\|\cdot\|_{L_2(Q_n)}$, respectively, arise naturally from these inner products, defined by

$$\|m\|_2^2 = \langle m, m \rangle_2 = \int_D m(x)^2 dQ(x), \quad \text{and} \quad \|m\|_n^2 = \langle m, m \rangle_n = \frac{1}{n} \sum_{i=1}^n m(X^i)^2.$$

Henceforth, we will abbreviate subscripts when using these norms and inner products, writing $\|\cdot\|_2$ and $\|\cdot\|_n$ for the $L_2(Q)$ and $L_2(Q_n)$ norms, respectively, and similarly for the inner products. This abbreviated notation omits the underlying distribution Q ; thus, unless explicitly stated otherwise, the underlying distribution should always be interpreted as the distribution of the input points. We will often call $\|\cdot\|_2$ the L_2 norm and $\|\cdot\|_n$ the empirical norm, and similarly for inner products.

In what follows, of particular interest will be the case when $D = [0, 1]^d$, and $m : [0, 1]^d \rightarrow \mathbb{R}$ is an additive function, of the form

$$m = \sum_{j=1}^d m_j,$$

which we write to mean $m(x) = \sum_{j=1}^d m_j(x_j)$. In a slight abuse of notation (overload of notation), for each $j = 1, \dots, d$, we will abbreviate the $L_2(Q_j)$ norm by $\|\cdot\|_2$, where Q_j is the j th marginal of Q , and will also abbreviate $L_2(Q_{jn})$ norm by $\|\cdot\|_n$, where Q_{jn} is the empirical distribution of X_j^i , $i = 1, \dots, n$. We will use similar abbreviations for the inner products.

A few more general definitions are in order. We denote the L_∞ norm, also called the sup norm, of a function $f : D \rightarrow \mathbb{R}$ by $\|f\|_\infty = \text{ess sup}_{z \in D} |f(z)|$. For a functional ν , acting on functions from D to \mathbb{R} , we write $B_\nu(\delta)$ for the ν -ball of radius $\delta > 0$, i.e., $B_\nu(\delta) = \{f : \nu(f) \leq \delta\}$. We abbreviate $B_n(\delta)$ for the $\|\cdot\|_n$ -ball of radius δ , $B_2(\delta)$ for the $\|\cdot\|_2$ -ball of radius δ , and $B_\infty(\delta)$ for the $\|\cdot\|_\infty$ -ball of radius δ . We will use these concepts fluidly, without explicit reference to the domain D (or its dimensionality), as the meaning should be clear from the context.

Lastly, for a set S and norm $\|\cdot\|$, we define the covering number $N(\delta, \|\cdot\|, S)$ to be the smallest number of $\|\cdot\|$ -balls of radius δ to cover S , and the packing number $M(\delta, \|\cdot\|, S)$ to be the largest number of disjoint $\|\cdot\|$ -balls of radius δ that are contained in S . We call $\log N(\delta, \|\cdot\|, S)$ the entropy number.

3.2 Error bounds for a fixed dimension d

We consider error bounds for the generic roughness-penalized estimator defined as a solution of

$$\begin{aligned} \min_{f_j \in \mathcal{S}_j, j=1, \dots, d} \quad & \frac{1}{2} \sum_{i=1}^n \left(Y^i - \bar{Y} - \sum_{j=1}^d f_j(X_j^i) \right)^2 + \lambda \sum_{j=1}^d J(f_j) \\ \text{subject to} \quad & \sum_{i=1}^n f_j(X_j^i) = 0, \quad j = 1, \dots, d, \end{aligned} \tag{22}$$

where \mathcal{S}_j , $j = 1, \dots, d$ are univariate function spaces, and J is a regularizer that acts on univariate functions. We assume in this subsection that the dimension d of the input space is fixed, i.e., it does not grow with n . Before stating our main result in this setting, we list and briefly discuss our other assumptions. First, we give our assumptions on the data generation process.

Assumption A1 (Continuous input distribution). *The input points X^i , $i = 1, \dots, n$ are i.i.d. from a continuous distribution Q supported on $[0, 1]^d$.*

Assumption B1 (Generic regression model, with sub-Gaussian errors). *The responses Y^i , $i = 1, \dots, n$ follow the model*

$$Y^i = \mu + f_0(X^i) + \epsilon^i, \quad i = 1, \dots, n,$$

with overall mean $\mu \in \mathbb{R}$, where $\sum_{i=1}^n f_0(X^i) = 0$ for identifiability. The errors ϵ^i , $i = 1, \dots, n$ are uniformly sub-Gaussian and have mean zero, i.e.,

$$\mathbb{E}(\epsilon) = 0, \quad \text{and} \quad \mathbb{E}[\exp(v^T \epsilon)] \leq \exp(\sigma^2 \|v\|_2^2 / 2), \quad \text{for all } v \in \mathbb{R}^n,$$

for a constant $\sigma > 0$. The errors and input points are independent.

Assumptions A1 and B1 are very weak. We do not place any specific smoothness or additivity conditions on the underlying regression function f_0 , as our error bound in Theorem 1 will involve the error of the closest additive function to f_0 , whose components lie in the given function spaces \mathcal{S}_j , $j = 1, \dots, d$.

Next, we present our assumptions on the regularizer J . We write $\|\cdot\|_{Z_n}$ for the empirical norm defined over a set of univariate points $Z_n = \{z^1, \dots, z^n\} \subseteq [0, 1]$, i.e., $\|g\|_{Z_n}^2 = \frac{1}{n} \sum_{i=1}^n g^2(z^i)$.

Assumption C1 (Seminorm regularizer, null space of polynomials). *The regularizer J is a seminorm, and its domain is contained in the space of k times weakly differentiable functions, for an integer $k \geq 0$. Furthermore, its null space contains all k th order polynomials, i.e., $J(g) = 0$ for all $g(t) = t^\ell$, $\ell = 0, \dots, k$.*

Assumption C2 (Relative boundedness of derivatives). *There is a constant $L > 0$ such that $\text{ess sup}_{t \in [0, 1]} g^{(k)}(t) - \text{ess inf}_{t \in [0, 1]} g^{(k)}(t) \leq L$ for $g \in B_J(1)$ (with $g^{(k)}$ the k th weak derivative of g).*

Assumption C3 (Entropy bound). *There are constants $0 < w < 2$ and $K > 0$ such that*

$$\sup_{Z_n = \{z^1, \dots, z^n\} \subseteq [0, 1]} \log N(\delta, \|\cdot\|_{Z_n}, B_J(1) \cap B_\infty(1)) \leq K \delta^{-w}.$$

Assumptions C1, C2, C3 on the regularizer J are not strong, and are satisfied by various common regularizers, e.g., $J(g) = [\int_0^1 (g^{(k+1)}(t))^2 dt]^{1/2}$ or $J(g) = \text{TV}(g^{(k)})$ (the latter studied shortly).

We now state our main result in the fixed d case, which is proved in Appendix A.5, A.6.

Theorem 1. Assume [A1](#), [B1](#) on the data distribution, and assume [C1](#), [C2](#), [C3](#) on the seminorm J . Also, assume that the dimension d of the input space is fixed. Let $C_n \geq 1$ be an arbitrary sequence. There exist constants $c_1, c_2, c_3, n_0 > 0$, that depend only on d, σ, k, L, K, w , such that for all $c \geq c_1$, $n \geq n_0$, and tuning parameter values $\lambda \geq cn^{w/(2+w)}C_n^{-(2-w)/(2+w)}$, any solution in [\(22\)](#) satisfies

$$\left\| \sum_{j=1}^d \hat{f}_j - f_0 \right\|_n^2 \leq \left\| \sum_{j=1}^d \tilde{f}_j - f_0 \right\|_n^2 + \frac{6\lambda}{n} \max \left\{ C_n, \sum_{j=1}^d J(\tilde{f}_j) \right\}, \quad (23)$$

with probability at least $1 - \exp(-c_2c) - \exp(-c_3\sqrt{n})$, simultaneously over all $\tilde{f} = \sum_{j=1}^d \tilde{f}_j$, feasible for the problem [\(22\)](#), such that $\|\tilde{f} - f_0\|_n \leq \max\{C_n, \sum_{j=1}^d J(\tilde{f}_j)\}$.

Remark 4 (Error bound for additive, J -smooth f_0). Assume $f_0 = \sum_{j=1}^d f_{0j}$, where $f_{0j} \in \mathcal{S}_j$, $j = 1, \dots, d$, and $\sum_{j=1}^d J(f_{0j}) \leq C_n$. Letting $\tilde{f} = f_0$, the approximation error term in [\(23\)](#) (the first term on the right-hand side) is zero, and for $\lambda = cn^{w/(2+w)}C_n^{-(2-w)/(2+w)}$, the result in the theorem reads

$$\left\| \sum_{j=1}^d \hat{f}_j - \sum_{j=1}^d f_{0j} \right\|_n^2 \leq 6cn^{-2/(2+w)}C_n^{2w/(2+w)}, \quad (24)$$

with probability at least $1 - \exp(-c_2c) - \exp(-c_3\sqrt{n})$. As we will see in the minimax lower bound in [Theorem 3](#) (plugging in $c_n = C_n/d$, and taking d to be a constant), the rate $n^{-2/(2+w)}C_n^{2w/(2+w)}$ is optimal for such a class of functions.

Remark 5 (Distance to best additive, J -smooth approximation of f_0). The arguments used to establish the oracle-type inequality [\(23\)](#) also imply a result on the empirical norm error between \hat{f} and the best additive approximation of f_0 . To be precise, let $(f_1^{\text{best}}, \dots, f_d^{\text{best}})$ denote a solution in the population-level problem

$$\begin{aligned} \min_{f_j \in \mathcal{S}_j, j=1, \dots, d} \quad & \frac{1}{2} \sum_{i=1}^n \left(f_0(X^i) - \sum_{j=1}^d f_j(X_j^i) \right)^2 + \frac{\lambda}{2} \sum_{j=1}^d J(f_j) \\ \text{subject to} \quad & \sum_{i=1}^n f_j(X_j^i) = 0, \quad j = 1, \dots, d. \end{aligned} \quad (25)$$

We note that [\(25\)](#) has “half” of the regularization of problem [\(22\)](#), as it uses a penalty parameter of $\lambda/2$ versus λ . We can think of $f^{\text{best}} = \sum_{j=1}^d f_j^{\text{best}}$ as the best additive, J -smooth approximation of f_0 , where λ as usual controls the level of smoothness. The following is a consequence of the proof of [Theorem 1](#), verified in [Appendix A.7](#): assume that $\|f^{\text{best}} - f_0\|_n \leq \max\{C_n, \sum_{j=1}^d J(f_j^{\text{best}})\}$ almost surely (with respect to Q), for sufficiently large λ ; then any solution in [\(22\)](#) satisfies for all $c \geq c_1$, $n \geq n_0$, and $\lambda \geq cn^{w/(2+w)}C_n^{-(2-w)/(2+w)}$,

$$\left\| \sum_{j=1}^d \hat{f}_j - \sum_{j=1}^d f_j^{\text{best}} \right\|_n^2 \leq \frac{6\lambda}{n} \max \left\{ C_n, \sum_{j=1}^d J(f_j^{\text{best}}) \right\}, \quad (26)$$

with probability at least $1 - \exp(-c_2c) - \exp(-c_3\sqrt{n})$, where as before $c_1, c_2, c_3, n_0 > 0$ are constants that depend only on d, σ, k, L, K, w . Notably, the right-hand side in the bound [\(26\)](#) does not depend on the approximation error; in particular, we do not even require $\|f^{\text{best}} - f_0\|_n$ to converge to zero. This is analogous to classical results from [Stone \[1985\]](#).

We examine a special case of the generic problem [\(22\)](#) when the regularizer is $J(g) = \text{TV}(g^{(k)})$, and derive implications of the above [Theorem 1](#) for additive locally regression adaptive splines and additive trend filtering, corresponding to different choices of the function classes \mathcal{S}_j , $j = 1, \dots, d$ in [\(22\)](#). We must introduce an additional (weak) assumption on the input distribution, for the results on restricted locally adaptive regression splines and trend filtering.

Assumption A2 (Bounded input density). *The density of the input distribution Q is bounded below by a constant $b_0 > 0$.*

Here is our result for additive locally adaptive splines and additive trend filtering. The proof is given in Appendix A.8, A.9.

Corollary 1. *Assume A1, B1 on the data distribution. Also, assume that the dimension d of the input space is fixed, and that the underlying regression function is additive, $f_0 = \sum_{j=1}^d f_{0j}$, where the components f_{0j} , $j = 1, \dots, d$ are k times weakly differentiable, such that $\sum_{j=1}^d \text{TV}(f_{0j}^{(k)}) \leq C_n$ for a sequence $C_n \geq 1$. For $J(g) = \text{TV}(g^{(k)})$, Assumptions C1, C2, C3 hold with $L = 1$ and $w = 1/(k+1)$. Furthermore, the following is true of the estimator defined by problem (22).*

- (a) *Let \mathcal{S}_j be the set of all k times weakly differentiable functions, for each $j = 1, \dots, d$. There are constants $c_1, c_2, c_3, n_0 > 0$, depending only on d, σ, k , such that for all $c \geq c_1$ and $n \geq n_0$, any solution in the additive locally adaptive regression spline problem (22), with tuning parameter value $\lambda = cn^{1/(2k+3)}C_n^{-(2k+1)/(2k+3)}$, satisfies*

$$\left\| \sum_{j=1}^d \hat{f}_j - \sum_{j=1}^d f_{0j} \right\|_n^2 \leq cn^{-(2k+2)/(2k+3)} C_n^{2/(2k+3)}, \quad (27)$$

with probability at least $1 - \exp(-c_2c) - \exp(-c_3\sqrt{n})$.

- (b) *Let $\mathcal{S}_j = \mathcal{G}_j$, the set of k th degree splines with knots in the set T_j in (21), for $j = 1, \dots, d$, and assume A2 on the input density. Then there are constants $c_1, c_2, c_3, n_0 \geq 0$, that depend only on d, b_0, σ, k , such that for all $c \geq c_1$ and $n(\log n)^{-(1+1/k)} \geq n_0 C_n^{(2k+2)/(2k^2+2k-1)}$, any solution in the restricted additive locally adaptive spline problem (22), with $\lambda = cn^{1/(2k+3)}C_n^{-(2k+1)/(2k+3)}$, satisfies the same result in (27), with probability at least $1 - \exp(-c_2c) - c_3/n$.*
- (c) *Let $\mathcal{S}_j = \mathcal{H}_j$, the set of k th degree falling factorial functions defined over X_j (the j th dimension of inputs), for $j = 1, \dots, d$, and assume A2. Then there exist constants $c_1, c_2, c_3, n_0 > 0$, that depend only on d, b_0, σ, k , such that for all $c \geq c_1$ and $n(\log n)^{-(2k+3)} \geq n_0 C_n^{4k+4}$, any solution in the additive trend filtering problem (22), with $\lambda = cn^{1/(2k+3)}C_n^{-(2k+1)/(2k+3)}$, satisfies (27), with probability at least $1 - \exp(-c_2c) - c_3/n$.*

Remark 6 (Spline and falling factorial approximants). *For part (a) of the corollary, the approximation error (the first term on the right-hand side) in (24) is zero by definition, and we need only verify Assumptions C1, C2, C3 for the regularizer $J(g) = \text{TV}(g^{(k)})$. Parts (b) and (c) require control over the approximation error, because the underlying regression function $f_0 = \sum_{j=1}^d f_{0j}$ need not have components that lie in the chosen function spaces \mathcal{S}_j , $j = 1, \dots, d$. To be clear: for $k = 0$ or $k = 1$, as discussed in Section 2.5, all three problems considered in parts (a), (b), (c) are equivalent; hence parts (b) and (c) really only concern the case $k \geq 2$. For both of these parts, we control the approximation error by controlling the univariate approximation error and then applying the triangle inequality. For part (b), we use a special spline quasi-interpolant from Proposition 7 in Mammen and van de Geer [1997] (who in turn construct this using results from de Boor [1978]); for part (c), we develop a new falling factorial approximant that may be of independent interest.*

3.3 Error bounds for a growing dimension d

In this subsection, we allow the input dimension d to grow with the sample size n . To keep our analysis as clean as possible, we consider a constrained version of the problem (22), namely

$$\begin{aligned} \min_{f_j \in \mathcal{S}_j, j=1, \dots, d} \quad & \frac{1}{2} \sum_{i=1}^n \left(Y^i - \bar{Y} - \sum_{j=1}^d f_j(X_j^i) \right)^2 \\ \text{subject to} \quad & \sum_{i=1}^n f_j(X_j^i) = 0, \quad J(f_j) \leq \delta, \quad j = 1, \dots, d, \end{aligned} \quad (28)$$

for a tuning parameter $\delta > 0$. (The penalized problem (22) can also be analyzed in the setting of growing d , but we find that the analysis is messier and requires more assumptions in order to obtain the same results.) Instead of A1, we now use the following assumption in the input distribution.

Assumption A3 (Product input distribution). *The input points $X^i, i = 1, \dots, n$ are i.i.d. from a continuous distribution Q supported on $[0, 1]^d$, that decomposes as $Q = Q_1 \times \dots \times Q_d$, where the density of each Q_j is lower and upper bounded by constants $b_1, b_2 > 0$, for $j = 1, \dots, d$.*

Assumption A3 is fairly restrictive, since it requires the input distribution Q to be independent across dimensions of the input space. The reason we use this assumption: when $Q = Q_1 \times \dots \times Q_d$, additive functions enjoy a key decomposability property in terms of the (squared) L_2 norm defined with respect to Q . In particular, if $m = \sum_{j=1}^d m_j$ has components with L_2 mean zero, denoted by $\bar{m}_j = \int_0^1 m_j(x_j) dQ_j(x_j) = 0, j = 1, \dots, d$, then we have

$$\left\| \sum_{j=1}^d m_j \right\|_2^2 = \sum_{j=1}^d \|m_j\|_2^2. \quad (29)$$

This is explained by the fact that each pair of components m_j, m_ℓ with $j \neq \ell$ are orthogonal with respect to the L_2 inner product, since

$$\langle m_j, m_\ell \rangle_2 = \int_{[0,1]^2} m_j(x_j) m_\ell(x_\ell) dQ_j(x_j) dQ_\ell(x_\ell) = \bar{m}_j \bar{m}_\ell = 0.$$

The above orthogonality, and thus the decomposability property in (29), is only true because of the product form $Q = Q_1 \times \dots \times Q_d$. Such decomposability is not generally possible with the empirical norm (the inner products between components do not vanish even if all empirical means are zero). In the proof of Theorem 2, we move from considering the empirical norm of the error vector to the L_2 norm, in order to leverage the property in (29), which eventually leads to an error rate that has a linear dependence on the dimension d . In the absence of L_2 decomposability, the same error rate can be achieved with a weaker incoherence bound, as in (34); see Remark 9 after the theorem.

We now state our main result in the growing d case, whose proof is in Appendix A.10, A.11.

Theorem 2. *Assume A3, B1 on the data distribution, and assume C1, C2, C3 on the seminorm J . Let $\delta \geq 1$ be arbitrary. There are constants $c_1, c_2, c_3, n_0 > 0$, that depend only on $b_1, b_2, \sigma, k, L, K, w$, such that for all $c \geq c_1$ and $n \geq n_0(d\delta)^{1+w/2}$, any solution in (28) satisfies both*

$$\left\| \sum_{j=1}^d \hat{f}_j - f_0 \right\|_n^2 \leq \left\| \sum_{j=1}^d \tilde{f}_j - f_0 \right\|_n^2 + cdn^{-2/(2+w)}\delta, \quad (30)$$

$$\left\| \sum_{j=1}^d \hat{f}_j - f_0 \right\|_2^2 \leq 2 \left\| \sum_{j=1}^d \tilde{f}_j - f_0 \right\|_2^2 + 24 \left\| \sum_{j=1}^d \tilde{f}_j - f_0 \right\|_n^2 + cdn^{-2/(2+w)}\delta^2, \quad (31)$$

with probability at least $1 - \exp(-c_2c) - c_3/n$, simultaneously over all functions $\tilde{f} = \sum_{j=1}^d \tilde{f}_j$, feasible for the problem (28).

Remark 7 (Error bound for additive, J -smooth f_0). Assume $f_0 = \sum_{j=1}^d f_{0j}$, where $f_{0j} \in \mathcal{S}_j$ and $J(f_{0j}) \leq c_n$, $j = 1, \dots, d$, for a sequence $c_n \geq 1$. Letting $\tilde{f} = f_0$, and $\delta = c_n$, the results in (30), (31) translate to

$$\left\| \sum_{j=1}^d \hat{f}_j - \sum_{j=1}^d f_{0j} \right\|_n^2 \leq cdn^{-2/(2+w)}c_n, \quad \text{and} \quad \left\| \sum_{j=1}^d \hat{f}_j - \sum_{j=1}^d f_{0j} \right\|_2^2 \leq cdn^{-2/(2+w)}c_n^2, \quad (32)$$

with probability at least $1 - \exp(-c_2c) - c_3/n$, provided that $n \geq n_0(dc_n)^{1+w/2}$. From the minimax lower bound in Theorem 3, we can see that the optimal rate for such a class of functions is in fact $dn^{-2/(2+w)}c_n^{2w/(2+w)}$, which reveals that the rates in (32) are tight when c_n is a constant, but not when c_n grows with n . It is worth noting that the dependence of the bounds on c_n in Theorem 2 (and hence in (32)) can be improved to have the optimal scaling of $c_n^{2w/(2+w)}$ by assuming that f_0 is sup norm bounded, and additionally placing a sup norm bound on the components in (28). This feels like an unnecessary restriction, so we prefer to present results without it, as in Theorem 2 (and (32)).

Remark 8 (Distance to best additive, J -smooth approximation of f_0). A consequence of the proof of (30) is a bound on the empirical norm error between \hat{f} and the best additive approximation of f_0 . To be precise, let $f^{\text{best}} = \sum_{j=1}^d f_j^{\text{best}}$ minimize $\left\| \sum_{j=1}^d f_j - f_0 \right\|_n^2$ over all additive functions $f = \sum_{j=1}^d f_j$ feasible for problem (28). Then following directly from (79) in the proof of Theorem 2, we have for all $c \geq c_1$ and $n \geq n_0(d\delta)^{1+w/2}$,

$$\left\| \sum_{j=1}^d \hat{f}_j - \sum_{j=1}^d f_j^{\text{best}} \right\|_n^2 \leq cdn^{-2/(2+w)}\delta, \quad (33)$$

with probability at least $1 - \exp(-c_2c) - c_3/n$, where again $c_1, c_2, c_3, n_0 > 0$ are constants that depend on $b_1, b_2, \sigma, k, L, K, w$. Just as we saw in fixed d case, the right-hand side in (33) does not depend on the approximation error $\|f^{\text{best}} - f_0\|_n$, which is analogous to classical results from Stone [1985].

Remark 9 (L_2 decomposability and incoherence). The decomposability property in (29) is critical in obtaining the sharp (linear) dependence on d in the error rates (30), (31). However, it is worth noting that all that is needed in the proof is in fact a lower bound of the form

$$\left\| \sum_{j=1}^d m_j \right\|_2^2 \geq \phi_0 \sum_{j=1}^d \|m_j\|_2^2, \quad (34)$$

for a constant $\phi_0 > 0$, rather than an equality, as in (29). The above is an incoherence condition that can hold for nonproduct distributions Q , over an appropriate class of functions (additive functions with smooth components), provided that the correlations between components of Q are not too large. See Meier et al. [2009], van de Geer [2014] for similar incoherence conditions.

Next we present our results for additive locally adaptive regression splines (both unrestricted and restricted variants) and additive trend filtering. The proof is in Appendix A.12.

Corollary 2. Assume A3, B1 on the data distribution. Also, assume that the underlying regression function is additive, $f_0 = \sum_{j=1}^d f_{0j}$, where the components f_{0j} , $j = 1, \dots, d$ are k times weakly differentiable, such that $\text{TV}(f_{0j}^{(k)}) \leq c_n$, $j = 1, \dots, d$, for a sequence $c_n \geq 1$. Then for $J(g) = \text{TV}(g^{(k)})$, the following is true of the estimator defined by problem (28).

- (a) Let \mathcal{S}_j be the space of all k times weakly differentiable functions, for each $j = 1, \dots, d$. There exist constants $c_1, c_2, c_3, n_0 > 0$, that depend only on b_1, b_2, σ, k , such that for all $c \geq c_1$ and

$n \geq n_0(dc_n)^{(2k+3)/(2k+2)}$, any solution in the constrained-form additive locally adaptive spline problem (28), with tuning parameter $\delta = c_n$, satisfies

$$\left\| \sum_{j=1}^d \hat{f}_j - \sum_{j=1}^d f_{0j} \right\|_n^2 \leq cdn^{-(2k+2)/(2k+3)}c_n, \quad \text{and} \quad \left\| \sum_{j=1}^d \hat{f}_j - \sum_{j=1}^d f_{0j} \right\|_2^2 \leq cdn^{-(2k+2)/(2k+3)}c_n^2, \quad (35)$$

with probability at least $1 - \exp(-c_2c) - c_3/n$.

(b) Let $\mathcal{S}_j = \mathcal{G}_j$, the set of k th degree splines with knots in the set T_j in (21), for $j = 1, \dots, d$. There exist constants $c_1, c_2, c_3, n_0 > 0$, that depend only on b_1, b_2, σ, k , such that for $c \geq c_1$ and $n \geq (dc_n)^{(2k+3)/(2k+2)}$, any solution in the constrained-form restricted additive locally adaptive spline problem (28), with tuning parameter $\delta = a_k c_n$, where $a_k \geq 1$ is a constant that depends only on k , satisfies (35), with probability at least $1 - \exp(-c_2c) - c_3d/n$.

(c) Let $\mathcal{S}_j = \mathcal{H}_j$, the set of k th degree falling factorial functions defined over X_j (the j th dimension of input points), for $j = 1, \dots, d$. Then there are constants $c_1, c_2, c_3, n_0 > 0$, depending only on b_1, b_2, σ, k , such that for all $c \geq c_1$ and $n \geq n_0(dc_n)^{(2k+3)/(2k+2)}$, any solution in the constrained-form additive trend filtering problem (28), with tuning parameter $\delta = a_k c_n$, where $a_k \geq 1$ is a constant depending only on k , satisfies (35), with probability at least $1 - \exp(-c_2c) - c_3d/n$.

3.4 Minimax lower bounds

We consider minimax lower bounds for estimation over the class of additive functions whose components are smooth with respect to the seminorm J . We allow the dimension d to grow with n . As for the data distribution, we will use the following assumptions in place of A1, A2, A3, B1.

Assumption A4 (Uniform input distribution). The inputs X^i , $i = 1, \dots, n$ are i.i.d. from the uniform distribution on $[0, 1]^d$.

Assumption B2 (Additive model, Gaussian errors). The responses Y^i , $i = 1, \dots, n$ follow

$$Y^i = \mu + \sum_{j=1}^d f_{0j}(X_j^i) + \epsilon^i, \quad i = 1, \dots, n,$$

with mean $\mu \in \mathbb{R}$, where $\int_{[0,1]^d} f_0(x) dx = 0$ for identifiability. The errors ϵ^i , $i = 1, \dots, n$ are i.i.d. $N(0, \sigma^2)$, for some constant $\sigma > 0$. The errors and input points are independent.

For the regularizer J , assumed to satisfy Assumptions C1, C2, we will replace Assumption C3 by the following assumption, on the log packing and log covering (entropy) numbers.

Assumption C4 (Matching packing and covering number bounds). There exist constants $0 < w < 2$ and $K_1, K_2 > 0$ such that

$$\begin{aligned} \log M(\delta, \|\cdot\|_2, B_J(1) \cap B_\infty(1)) &\geq K_1 \delta^{-w}, \\ \log N(\delta, \|\cdot\|_2, B_J(1) \cap B_\infty(1)) &\leq K_2 \delta^{-w}. \end{aligned}$$

(To be clear, here $\|\cdot\|_2$ is the L_2 norm defined with respect to the uniform distribution on $[0, 1]^d$.)

Let us introduce the notation

$$B_J^d(\delta) = \left\{ \sum_{j=1}^d f_j : J(f_j) \leq \delta, j = 1, \dots, d \right\},$$

Now we state our main minimax lower bound. The proof is given in Appendix A.13, A.14.

Theorem 3. Assume *A4*, *B2* on the data distribution, and *C1*, *C2*, *C4* on the seminorm J . Then there exist constants $c_0, n_0 > 0$, that depend only on $\sigma, k, L, K_1, K_2, w$, such that for all $c_n \geq 1$ and $n \geq n_0 d^{1+w/2} c_n^w$, we have

$$\inf_{\hat{f}} \sup_{f_0 \in B_J^d(c_n)} \mathbb{E} \|\hat{f} - f_0\|_2^2 \geq c_0 d n^{-2/(2+w)} c_n^{2w/(2+w)}. \quad (36)$$

When we choose $J(g) = \text{TV}(g^{(k)})$ as our regularizer, the additive function class $B_J^d(\delta)$ becomes

$$\mathcal{F}_k^d(\delta) = \left\{ \sum_{j=1}^d f_j : \text{TV}(f_j^{(k)}) \leq \delta, j = 1, \dots, d \right\},$$

and Theorem 3 implies the following result, whose proof is in Appendix A.15.

Corollary 3. Assume *A4*, *B2* on the data distribution. Assume further that $f_{0j}, j = 1, \dots, d$ are k times weakly differentiable. Then there are constants $c_0, n_0 > 0$, that depend only on σ, k , such that for all $c_n \geq 1$ and $n \geq n_0 d^{(2k+3)/(2k+2)} c_n^{1/(k+1)}$,

$$\inf_{\hat{f}} \sup_{f_0 \in \mathcal{F}_k^d(c_n)} \mathbb{E} \|\hat{f} - f_0\|_2^2 \geq c_0 d n^{-(2k+2)/(2k+3)} c_n^{2/(2k+3)}. \quad (37)$$

Remark 10 (Optimality for a fixed dimension d). For a fixed d , the estimator defined by (22) is minimax rate optimal over the class of additive functions f_0 such that $\sum_{j=1}^d J(f_{0j}) \leq C_n$. To see this, note that such a class of functions contains $B_J^d(C_n/d)$, therefore plugging $c_n = C_n/d$ into the right-hand side in (36) yields a lower bound rate of $n^{-2/(2+w)} C_n^{2w/(2+w)}$, which matches the upper bound rate in (24).

Furthermore, when $J(g) = \text{TV}(g^{(k)})$, the lower bound rate given by plugging $c_n = C_n/d$ into the right-hand side in (37) is $n^{-(2k+2)/(2k+3)} C_n^{2/(2k+3)}$, matching the upper bound rate in (27). Hence additive locally adaptive regression splines, restricted additive locally adaptive regression splines, and additive trend filtering all achieve the minimax rate over the space of additive functions f_0 such that $\sum_{j=1}^d \text{TV}(f_{0j}^{(k)}) \leq C_n$.

Remark 11 (Optimality for a growing dimension d). For growing d , the estimator defined by (28) is minimax rate optimal over the class of additive functions f_0 such that $J(f_{0j}) \leq c, j = 1, \dots, d$, where $c > 0$ is a constant. This is verified by noting that the lower bound rate of $d n^{-2/(2+w)}$ in (36) matches the upper bound rates in (30), (31).

When $J(g) = \text{TV}(g^{(k)})$, and again, $c_n = c$ (a constant), the lower bound rate of $d n^{-(2k+2)/(2k+3)}$ in (37) matches the upper bound rates in (35). Thus additive locally adaptive regression splines, restricted additive locally adaptive regression splines, and additive trend filtering all attain the minimax rate over the space of additive functions f_0 with $\text{TV}(f_{0j}^{(k)}) \leq c, j = 1, \dots, d$.

For growing c_n , we note that the upper bounds in (32) and (35) have an inflated dependence on c_n , compared to (36) and (37). It turns out that the latter (lower bounds) are tight, and the former (upper bounds) are loose. The upper bounds can be tightened under further boundedness assumptions (see Remark 7).

Remark 12 (Suboptimality of additive linear smoothers). Seminal theory from Donoho and Johnstone [1998] on minimax linear rates over Besov spaces shows that, under Assumption B2, and with the inputs $X^i, i = 1, \dots, n$ being now nonrandom and occurring over the regular d -dimensional lattice $\{1/N, 2/N, \dots, 1\}^d \subseteq [0, 1]^d$ with $N = n^{1/d}$, we have

$$\inf_{\hat{f} \text{ additive linear}} \sup_{f_0 \in \mathcal{F}_k^d(c_n)} \mathbb{E} \|\hat{f} - f_0\|_2^2 \geq c_0 d n^{-(2k+1)/(2k+2)} c_n^{2/(2k+2)}, \quad (38)$$

for all $n \geq n_0$, where $c_0, n_0 > 0$ are constants, depending only on σ, k . On the left-hand side in (38) the infimum is taken over all additive linear smoothers, i.e., estimators $\hat{f} = \sum_{j=1}^d \hat{f}_j$ such that each

Algorithm 1 Backfitting for additive trend filtering

Given responses $Y^i \in \mathbb{R}$ and input points $X^i \in \mathbb{R}^d$, $i = 1, \dots, n$.

1. Set $t = 0$ and initialize $\theta_j^{(0)} = 0$, $j = 1, \dots, d$.
 2. For $t = 1, 2, 3, \dots$ (until convergence):
 - a. For $j = 1, \dots, d$:
 - (i) $\theta_j^{(t)} = \text{TF}_\lambda \left(Y - \bar{Y} \mathbb{1} - \sum_{\ell < j} \theta_\ell^{(t)} - \sum_{\ell > j} \theta_\ell^{(t-1)}, X_j \right)$
 - (ii) (Optional) $\theta_j^{(t)} = \theta_j^{(t)} - \frac{1}{n} \mathbb{1}^T \theta_j^{(t)}$
 3. Return $\hat{\theta}_j$, $j = 1, \dots, d$ (parameters $\theta_j^{(t)}$, $j = 1, \dots, d$ at convergence).
-

component \hat{f}_j is a linear smoother, for $j = 1, \dots, d$. The additive linear smoother lower bound (38) is verified in Appendix A.16.

For a fixed d , we can see that all additive linear smoothers—e.g., additive smoothing splines, additive kernel smoothing estimators, additive RKHS estimators, etc.—are suboptimal over the class of additive functions f_0 with $\sum_{j=1}^d \text{TV}(f_{0j}^{(k)}) \leq C_n$, as the optimal linear rate in (38) (set $c_n = C_n/d$) is $n^{-(2k+1)/(2k+2)} C_n^{2/(2k+2)}$, slower than the optimal rate $n^{-(2k+2)/(2k+3)} C_n^{2/(2k+2)}$ of additive locally adaptive regression splines and additive trend filtering in (27).

For growing d , and $c_n = c$ (a constant), we also see that additive linear smoothers are suboptimal over the class of additive functions f_0 such that $\text{TV}(f_{0j}^{(k)}) \leq c$, $j = 1, \dots, d$, as the optimal linear rate in (38) is $dn^{-(2k+1)/(2k+2)}$, slower than the optimal rate $dn^{-(2k+2)/(2k+3)}$ of additive locally adaptive regression splines and additive trend filtering in (35).

4 Backfitting and the dual

We now examine computational approaches for the additive trend filtering problem (7). This is a convex optimization problem, and many standard approaches can be applied. For its simplicity and its ubiquity in additive modeling, we focus on the backfitting algorithm in particular.

4.1 Backfitting

The backfitting approach for problem (7) is described in Algorithm 1. We write $\text{TF}_\lambda(r, X_j)$ for the univariate trend filtering fit, with a tuning parameter $\lambda > 0$, to a response vector $r = (r^1, \dots, r^n) \in \mathbb{R}^n$ over an input vector $X_j = (X_j^1, \dots, X_j^n) \in \mathbb{R}^n$. In words, the algorithm cycles over $j = 1, \dots, d$, and at each step updates the estimate for component j by applying univariate trend filtering to the j th partial residual (i.e., the current residual excluding component j). Centering in Step 2b part (ii) is optional, because the fit $\text{TF}_\lambda(r, X_j)$ will have mean zero whenever r has mean zero, but centering can still be performed for numerical stability. In general, the efficiency of backfitting hinges on the efficiency of the univariate smoother employed; to implement Algorithm 1 in practice we can use fast interior point methods [Kim et al., 2009] or fast operator splitting methods [Ramdas and Tibshirani, 2016] for univariate trend filtering, both of which result in efficient empirical performance.

Algorithm 1 is equivalent to block coordinate descent (BCD), also called exact blockwise minimization, applied to problem (7) over the coordinate blocks θ_j , $j = 1, \dots, d$. A general treatment of BCD is given in Tseng [2001], who shows that for a convex criterion that decomposes into smooth plus separable terms, as does that in (7), all limit points of the sequence of iterates produced by

BCD are optimal solutions. A recent wave of work from the optimization community gives refined convergence analyses for coordinate descent (or its variants) in particular settings. We do not pursue the implications of this work for our problem; our interest here is primarily in developing a connection between BCD for problem (7) and alternating projections in its dual problem (18), which is the topic of the next subsection.

4.2 Dual alternating projections

Using the additive trend filtering problem (7) and its dual (18), related by the transformation (19), we see that for any dimension $j = 1, \dots, d$, the univariate trend filtering fit with response vector $r = (r^1, \dots, r^n)$ and input vector $X_j = (X_j^1, \dots, X_j^n)$ can be written as

$$\text{TF}_\lambda(r, X_j) = (\text{Id} - \Pi_{U_j})(r), \quad (39)$$

where $U_j = \{S_j D_j^T v_j : \|u\|_\infty \leq \lambda\}$, and recall, we abbreviate $D_j = D^{(X_j, k+1)}$. The backfitting approach in Algorithm 1 can be viewed (ignoring the optional centering step) as performing the updates, for $t = 1, 2, 3, \dots$,

$$\theta_j^{(t)} = (\text{Id} - \Pi_{U_j}) \left(Y - \bar{Y} \mathbb{1} - \sum_{\ell < j} \theta_\ell^{(t)} - \sum_{\ell > j} \theta_\ell^{(t-1)} \right), \quad j = 1, \dots, d, \quad (40)$$

or, reparametrized in terms of the primal-dual relationship $u = Y - \bar{Y} \mathbb{1} - \sum_{j=1}^d \theta_j$ in (19),

$$\begin{aligned} u_0^{(t)} &= Y - \bar{Y} \mathbb{1} - \sum_{j=1}^d \theta_j^{(t-1)}, \\ u_j^{(t)} &= \Pi_{U_j} (u_{j-1}^{(t)} + \theta_j^{(t-1)}), \quad j = 1, \dots, d, \\ \theta_j^{(t)} &= \theta_j^{(t-1)} + u_{j-1}^{(t)} - u_j^{(t)}, \quad j = 1, \dots, d. \end{aligned} \quad (41)$$

Thus the backfitting algorithm for (7), as expressed above in (41), is seen to be a particular type of *alternating projections* method applied to the dual problem (18), cycling through projections onto U_j , $j = 1, \dots, d$. Interestingly, as opposed to the classical alternating projections approach, which would repeatedly project the current iterate $u_{j-1}^{(t)}$ onto U_j , $j = 1, \dots, d$, the steps in (41) repeatedly project an “offset” version $u_{j-1}^{(t)} + \theta_j^{(t-1)}$ of the current iterate, for $j = 1, \dots, d$ (this corresponds to running univariate trend filtering on the current residual, in the iterations (40)).

There is a considerable literature on alternating projections in optimization, see, e.g., [Bauschke and Borwein \[1996\]](#) for a review. Many alternating projections algorithms can be derived from the perspective of an operator splitting technique, e.g., the alternative direction method of multipliers (ADMM). Indeed, the steps in (41) appear very similar to those from an ADMM algorithm applied to the dual (18), if we think of the “offset” variables θ_j , $j = 1, \dots, d$ in the iterations (41) as *dual* variables in the dual problem (18) (i.e., if we think of the primal variables θ_j , $j = 1, \dots, d$ as dual variables in the dual problem (18)). This connection inspires a new parallel version of backfitting, presented in the next subsection.

4.3 Parallelized backfitting

We have seen that backfitting is a special type of alternating projections algorithm, applied to the dual problem (18). For set intersection problems (where we seek a point in the intersection of given closed, convex sets), the optimization literature offers a variety of *parallel projections* methods (in contrast to alternating projections methods) that are provably convergent. One such method can be

derived using ADMM (e.g., see Section 5.1 of [Boyd et al. \[2011\]](#)), and a similar construction can be used for the dual problem (18). We first rewrite this problem as

$$\begin{aligned} \min_{u_0, u_1, \dots, u_d \in \mathbb{R}^n} \quad & \frac{1}{2} \|Y - \bar{Y} \mathbb{1} - u_0\|_2^2 + \sum_{j=1}^d I_{U_j}(u_j) \\ \text{subject to} \quad & u_0 = u_1, u_0 = u_2, \dots, u_0 = u_d, \end{aligned} \quad (42)$$

where we write I_S for the indicator function of a set S (equal to 0 on S , and ∞ otherwise). Then we define the augmented Lagrangian, for an arbitrary $\rho > 0$, as

$$L_\rho(u_0, u_1, \dots, u_d, \gamma_1, \dots, \gamma_d) = \frac{1}{2} \|Y - \bar{Y} \mathbb{1} - u_0\|_2^2 + \sum_{j=1}^d \left(I_{U_j}(u_j) + \frac{\rho}{2} \|u_0 - u_j + \gamma_j\|_2^2 - \frac{\rho}{2} \|\gamma_j\|_2^2 \right).$$

The ADMM steps for problem (42) are now given by repeating, for $t = 1, 2, 3, \dots$,

$$\begin{aligned} u_0^{(t)} &= \frac{1}{\rho d + 1} \left(Y - \bar{Y} \mathbb{1} + \rho \sum_{j=1}^d (u_j^{(t-1)} - \gamma_j^{(t-1)}) \right) \\ u_j^{(t)} &= \Pi_{U_j}(u_0^{(t)} + \gamma_j^{(t-1)}), \quad j = 1, \dots, d \\ \gamma_j^{(t)} &= \gamma_j^{(t-1)} + u_0^{(t)} - u_j^{(t)}, \quad j = 1, \dots, d. \end{aligned} \quad (43)$$

Now compare (43) to (41)—the key difference is that in (43), the updates to u_j , $j = 1, \dots, d$, i.e., the projections onto U_j , $j = 1, \dots, d$, completely decouple and can hence be performed *in parallel*. Run properly, this could provide a large speedup over the sequential projections in (41).

Of course, for our current study, the dual problem (42) is really only interesting insofar as it is connected to the additive trend filtering problem (7). Fortunately, the parallel projections algorithm (43) maintains a very useful connection to the primal problem (7): $\rho \hat{\gamma}_j$, $j = 1, \dots, d$ (i.e., the scaled iterates $\rho \gamma_j^{(t)}$, $j = 1, \dots, d$ at convergence) are optimal for the additive trend filtering problem (7). This is simply because the dual of the dual problem (42) is the additive trend filtering problem (7) (therefore $\rho \gamma_j$, $j = 1, \dots, d$, which are dual to the constraints in (42), are equivalent to the primal parameters θ_j , $j = 1, \dots, d$ in (7)). We state this next as a theorem, and transcribe the iterations in (43) into an equivalent primal form, in Algorithm 2. For details, see Appendix A.17.

Theorem 4. *Initialized arbitrarily, the ADMM steps (43) produce parameters $\hat{\gamma}_j$, $j = 1, \dots, d$ (i.e., the iterates $\gamma_j^{(t)}$, $j = 1, \dots, d$ at convergence) such that the scaled parameters $\rho \hat{\gamma}_j$, $j = 1, \dots, d$ solve additive trend filtering (7). Further, the outputs $\hat{\theta}_j$, $j = 1, \dots, d$ of Algorithm 2 solve additive trend filtering (7).*

Written in primal form, we see that the parallel backfitting approach in Algorithm 2 differs from what may be considered the “naive” approach to parallelizing the usual backfitting iterations in Algorithm 1. Consider $\rho = 1$. If we were to replace Step 2a in Algorithm 2 with $u_0^{(t)} = r^{(t-1)}$, the full residual

$$r^{(t-1)} = Y - \bar{Y} \mathbb{1} - \sum_{j=1}^d \theta_j^{(t-1)},$$

then the update steps for $\theta_j^{(t)}$, $j = 1, \dots, d$ that follow would be just given by applying univariate trend filtering to each partial residual (without sequentially updating the partial residuals between trend filtering runs). This naive parallel method has no convergence guarantees, and can fail even in simple practical examples to produce optimal solutions. Importantly, Algorithm 2 does not take $u_0^{(t)}$ to be the full residual, but as Step 2a shows, uses a less greedy choice: it basically takes $u_0^{(t)}$ to be a convex combination of the residual $r^{(t-1)}$ and its previous value $u_0^{(t-1)}$, with higher weight on

Algorithm 2 Parallel backfitting for additive trend filtering

Given responses $Y^i \in \mathbb{R}$, input points $X^i \in \mathbb{R}^d$, $i = 1, \dots, n$, and $\rho > 0$.

1. Initialize $u_0^{(0)} = 0$, $\theta_j^{(0)} = 0$ and $\theta_j^{(-1)} = 0$ for $j = 1, \dots, d$.
2. For $t = 1, 2, 3, \dots$ (until convergence):

- a. $u_0^{(t)} = \frac{1}{\rho d + 1} \left(Y - \bar{Y} \mathbb{1} - \sum_{j=1}^d \theta_j^{(t-1)} \right) + \frac{\rho d}{\rho d + 1} \left(u_0^{(t-1)} + \frac{1}{\rho d} \sum_{j=1}^d (\theta_j^{(t-2)} - \theta_j^{(t-1)}) \right)$

- b. For $j = 1, \dots, d$ (in parallel):

- (i) $\theta_j^{(t)} = \rho \cdot \text{TF}_\lambda(u_0^{(t)} + \theta_j^{(t-1)}) / \rho, X_j$
- (ii) (Optional) $\theta_j^{(t)} = \theta_j^{(t-1)} - \frac{1}{n} \mathbb{1}^T \theta_j^{(t)}$

3. Return $\hat{\theta}_j$, $j = 1, \dots, d$ (parameters $\theta_j^{(t)}$, $j = 1, \dots, d$ at convergence).
-

the latter. The subsequent parallel updates for $\theta_j^{(t)}$, $j = 1, \dots, d$ are still given by univariate trend filtering fits, and though these steps do not exactly use partial residuals (since $u_0^{(t)}$ is not exactly the full residual), they are guaranteed to produce additive trend filtering solutions upon convergence (as per Theorem 4). An example of cyclic versus parallelized backfitting is given in Appendix A.18.

5 Experiments

Through empirical experiments, we examine the performance of additive trend filtering relative to additive smoothing splines. We also examine the efficacy of cross-validation for choosing the tuning parameter λ , as well as the use of multiple tuning parameters (i.e., a separate parameter λ_j for each component $j = 1, \dots, d$). All experiments were performed in R. For the univariate trend filtering solver, we used the `trendfilter` function in the `glmgen` package, which is an implementation of the fast ADMM algorithm given in Ramdas and Tibshirani [2016]. For the univariate smoothing spline solver, we used the `smooth.spline` function in base R.

5.1 Simulated heterogeneously-smooth data

We sampled $n = 2500$ input points in $d = 10$ dimensions, by assigning the inputs along each dimension $X_j = (X_j^1, \dots, X_j^n)$ to be a different permutation of the equally spaced points $(1/n, 2/n, \dots, 1)$, for $j = 1, \dots, 10$. For the componentwise trends, we examined sinusoids with Doppler-like spatially-varying frequencies:

$$g_{0j}(x_j) = \sin \left(\frac{2\pi}{(x_j + 0.1)^{j/10}} \right), \quad j = 1, \dots, 10.$$

We then defined the component functions as $f_{0j} = a_j g_{0j} - b_j$, $j = 1, \dots, d$, where a_j, b_j were chosen so that f_{0j} had empirical mean zero and empirical norm $\|f_{0j}\|_n = 1$, for $j = 1, \dots, d$. The responses were generated according to $Y^i \stackrel{\text{i.i.d.}}{\sim} N(\sum_{j=1}^d f_{0j}(X_j^i), \sigma^2)$, $i = 1, \dots, 2500$. By construction, in this setup, there is considerable heterogeneity in the levels of smoothness both within and between the component functions.

The left panel of Figure 4 shows a comparison of the MSE curves from additive trend filtering in (7) (of quadratic order, $k = 2$) and additive smoothing splines in (1) (of cubic order). We set σ^2 in the generation of the responses so that the signal-to-noise ratio (SNR) was $\|f_0\|_n^2 / \sigma^2 = 4$, where $f_0 = \sum_{j=1}^d f_{0j}$. The two methods (additive trend filtering and additive smoothing splines) were each




Figure 4: The left panel shows the MSE curves for additive trend filtering (7) (of quadratic order) and additive smoothing splines (1) (of cubic order), computed over 10 repetitions from the heterogeneous smoothness simulation with $n = 2500$ and $d = 10$, described in Section 5.1, where the SNR is set to 4. Vertical segments denote ± 1 standard deviations. The right panel displays the best-case MSE for each method (the minimum MSE over its regularization path), in a problem setup with $n = 1000$ and $d = 6$, as the signal-to-noise ratio (SNR) varies from 0.7 to 16, in equally spaced values on the log scale.

allowed their own sequence of tuning parameter values, and results were averaged over 10 repetitions from the simulation setup described above. As we can see, additive trend filtering achieves a better minimum MSE along its regularization path, and does so at a less complex model (lower df).

The right panel of Figure 4 shows the best-case MSEs for additive trend filtering and additive smoothing splines (i.e., the minimum MSE over their regularization paths) as the noise level σ^2 is varied so that the SNR ranges from 0.7 to 1.6, in equally spaced values on the log scale. The results were again averaged over 10 repetitions of data drawn from a simulation setup essentially the same as the one described above, except that we considered a smaller problem size, with $n = 1000$ and $d = 6$. The plot reveals that additive trend filtering performs increasingly well (in comparison to additive smoothing splines) as the SNR grows—not surprising, as here it is able to better capture the heterogeneity in the component functions.

Lastly, in Appendix A.19, we present results from an experimental setup mimicking that in this subsection, except with the component functions f_{0j} , $j = 1, \dots, d$ having homogeneous smoothness throughout. Here additive trend filtering and additive smoothing splines perform very similarly.

5.2 Cross-validation and multiple tuning parameters

Sticking to the simulation setup from the last subsection, but at the smaller problem size, $n = 1000$ and $d = 6$ (used to produce the right panel of Figure 4), we study in the left panel of Figure 5 the use of 5-fold cross-validation (CV) to select the tuning parameter λ for additive trend filtering and additive smoothing splines. Displayed are the resulting MSE curves as the SNR varies from 0.7 to 16. Also shown on the same plot are the oracle MSE curves (which are the same as those the right panel of Figure 4), in which λ has been chosen to minimize the MSE for each method. We can see that the performance of each method degrades using CV, but not by much.

In the right panel of the figure, we examine the use of multiple tuning parameters for additive

smoothing splines and additive trend filtering, i.e., replacing the penalties in (1) and (7) by

$$\sum_{j=1}^d \lambda_j \theta_j^T Q_j \theta_j \quad \text{and} \quad \sum_{j=1}^d \lambda_j \|D^{(X_j, k+1)} S_j \theta_j\|_1,$$

respectively, which means we would now have d tuning parameters λ_j , $j = 1, \dots, d$. When the function we are estimating has different amounts of smoothness along different dimensions, we have argued (and seen through examples) that additive trend filtering—using only a single tuning parameter λ —can accommodate these differences, at least somewhat, thanks to its locally adaptive nature. But, when the difference in smoothness across dimensions is drastic enough, it may be worthwhile to use multiple tuning parameters.

When d is moderate (even just for $d = 6$), cross-validation over a d -dimensional grid of values for λ_j , $j = 1, \dots, d$ can be prohibitive. However, as pointed out by a referee of this article, there has been a considerable amount of work dedicated to this problem by authors studying additive models built from splines (or other linear smoothers), e.g., Gu and Wahba [1991], Wood [2000], Fahrmeir and Lang [2001], Ruppert et al. [2003], Wood [2004], Kim and Gu [2004], Rue et al. [2009], Wood [2011], Wood et al. [2015, 2016]. Many of these papers use an efficient computational approach based on restricted maximum likelihood (REML) for selecting λ_j , $j = 1, \dots, d$; see also Wood [2017] for a nice introduction and description of this approach. Unfortunately, as far as we see it, REML does not easily apply to additive trend filtering.

We thus use the following simple approach for multiple tuning parameter selection: within each backfitting loop, for each component $j = 1, \dots, d$, we use (univariate) CV to choose λ_j . While this does not solve a particular convex optimization problem, and is not guaranteed to converge in general, we have found it to work quite well in practice. The right panel of Figure 5 compares the performance of this so-called backfit-CV tuning to the oracle, that chooses just a single tuning parameter. Both additive trend filtering and additive smoothing splines are seen to improve with d tuning parameters, tuned by backfit-CV, in comparison to the oracle choice of tuning parameter. Interestingly, we also see that additive smoothing splines with d tuning parameters performs on par with additive trend filtering with the oracle choice of tuning parameter. (In this example, REML tuning for additive smoothing splines—as implemented by the `mgcv` R package—performed a bit worse than backfit-CV tuning, and so we only show results from the latter.)

6 Discussion

We have studied additive models built around the univariate trend filtering estimator, i.e., defined by penalizing according to the sum of ℓ_1 norms of discrete derivatives of the component functions. We examined basic properties of these additive models, such as extrapolation of the fitted values to a d -dimensional surface, uniqueness of the component fits, and characterization of effective degrees of freedom of the fit. When the underlying regression function is additive, with components whose k th derivatives are of bounded variation, we derived error rates for k th order additive trend filtering: $n^{-(2k+2)/(2k+3)}$ for a fixed input dimension d (under weak assumptions), and $dn^{-(2k+2)/(2k+3)}$ for a growing dimension d (under stronger assumptions). We showed these rates are sharp by establishing matching minimax lower bounds, and showed that additive linear smoothers (e.g., additive smoothing splines) can at best achieve a rate of $n^{-(2k+1)/(2k+2)}$ for fixed d , and $dn^{-(2k+1)/(2k+2)}$ for growing d , over the same function class. Finally, on the computational side, we devised a provably convergent parallel backfitting algorithm for additive trend filtering.

It is worth noting that our parallel backfitting method is not specific to additive trend filtering, but it can be embedded in a more general parallel coordinate descent framework [Tibshirani, 2017]. It would be interesting to follow up on this framework to see what practical benefits it can provide for additive models in particular.

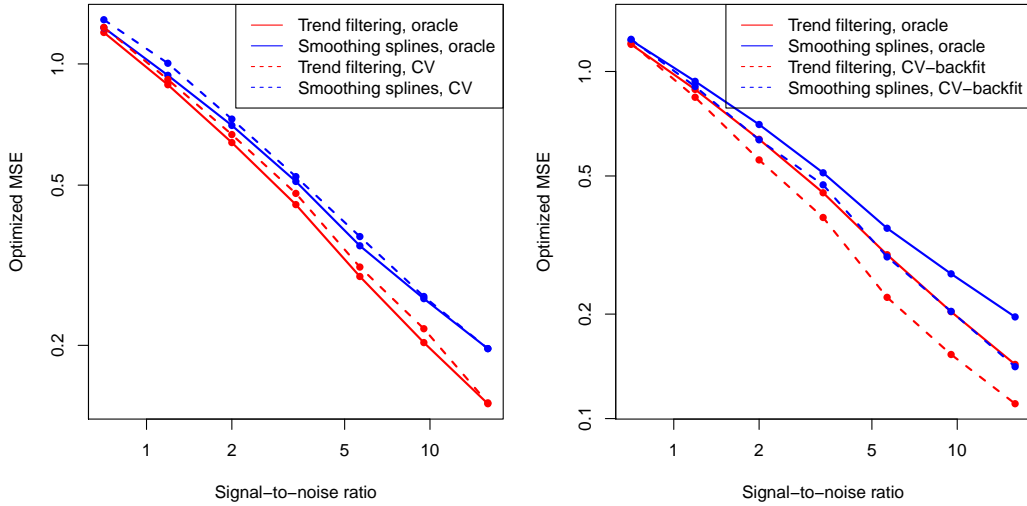


Figure 5: Both panels display results from the same simulation setup as that in the right panel of Figure 4. The left panel shows MSE curves when the estimators are tuned by 5-fold cross-validation (CV), and also by the oracle (reflecting the minimum possible MSE). The right panel displays MSE curves when we allow each estimator to have d tuning parameters, tuned by a hybrid backfit-CV method explained in the text, versus the oracle MSE curves for a single tuning parameter.

A natural extension of our work is to consider the high-dimensional case, where d is comparable or possibly even much larger than n , and we fit a *sparse additive model* by employing an additional sparsity penalty in problem (7). Another natural extension is to consider responses $Y^i|X^i, i = 1, \dots, n$ from an exponential family distribution, and we fit a *generalized additive model* by changing the loss in (7). After we completed an initial version of this paper, both extensions have been pursued: [Tan and Zhang \[2017\]](#) develop a suite of error bounds for sparse additive models, with various form of penalties (which include total variation on derivatives of components); and [Haris et al. \[2018\]](#) give comprehensive theory for sparse generalized additive models, with various types of penalties (which again include total variation on derivatives of components).

Acknowledgements

We are very thankful to Garvesh Raskutti for his generous help and insight on various issues, and Martin Wainwright for generously sharing his unpublished book with us. We are also grateful to an anonymous referee whose thoughtful comments improved our paper.

A Appendix

A.1 Fast extrapolation

We discuss extrapolation using the fitted functions \hat{f}_j , $j = 1, \dots, d$ from additive trend filtering (10), as in (13). We must compute the coefficients $\hat{\alpha}_j = (\hat{a}_j, \hat{b}_j)$ whose block form is given in (14), (15). Clearly, the computation of \hat{b}_j in (15) requires $O(n)$ operations (owing to the bandedness of $D^{(X_j, k+1)}$, and treating k as a constant). As for \hat{a}_j in (14), it can be seen from the structure of $C^{(X_j, k+1)}$ as described in Wang et al. [2014] that

$$(\hat{a}_j)_1 = (S_j \hat{\theta}_j)_1,$$

$$(\hat{a}_j)_\ell = \frac{1}{(\ell-1)!} \left[\text{diag} \left(\frac{1}{X_j^\ell - X_j^1}, \dots, \frac{1}{X_j^n - X_j^{n-\ell+1}} \right) D^{(X_j, \ell-1)} S_j \hat{\theta}_j \right]_1, \quad \ell = 2, \dots, k+1,$$

which takes only $O(1)$ operations (again treating k as constant, and now using the bandedness of each $D_j^{(X_j, \ell-1)}$, $\ell = 2, \dots, k+1$). In total then, computing the coefficients $\hat{\alpha}_j = (\hat{a}_j, \hat{b}_j)$ requires $O(n)$ operations, and computing $\hat{\alpha} = (\hat{\alpha}_1, \dots, \hat{\alpha}_d)$ requires $O(nd)$ operations.

After having computed $\hat{\alpha} = (\hat{\alpha}_1, \dots, \hat{\alpha}_d)$, which only needs to be done once, a prediction at a new point $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ with the additive trend filtering fit \hat{f} is given by

$$\hat{f}(x) = \bar{Y} + \sum_{j=1}^d \sum_{\ell=1}^n \hat{\alpha}_j^\ell h_\ell^{(X_j)}(x_j),$$

This requires $O(d + \sum_{j=1}^d \sum_{\ell=k+2}^n \mathbb{1}\{\hat{\alpha}_j^\ell \neq 0\})$ operations, utilizing the sparsity of the components in $\hat{\alpha}$ not associated with the polynomial basis functions.

A.2 Proof of Lemma 2

We begin by eliminating the constraint in the additive trend filtering problem (7), rewriting it as

$$\min_{\theta_1, \dots, \theta_d \in \mathbb{R}^n} \frac{1}{2} \left\| MY - \sum_{j=1}^d M \theta_j \right\|_2^2 + \lambda \sum_{j=1}^d \| D^{(X_j, k+1)} S_j M \theta_j \|_1,$$

where $M = I - \mathbb{1} \mathbb{1}^T / n$. Noting that $D^{(X_j, k+1)} \mathbb{1} = 0$ for $j = 1, \dots, d$, we can replace the penalty term above by $\sum_{j=1}^d \| D^{(X_j, k+1)} S_j \theta_j \|_1$. Reparametrizing using the falling factorial basis, as in Lemma 1, yields the problem

$$\min_{a \in \mathbb{R}^{(k+1)d}, b \in \mathbb{R}^{(n-k-1)d}} \frac{1}{2} \left\| MY - M \sum_{j=1}^d P_j a_j - M \sum_{j=1}^d K_j b_j \right\|_2^2 + \lambda k! \sum_{j=1}^d \| b_j \|_1,$$

where we have used the abbreviation $P_j = P^{(X_j, k)}$ and $K_j = K^{(X_j, k)}$, as well as the block representation $\alpha_j = (a_j, b_j) \in \mathbb{R}^{(k+1)} \times \mathbb{R}^{(n-k-1)}$, for $j = 1, \dots, d$. Since each P_j , $j = 1, \dots, d$ has $\mathbb{1}$ for its first column, the above problem is equivalent to

$$\min_{a \in \mathbb{R}^{kd}, b \in \mathbb{R}^{(n-k-1)d}} \frac{1}{2} \left\| MY - M \sum_{j=1}^d \tilde{P}_j a_j - M \sum_{j=1}^d K_j b_j \right\|_2^2 + \lambda k! \sum_{j=1}^d \| b_j \|_1,$$

where \tilde{P}_j denotes P_j with the first column removed, for $j = 1, \dots, d$. To be clear, solutions in the above problem and the original trend filtering formulation (7) are related by

$$\hat{\theta}_j = \tilde{P}_j \hat{a}_j + K_j \hat{b}_j, \quad j = 1, \dots, d.$$

Furthermore, we can see that $\hat{a} = (\hat{a}_1, \dots, \hat{a}_d)$ solves

$$\min_{a \in \mathbb{R}^{kd}} \frac{1}{2} \left\| \left(MY - M \sum_{j=1}^d K_j \hat{b}_j \right) - \tilde{P}a \right\|_2^2, \quad (44)$$

where \tilde{P} is as defined in (16), and $\hat{b} = (\hat{b}_1, \dots, \hat{b}_d)$ solves

$$\min_{b \in \mathbb{R}^{(n-k-1)d}} \frac{1}{2} \left\| UU^T MY - UU^T M \sum_{j=1}^d K_j b_j \right\|_2^2 + \lambda k! \|b\|_1,$$

where UU^T is the projection orthogonal to the column space of \tilde{P} , i.e., it solves

$$\min_{b \in \mathbb{R}^{(n-k-1)d}} \frac{1}{2} \left\| U^T MY - \tilde{K}b \right\|_2^2 + \lambda k! \|b\|_1, \quad (45)$$

where \tilde{K} is as in (17). Since problem (45) is a standard lasso problem, existing results on the lasso (e.g., Tibshirani [2013]) imply that the solution \hat{b} is unique whenever \tilde{K} has columns in general position. This proves the first part of the lemma. For the second part of the lemma, note that the solution \hat{a} in the least squares problem (44) is just given by the regression of $MY - M \sum_{j=1}^d K_j \hat{b}_j$ onto \tilde{P} , which is unique whenever \tilde{P} has full column rank. This completes the proof.

A.3 Derivation of additive trend filtering dual

As in the proof of Lemma 2, we begin by rewriting the problem (7) as

$$\min_{\theta_1, \dots, \theta_d \in \mathbb{R}^n} \frac{1}{2} \left\| MY - \sum_{j=1}^d M\theta_j \right\|_2^2 + \lambda \sum_{j=1}^d \|D_j S_j M\theta_j\|_1,$$

where $M = I - \mathbb{1}\mathbb{1}^T/n$. Then, we reparametrize the above problem,

$$\begin{aligned} \min_{\substack{\theta_1, \dots, \theta_d \in \mathbb{R}^n \\ w \in \mathbb{R}^n, z \in \mathbb{R}^{md}}} & \frac{1}{2} \|MY - w\|_2^2 + \lambda \sum_{j=1}^d \|z_j\|_1 \\ \text{subject to} & \quad w = \sum_{j=1}^d M\theta_j, \quad z_j = D_j S_j M\theta_j, \quad j = 1, \dots, d, \end{aligned}$$

and form the Lagrangian

$$L(\theta, w, z, u, v) = \frac{1}{2} \|MY - w\|_2^2 + \lambda \sum_{j=1}^d \|z_j\|_1 + u^T \left(w - \sum_{j=1}^d M\theta_j \right) + \sum_{j=1}^d v_j^T (D_j S_j M\theta_j - z_j).$$

Minimizing the Lagrangian L over all θ, z yields the dual problem

$$\begin{aligned} \max_{\substack{u \in \mathbb{R}^n \\ v_1, \dots, v_d \in \mathbb{R}^m}} & \frac{1}{2} \|MY\|_2^2 - \frac{1}{2} \|MY - u\|_2^2 \\ \text{subject to} & \quad u = S_j D_j^T v_j, \quad \|v_j\|_\infty \leq \lambda, \quad j = 1, \dots, d. \end{aligned}$$

The claimed dual problem (18) is just the above, rewritten in an equivalent form.

A.4 Proof of Lemma 3

We first eliminate the equality constraint in (7), rewriting this problem, as was done in the proof of Lemma 2, as

$$\min_{\theta_1, \dots, \theta_d \in \mathbb{R}^d} \frac{1}{2} \left\| MY - \sum_{j=1}^d M\theta_j \right\|_2^2 + \lambda \sum_{j=1}^d \|D_j S_j \theta_j\|_1,$$

where $M = I - \mathbb{1}\mathbb{1}^T/n$, and $D_j = D^{(X_j, k+1)}$, $j = 1, \dots, d$. This is a generalized lasso problem with a design matrix $T \in \mathbb{R}^{n \times nd}$ that has d copies of M stacked along its columns, and a penalty matrix $D \in \mathbb{R}^{nd \times nd}$ that is block diagonal in the blocks D_j , $j = 1, \dots, d$. Applying Theorem 3 of Tibshirani and Taylor [2012], we see that

$$\text{df}(T\hat{\theta}) = \mathbb{E}[\dim(T\text{null}(D_{-A}))],$$

where $A = \text{supp}(D\hat{\theta})$, and where D_{-A} denotes the matrix D with rows removed that correspond to the set A . The conditions for uniqueness in the lemma now precisely imply that

$$\dim(T\text{null}(D_{-A})) = \left(\sum_{j=1}^d |A_j| \right) + kd,$$

where A_j denotes the subset of A corresponding to the block of rows occupied by D_j , and $|A_j|$ its cardinality, for $j = 1, \dots, d$. This can be verified by transforming to the basis perspective as utilized in the proofs of Lemmas 1 and 2. The desired result is obtained by noting that, for $j = 1, \dots, d$, the component $\hat{\theta}_j$ exhibits a knot for each element in A_j .

A.5 Preliminaries for the proof of Theorem 1

Before the proof of Theorem 1, we collect important preliminary results. We start with a result on orthonormal polynomials. We thank Dejan Slepcev for his help with the next lemma.

Lemma 4. *Given an integer $\kappa \geq 0$, and a continuous measure Λ on $[0, 1]$, whose Radon-Nikodym derivative λ is bounded below and above by constants $b_1, b_2 > 0$, respectively. Denote by $\phi_0, \dots, \phi_\kappa$ an orthonormal basis for the space of polynomials of degree κ on $[0, 1]$, given by running the Gram-Schmidt procedure on the polynomials $1, t, \dots, t^\kappa$, with respect to the $L_2(\Lambda)$ inner product. Hence, for $\ell = 0, \dots, \kappa$, ϕ_ℓ is an ℓ th degree polynomial, orthogonal (in $L_2(\Lambda)$) to all polynomials of degree less than ℓ , and we denote its leading coefficient by $a_\ell > 0$. Now define, for $t \in [0, 1]$:*

$$\begin{aligned} \Phi_{\kappa,0}(t) &= \phi_\kappa(t)\lambda(t), \\ \Phi_{\kappa,\ell+1}(t) &= \int_0^t \Phi_{\kappa,\ell}(u) du, \quad \ell = 0, \dots, \kappa. \end{aligned}$$

Then the following two relations hold:

$$\Phi_{\kappa,\ell}(1) = \begin{cases} 0 & \text{for } \ell = 1, \dots, \kappa, \\ \frac{(-1)^\kappa}{a_\kappa \kappa!} & \text{for } \ell = \kappa + 1, \end{cases} \quad (46)$$

and

$$a_\kappa \kappa! |\Phi_{\kappa,\kappa}(t)| \leq \binom{2\kappa}{\kappa} \sqrt{\frac{b_2}{b_1}}, \quad t \in [0, 1]. \quad (47)$$

Proof. First, we use induction to show that for $t \in [0, 1]$,

$$\Phi_{\kappa,\ell}(t) = \int_0^t \phi_\kappa(u) \frac{(t-u)^{\ell-1}}{(\ell-1)!} \lambda(u) du, \quad \ell = 1, \dots, \kappa + 1. \quad (48)$$

This statement holds for $\ell = 1$ by definition of $\Phi_{\kappa,0}, \Phi_{\kappa,1}$. Assume it holds at some $\ell > 1$. Then

$$\begin{aligned}\Phi_{\kappa,\ell+1}(t) &= \int_0^t \int_0^u \phi_{\kappa}(v) \frac{(u-v)^{\ell-1}}{(\ell-1)!} \lambda(v) dv du \\ &= \int_0^t \phi_{\kappa}(v) \left(\int_v^t \frac{(u-v)^{\ell-1}}{(\ell-1)!} du \right) \lambda(v) dv \\ &= \int_0^t \phi_{\kappa}(v) \frac{(t-v)^{\ell}}{\ell!} \lambda(v) dv,\end{aligned}$$

where we used inductive hypothesis in the first line and Fubini's theorem in the second line, which completes the inductive proof.

Now, the relation in (48) shows that $\Phi_{\kappa,\ell}(1)$ is the $L_2(\Lambda)$ inner product of ϕ_{κ} and an $(\ell-1)$ st degree polynomial, for $\ell = 1, \dots, \kappa$. As ϕ_{κ} is orthogonal to all polynomials of degree less than κ , we have $\Phi_{\kappa,\ell}(1) = 0$, $\ell = 1, \dots, \kappa$. For $\ell = \kappa + 1$, note that this same orthogonality along with (48) also shows

$$\Phi_{\kappa,\kappa+1}(1) = \left\langle \phi_{\kappa}, \frac{(-1)^{\kappa}}{a_{\kappa}\kappa!} \phi_{\kappa} \right\rangle_2 = \frac{(-1)^{\kappa}}{a_{\kappa}\kappa!}.$$

where $\langle \cdot, \cdot \rangle_2$ is the $L_2(\Lambda)$ inner product. This establishes the statement in (46).

As for (47), note that if $\kappa = 0$ then the statement holds, because the left-hand side is 1 and the right-hand side is always larger than 1. Hence consider $\kappa \geq 1$. From (48), we have, for any $t \in [0, 1]$,

$$\begin{aligned}|\Phi_{\kappa,\kappa}(t)| &\leq \int_0^t |\phi_{\kappa}(u)| \frac{(t-u)^{\kappa-1}}{(\kappa-1)!} \lambda(u) du \\ &\leq \left(\int_0^t \phi_{\kappa}^2(u) \lambda(u) du \right)^{1/2} \left(\int_0^t \frac{(t-u)^{2\kappa-2}}{(\kappa-1)!^2} \lambda(u) du \right)^{1/2} \\ &\leq \frac{\sqrt{b_2}}{(\kappa-1)! \sqrt{2\kappa-1}},\end{aligned}\tag{49}$$

where in the second line we used Cauchy-Schwartz, and in the third line we used the fact that ϕ_{κ} has unit norm, and the upper bound b_2 on λ . Next we bound a_{κ} . Let p be the projection of x^{κ} onto the space of polynomials of degree $\kappa-1$, with respect to the $L_2(\Lambda)$ inner product. Then we have $\phi_{\kappa} = (x^{\kappa} - p)/\|x^{\kappa} - p\|_2$, thus its leading coefficient is $a_{\kappa} = 1/\|x^{\kappa} - p\|_2$, where $\|\cdot\|_2$ is the $L_2(\Lambda)$ norm. Consider

$$\begin{aligned}\|x^{\kappa} - p\|_2 &\geq \sqrt{b_1} \left(\int_0^1 (x^{\kappa} - p)^2(t) dt \right)^{1/2} \\ &\geq \sqrt{b_1} \left(\int_0^1 P_{\kappa}^2(t) dt \right)^{1/2} = \frac{\sqrt{b_1}}{\sqrt{2\kappa+1} \binom{2\kappa}{\kappa}}.\end{aligned}\tag{50}$$

In the first line we used the lower bound b_1 on λ . In the second we used the fact the Legendre polynomial P_{κ} of degree κ , shifted to $[0, 1]$ but unnormalized, is the result from projecting out $1, t, \dots, t^{\kappa-1}$ from t^{κ} , with respect to the uniform measure. In the last step we used the fact that P_{κ} has norm $1/(\sqrt{2\kappa+1} \binom{2\kappa}{\kappa})$. Combining (49) and (50) gives the result (47). \square

Remark 13 (Special case: uniform measure and Rodrigues' formula). *In the case of $\Lambda = U$, the uniform measure on $[0, 1]$, we can just take $\phi_0, \dots, \phi_{\kappa}$ to be the Legendre polynomials, shifted to $[0, 1]$ and normalized appropriately. Invoking Rodrigues' formula to express these functions,*

$$\phi_{\ell}(t) = \frac{\sqrt{2\ell+1}}{\ell!} \frac{d^{\ell}}{dt^{\ell}} (t^2 - t)^{\ell}, \quad \ell = 0, \dots, \kappa,$$

the results in Lemma 4 can be directly verified.

We use Lemma 4 to construct a sup norm bound on functions in $B_J(1)$ that are orthogonal (in $L_2(\Lambda)$) to all polynomials of degree k . We again Dejan Slepcev for his help with the next lemma.

Lemma 5. *Given an integer $k \geq 0$, and a continuous measure Λ on $[0, 1]$, whose Radon-Nikodym derivative λ is bounded below and above by constants $b_1, b_2 > 0$, respectively. Let J be a functional satisfying Assumptions C1 and C2, for a constant $L > 0$. There is a constant $R_0 > 0$, that depends only on k, b_1, b_2, L , such that*

$$\|g\|_\infty \leq R_0, \quad \text{for all } g \in B_J(1), \text{ such that } \langle g, p \rangle_2 = 0 \text{ for all polynomials } p \text{ of degree } k,$$

where $\langle \cdot, \cdot \rangle_2$ denotes the $L_2(\Lambda)$ inner product.

Proof. Fix an arbitrary function $g \in B_J(1)$, orthogonal (in $L_2(\Lambda)$) to all polynomials of degree k . Using integration by parts, and repeated application of Lemma 4, we have

$$0 = a_\ell \ell! \cdot \langle g, \phi_\ell \rangle_2 = \int_0^1 g^{(\ell)}(t) w_\ell(t) dt, \quad \ell = 0, \dots, k, \quad (51)$$

where $w_\ell(t) = (-1)^\ell a_\ell \ell! \Phi_{\ell, \ell}(t)$, $\ell = 0, \dots, k$, and by properties (46), (47) of Lemma 4,

$$\int_0^1 w_\ell(t) dt = 1, \quad \int_0^1 |w_\ell(t)| dt \leq \binom{2\ell}{\ell} \sqrt{\frac{b_2}{b_1}}, \quad \ell = 0, \dots, k. \quad (52)$$

Now, we will prove the following by induction:

$$\|g^{(\ell)}\|_\infty \leq L \left(\frac{b_2}{b_1} \right)^{(k-\ell+1)/2} \prod_{i=\ell}^k \binom{2i}{i}, \quad \ell = 0, \dots, k. \quad (53)$$

Starting at $\ell = k$, the statement holds because, using (51), for almost every $t \in [0, 1]$,

$$\begin{aligned} |g^{(k)}(t)| &= \left| g^{(k)}(t) - \int_0^1 g^{(k)}(u) w_k(u) du \right| \\ &= \left| \int_0^1 \left(g^{(k)}(t) - g^{(k)}(u) \right) w_k(u) du \right| \\ &\leq L \binom{2k}{k} \sqrt{\frac{b_2}{b_1}}, \end{aligned}$$

where in the second line we used the fact that the weight function integrates to 1 from (52), and in the third we used Assumption C2 and the upper bound on the integrated absolute weights from (52). Assume the statement holds at some $\ell < k$. Then again by (51), (52), for almost every $t \in [0, 1]$,

$$\begin{aligned} |g^{(\ell-1)}(t)| &= \left| \int_0^1 \left(g^{(\ell-1)}(t) - g^{(\ell-1)}(u) \right) w_{\ell-1}(u) du \right| \\ &\leq \left(\operatorname{ess\,sup}_{0 \leq u < v \leq 1} |g^{(\ell-1)}(v) - g^{(\ell-1)}(u)| \right) \binom{2\ell-2}{\ell-1} \sqrt{\frac{b_2}{b_1}} \\ &= \left(\operatorname{ess\,sup}_{0 \leq u < v \leq 1} \left| \int_u^v g^{(\ell)}(s) ds \right| \right) \binom{2\ell-2}{\ell-1} \sqrt{\frac{b_2}{b_1}} \\ &\leq L \left(\frac{b_2}{b_1} \right)^{(k-\ell+2)/2} \prod_{i=\ell-1}^k \binom{2i}{i}, \end{aligned}$$

the last line using $\operatorname{ess\,sup}_{0 \leq u < v \leq 1} \left| \int_u^v g^{(\ell)}(s) ds \right| \leq \|g^{(\ell)}\|_\infty$ and the inductive hypothesis. This verifies (53). Taking $\ell = 0$ in (53) and defining $R_0 = L(b_2/b_1)^{(k+1)/2} \prod_{i=0}^k \binom{2i}{i}$ proves the lemma. \square

We study the minimum eigenvalue of the (uncentered) empirical covariance matrix of a certain basis for additive k th degree polynomials in \mathbb{R}^{kd} . We thank Mathias Drton for his help with part (a) of the next lemma.

Lemma 6. *Let X^i , $i = 1, \dots, n$ denote an i.i.d. sample from a continuous distribution Q on $[0, 1]^d$. For an integer $k \geq 0$, let $V \in \mathbb{R}^{n \times kd}$ be a matrix whose i th row is given by*

$$V^i = \left(X_1^i, (X_1^i)^2, \dots, (X_1^i)^k, \dots, X_d^i, (X_d^i)^2, \dots, (X_d^i)^k \right) \in \mathbb{R}^{kd}, \quad (54)$$

for $i = 1, \dots, n$. Let

$$\nu_n^2 = \lambda_{\min} \left(\frac{1}{n} V^T V \right), \quad \text{and} \quad \nu_0^2 = \lambda_{\min} \left(\frac{1}{n} \mathbb{E}[V^T V] \right),$$

where $\lambda_{\min}(\cdot)$ denotes the minimum eigenvalue of its argument. Assuming that $n \geq kd$, the following properties hold:

(a) $\nu_n > 0$, almost surely with respect to Q ;

(b) $\nu_0 > 0$;

(c) for any $0 \leq t \leq 1$, $\mathbb{P}(\nu_n^2 > t\nu_0^2)$ with probability at least $1 - d \exp \left(- \frac{(1-t)^2 \nu_0 n}{2(kd)^2} \right)$.

Proof. For part (a), if the claim holds for $n = kd$, then it holds for all $n > kd$, so we may assume without a loss of generality that $n = kd$. Note that the determinant of $V \in \mathbb{R}^{n \times kd}$ is a polynomial function, call it $q(X)$, of the elements X_j^i , $i = 1, \dots, n$, $j = 1, \dots, d$. By Lemma 1 of Okamoto [1973], the roots of any polynomial—that is not identically zero—form a set of Lebesgue measure zero. To check that the polynomial q in question is not identically zero, it suffices to show that it is nonzero at a single realization of X . To this end, consider an input matrix defined by

$$X = \begin{bmatrix} \alpha_1 I \\ \vdots \\ \alpha_k I \end{bmatrix} \in \mathbb{R}^{n \times kd},$$

the rowwise concatenation of $\alpha_\ell I \in \mathbb{R}^{d \times d}$, $\ell = 1, \dots, k$. By the blockwise Vandermonde structure of the corresponding basis matrix V , we have that $q(X) \neq 0$ provided the coefficients α_ℓ , $\ell = 1, \dots, k$ are all distinct. Therefore q is not identically zero, and with respect to the continuous distribution Q , the determinant of V is nonzero, i.e., $\nu_n > 0$, almost surely.

For part (b), given any $a \in \mathbb{R}^{kd}$ with $a \neq 0$, we know that $Va \neq 0$ almost surely, since $\nu_n > 0$ almost surely, by part (a). Thus

$$a^T \mathbb{E}[V^T V] a = \mathbb{E} \|Va\|_2^2 > 0,$$

which proves that $\nu_0 > 0$.

Part (c) is an application of a matrix Chernoff bound from Tropp [2012]. In order to apply this result, we must obtain an almost sure upper bound R on $\lambda_{\max}(V^i(V^i)^T)$, with V^i as in (54) and $\lambda_{\max}(\cdot)$ denoting the maximum eigenvalue of its argument. This follows as

$$\lambda_{\max}(V^i(V^i)^T) \leq \sum_{j=1}^{kd} \sum_{\ell=1}^{kd} (V_j^i V_\ell^i)^2 \leq (kd)^2,$$

as each component of V^i has absolute magnitude at most 1 (recalling that Q is supported on $[0, 1]^d$). Taking $R = (kd)^2$ and applying Corollary 5.2 of Tropp [2012] (to be specific, applying the form of the Chernoff bound given in Remark 5.3 of this paper) gives the result. \square

The next lemma pertains to the additive function space

$$\mathcal{M}_n(\delta) = \left\{ \sum_{i=1}^d m_j : \sum_{j=1}^d J(m_j) \leq \delta, \text{ and } \langle m_j, 1 \rangle_n = 0, j = 1, \dots, d \right\}. \quad (55)$$

We give a sup norm bound on the components of functions in $\mathcal{M}_n(1) \cap B_n(\rho)$. The proof combines Lemmas 5 and 6, and uses a general strategy that follows the arguments given in Example 2.1(ii) of van de Geer [1990].

Lemma 7. *Let X^i , $i = 1, \dots, n$ denote an i.i.d. sample from a continuous distribution Q on $[0, 1]^d$, and let J be a seminorm satisfying Assumptions C1 and C2. There are constants $R_1, R_2, c_0, n_0 > 0$, depending only on d, k, L , such that for all $\rho > 0$ and $n \geq n_0$,*

$$\|m_j\|_\infty \leq R_1\rho + R_2, \text{ for all } j = 1, \dots, d \text{ and } \sum_{j=1}^d m_j \in \mathcal{M}_n(1) \cap B_n(\rho),$$

with probability at least $1 - \exp(-c_0n)$, where $\mathcal{M}_n(1)$ is the function space in (55).

Proof. Fix an arbitrary $m = \sum_{j=1}^d m_j \in \mathcal{M}_n(1) \cap B_n(\rho)$. For each $j = 1, \dots, d$, decompose

$$m_j = \langle m_j, 1 \rangle_n + p_j + g_j,$$

where p_j is a polynomial of degree k such that $\langle p_j, 1 \rangle_n = 0$, and g_j is orthogonal to all polynomials of degree k , with respect to the $L_2(U)$ inner product, with U the uniform distribution on $[0, 1]$; in fact, by definition of $\mathcal{M}_n(1)$, we know that $\langle m_j, 1 \rangle_n = 0$ so

$$m_j = p_j + g_j.$$

By the triangle inequality and Lemma 5 applied to the measure $\Lambda = U$ (whose density is of course lower and upper bounded with $b_1 = b_2 = 1$), we have, for each $j = 1, \dots, d$,

$$\left\| \sum_{j=1}^d g_j \right\|_\infty \leq \sum_{j=1}^d \|g_j\|_\infty \leq R_0 \sum_{j=1}^d J(g_j) \leq R_0, \quad (56)$$

where $R_0 > 0$ is the constant from Lemma 5, and we have used $J(m_j) = J(g_j)$, for $j = 1, \dots, d$, as the null space of J contains k th degree polynomials.

The triangle inequality and (56) now imply

$$\|p\|_n \leq \|m\|_n + \|g\|_n \leq \rho + R_0. \quad (57)$$

Write

$$p(x) = \sum_{j=1}^d \sum_{\ell=1}^k \alpha_{j\ell} x_j^\ell, \text{ for } x \in [0, 1]^d,$$

for some coefficients $\alpha_{j\ell}$, $j = 1, \dots, \ell = 1, \dots, k$. For $V \in \mathbb{R}^{n \times kd}$ the basis matrix as in Lemma 6, and $\alpha = (\alpha_{11}, \dots, \alpha_{1k}, \dots, \alpha_{d1}, \dots, \alpha_{dk}) \in \mathbb{R}^{kd}$, we have

$$\|p\|_n = \frac{1}{\sqrt{n}} \|V\alpha\|_2.$$

Furthermore, noting

$$\|p\|_n \geq \sqrt{\lambda_{\min}\left(\frac{1}{n} V^T V\right)} \|\alpha\|_2,$$

we have

$$\|\alpha\|_2 \leq \frac{\rho + R_0}{\nu_n},$$

where $\nu_n^2 = \lambda_{\min}(V^T V/n)$, as in Lemma 6, and we have used the upper bound in (57). Using part (c) of Lemma 6, with $t = 1/2$, we have

$$\|\alpha\|_2 \leq \frac{2(\rho + R_0)}{\nu_0},$$

with probability at least $1 - d \exp(-\nu_0 n / (8(kd)^2))$, where $\nu_0^2 = \lambda_{\min}(\mathbb{E}[V^T V]/n)$, as in Lemma 6. Therefore, using the triangle inequality and the fact that Q is supported on $[0, 1]^d$, we have for each $j = 1, \dots, d$, and any $x_j \in [0, 1]$,

$$|p_j(x_j)| \leq \sum_{\ell=1}^k |\alpha_{j\ell}| \leq \|\alpha\|_1 \leq \frac{2\sqrt{kd}(\rho + R_0)}{\nu_0},$$

with probability at least $1 - d \exp(-\nu_0 n / (8(kd)^2))$. Finally, for each $j = 1, \dots, d$, using the triangle inequality, and the sup norm bound from Lemma 5 once again,

$$\|m_j\|_\infty \leq \|p_j\|_\infty + \|g_j\|_\infty \leq \frac{2\sqrt{kd}(\rho + R_0)}{\nu_0} + R_0,$$

with probability $1 - d \exp(-\nu_0 n / (8(kd)^2))$, completing the proof. \square

We give a simple bound on the entropy of an arbitrary sum of sets in terms of the entropies of the original sets.

Lemma 8. *Given sets S_1, \dots, S_d and a norm $\|\cdot\|$, it holds that*

$$\log N(\delta, \|\cdot\|, S_1 + \dots + S_d) \leq \sum_{j=1}^d \log N(\delta/d, \|\cdot\|, S_j).$$

Proof. For $j = 1, \dots, d$, suppose that S_j can be covered in N_j balls of radius δ/d , with centers at $s_j^1, \dots, s_j^{N_j}$. Take an arbitrary $s \in S_1 + \dots + S_d$, and write $s = \sum_{j=1}^d s_j$, with $s_j \in S_j$, $j = 1, \dots, d$. For each $j = 1, \dots, d$, there is some $s_j^{\ell_j}$ such that $\|s_j - s_j^{\ell_j}\| \leq \delta/d$, and so by the triangle inequality

$$\left\| \sum_{j=1}^d s_j - \sum_{j=1}^d s_j^{\ell_j} \right\| \leq \delta.$$

That is, we have shown that $\prod_{j=1}^d N_j$ balls of radius δ with centers at

$$\sum_{j=1}^d s_j^{\ell_j}, \quad \text{for } (\ell_1, \dots, \ell_d) \in \{1, \dots, N_1\} \times \dots \times \{1, \dots, N_d\},$$

cover S . This completes the proof. \square

The next result represents our main tool from empirical process theory that will be used in the proof of Theorem 1. It is essentially an application of Lemma 3.5 in van de Geer [1990] (see also van de Geer [2000]).

Lemma 9. Let X^i , $i = 1, \dots, n$ denote an i.i.d. sample from a continuous distribution Q on $[0, 1]^d$. Let ϵ^i , $i = 1, \dots, n$ be uniformly sub-Gaussian random variates that have variance proxy $\sigma^2 > 0$ and are independent of X^i , $i = 1, \dots, n$. Let J be a seminorm satisfying Assumptions C1, C2, C3, and let $\rho > 0$ be arbitrary. Then there are constants $c_1, c_2, c_3, n_0 > 0$, depending only on $d, \sigma, k, L, K, w, \rho$, such that for all $c \geq c_1$ and $n \geq n_0$,

$$\sup_{m \in \mathcal{M}_n(1) \cap B_n(\rho)} \frac{\frac{1}{n} \sum_{i=1}^n \epsilon^i m(X^i)}{\|m\|_n^{1-w/2}} \leq \frac{c}{\sqrt{n}},$$

with probability at least $1 - \exp(-c_2 c^2) - \exp(-c_3 n)$.

Proof. Let Ω_1 denote the event on which the conclusion in Lemma 7 holds, which has probability at least $1 - \exp(-c_3 n)$ for $n \geq n_1$, for constants $c_3, n_1 > 0$. Also let $R_0 = R_1 \rho + R_2$, where $R_1, R_2 > 0$ are the constants from the lemma. Denote

$$B_\infty^d(\delta) = \left\{ \sum_{j=1}^d f_j : \|f_j\|_\infty \leq \delta, j = 1, \dots, d \right\}.$$

On Ω_1 , consider

$$\log N(\delta, \|\cdot\|_n, \mathcal{M}_n(1) \cap B_n(\rho)) \leq \log N(\delta, \|\cdot\|_n, \mathcal{M}_n(1) \cap B_\infty^d(R_0)) \quad (58)$$

$$\leq \sum_{j=1}^d \log N(\delta/d, \|\cdot\|_n, B_J(1) \cap B_\infty(R_0)) \quad (59)$$

$$\leq \sum_{j=1}^d \log N(\delta/(R_0 d), \|\cdot\|_n, B_J(1) \cap B_\infty(1)) \quad (60)$$

$$\leq K d^{1+w} (R_0)^w \delta^{-w}. \quad (61)$$

The first inequality (58) uses the sup norm bound from Lemma 7; the second inequality (59) uses

$$\mathcal{M}_n(1) \cap B_\infty^d(R_0) \subseteq \left\{ \sum_{j=1}^d m_j : m_j \in B_J(1) \cap B_\infty(R_0), j = 1, \dots, d \right\},$$

and applies Lemma 8 to the space on the right-hand side above. The third inequality (60) just uses the fact we may assume $R_0 \geq 1$, without a loss of generality; and the last inequality (61) uses Assumption C3. The entropy bound established in (61) allows us to apply Lemma 3.5 van de Geer [1990] (see also Lemma 8.4 in van de Geer [2000]), which gives constants $c_1, c_2, n_2 > 0$, depending only on d, σ, k, R_0, K, w , such that for all $c \geq c_1$ and $n \geq n_1$,

$$\sup_{m \in \mathcal{M}_n(1) \cap B_n(\rho)} \frac{\frac{1}{\sqrt{n}} \sum_{i=1}^n \epsilon^i m(X^i)}{\|m\|_n^{1-w/2}} \leq c$$

on an event Ω_2 with probability at least $1 - \exp(-c_2 c^2)$. The desired result in the lemma therefore holds for all $c \geq c_1$ and $n \geq n_0 = \max\{n_1, n_2\}$, on $\Omega_1 \cap \Omega_2$. \square

We finish with two simple results, on shifting around exponents in sums and products.

Lemma 10. For any $a, b \geq 0$, and any $0 < q < 1$,

$$(a + b)^q \leq a^q + b^q.$$

Proof. The function $f(t) = (1+t)^q - (1+t^q)$ has derivative $f'(t) = q(1+t)^{q-1} - qt^{q-1} < 0$ for all $t > 0$, and so $f(t) < f(0) = 0$ for all $t > 0$. Plugging in $t = a/b$ and rearranging gives the claim. \square

Lemma 11. For any $a, b \geq 0$, and any w ,

$$ab^{1-w/2} \leq a^{1/(1+w/2)}b + a^{2/(1+w/2)}.$$

Proof. Note that either $ab^{1-w/2} \leq a^{1/(1+w/2)}b$ or $ab^{1-w/2} \geq a^{1/(1+w/2)}b$, and in the latter case we get $b \leq a^{1/(1+w/2)}$, so $ab^{1-w/2} \leq a^{2/(1+w/2)}$. \square

A.6 Proof of Theorem 1

This proof roughly follows the ideas in the proof of Theorem 9 in [Mammen and van de Geer \[1997\]](#), though it differs in a few key ways. We use $c > 0$ to denote a constant that will multiply our final estimation error bound; it will also control the probability with which our final result holds. Some steps will only hold for sufficiently large n , but we do not always make this explicit. Lastly, we will occasionally abuse our notation for the empirical norms and empirical inner products by using them with vector arguments, to be interpreted in the appropriate sense (e.g., $\langle m, v \rangle_n = \frac{1}{n} \sum_{i=1}^n v^i m(X^i)$ for a function m and vector $v \in \mathbb{R}^n$).

We break down the presentation of our proof into mini sections for readability.

Deriving a basic inequality. Denote by $\hat{f} = \sum_{j=1}^d \hat{f}_j$ the total additive fit in (22). Let \mathcal{S} denote feasible set in (22). For any $f \in \mathcal{S}$, note that by orthogonality,

$$\|Y - \bar{Y}\mathbb{1} - f\|_n^2 = \|(f_0 + \epsilon - \bar{\epsilon}\mathbb{1}) - f\|_n^2 + (\bar{\epsilon})^2$$

where $\bar{\epsilon} = \frac{1}{n} \sum_{i=1}^n \epsilon^i$. Therefore \hat{f} must also be optimal for the problem

$$\min_{f \in \mathcal{S}} \frac{1}{2} \|W - f\|_n^2 + \lambda_n J_d(f),$$

where $W^i = f_0(X^i) + \epsilon^i - \bar{\epsilon}$, $i = 1, \dots, n$, and we denote $\lambda_n = \lambda/n$ and $J_d(f) = \sum_{j=1}^d J(f_j)$. Standard arguments (from first-order optimality) show that any solution \hat{f} in the above satisfies

$$\langle W - \hat{f}, \tilde{f} - \hat{f} \rangle_n \leq \lambda_n (J_d(\tilde{f}) - \lambda_n J_d(\hat{f})),$$

for any feasible $\tilde{f} = \sum_{j=1}^d \tilde{f}_j \in \mathcal{S}$. Expanding the definition of W and rearranging gives

$$\langle \hat{f} - f_0, \hat{f} - \tilde{f} \rangle_n \leq \langle \epsilon - \bar{\epsilon}\mathbb{1}, \hat{f} - \tilde{f} \rangle_n + \lambda_n (J_d(\tilde{f}) - \lambda_n J_d(\hat{f})).$$

Using the polarization identity $\langle a, b \rangle = \frac{1}{2} (\|a\|^2 + \|b\|^2 - \|a - b\|^2)$ for an inner product $\langle \cdot, \cdot \rangle$ and its corresponding norm $\|\cdot\|$,

$$\|\hat{f} - f_0\|_n^2 + \|\hat{f} - \tilde{f}\|_n^2 \leq 2\langle \epsilon - \bar{\epsilon}\mathbb{1}, \hat{f} - \tilde{f} \rangle_n + 2\lambda_n (J_d(\tilde{f}) - \lambda_n J_d(\hat{f})) + \|\tilde{f} - f_0\|_n^2.$$

Abbreviating $\hat{\Delta} = \hat{f} - \tilde{f}$, $\hat{J} = J_d(\hat{f})$, and $\tilde{J} = J_d(\tilde{f})$, and using $\langle \bar{\epsilon}\mathbb{1}, \hat{\Delta} \rangle = 0$, this becomes

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq 2\langle \epsilon, \hat{\Delta} \rangle_n + 2\lambda_n (\tilde{J} - \hat{J}) + \|\tilde{f} - f_0\|_n^2, \quad (62)$$

which is our basic inequality. In what follows, we will restrict our attention to feasible \tilde{f} such that $\|\tilde{f} - f_0\|_n \leq \max\{C_n, \tilde{J}\}$.

Localizing the error vector. We prove that $\hat{\Delta}$ is appropriately bounded in the empirical norm. By the tail bound for quadratic forms of sub-Gaussian random variates in Theorem 2.1 of [Hsu et al. \[2012\]](#), for all $t > 0$,

$$\mathbb{P}\left(\|\epsilon\|_n^2 > \sigma^2 \left(1 + \frac{2\sqrt{t}}{\sqrt{n}} + \frac{2t}{n}\right)\right) \leq e^{-t},$$

and hence taking $t = \sqrt{n}$,

$$\|\epsilon\|_n^2 \leq 5\sigma^2,$$

on an event Ω_1 with probability at least $1 - \exp(-\sqrt{n})$. Thus returning to the basic inequality (62), using the Cauchy-Schwartz inequality, and the above bound, we have

$$\|\hat{\Delta}\|_n^2 \leq 2\sqrt{5}\sigma\|\hat{\Delta}\|_n + 2\lambda_n\tilde{J} + \|\tilde{f} - f_0\|_n^2,$$

on Ω_1 . This is a quadratic inequality of the form $x^2 \leq bx + c$ in $x = \|\hat{\Delta}\|_n$, so we can upper bound x by the larger of the two roots, $x \leq (b + \sqrt{b^2 + 4c})/2 \leq b + \sqrt{c}$, i.e.,

$$\|\hat{\Delta}\|_n \leq 2\sqrt{5}\sigma + \sqrt{2\lambda_n\tilde{J} + \|\tilde{f} - f_0\|_n^2},$$

on Ω_1 . Abbreviating $J^* = \max\{C_n, \tilde{J}\}$, and using $J^* \geq 1$ (as $C_n \geq 1$ by assumption),

$$\|\hat{\Delta}\|_n \leq J^* \left(2\sqrt{5}\sigma + \sqrt{2\lambda_n + \|\tilde{f} - f_0\|_n^2 / (J^*)^2} \right),$$

on Ω_1 . Recalling $\|\tilde{f} - f_0\|_n \leq J^*$, and using the fact that $\lambda_n = o(1)$ for our eventual choice of λ_n , we have that for sufficiently large n ,

$$\|\hat{\Delta}\|_n \leq J^*(2\sqrt{5}\sigma + \sqrt{2}), \quad (63)$$

on Ω_1 .

Bounding the sub-Gaussian complexity term. We focus on the first term on the right-hand side in (62), i.e., the sub-Gaussian complexity term. Let $m = \hat{\Delta}/(\hat{J} + J^*)$. By construction, we have $J(m) \leq 1$, and from (63), we have $\|m\|_n \leq 2\sqrt{5}\sigma + \sqrt{2}$ on Ω_1 . Then, applying Lemma 9, with the choice $\rho = 2\sqrt{5}\sigma + \sqrt{2}$, we see that there are constants $c_1, c_2, c_3 > 0$ such that for all $c \geq c_1$,

$$\frac{2\langle \epsilon, m \rangle_n}{\|m\|_n^{1-w/2}} \leq \frac{c}{\sqrt{n}},$$

on $\Omega_1 \cap \Omega_2$, where Ω_2 is an event with probability at least $1 - \exp(-c_2c^2) - \exp(-c_3n)$. Plugging this into (62) gives

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq \frac{c}{\sqrt{n}}(\hat{J} + J^*)^{w/2}\|\hat{\Delta}\|_n^{1-w/2} + 2\lambda_n(\tilde{J} - \hat{J}) + \|\tilde{f} - f_0\|_n^2,$$

on $\Omega_1 \cap \Omega_2$. By the inequality in Lemma 11, applied to the first term on the right-hand side above, with $a = n^{-1/2}(\hat{J} + J^*)^{w/2}$ and $b = \|\hat{\Delta}\|_n$,

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq cr_n(\hat{J} + J^*)^{w/(2+w)}\|\hat{\Delta}\|_n + cr_n^2(\hat{J} + J^*)^{2w/(2+w)} + 2\lambda_n(\tilde{J} - \hat{J}) + \|\tilde{f} - f_0\|_n^2,$$

on $\Omega_1 \cap \Omega_2$, where we abbreviate $r_n = n^{-1/(2+w)}$. Applying the simple inequality $2ab \leq a^2 + b^2$ to the first term on the right-hand side, with $a = cr_n(\hat{J} + J^*)^{w/(2+w)}$ and $b = \|\hat{\Delta}\|_n$, and subtracting $\|\hat{\Delta}\|_n^2/2$ from both sides,

$$\|\hat{f} - f_0\|_n^2 + \frac{1}{2}\|\hat{\Delta}\|_n^2 \leq \frac{3}{2}c^2r_n^2(\hat{J} + J^*)^{2w/(2+w)} + 2\lambda_n(\tilde{J} - \hat{J}) + \|\tilde{f} - f_0\|_n^2, \quad (64)$$

on $\Omega_1 \cap \Omega_2$ (where we have assumed without a loss of generality that $c \geq 1$).

Controlling the effect of the penalty terms. Now we handle the appearances of the achieved penalty term \hat{J} . First, set $\lambda_n \geq (3/4)c^2r_n^2/C_n^{(2-w)/(2+w)}$, and denote

$$a = \frac{3}{2}c^2r_n^2(\hat{J} + J^*)^{2w/(2+w)} + 2\lambda_n(\tilde{J} - \hat{J}).$$

Consider the case $\hat{J} \geq C_n$. Then $-1/C_n^{(2-w)/(2+w)} \geq -1/\hat{J}^{(2-w)/(2+w)}$, and

$$2\lambda_n(\tilde{J} - \hat{J}) \leq 2\lambda_n\tilde{J} - (3/2)c^2r_n^2\hat{J}^{2w/(2+w)},$$

thus, using the simple inequality in Lemma 10, we have $a \leq 4\lambda_nJ^*$. In the case $\hat{J} < C_n$, we have by Lemma 10 again,

$$a \leq \frac{3}{2}c^2r_n^2\left(C_n^{2w/(2+w)} + (J^*)^{2w/(2+w)}\right) + 2\lambda_n\tilde{J} \leq 6\lambda_nJ^*.$$

Therefore, altogether, we conclude that $a \leq 6\lambda_nJ^*$, and plugging this into (64) gives

$$\|\hat{f} - f_0\|_n^2 + \frac{1}{2}\|\hat{\Delta}\|_n^2 \leq 6\lambda_nJ^* + \|\tilde{f} - f_0\|_n^2,$$

on $\Omega_1 \cap \Omega_2$. The statement (23) as made in the theorem follows by dropping the nonnegative term $\|\hat{\Delta}\|_n^2/2$ from the left-hand side, and adjusting the constants $c, c_1, c_2, c_3 > 0$ as needed.

A.7 Proof of the best additive approximation bound in (26)

We follow the exact same arguments as in the proof of Theorem 1, up until the last part, in which we control the achieved penalty terms \hat{J} . Now we deviate from the previous arguments, slightly. Set $\lambda_n \geq (3/2)c^2r_n^2/C_n^{(2-w)/(2+w)}$, and denote

$$a = \frac{3}{2}c^2r_n^2(\hat{J} + J^*)^{2w/(2+w)} + \lambda_n(\tilde{J} - \hat{J}).$$

By the same logic as in the proof of Theorem 1, we have $a \leq 3\lambda_nJ^*$. Plugging this into (64) gives

$$\|\hat{f} - f_0\|_n^2 + \frac{1}{2}\|\hat{\Delta}\|_n^2 \leq 3\lambda_nJ^* + \lambda_n(\tilde{J} - \hat{J}) + \|\tilde{f} - f_0\|_n^2,$$

on $\Omega_1 \cap \Omega_2$. Rearranging,

$$\frac{1}{2}\|\hat{\Delta}\|_n^2 \leq 3\lambda_nJ^* + \left(\|\tilde{f} - f_0\|_n^2 + \lambda_n\tilde{J} - \|\hat{f} - f_0\|_n^2 - \lambda_n\hat{J}\right),$$

on $\Omega_1 \cap \Omega_2$. But, setting $\tilde{f} = f^{\text{best}}$, the bracketed term on the right-hand side above is nonpositive (by definition of f^{best} in (25)). This leads to (26), after adjusting $c, c_1, c_2, c_3 > 0$ as needed.

A.8 Preliminaries for the proof of Corollary 1

The following two lemmas will be helpful for the proof of Corollary 1.

Lemma 12. *Given $f = \sum_{j=1}^d f_j$, whose component functions are each k times weakly differentiable, there exists an additive spline approximant $\check{f} = \sum_{j=1}^d \check{f}_j$, where $\check{f}_j \in \mathcal{G}_j$, the set of k th order splines with knots in the set T_j defined in (21), for $j = 1, \dots, d$, such that*

$$(i) \text{TV}(\check{f}_j^{(k)}) \leq a_k \text{TV}(f_j^{(k)}), \text{ for } j = 1, \dots, d; \text{ and}$$

$$(ii) \|\check{f}_j - f_j\|_\infty \leq a_k W_{\max}^k \text{TV}(f_j^{(k)}), \text{ for } j = 1, \dots, d.$$

Above, $a_k \geq 1$ is a constant depending only on k , and we define $W_{\max} = \max_{j=1, \dots, d} W_j$, where

$$W_j = \max_{i=1, \dots, n-1} |X_j^{(i)} - X_j^{(i+1)}|, \quad j = 1, \dots, d.$$

When the input points are drawn from a distribution Q that satisfies Assumptions A1, A2, there are universal constants $c_0, n_0 > 0$ such that for $n \geq n_0$, we have $W_{\max} \leq (c_0/b_0) \log n/n$ with probability at least $1 - 2b_0d/n$, and so the bounds in (ii) become

$$\|\check{f}_j - f_j\|_\infty \leq \frac{c_0^k a_k}{b_0^k} \left(\frac{\log n}{n} \right)^k \text{TV}(f_j^{(k)}), \quad \text{for } j = 1, \dots, d, \quad (65)$$

with probability at least $1 - 2b_0d/n$.

Proof. Parts (i) and (ii) are simply a componentwise application of Proposition 7 in Mammen and van de Geer [1997]. In particular, from their result, we know that for $j = 1, \dots, d$, there is a k th degree spline function \check{f}_j whose knots lie in T_j in (21), with $\text{TV}(\check{f}_j^{(k)}) \leq a_k \text{TV}(f_j^{(k)})$ and

$$\|\check{f}_j - f_j\|_\infty \leq a_k W_j^k \text{TV}(f_j^{(k)}),$$

where $a_k \geq 1$ depends only on k . (This result follows from strong quasi-interpolating properties of spline functions, from de Boor [1978].) This proves parts (i) and (ii).

When we consider random inputs drawn from a distribution Q satisfying Assumptions A1, A2, the densities of the marginals Q_j , $j = 1, \dots, d$ will be bounded below by $b_0 > 0$, and thus there are universal constants $c_0, n_0 > 0$ such that for $n \geq n_0$, we have $W_j \leq (c_0/b_0) \log n/n$ with probability at least $1 - 2b_0/n$ (see, e.g., Lemma 5 in Wang et al. [2014]), for $j = 1, \dots, d$, and hence applying a union bound gives the result for W_{\max} . \square

Lemma 13. *Given $f = \sum_{j=1}^d f_j$, whose component functions are each k times weakly differentiable, there is an additive falling factorial approximant $\check{f} = \sum_{j=1}^d \check{f}_j$, where $\check{f}_j \in \mathcal{H}_j$, the set of k th order falling factorial functions defined over X_j^1, \dots, X_j^n , for each $j = 1, \dots, d$, such that*

$$(i) \text{TV}(\check{f}_j^{(k)}) \leq a_k \text{TV}(f_j^{(k)}), \text{ for } j = 1, \dots, d; \text{ and}$$

$$(ii) \|\check{f}_j - f_j\|_\infty \leq a_k (W_{\max}^k + 2k^2 W_{\max}) \text{TV}(f_j^{(k)}), \text{ for } j = 1, \dots, d.$$

Again, $a_k \geq 1$ is a constant depending only on k , and W_{\max} is as defined in Lemma 12. When the inputs are drawn from a distribution Q satisfying Assumptions A1, A2, the bound in (ii) become

$$\|\check{f}_j - f_j\|_\infty \leq a_k \left(\frac{c_0^k}{b_0^k} \left(\frac{\log n}{n} \right)^k + 2k^2 \frac{c_0}{b_0} \frac{\log n}{n} \right) \text{TV}(f_j^{(k)}), \quad \text{for } j = 1, \dots, d, \quad (66)$$

with probability at least $1 - 2b_0d/n$.

Proof. First we apply Lemma 12 to produce an additive spline approximant, call it $f^* = \sum_{j=1}^d f_j^*$, to the given $f = \sum_{j=1}^d f_j$. Next, we parametrize the spline component functions in a helpful way:

$$f_j^* = \sum_{\ell=1}^n \alpha_j^\ell g_{j\ell}, \quad j = 1, \dots, d.$$

where $\alpha_j^1, \dots, \alpha_j^n \in \mathbb{R}$ are coefficients and g_{j1}, \dots, g_{jn} are the truncated power basis functions over the knot set T_j defined in (21), and we write $g_{j\ell}(t) = t^{\ell-1}$, $\ell = 1, \dots, k$ without a loss of generality, for $j = 1, \dots, d$. It is not hard to check that $\text{TV}((f_j^*)^{(k)}) = \sum_{\ell=k+2}^n |\alpha_j^\ell|$, for $j = 1, \dots, d$.

We now define $\check{f} = \sum_{j=1}^d \check{f}_j$, our falling factorial approximant, to have component functions

$$\check{f}_j = \sum_{\ell=1}^{k+1} \alpha_j^\ell g_{j\ell} + \sum_{\ell=k+2}^n \alpha_j^\ell h_{j\ell}, \quad j = 1, \dots, d.$$

where h_{j1}, \dots, h_{jn} are the falling factorial basis functions defined over X_j^1, \dots, X_j^n , for $j = 1, \dots, d$. (Note that \check{f}_j preserves the polynomial part of f_j^* exactly, for $j = 1, \dots, d$.) Again, it is straightforward to check that $\text{TV}(\check{f}_j^{(k)}) = \sum_{\ell=k+2}^n |\alpha_j^\ell|$, for $j = 1, \dots, d$, i.e.,

$$\text{TV}(\check{f}_j^{(k)}) = \text{TV}((f_j^*)^{(k)}) \leq a_k \text{TV}(f_j^{(k)}), \quad \text{for } j = 1, \dots, d,$$

the inequality coming from part (i) of Lemma 12. This verifies part (i) of the current lemma. As for part (ii), we note that Lemma 4 of Wang et al. [2014] shows that

$$|h_{j\ell}(X_j^i) - g_{j\ell}(X_j^i)| \leq k^2 W_j, \quad \text{for } \ell = k+2, \dots, n, i = 1, \dots, n, j = 1, \dots, d,$$

where recall W_j is the maximum gap between sorted input points along the j th dimension, $j = 1, \dots, d$, as defined in Lemma 12. In fact, a straightforward modification of their proof can be used to strengthen this result to

$$\|h_{j\ell} - g_{j\ell}\|_\infty \leq 2k^2 W_j, \quad \text{for } \ell = k+2, \dots, n, j = 1, \dots, d,$$

which means that by Holder's inequality,

$$\|\check{f}_j - f_j^*\|_\infty \leq 2k^2 W_j \sum_{\ell=k+2}^n |\alpha_j^\ell| \leq 2k^2 a_k W_j \text{TV}(f_j^{(k)}) \quad \text{for } j = 1, \dots, d.$$

Then, by the triangle inequality,

$$\|\check{f}_j - f_j\|_\infty \leq \|\check{f}_j - f_j^*\|_\infty + \|f_j^* - f_j\|_\infty \leq a_k \left(W_{\max}^k + 2k^2 W_{\max} \right) \text{TV}(f_j^{(k)}), \quad \text{for } j = 1, \dots, d,$$

where we have used part (ii) of Lemma 12. This verifies part (ii) of the current lemma.

Lastly, for random inputs drawn from a distribution Q satisfying Assumptions A1, A2, the proof of (66) follows the same arguments as the proof of (65). \square

A.9 Proof of Corollary 1

We consider first the statement in part (a). We must check that Assumptions C1, C2, C3 hold for our choice of regularizer $J(g) = \text{TV}(g^{(k)})$, and then we can apply Theorem 1. Assumptions C1, C2 are immediate. As for Assumption C3, consider the univariate function class

$$\mathcal{W}_{k+1} = \left\{ f : \int_0^1 |f^{(k+1)}(t)| dt \leq 1, \|f\|_\infty \leq 1 \right\}.$$

The results in Birman and Solomyak [1967] imply that for any set $Z_n = \{z^1, \dots, z^n\} \subseteq [0, 1]$,

$$\log N(\delta, \|\cdot\|_{Z_n}, \mathcal{W}_{k+1}) \leq K \delta^{-1/(k+1)},$$

for a universal constant $K > 0$. As explained in Mammen [1991], Mammen and van de Geer [1997], this confirms that Assumption C3 holds for our choice of regularizer, with $w = 1/(k+1)$. Applying Theorem 1, with $\check{f} = f_0$, gives the result in (27).

For the statement in part (b), note first that we can consider $k \geq 2$ without a loss of generality, as pointed out in Remark 6 following the corollary. Using Lemma 12, can choose an additive spline approximant \check{f} to f_0 , with components $\check{f}_j \in \mathcal{G}_j$, $j = 1, \dots, d$. Define \check{f}_j to be the centered version of f_j , with zero empirical mean, $j = 1, \dots, d$. By the fact that centering does not change the penalty, and part (i) of the lemma, $\sum_{j=1}^d \text{TV}(\check{f}_j^{(k)}) \leq a_k \sum_{j=1}^d \text{TV}(f_{0j}^{(k)})$. Also, using the fact that centering cannot increase the empirical norm, the triangle inequality, and (65), we get that with probability

least $1 - 2b_0d/n$,

$$\begin{aligned} \left\| \sum_{j=1}^d \tilde{f}_j - \sum_{j=1}^d f_{0j} \right\|_n &\leq \left\| \sum_{j=1}^d \tilde{f}_j - \sum_{j=1}^d f_{0j} \right\|_n \\ &\leq \sum_{j=1}^d \|\tilde{f}_j - f_{0j}\|_\infty \\ &\leq \frac{c_0^k a_k}{b_0^k} \left(\frac{\log n}{n} \right)^k \sum_{j=1}^d \text{TV}(f_{0j}^{(k)}), \end{aligned}$$

When $\sum_{j=1}^d \text{TV}(f_{0j}^{(k)}) \leq C_n$, we see that $\|\tilde{f} - f_0\|_n$ is bounded by C_n for large enough n . This meets required condition for Theorem 1, by the above display, the approximation error in (23) satisfies

$$\left\| \sum_{j=1}^d \tilde{f}_j - \sum_{j=1}^d f_{0j} \right\|_n^2 \leq \left(\frac{c_0^k a_k}{b_0^k} \right)^2 \left(\frac{\log n}{n} \right)^{2k} C_n^2.$$

But when $n/(\log n)^{1+1/k} \geq n_0 C_n^{(2k+2)/(2k^2+2k-1)}$, the right-hand side above is upper bounded by $a_0 n^{-(2k+2)/(2k+3)} C_n^{2/(2k+3)}$, for a constant $a_0 > 0$. This establishes the result in (27) for restricted additive locally adaptive splines.

For the statement in part (c), we can again consider $k \geq 2$ without a loss of generality. Then the same arguments as given for part (b) apply here, but now we use Lemma 13 for the additive falling factorial approximant \tilde{f} to f_0 , and we require $n/(\log n)^{2k+3} \geq n_0 C_n^{4k+4}$ for the approximation error to be bounded by the estimation error.

A.10 Preliminaries for the proof of Theorem 2

Our first lemma is similar to Lemma 9, but concerns univariate functions. As in Lemma 9, this result relies on Lemma 3.5 in van de Geer [1990] (see also van de Geer [2000]).

Lemma 14. *Let ϵ^i , $i = 1, \dots, n$ be uniformly sub-Gaussian random variables having variance proxy $\sigma^2 > 0$. Let J be a seminorm satisfying Assumption C3, and let $\rho > 0$ be arbitrary. Then there exist constants $c_1, c_2, n_0 > 0$, depending only on σ, K, w, ρ , such that for all $c \geq c_1$ and $n \geq n_0$,*

$$\sup_{Z_n = \{z^1, \dots, z^n\} \subseteq [0,1]} \sup_{g \in B_J(1) \cap B_\infty(\rho)} \frac{\frac{1}{n} \sum_{i=1}^n \epsilon^i g(z^i)}{\|g\|_{Z_n}^{1-w/2}} \leq \frac{c}{\sqrt{n}},$$

with probability at least $1 - \exp(-c_2 c^2)$, where we write $\|\cdot\|_{Z_n}$ for the empirical norm defined over a set of univariate points $Z_n = \{z^1, \dots, z^n\} \subseteq [0, 1]$.

Proof. Assume without a loss of generality that $\rho \geq 1$. Note that for any $Z_n = \{z^1, \dots, z^n\} \subseteq [0, 1]$,

$$\log N(\delta, \|\cdot\|_{Z_n}, B_J(1) \cap B_\infty(\rho)) \leq K \rho^w \delta^{-w},$$

by Assumption C3. As the right-hand side in the above entropy bound does not depend on Z_n , we can apply Lemma 3.5 in van de Geer [1990] to get the desired uniform control over all subsets. \square

We give a coupling between the empirical and L_2 norms over $B_J(1) \cap B_\infty(\rho)$, using Theorem 14.1 in Wainwright [2017] (see also van de Geer [2000], Bartlett et al. [2005], Raskutti et al. [2012]).

Lemma 15. *Let z^i , $i = 1, \dots, n$ denote an i.i.d. sample from a distribution Λ on $[0, 1]$. Write $\|\cdot\|_2$ for the $L_2(\Lambda)$ norm, and $\|\cdot\|_n$ for the $L_2(\Lambda_n)$ norm. Let J satisfy Assumption C3, and let $\rho > 0$ be*

arbitrary. Then there are constants $c_1, c_2, c_3, n_0 > 0$, that depend only on K, w, ρ , such that for any $t \geq c_1 n^{-1/(2+w)}$ and $n \geq n_0$,

$$\left| \|g\|_n^2 - \|g\|_2^2 \right| \leq \frac{1}{2} \|g\|_2^2 + \frac{t^2}{2}, \quad \text{for all } g \in B_J(1) \cap B_\infty(\rho),$$

with probability at least $1 - c_2 \exp(-c_3 n t^2)$.

Proof. Abbreviate $\mathcal{F} = B_J(1) \cap B_\infty(\rho)$. We will analyze the local Rademacher complexity

$$\mathcal{R}(\mathcal{F} \cap B_2(t)) = \mathbb{E}_{z, \sigma} \left[\sup_{g \in \mathcal{F} \cap B_2(t)} \frac{1}{n} \left| \sum_{i=1}^n \sigma^i g(z^i) \right| \right],$$

the expectation being taken over i.i.d. draws $z^i, i = 1, \dots, n$ from Λ and i.i.d. Rademacher variables $\sigma^i, i = 1, \dots, n$, as usual. Define the critical radius $\tau_n > 0$ to be the smallest solution of the equation

$$\frac{\mathcal{R}(\mathcal{F} \cap B_2(t))}{t} = \frac{t}{\rho}.$$

We will prove $\tau_n \leq c_1 n^{-1/(2+w)}$ for a constant $c_1 > 0$. Applying Theorem 14.1 in [Wainwright \[2017\]](#) would then give the result.

In what follows, we will use $c > 0$ to denote a constant whose value may change from line to line (but does not depend on $z^i, i = 1, \dots, n$). Consider the empirical local Rademacher complexity

$$\mathcal{R}_n(\mathcal{F} \cap B_2(t)) = \mathbb{E}_\sigma \left[\sup_{g \in \mathcal{F} \cap B_2(t)} \frac{1}{n} \left| \sum_{i=1}^n \sigma^i g(z^i) \right| \right].$$

As we are considering $t \geq \tau_n$, Corollary 2.2 of [Bartlett et al. \[2005\]](#) gives

$$\mathcal{F} \cap B_2(t) \subseteq \mathcal{F} \cap B_n(\sqrt{2}t),$$

with probability at least $1 - 1/n$. Denote by \mathcal{E} the event that this occurs. Then on \mathcal{E} ,

$$\begin{aligned} \mathcal{R}_n(\mathcal{F} \cap B_2(t)) &\leq \mathbb{E}_\sigma \left[\sup_{g \in \mathcal{F} \cap B_n(\sqrt{2}t)} \frac{1}{n} \left| \sum_{i=1}^n \sigma^i g(z^i) \right| \right] \\ &\leq \frac{c}{\sqrt{n}} \int_0^{\sqrt{2}t} \sqrt{\log N(\delta, \|\cdot\|_n, \mathcal{F})} d\delta \\ &\leq \frac{c\sqrt{K}\rho^{w/2}}{\sqrt{n}} \int_0^{\sqrt{2}t} \delta^{-w/2} d\delta = \frac{c}{\sqrt{n}} t^{1-w/2}, \end{aligned}$$

where in second line we used Dudley's entropy integral [[Dudley, 1967](#)], and in the third line we used Assumption [C3](#). On \mathcal{E}^c , note that we have the trivial bound $\mathcal{R}_n(\mathcal{F} \cap B_2(t)) \leq \rho$. Therefore we can upper bound the local Rademacher complexity, splitting the expectation over \mathcal{E} and \mathcal{E}^c ,

$$\mathcal{R}(\mathcal{F} \cap B_2(t)) = \mathbb{E}_z \mathcal{R}_n(\mathcal{F} \cap B_2(t)) \leq \frac{c t^{1-w/2}}{\sqrt{n}} + \frac{\rho}{n} \leq \frac{c t^{1-w/2}}{\sqrt{n}},$$

where the second inequality holds when n is large enough, as we may assume $t \geq n^{-1/2}$ without a loss of generality. An upper bound on the critical radius τ_n is thus given by the solution of

$$\frac{c t^{-w/2}}{\sqrt{n}} = \frac{t}{\rho},$$

which is $t = c n^{-1/(2+w)}$. This completes the proof. \square

We extend Lemma 5 to give a uniform sup norm bound on the functions in $B_J(1) \cap B_2(\rho)$.

Lemma 16. *Assume the conditions of Lemma 5. Then there are constants $R_1, R_2 > 0$ that depend only on k, b_1, b_2, L , such that*

$$\|m\|_\infty \leq R_1\rho + R_2, \quad \text{for all } m \in B_J(1) \cap B_2(\rho).$$

Proof. For $m \in B_J(1) \cap B_2(\rho)$, decompose $m = p + g$ where p is a polynomial of degree k , and g is orthogonal to all polynomials of degree k with respect to the $L_2(\Lambda)$ inner product. By Lemma 5, we have $\|g\|_\infty \leq R_0$ for a constant $R_0 > 0$, and by the triangle inequality,

$$\|p\|_2 \leq \|m\|_2 + \|g\|_2 \leq \rho + R_0.$$

Now write

$$p(x) = \sum_{\ell=1}^{k+1} \alpha_\ell \phi_\ell(x), \quad \text{for } x \in [0, 1]^d,$$

where $\phi_\ell, \ell = 1, \dots, k+1$ are orthonormal polynomials on $[0, 1]$ with respect to the $L_2(\Lambda)$ inner product. Then $\|\alpha\|_2 = \|p\|_2 \leq \rho + R_0$, from the second to last display, and $\|\alpha\|_2 \leq \sqrt{k+1}(\rho + R_0)$, so for any $x \in [0, 1]$,

$$|p(x)| \leq \|\alpha\|_1 \max_{\ell=1, \dots, k+1} |\phi_\ell(x)| \leq c_k \sqrt{k+1}(\rho + R_0),$$

where $c_k = \max_{\ell=1, \dots, k+1} \|\phi_\ell\|_\infty$ is a constant that depends only on k, b_1 from Aptekarev et al. [2016]. Therefore

$$\|m\|_\infty \leq \|p\|_\infty + \|g\|_\infty \leq c_k \sqrt{k+1}(\rho + R_0) + R_0,$$

and defining $R_1, R_2 > 0$ appropriately, this is of the desired form, and completes the proof. \square

Our last two lemmas pertain to the function space

$$\mathcal{M}_2(\delta) = \left\{ \sum_{i=1}^d m_j : J(m_j) \leq \delta, \text{ and } \langle m_j, 1 \rangle_2 = 0, j = 1, \dots, d \right\}. \quad (67)$$

We derive a one-sided bound on the L_2 norm in terms of the empirical norm, over $\mathcal{M}_2(1)$. Our proof uses Theorem 14.2 in Wainwright [2017], which is a somewhat unique theorem, because it does not require a global sup norm bound on the function class in consideration (unlike many standard results of this type).

Lemma 17. *Let $X^i, i = 1, \dots, n$ denote an i.i.d. sample from a distribution Q on $[0, 1]^d$ satisfying Assumption A3, and let J satisfy Assumption C3. Then there are constants $c_1, c_2, c_3, n_0 > 0$, that depend only on b_1, b_2, k, L, K, w , such that for any $c_1 \sqrt{dn}^{-1/(2+w)} \leq t \leq 1$ and $n \geq n_0$,*

$$\|m\|_2^2 \leq 2\|m\|_n^2 + t^2, \quad \text{for all } m \in \mathcal{M}_2(1),$$

with probability at least $1 - c_2 \exp(-c_3 n t^2)$, where $\mathcal{M}_2(1)$ is the space in (67).

Proof. Let $m \in \mathcal{M}_2(1)$ with $\|m\|_2 \leq 1$. Then as $\|m\|_2^2 = \sum_{j=1}^d \|m_j\|_2^2$, it follows that $\|m_j\|_2 \leq 1, j = 1, \dots, d$, and by Lemma 16, we have $\|m_j\|_\infty \leq R_1 + R_2, j = 1, \dots, d$. From the calculation in Example 14.6 of Wainwright [2017], we have the property

$$\|m^2\|_2^2 \leq C^2 \|m\|_2^4, \quad \text{for all } m \in \mathcal{M}_2(1) \cap B_2(1),$$

where $C^2 = (R_1 + R_2)^2 + 6$. Abbreviating $\mathcal{F} = \mathcal{M}_2(1)$, we will study the local Rademacher complexity

$$\mathcal{R}(\mathcal{F} \cap B_2(t)) = \mathbb{E}_{z, \sigma} \left[\sup_{m \in \mathcal{F} \cap B_2(t)} \frac{1}{n} \left| \sum_{i=1}^n \sigma^i m(z^i) \right| \right],$$

and the associated critical radius $\tau_n > 0$, defined as usual to be the smallest solution of

$$\frac{\mathcal{R}(\mathcal{F} \cap B_2(t))}{t} = \frac{t}{C}.$$

We will establish $\tau_n \leq c_1 \sqrt{dn}^{-1/(2+w)}$ for a constant $c_1 > 0$. Applying Theorem 14.2 in [Wainwright \[2017\]](#) would then give the result.

For the rest of the proof, we will use $c > 0$ for a constant whose value may change from line to line; also, many statements will hold for large enough n , but this will not always be made explicit. Fix some $0 < t \leq 1$. By L_2 orthogonality of the components of functions in \mathcal{F} ,

$$\begin{aligned} \sup_{m \in \mathcal{F} \cap B_2(t)} \left| \frac{1}{n} \sum_{i=1}^n \sigma^i m(X^i) \right| &\leq \sup_{\|\beta\|_2 \leq t} \sup_{\substack{m_j \in B_J(1) \cap B_2(|\beta_j|), \\ j=1, \dots, d}} \left| \sum_{i=1}^n \sigma^i \sum_{j=1}^d m_j(X_j^i) \right| \\ &\leq \sup_{\|\beta\|_2 \leq t} \sum_{j=1}^d \sup_{m_j \in B_J(1) \cap B_2(|\beta_j|)} \left| \frac{1}{n} \sum_{i=1}^n \sigma^i m_j(X_j^i) \right|. \end{aligned}$$

We now bound the inner supremum above, for an arbitrary $j = 1, \dots, d$. Denote by τ_{nj} the critical radius of $B_J(1) \cap B_2(|\beta_j|)$, denote $r_n = n^{-1/(2+w)}$, and define the abbreviation $a \vee b = \max\{a, b\}$. Observe

$$\begin{aligned} &\sup_{m_j \in B_J(1) \cap B_2(|\beta_j|)} \left| \frac{1}{n} \sum_{i=1}^n \sigma^i m_j(X_j^i) \right| \\ &\leq c \left(\mathcal{R}_n(B_J(1) \cap B_2(|\beta_j|)) + \sqrt{\frac{\log n}{n}} \left(\sup_{m_j \in B_J(1) \cap B_2(|\beta_j|)} \|m_j\|_n \right) \right) \\ &\leq c \left(\mathcal{R}(B_J(1) \cap B_2(|\beta_j|)) + \frac{\log n}{n} + \sqrt{\frac{\log n}{n}} \left(\sup_{m_j \in B_J(1) \cap B_2(|\beta_j|)} \|m_j\|_n \right) \right) \\ &\leq c \left(\mathcal{R}(B_J(1) \cap B_2(|\beta_j|)) + \frac{\log n}{n} + \sqrt{\frac{\log n}{n}} \sqrt{2} (|\beta_j| \vee \tau_{nj}) \right) \\ &\leq c \left(\frac{|\beta_j|^{1-w/2}}{\sqrt{n}} + \frac{\log n}{n} + (|\beta_j| \vee \tau_{nj}) \sqrt{\frac{\log n}{n}} \right) \\ &\leq c \left(\frac{|\beta_j|^{1-w/2}}{\sqrt{n}} + (|\beta_j| \vee r_n) \sqrt{\frac{\log n}{n}} \right). \end{aligned}$$

The first three inequalities above hold with probability at least $1 - 1/3n^2$ each. The first inequality is by Theorem 3.6 in [Wainwright \[2017\]](#) (see also Example 3.9 in [Wainwright \[2017\]](#)); the second and third are by Lemma A.4 and Lemma 3.6 in [Bartlett et al. \[2005\]](#), respectively. The fourth upper bounds the local Rademacher complexity of $B_J(1) \cap B_2(|\beta_j|)$, and the fifth upper bounds the critical radius τ_{nj} of this class, both following the proof of Lemma 15 (recall, the functions in $B_J(1) \cap B_2(|\beta_j|)$ have a uniform sup norm bound of $\rho = R_1 + R_2$, by Lemma 16). The last step also uses $\log n/n \leq r_n \sqrt{\log n/n}$ for n sufficiently large. The final result of the above display holds with probability at least $1 - 1/n^2$; by a union bound, it holds with probability at least $1 - d/n^2$ simultaneously over $j = 1, \dots, d$. Call this event \mathcal{E} . Then on \mathcal{E} ,

$$\begin{aligned} \sup_{m \in \mathcal{F} \cap B_2(t)} \left| \frac{1}{n} \sum_{i=1}^n \sigma^i m(X^i) \right| &\leq c \sum_{j=1}^d \left(\frac{|\beta_j|^{1-w/2}}{\sqrt{n}} + (|\beta_j| \vee r_n) \sqrt{\frac{\log n}{n}} \right) \\ &\leq c \left(\frac{d^{(2+w)/4} t^{1-w/2}}{\sqrt{n}} + \sqrt{\frac{d \log n}{n}} t + dr_n^2 \right). \end{aligned} \tag{68}$$

In the second line, we use Holder's inequality $a^T b \leq \|a\|_p \|b\|_q$ for the first term, with $p = 4/(2+w)$ and $q = 4/(2-w)$; we use $a \vee b \leq a + b$ for the second term, along the bound $\|\beta\|_1 \leq \sqrt{dt}$, and the fact that $r_n \sqrt{\log n/n} \leq r_n^2$ for large enough n .

Meanwhile, on \mathcal{E}^c , we can apply the simple bound $\|m\|_\infty \leq \sum_{j=1}^d \|m_j\|_\infty \leq \rho d$ for functions in $\mathcal{F} \cap B_2(t)$, where $\rho = R_1 + R_2$ (owing to Lemma 16), and thus

$$\sup_{m \in \mathcal{F} \cap B_2(t)} \frac{1}{n} \left| \sum_{i=1}^n \sigma^i m(X^i) \right| \leq \rho d. \quad (69)$$

Splitting the expectation defining the local Rademacher complexity over $\mathcal{E}, \mathcal{E}^c$, and using (68), (69),

$$\begin{aligned} \mathcal{R}(\mathcal{F} \cap B_2(t)) &= \mathbb{E}_{X, \sigma} \left[\sup_{m \in \mathcal{F} \cap B_2(t)} \frac{1}{n} \left| \sum_{i=1}^n \sigma^i m(X^i) \right| \right] \\ &\leq c \left(\frac{d^{(2+w)/4} t^{1-w/2}}{\sqrt{n}} + \sqrt{\frac{d \log n}{n} t + d r_n^2} \right) + \frac{\rho d^2}{n^2}. \end{aligned} \quad (70)$$

It can be easily verified that for $t = c\sqrt{d}r_n^2$, the upper bound in (70) is at most t^2/C . Therefore this is an upper bound on the critical radius of \mathcal{F} , which completes the proof. \square

Lastly, we bound the gap in the empirical and L_2 means of functions in $\mathcal{M}_2(1) \cap B_2(t)$, for small enough t . The proof uses Theorem 2.1 in Bartlett et al. [2005].

Lemma 18. *Let $X^i, i = 1, \dots, n$ denote an i.i.d. sample from a distribution Q on $[0, 1]^d$ satisfying Assumption A3, and let J satisfy Assumption C3. There are constants $c_0, n_0 > 0$, that depend only on b_1, b_2, k, L, K, w , such that for any $0 < t \leq 1$ and $n \geq n_0$,*

$$|\langle m, 1 \rangle_n - \langle m, 1 \rangle_2| \leq c_0 \left(\frac{d^{(2+w)/4} t^{1-w/2}}{\sqrt{n}} + \sqrt{\frac{d \log n}{n} t + d n^{-2/(2+w)}} \right), \quad \text{for all } m \in \mathcal{M}_2(1) \cap B_2(t),$$

with probability at least $1 - 1/n$, where $\mathcal{M}_2(1)$ is the space in (67).

Proof. This follows by combining the local Rademacher bound in (70) from the proof of Lemma 17 with Theorem 2.1 in Bartlett et al. [2005], and simplifying by keeping the dominant terms for large enough n . \square

A.11 Proof of Theorem 2

At a high-level, the difference between this proof and that of Theorem 1 is that here we do not try to directly control the sub-Gaussian complexity term (as this would lead to a poor dependence on the dimension d). Instead, we reduce the problem to controlling univariate sub-Gaussian complexities, and then assemble the result using ties between the empirical and L_2 norms, and the decomposition property (29). We will use the same general notation as in the proof of Theorem 1: $c > 0$ denotes a constant that will multiply our final bound, and will control the probability with which the final result holds; we will use the empirical norms and inner products with vector arguments, to be interpreted appropriately; we use the abbreviations $r_n, \hat{\Delta}$, and so on. Finally, in many lines that follow, we will redefine c by absorbing constant factors into it, without explicit notice.

The same arguments that led us to (62) yield the basic inequality

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq 2\langle \epsilon, \hat{\Delta} \rangle_n + \|\tilde{f} - f_0\|_n^2 = 2 \sum_{j=1}^d \langle \epsilon, \hat{\Delta}_j \rangle_n + \|\tilde{f} - f_0\|_n^2, \quad (71)$$

where we write $\hat{\Delta} = \sum_{j=1}^d \hat{\Delta}_j$.

Bounding the sub-Gaussian complexity terms. We now bound the univariate sub-Gaussian complexity terms, appearing in the sum on the right-hand side in (71). For $j = 1, \dots, d$, define $g_j = \hat{\Delta}_j / (2\delta + \|\hat{\Delta}_j\|_n)$, and note that by construction $J(g_j) \leq 1$ and $\|g_j\|_n \leq 1$. By Lemma 7, there are constants $c_0, R > 0$ such that $\|g_j\|_\infty \leq R$ on an event whose probability is at least $1 - \exp(-c_0 n)$. Thus by Lemma 14, there are constants $c_1, c_2 > 0$ such that for all $c \geq c_1$,

$$\frac{2\langle \epsilon, g_j \rangle_n}{\|g_j\|_n^{1-w/2}} \leq \frac{c}{\sqrt{n}}, \quad \text{for all } j = 1, \dots, d,$$

on an event Ω_1 with probability at least $1 - \exp(-c_0 n) - \exp(-c_2 c^2)$. Plugging this into (71) gives

$$\begin{aligned} \|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 &\leq \frac{c}{\sqrt{n}} \sum_{j=1}^d (2\delta + \|\hat{\Delta}_j\|_n)^{w/2} \|\hat{\Delta}_j\|_n^{1-w/2} + \|\tilde{f} - f_0\|_n^2, \\ &\leq \frac{c\delta^{w/2}}{\sqrt{n}} \sum_{j=1}^d \|\hat{\Delta}_j\|_n^{1-w/2} + \frac{c}{\sqrt{n}} \sum_{j=1}^d \|\hat{\Delta}_j\|_n + \|\tilde{f} - f_0\|_n^2, \end{aligned} \quad (72)$$

on Ω_1 , where we used Lemma 10 in the second inequality.

Converting empirical norms into L_2 norms. For each $j = 1, \dots, d$, let $\bar{\Delta}_j = \langle \hat{\Delta}_j, 1 \rangle_2$ be the L_2 mean of $\hat{\Delta}_j$, and $\tilde{\Delta}_j = \hat{\Delta}_j - \bar{\Delta}_j$ the L_2 centered version of $\hat{\Delta}_j$. Note that, for each $j = 1, \dots, d$, we have by empirical orthogonality $\|\tilde{\Delta}_j\|_n^2 = \|\hat{\Delta}_j\|_n^2 - |\bar{\Delta}_j|^2$, which implies $\|\hat{\Delta}_j\|_n \leq \|\tilde{\Delta}_j\|_n$. Applying this to upper bound the right-hand side in (72) gives

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq \frac{c\delta^{w/2}}{\sqrt{n}} \sum_{j=1}^d \|\tilde{\Delta}_j\|_n^{1-w/2} + \frac{c}{\sqrt{n}} \sum_{j=1}^d \|\tilde{\Delta}_j\|_n + \|\tilde{f} - f_0\|_n^2, \quad (73)$$

on Ω_1 . We bound each empirical norm in the sum on the right-hand side in (73) by its L_2 norm counterpart. Now, for each $j = 1, \dots, d$, define $g_j = \tilde{\Delta}_j / (2\delta + \|\tilde{\Delta}_j\|_2)$. Since $J(g_j) \leq 1$ and $\|g_j\|_2 \leq 1$, by Lemma 16, there is a constant $R > 0$ such that $\|g_j\|_\infty \leq R$. We can hence apply Lemma 15 to the measure $\Lambda = Q_j$, which gives constants $c_3, c_4, c_5 > 0$ such that

$$\|g_j\|_n \leq \sqrt{\frac{3}{2}} \|g_j\|_2 + c_3 r_n, \quad \text{for all } j = 1, \dots, d,$$

on an event Ω_2 with probability at least $1 - c_4 d \exp(-c_5 n r_n^2)$, where recall $r_n = n^{-1/(2+w)}$, i.e.,

$$\|\tilde{\Delta}_j\|_n \leq 2\sqrt{\frac{3}{2}} \|\tilde{\Delta}_j\|_2 + 2c_3 r_n \delta, \quad \text{for all } j = 1, \dots, d,$$

on Ω_2 , where we assume n is large enough so that $c_3 r_n \leq \sqrt{3/2}$. Returning to (73), and using the simple inequality in Lemma 10, we have

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq \frac{c\delta^{w/2}}{\sqrt{n}} \sum_{j=1}^d \|\tilde{\Delta}_j\|_2^{1-w/2} + \frac{c}{\sqrt{n}} \sum_{j=1}^d \|\tilde{\Delta}_j\|_2 + c d r_n^2 \delta + \|\tilde{f} - f_0\|_n^2, \quad (74)$$

on $\Omega_1 \cap \Omega_2$.

Invoking L_2 decomposability. We recall the key L_2 decomposition property (29), of additive functions with L_2 mean zero components. Using Holder's inequality $a^T b \leq \|a\|_p \|b\|_q$ to bound the

first sum on the right-hand side in (74), with $p = 4/(2+w)$ and $q = 4/(2-w)$, and Cauchy-Schwartz to bound the second sum in (74), we get

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq \frac{cd^{(2+w)/4}\delta^{w/2}}{\sqrt{n}} \|\tilde{\Delta}\|_2^{1-w/2} + c\sqrt{\frac{d}{n}} \|\tilde{\Delta}\|_2 + cdr_n^2\delta + \|\tilde{f} - f_0\|_n^2, \quad (75)$$

on $\Omega_1 \cap \Omega_2$, where we denote $\tilde{\Delta} = \sum_{j=1}^d \tilde{\Delta}_j$.

Converting back to empirical norm. We bound the L_2 norm of the centered error vector on the right-hand side in (75) with its empirical norm counterpart. By Lemma 17 applied to $m = \tilde{\Delta}/(2\delta)$, provided n is large enough so that $c_6\sqrt{dr_n} \leq 1$, there are constants $c_6, c_7, c_8 > 0$ such that

$$\|\tilde{\Delta}\|_2 \leq \sqrt{2}\|\tilde{\Delta}\|_n + 2c_6\sqrt{dr_n}\delta, \quad (76)$$

on an event Ω_3 with probability at least $1 - c_7 \exp(-c_8 dnr_n^2)$. Plugging this into the right-hand side in (75), and using Lemma 10, we have

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq \frac{cd^{(2+w)/4}\delta^{w/2}}{\sqrt{n}} \|\tilde{\Delta}\|_n^{1-w/2} + c\sqrt{\frac{d}{n}} \|\tilde{\Delta}\|_n + cdr_n^2\delta + \|\tilde{f} - f_0\|_n^2,$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3$. Using Lemma 11 on the first term above, with $a = d^{(2+w)/4}\delta^{w/2}/\sqrt{n}$ and $b = \|\tilde{\Delta}\|_n$, and simplifying, gives

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq c\sqrt{dr_n}\sqrt{\delta}\|\tilde{\Delta}\|_n + cdr_n^2\delta + \|\tilde{f} - f_0\|_n^2, \quad (77)$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3$.

Deriving an empirical norm error bound. Note that in (77), we have $\|\hat{\Delta}\|_n$ on the left-hand side and $\|\tilde{\Delta}\|_n$ on the right-hand side, where $\tilde{\Delta} = \hat{\Delta} - \bar{\Delta}$ is the centered error vector, and we are abbreviating $\bar{\Delta} = \sum_{j=1}^d \bar{\Delta}_j$. We seek to bound $|\bar{\Delta}|$. Define $t = c_6\sqrt{dr_n}$, where c_6 is the constant in (76), and define

$$m = \frac{t\tilde{\Delta}/(2\delta)}{\sqrt{2}\|\tilde{\Delta}\|_n/(2\delta) + t}.$$

Note that $J(m_j) \leq J(\tilde{\Delta}_j)/(2\delta) \leq 1$, for $j = 1, \dots, d$, by construction, and also

$$\|m\|_2 = \frac{t\|\tilde{\Delta}\|_2/(2\delta)}{\sqrt{2}\|\tilde{\Delta}\|_n/(2\delta) + t} \leq t,$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3$, recalling (76). By Lemma 18 applied to m , provided n is large enough such that $t = c_6\sqrt{dr_n} \leq 1$, there is a constant $c_9 > 0$ such that $|\langle m, 1 \rangle_n| \leq c_9 t^2$ on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$, where Ω_4 is an event with probability at least $1 - 1/n$, i.e.,

$$|\langle 1, \tilde{\Delta} \rangle_n|/(2\delta) \leq c_9 t(\sqrt{2}\|\tilde{\Delta}\|_n/(2\delta) + t),$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$, i.e.,

$$|\bar{\Delta}| \leq \sqrt{2}c_9 t \|\tilde{\Delta}\|_n + 2c_9 t^2 \delta,$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$. Thus, by empirical orthogonality,

$$\|\tilde{\Delta}\|_n^2 = \|\hat{\Delta}\|_n^2 + |\bar{\Delta}|^2 \leq \|\hat{\Delta}\|_n^2 + 2(\sqrt{2}c_9 t)^2 \|\tilde{\Delta}\|_n^2 + 2(2c_9 t^2 \delta)^2,$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$, and assuming n is large enough so that $2(\sqrt{2}c_9 t)^2 \leq 1/2$ and $2(2c_9)^2 t^2 \delta \leq 1$, this becomes

$$\frac{1}{2}\|\tilde{\Delta}\|_n^2 \leq \|\hat{\Delta}\|_n^2 + t^2 \delta, \quad (78)$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$. Using this on the right-hand side in (77) gives

$$\|\hat{f} - f_0\|_n^2 + \|\hat{\Delta}\|_n^2 \leq c\sqrt{dr_n}\sqrt{\delta}\|\hat{\Delta}\|_n + cdr_n^2\delta + \|\tilde{f} - f_0\|_n^2,$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$. Using the simple inequality $2ab \leq a^2 + b^2$ on the first term on the right-hand side above, with $a = c\sqrt{dr_n}\sqrt{\delta}$ and $b = \|\hat{\Delta}\|_n$, gives

$$\|\hat{f} - f_0\|_n^2 + \frac{1}{2}\|\hat{\Delta}\|_n^2 \leq \|\tilde{f} - f_0\|_n^2 + c^2dr_n^2\delta, \quad (79)$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$. The empirical norm result in (30) in the theorem follows by dropping the nonnegative term $\|\hat{\Delta}\|_n^2/2$ from the left-hand side, and adjusting the constants $c, c_1, c_2, c_3 > 0$ as needed.

Deriving an L_2 norm error bound. Note that (79) also implies

$$\frac{1}{2}\|\hat{\Delta}\|_n^2 \leq \|\tilde{f} - f_0\|_n^2 + c^2dr_n^2\delta,$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$. Recalling (78), this gives

$$\|\tilde{\Delta}\|_n^2 \leq 4\|\tilde{f} - f_0\|_n^2 + c^2dr_n^2\delta, \quad (80)$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$. By L_2 orthogonality,

$$\begin{aligned} \|\hat{\Delta}\|_2^2 &= \|\tilde{\Delta}\|_2^2 + |\bar{\Delta}|^2 \\ &\leq 3\|\tilde{\Delta}\|_n^2 + t^2\delta^2 \\ &\leq 12\|\tilde{f} - f_0\|_n^2 + c^2dr_n^2\delta^2, \end{aligned}$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$, where in the second line we used (76) and $|\bar{\Delta}| \leq \|\tilde{\Delta}\|_n$, and in the third line we used (80). Finally,

$$\|\hat{f} - f_0\|_2^2 \leq 2\|\hat{f} - \tilde{f}\|_2^2 + 2\|\tilde{f} - f_0\|_2^2 \leq 24\|\tilde{f} - f_0\|_n^2 + 2\|\tilde{f} - f_0\|_2^2 + c^2dr_n^2\delta^2,$$

on $\Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$. The L_2 norm result in (31) in the theorem follows by simply adjusting the constants $c, c_1, c_2, c_3 > 0$ as needed.

A.12 Proof of Corollary 2

The proof of the statement in part (a) is exactly as in the proof of part (a) in Corollary 1.

For part (b), we can consider $k \geq 2$ without a loss of generality, and start with an additive spline approximant \tilde{f} to f_0 from Lemma 12. Let \tilde{f} denote the result of centering each component of \tilde{f} to have zero empirical mean. Then $\text{TV}(\tilde{f}_j^{(k)}) \leq a_k c_n = \delta$, $j = 1, \dots, d$, and just as in the proof of part (b) in Corollary 1, letting $\|\cdot\|$ denote either the empirical or L_2 norm, we have

$$\left\| \sum_{j=1}^d \tilde{f}_j - \sum_{j=1}^d f_{0j} \right\|^2 \leq \left(\frac{c_0^k a_k}{b_0^k} \right)^2 \left(\frac{\log n}{n} \right)^{2k} d^2 c_n^2.$$

But when $n \geq n_0(dc_n)^{(2k+3)/(2k+2)}$, the right-hand side above is bounded by $a_0 dn^{-(2k+2)/(2k+3)} c_n$ for a constant $a_0 > 0$, which shows the approximation error terms in (30), (31) are of the desired order. This proves the desired result for restricted locally adaptive splines.

For part (c), we follow the same arguments, the only difference being that we construct a falling factorial approximant \check{f} to f_0 from Lemma 13.

A.13 Preliminaries for the proof of Theorem 3

The next two results in this subsection are helper lemmas for the last lemma.

Lemma 19. *Let J be a functional that satisfies Assumptions C1, C2, C4. Then there are constants $\tilde{K}_1, \tilde{\delta}_1 > 0$, that depend only on k, L, K_1, w , such that for all $0 < \delta \leq \tilde{\delta}_1$,*

$$\log M\left(\delta, \|\cdot\|_2, \Pi_k^\perp(B_J(1))\right) \geq \tilde{K}_1 \delta^{-w},$$

where $\|\cdot\|_2$ is the $L_2(U)$ norm, with U the uniform distribution on $[0, 1]$, and Π_k^\perp is defined by

$$\Pi_k^\perp(g) = g - \Pi_k(g), \quad \text{where } \Pi_k(g) = \operatorname{argmin}_{p \in \mathcal{P}_k} \|g - p\|_2,$$

with \mathcal{P}_k denoting the space of polynomials of degree k . In other words, Π_k^\perp is the projection operator onto the space orthogonal (in $L_2(U)$) to the polynomials of degree k .

Proof. Let $R_0 > 0$ be the constant from Lemma 5, when we take $\Lambda = U$. Note that

$$B_J(1) \cap B_\infty(R_0) = \Pi_k^\perp(B_J(1)) + (\mathcal{P}_k \cap B_\infty(R_0)). \quad (81)$$

In general, for $S = S_1 + S_2$ and a norm $\|\cdot\|$, observe that, from basic relationships between covering and packing numbers,

$$M(4\delta, \|\cdot\|, S) \leq N(2\delta, \|\cdot\|, S) \leq N(\delta, \|\cdot\|, S_1)N(\delta, \|\cdot\|, S_2) \leq M(\delta, \|\cdot\|, S_1)N(\delta, \|\cdot\|, S_2),$$

so that

$$\log M(\delta, \|\cdot\|, S_1) \geq \log \frac{M(4\delta, \|\cdot\|, S)}{N(\delta, \|\cdot\|, S_2)}.$$

Applying this to our decomposition in (81),

$$\begin{aligned} \log M\left(\delta, \|\cdot\|_2, \Pi_k^\perp(B_J(1))\right) &\geq \log \frac{M(4\delta, \|\cdot\|_2, B_J(1) \cap B_\infty(R_0))}{N(\delta, \|\cdot\|_2, \mathcal{P}_k \cap B_\infty(R_0))} \\ &\geq K_1 R_0^w 4^{-w} \delta^{-w} - A(k+1) \log(1/\delta), \end{aligned}$$

where in the second inequality we used Assumption C4 (assuming without a loss of generality that $R_0 \geq 1$), and a well-known entropy bound for a finite-dimensional ball (e.g., Mammen [1991]), with $A > 0$ being a constant that depends only on R_0 . For small enough δ , the right-hand side above is of the desired order, and this completes the proof. \square

Lemma 20. *Let $d, M > 0$ be integers, and $I = \{1, \dots, M\}$. Denote by $H(u, v) = \sum_{j=1}^d 1\{u_j \neq v_j\}$ the Hamming distance between $u, v \in I^d$. Then there is a subset $S \subseteq I^d$ with $|S| \geq (M/4)^{d/2}$ such that $H(u, v) \geq d/2$ for any $u, v \in S$.*

Proof. Let $\Omega_0 = I^d$, $u_0 = (1, \dots, 1) \in \Omega_0$. For $j = 0, 1, \dots$, recursively define

$$\Omega_{j+1} = \{u \in \Omega_j : H(u, u_j) > a = \lceil d/2 \rceil\},$$

where u_{j+1} is arbitrarily chosen from Ω_{j+1} . The procedure is stopped when Ω_{j+1} is empty; denote the last set defined in this procedure by Ω_E , and denote $S = \{u_0, \dots, u_E\}$. For $0 \leq i, j \leq E$, by construction, $H(u_i, u_j) > a$. For $j = 0, \dots, E$,

$$\begin{aligned} n_j &= |\Omega_j - \Omega_{j+1}| = |\{u \in \Omega_j : H(u, u_j) \leq a\}| \\ &\leq |\{u \in I^d : H(u, u_j) \leq a\}| \\ &= \binom{d}{d-a} M^a \end{aligned}$$

The last step is true because we can choose $d - a$ positions in which u matches u_j in $\binom{d}{d-a}$ ways, and the rest of the a positions can be filled arbitrarily in M ways. Also note $M^d = n_0 + \dots + n_E$. Therefore

$$M^d \leq (E + 1) \binom{d}{d-a} M^a,$$

which implies

$$E + 1 \geq \frac{M^{d-a}}{\binom{d}{d-a}} \geq \frac{M^{d-a}}{2^d} \geq (M/4)^{d/2}.$$

□

The lemma below gives a key technical result used in the proof of Theorem 3.

Lemma 21. *Let J be a functional that satisfies Assumptions C1, C2, C4. Then there are constants $\bar{K}_1, \bar{\delta}_1 > 0$, that depend only on $w, \bar{K}_1, \bar{\delta}_1$, where $\bar{K}_1, \bar{\delta}_1 > 0$ are the constants from Lemma 19, such that for all $0 < \delta \leq \bar{\delta}_1$,*

$$\log M \left(\delta, \|\cdot\|_2, \Pi_{k,d}^\perp(B_J^d(1)) \right) \geq \bar{K}_1 d^{1+w/2} \delta^{-w},$$

where $\|\cdot\|_2$ is the $L_2(U)$ norm, with U the uniform distribution on $[0, 1]^d$, and $\Pi_{k,d}^\perp$ is defined by

$$\Pi_{k,d}^\perp(g) = g - \Pi_{k,d}(g), \quad \text{where } \Pi_{k,d}(g) = \operatorname{argmin}_{p \in \mathcal{P}_{k,d}} \|g - p\|_2,$$

and $\mathcal{P}_{k,d}$ contains all functions of the form $p(x) = \sum_{j=1}^d p_j(x_j)$, for polynomials p_j , $j = 1, \dots, d$ of degree k . In other words, $\Pi_{k,d}^\perp$ is the projection operator onto the space orthogonal (in $L_2(U)$) to the space $\mathcal{P}_{k,d}$ of additive polynomials of degree k .

Proof. It is easy to check that the decomposability property of the $L_2(U)$ norm, in (29), implies a certain decomposability of the $L_2(U)$ projection operators $\Pi_{k,d}, \Pi_{k,d}^\perp$ over additive functions:

$$\Pi_{k,d} \left(\sum_{j=1}^d m_j \right) = \sum_{j=1}^d \Pi_k(m_j), \quad \Pi_{k,d}^\perp \left(\sum_{j=1}^d m_j \right) = \sum_{j=1}^d \Pi_k^\perp(m_j),$$

where Π_k, Π_k^\perp are projection operators onto \mathcal{P}_k and its orthocomplement, respectively, as defined in Lemma 19. The decomposability result for $\Pi_{k,d}^\perp$ in particular implies that

$$\Pi_{k,d}^\perp(B_J^d(1)) = \left\{ \sum_{j=1}^d f_j : f_j \in \Pi_k^\perp(B_J(1)), j = 1, \dots, d \right\}. \quad (82)$$

Abbreviate $M = M(\delta/\sqrt{d/2}, \|\cdot\|_2, \Pi_k^\perp(B_J(1)))$. By Lemma 19, we have for small enough δ ,

$$\log M \geq \tilde{K}_1 2^{-w/2} d^{w/2} \delta^{-w}.$$

Now let g_1, \dots, g_M denote a $(\delta/\sqrt{d/2})$ -packing of $\Pi_k^\perp(B_J(1))$. Let $I = \{1, \dots, M\}$, and for $u \in I^d$, define $f_u \in \Pi_k^\perp(B_J^d(1))$ by

$$f_u = \sum_{j=1}^d g_{u_j},$$

i.e., f_u is an additive function with components g_{u_j} , $j = 1, \dots, d$. If the Hamming distance between indices u, v satisfies $H(u, v) \geq d/2$, then

$$\|f_u - f_v\|_2^2 = \sum_{j=1}^d \|g_{u_j} - g_{v_j}\|_2^2 \geq H(u, v) \frac{\delta^2}{d/2} \geq \delta^2,$$

where we have again used the $L_2(U)$ decomposability property in (29). Thus, it is sufficient to find a subset S of I^d such that $u, v \in S \Rightarrow H(u, v) \geq d/2$. By Lemma 20, we can choose such an S with $|S| \geq (M/4)^{d/2}$. For small enough δ , such that $M \geq 16$, this gives the desired result because

$$\log |S| \geq \frac{d}{2} \log \frac{M}{4} \geq \frac{d}{4} \log M \geq \tilde{K}_1 2^{-w/2-2} d^{1+w/2} \delta^{-w}.$$

□

A.14 Proof of Theorem 3

Clearly, by orthogonality, for any functions \hat{f}, f_0 ,

$$\|\hat{f} - f_0\|_2^2 = \|\Pi_{k,d}(\hat{f}) - \Pi_{k,d}(f_0)\|_2^2 + \|\Pi_{k,d}^\perp(\hat{f}) - \Pi_{k,d}^\perp(f_0)\|_2^2 \geq \|\Pi_{k,d}^\perp(\hat{f}) - \Pi_{k,d}^\perp(f_0)\|_2^2,$$

where $\Pi_{k,d}, \Pi_{k,d}^\perp$ are projection operators onto $\mathcal{P}_{k,d}$ and its orthocomplement, respectively, defined in Lemma 19. Thus it suffices to consider the minimax error over $\Pi_{k,d}^\perp(B_J^d(c_n))$.

First, we lower bound the packing number and upper bound the covering number of the class $\Pi_{k,d}^\perp(B_J^d(c_n))$. The upper bound is more straightforward:

$$\begin{aligned} \log N\left(\epsilon, \|\cdot\|_2, \Pi_{k,d}^\perp(B_J^d(c_n))\right) &= \log N\left(\epsilon/c_n, \|\cdot\|_2, \Pi_{k,d}^\perp(B_J^d(1))\right) \\ &\leq \sum_{j=1}^d \log N\left(\epsilon/(c_n \sqrt{d}), \|\cdot\|_2, \Pi_k^\perp(B_J(1))\right) \\ &\leq K_2 c_n^w d^{1+w/2} \epsilon^{-w}. \end{aligned} \tag{83}$$

The second inequality follows from property (82) in the proof of Lemma 21 and similar arguments to those in the proof of Lemma 8—except that we leverage the decomposability of the L_2 norm, as in (29), instead of using the triangle inequality. The third inequality follows from Assumption C4.

The lower bound is less straightforward, and is given by Lemma 21:

$$\begin{aligned} \log M\left(\delta, \|\cdot\|_2, \Pi_{k,d}^\perp(B_J^d(c_n))\right) &= \log M\left(\delta/c_n, \|\cdot\|_2, \Pi_{k,d}^\perp(B_J^d(1))\right) \\ &\geq \bar{K}_1 c_n^w d^{1+w/2} \delta^{-w}. \end{aligned} \tag{84}$$

We note that (84) holds for $0 < \delta \leq \bar{\delta}_1$, where $\bar{\delta}_1 > 0$ is the constant from Lemma 21.

Now, following the strategy in Yang and Barron [1999], we use these bounds on the packing and covering numbers, along with Fano's inequality, to establish the desired result. Let f_1, f_2, \dots, f_M be a δ_n -packing of $\Pi_{k,d}^\perp(B_J(c_n))$, for $\delta_n > 0$ to be specified later. Fix an arbitrary estimator \hat{f} , and let

$$\hat{Z} = \operatorname{argmin}_{j \in \{1, \dots, M\}} \|\hat{f} - f_j\|_2.$$

We will use $P_{X,f}$ and $\mathbb{E}_{X,f}$ to denote the probability and expectation operators, respectively, over i.i.d. draws $X^i \sim U$, $i = 1, \dots, n$ (where U is the uniform distribution on $[0, 1]^d$), and i.i.d. draws

$Y^i | X^i \sim N(f(X^i), \sigma^2)$, $i = 1, \dots, n$. Then

$$\begin{aligned}
& \sup_{f_0 \in \Pi_{k,d}^\perp(B_J^d(c_n))} \mathbb{E}_{X, f_0} \|\hat{f} - f_0\|_2^2 \geq \sup_{f_0 \in \{f_1, \dots, f_M\}} \mathbb{E}_{X, f_0} \|\hat{f} - f_0\|_2^2 \\
& \geq \frac{1}{M} \mathbb{E}_X \sum_{j=1}^M \mathbb{E}_{f_j} \|\hat{f} - f_j\|_2^2 \\
& = \frac{1}{M} \mathbb{E}_X \sum_{j=1}^M \left(\mathbb{P}_{f_j}(\hat{Z} \neq j) \mathbb{E}_{f_j} (\|\hat{f} - f_j\|_2^2 | \hat{Z} \neq j) + \mathbb{P}_{f_j}(\hat{Z} = j) \mathbb{E}_{f_j} (\|\hat{f} - f_j\|_2^2 | \hat{Z} = j) \right) \\
& \geq \frac{1}{M} \mathbb{E}_X \sum_{j=1}^M \mathbb{P}_{f_j}(\hat{Z} \neq j) \mathbb{E}_{f_j} (\|\hat{f} - f_j\|_2^2 | \hat{Z} \neq j) \\
& \geq \frac{1}{M} \mathbb{E}_X \sum_{j=1}^M \mathbb{P}_{f_j}(\hat{Z} \neq j) \frac{\delta_n^2}{4}, \tag{85}
\end{aligned}$$

where in the last inequality we have used the fact that if $\hat{Z} \neq j$, then \hat{f} must be at least $\delta_n/2$ away from f_j , for each $j = 1, \dots, M$.

Abbreviate q_j for the distribution P_{f_j} , $j = 1, \dots, M$, and define the mixture $\bar{q} = \frac{1}{M} \sum_{j=1}^M q_j$. By Fano's inequality,

$$\frac{1}{M} \mathbb{E}_X \sum_{j=1}^M \mathbb{P}_{f_j}(\hat{Z} \neq j) \geq 1 - \frac{\frac{1}{M} \sum_{j=1}^M \mathbb{E}_X \text{KL}(q_j \| \bar{q}) + \log 2}{\log M}, \tag{86}$$

where $\text{KL}(P_1 \| P_2)$ denotes the Kullback-Leibler (KL) divergence between distributions P_1, P_2 . Let g_1, g_2, \dots, g_N be an ϵ_n -covering of $\Pi_{k,d}^\perp(B_J^d(c_n))$, for $\epsilon_n > 0$ to be determined shortly. Abbreviate s_ℓ for the distribution P_{g_ℓ} , $\ell = 1, \dots, N$, and $\bar{s} = \frac{1}{N} \sum_{\ell=1}^N s_\ell$. Also, write $p(N(f(X), \sigma^2 I))$ for the density of a $N(f(X), \sigma^2 I)$ random variable, where $f(X) = (f(X^1), \dots, f(X^n)) \in \mathbb{R}^n$. Then

$$\begin{aligned}
\frac{1}{M} \sum_{j=1}^M \mathbb{E}_X \text{KL}(q_j \| \bar{q}) & \leq \frac{1}{M} \sum_{j=1}^M \mathbb{E}_X \text{KL}(q_j \| \bar{s}) \\
& = \frac{1}{M} \sum_{j=1}^M \mathbb{E}_{X, f_j} \log \frac{p(N(f_j(X), \sigma^2 I))}{\frac{1}{N} \sum_{\ell=1}^N p(N(g_\ell(X), \sigma^2 I))} \\
& \leq \frac{1}{M} \sum_{j=1}^M \left(\log N + \mathbb{E}_X \min_{\ell=1, \dots, N} \text{KL}(q_j \| s_\ell) \right) \\
& \leq \frac{1}{M} \sum_{j=1}^M \left(\log N + \frac{n\epsilon_n^2}{2\sigma^2} \right) \\
& \leq K_2 c_n^w d^{1+w/2} \epsilon_n^{-w} + \frac{n\epsilon_n^2}{2\sigma^2}. \tag{87}
\end{aligned}$$

In the first line above, we used the fact that $\sum_{j=1}^M \text{KL}(q_j \| \bar{q}) \leq \sum_{j=1}^M \text{KL}(q_j \| s)$ for any other distribution s ; in the second and third, we explicitly expressed and manipulated the definition of KL divergence; in the fourth, we used $\text{KL}(q_j \| s_\ell) = \|f_j(X) - g_\ell(X)\|_2^2 / (2\sigma^2)$, and for each j , there is at least one ℓ such that $\mathbb{E}_X \|f_j(X) - g_\ell(X)\|_2^2 = \|f_j - g_\ell\|_2^2 \leq \epsilon_n^2$; in the fifth line, we used the entropy bound from (83). Minimizing (87) over $\epsilon_n > 0$ gives

$$\frac{1}{M} \sum_{j=1}^M \mathbb{E}_X \text{KL}(q_j \| \bar{q}) \leq \bar{K}_2 d n^{w/(2+w)} c_n^{2w/(2+w)},$$

for a constant $\bar{K}_2 > 0$. Returning to Fano's inequality (85), (86), we see that a lower bound on the minimax error is

$$\frac{\delta_n^2}{4} \left(1 - \frac{\bar{K}_2 d n^{w/(2+w)} c_n^{2w/(2+w)} + \log 2}{\log M} \right),$$

Therefore, a lower bound on the minimax error is $\delta_n^2/8$, for any $\delta_n > 0$ such that

$$\log M \geq 2\bar{K}_2 d n^{w/(2+w)} c_n^{2w/(2+w)} + 2 \log 2,$$

and for large enough n , the first term on the right-hand side above will be larger than $2 \log 2$, so it suffices to have

$$\log M \geq 4\bar{K}_2 d n^{w/(2+w)} c_n^{2w/(2+w)}. \quad (88)$$

Set $\delta_n = (\bar{K}_1/4\bar{K}_2)^{1/w} \sqrt{d} n^{-1/(2+w)} c_n^{w/(2+w)}$. Provided that $\delta_n \leq \bar{\delta}_1$, our log packing bound (84) is applicable, and ensures that (88) will be satisfied. This completes the proof.

A.15 Proof of Corollary 3

We only need to check Assumption C4 for $J(g) = \text{TV}(g^{(k)})$, $w = 1/(k+1)$, and then we can apply Theorem 3. As before, the entropy bound upper bound is implied by results in Birman and Solomyak [1967] (see Mammen [1991] for an explanation and discussion). The packing number lower bound is verified as follows. For f a $(k+1)$ times weakly differentiable function on $[0, 1]$,

$$\text{TV}(f^{(k)}) = \int_0^1 |f^{(k+1)}(t)| dt \leq \left(\int_0^1 |f^{(k+1)}(t)|^2 dt \right)^{1/2}.$$

Hence

$$\left\{ f : \text{TV}(f^{(k)}) \leq 1, \|f\|_\infty \leq 1 \right\} \supseteq \left\{ f : \int_0^1 |f^{(k+1)}(t)|^2 dt \leq 1, \|f\|_\infty \leq 1 \right\}.$$

Results in Kolmogorov and Tikhomirov [1959] imply that the space on the right-hand side satisfies the desired log packing number lower bound. This proves the result.

A.16 Proof of the linear smoother lower bound in (38)

We may assume without a loss of generality that each f_{0j} , $j = 1, \dots, d$ has L_2 mean zero (since f_0 does). By the decomposability property of the L_2 norm over additive functions with L_2 mean zero components, as in (29), we have for any additive linear smoother $\hat{f} = \sum_{j=1}^d \hat{f}_j$,

$$\|\hat{f} - f_0\|_2^2 = \left(\sum_{j=1}^d \bar{f}_j \right)^2 + \sum_{j=1}^d \|(\hat{f}_j - \bar{f}_j) - f_{0j}\|_2^2$$

where \bar{f}_j denotes the L_2 mean of \hat{f}_j , $j = 1, \dots, d$. Note that the estimator $\hat{f}_j - \bar{f}_j$ is itself a linear smoother, for each $j = 1, \dots, d$, since if we write $\hat{f}_j(x_j) = w_j(x_j)^T Y$ for a weight function w_j over $x_j \in [0, 1]$, then $\hat{f}_j(x_j) - \bar{f}_j = \tilde{w}_j(x_j)^T Y$ for a weight function $\tilde{w}_j(x_j) = w_j(x_j) - \int_0^1 w_j(t) dt$. This, and the last display, imply that

$$\inf_{\hat{f} \text{ additive linear}} \sup_{f_0 \in \mathcal{F}_k^d(c_n)} \mathbb{E} \|\hat{f} - f_0\|_2^2 = \sum_{j=1}^d \inf_{\hat{f}_j \text{ linear}} \sup_{f_0 \in \mathcal{F}_k^d(c_n)} \mathbb{E} \|\hat{f}_j(Y) - f_{0j}\|_2^2. \quad (89)$$

Now fix an arbitrary $j = 1, \dots, d$, and consider the j th term in the sum on the right-hand side above. Here we are looking at a linear smoother \hat{f}_j fit to data

$$Y^i = \mu + f_{0j}(X_j^i) + \sum_{\ell \neq j} f_{0\ell}(X_\ell^i) + \epsilon^i, \quad i = 1, \dots, n. \quad (90)$$

which depends on the components $f_{0\ell}$, for $\ell \neq j$. This is why the supremum in the j th term of the sum on the right-hand side in (89) must be taken over $f_0 \in \mathcal{F}_k^d(c_n)$, rather than $f_{0j} \in \mathcal{F}_k(c_n)$. Our notation $\hat{f}_j(Y)$ is used as a reminder to emphasize the dependence on the full data vector in (90).

A simple reformulation, by appropriate averaging over the lattice, helps untangle this supremum. Write $\hat{f}_j(x_j) = w_j(x_j)^T Y$ for a weight function w_j over $x_j \in [0, 1]$, and for each $v = 1, \dots, N$, let I_j^v be the set of indices i such that $X_j^i = v/N$. Also let

$$\bar{Y}_j^v = \frac{1}{N^{d-1}} \sum_{i \in I_j^v} Y^i, \quad v = 1, \dots, N,$$

and $\bar{Y}_j = (\bar{Y}_j^1, \dots, \bar{Y}_j^N) \in \mathbb{R}^N$. Then note that we can also write $\hat{f}_j(x_j) = \bar{w}_j(x_j)^T \bar{Y}_j$ for a suitably defined weight function \bar{w}_j , i.e., note that we can think of \hat{f}_j as a linear smoother fit to data \bar{Y}_j , whose components follow the distribution

$$\bar{Y}_j^v = \mu_j + f_{0j}(v/N) + \bar{\epsilon}_j^v, \quad v = 1, \dots, N, \quad (91)$$

where we let $\mu_j = \mu + \frac{1}{N} \sum_{\ell \neq j} \sum_{u=1}^n f_{0\ell}(u/N)$, and $\bar{\epsilon}_j^v$, $v = 1, \dots, n$ are i.i.d. $N(0, \sigma^2/N^{d-1})$. Recalling that $f_{0j} \in \mathcal{F}_k(c_n)$, we are in a position to invoke univariate minimax results from [Donoho and Johnstone \[1998\]](#). As shown in Section 5.1 of [Tibshirani \[2014\]](#), the space $\mathcal{F}_k(c_n)$ contains the Besov space $B_{1,1}^{k+1}(c'_n)$, for a radius c'_n that differs from c_n only by a constant factor. Therefore, by Theorem 1 of [Donoho and Johnstone \[1998\]](#) on the minimax risk of linear smoothers fit to data from the model (91), we see that for large enough N and a constant $c_0 > 0$,

$$\begin{aligned} \inf_{\hat{f}_j \text{ linear}} \sup_{f_{0j} \in \mathcal{F}_k(c_n)} \mathbb{E} \|\hat{f}_j(\bar{Y}_j) - f_{0j}\|_2^2 &\geq c_0 (c_n N^{(d-1)/2})^{2/(2k+2)} \frac{N^{-(2k+1)/(2k+2)}}{N^{d-1}} \\ &= c_0 N^{-d(2k+1)/(2k+2)} c_n^{2/(2k+2)} \\ &= c_0 n^{-(2k+1)/(2k+2)} c_n^{2/(2k+2)}. \end{aligned} \quad (92)$$

As we have reduced the lower bound to the minimax risk of linear smoothers over a Besov ball, we can see that the same result (92) indeed holds simultaneously over all $j = 1, \dots, d$. Combining this with (89) gives the desired result (38).

A.17 Proof of Theorem 4 and derivation details for Algorithm 2

We show that the dual of (42) is equivalent to the additive trend filtering problem (7), and further, the Lagrange multipliers corresponding to the constraints $u_0 = u_j$, for $j = 1, \dots, d$, are equivalent to the primal variables θ_j , $j = 1, \dots, d$. Let $M = I - \mathbb{1}\mathbb{1}^T/n$, and rewrite problem (42) as

$$\begin{aligned} \min_{u_0, u_1, \dots, u_d \in \mathbb{R}^n} \quad & \frac{1}{2} \|MY - Mu_0\|_2^2 + \sum_{j=1}^d I_{U_j}(u_j) \\ \text{subject to} \quad & Mu_0 = Mu_1, Mu_0 = Mu_2, \dots, Mu_0 = Mu_d, \end{aligned}$$

We can write the Lagrangian of this problem as

$$L(u_0, u_1, \dots, u_d, \theta_1, \dots, \theta_d) = \frac{1}{2} \|MY - Mu_0\|_2^2 + \sum_{j=1}^d I_{U_j}(u_j) + \sum_{j=1}^d \theta_j^T M(u_0 - u_j).$$

and we want to minimize this over u_0, \dots, u_d to form the dual of (42). This gives

$$\max_{\theta_1, \dots, \theta_d \in \mathbb{R}^n} \frac{1}{2} \|MY\|_2^2 - \frac{1}{2} \left\| MY - \sum_{j=1}^d M\theta_j \right\|_2^2 - \sum_{j=1}^d \left(\max_{u_j \in U_j} u_j^T M\theta_j \right). \quad (93)$$

We use the fact that the support function of U_j is just ℓ_1 penalty composed with $S_j D_j$ (invoking the duality between ℓ_∞ and ℓ_1 norms),

$$\max_{u_j \in U_j} u_j^T M \theta_j = \max_{\|v_j\| \leq \lambda} v_j^T D_j S_j M \theta_j = \lambda \|D_j S_j M \theta_j\|_1,$$

where recall we abbreviate $D_j = D^{(X_j, k+1)}$, for $j = 1, \dots, d$, and this allows us to rewrite the above problem (93) as

$$\min_{\theta_1, \dots, \theta_d \in \mathbb{R}^n} \frac{1}{2} \left\| M Y - \sum_{j=1}^d M \theta_j \right\|_2^2 + \lambda \sum_{j=1}^d \|D_j S_j M \theta_j\|_1,$$

which is precisely the same as the original additive trend filtering problem (7).

This realization has important consequences. In the ADMM iterations (43), the scaled parameters $\rho \gamma_j$, $j = 1, \dots, d$ correspond to dual variables θ_j , $j = 1, \dots, d$ in problem (42), which from the above calculation, are precisely primal variables in (7). Under weak conditions, ADMM is known to produce convergent dual iterates, e.g., Section 3.2 of [Boyd et al. \[2011\]](#) shows that if (i) the criterion is a sum of closed, convex functions and (ii) strong duality holds, then the dual iterates from ADMM converge to optimal dual solutions. (Convergence of primal iterates requires stronger assumptions.) Our problem (42) satisfies these two conditions, and so for the ADMM algorithm outlined in (43), the scaled iterates $\rho \gamma_j^{(t)}$, $j = 1, \dots, d$ converge to optimal solutions in the dual of (42), i.e., optimal solutions in the additive trend filtering problem (7). This proves the first part of the theorem.

As for the second part of the theorem, it remains to show that Algorithm 2 is equivalent to the ADMM iterations (41). This follows by notationally swapping γ_j , $j = 1, \dots, d$ for θ_j / ρ , $j = 1, \dots, d$, rewriting the updates

$$\theta_j^{(t)} / \rho = u_0^{(t)} + \theta_j^{(t-1)} / \rho - u_j^{(t)}, \quad j = 1, \dots, d,$$

as

$$\theta_j^{(t)} = \rho \cdot \text{TF}_\lambda(u_0^{(t)} + \theta_j^{(t-1)} / \rho, X_j), \quad j = 1, \dots, d,$$

using (39), and lastly, eliminating u_j , $j = 1, \dots, d$ from the u_0 update by solving for these variables in terms of θ_j , $j = 1, \dots, d$, i.e., by using

$$u_j^{(t-1)} = u_0^{(t-1)} + \theta_j^{(t-2)} / \rho - \theta_j^{(t-1)} / \rho, \quad j = 1, \dots, d.$$

A.18 Cyclic versus parallel backfitting

We compare the performances of the usual cyclic backfitting method in Algorithm 1 to the parallel version in Algorithm 2, on a simulated data set generated as in Section 5.1, except with $n = 2000$ and $d = 24$. We computed the additive trend filtering estimate (7) (of quadratic order), at a fixed value of λ lying somewhere near the middle of the regularization path, by running the cyclic and parallel backfitting algorithms until each obtained a suboptimality of 10^{-8} in terms of the achieved criterion value (the optimal criterion value here was determined by running Algorithm 1 for a very large number of iterations). We used simply $\rho = 1$ in Algorithm 2.

Figure 6 shows the progress of the two algorithms, plotting the suboptimality of the criterion value across the iterations. The two panels, left and right, differ in how iterations are counted for the parallel method. On the left, one full cycle of d component updates is counted as one iteration for the parallel method—this corresponds to running the parallel algorithm in “naive” serial mode, where each component update is actually performed in sequence. On the right, d full cycles of d component updates is counted as one iteration for the parallel method—this corresponds to running the parallel algorithm in an “ideal” parallel mode with d parallel processors. In both panels, one full cycle of d component updates is counted as one iteration for the cyclic method. We see that, if parallelization is fully utilized, the parallel method cuts down the iteration cost by about a factor of 2, compared to the cyclic method. We should expect these computational gains to be even larger as the number of components d grows.




Figure 6: *Suboptimality in criterion value versus iteration number for the cyclic (Algorithm 1) and parallel (Algorithm 2) backfitting methods, on a synthetic data set with $n = 2000$ and $d = 24$. On the left, iterations for the parallel method are counted as if “ideal” parallelization is used, where the d component updates are performed by d processors, at the total cost of one update, and on the right, iterations for the parallel method are counted as if “naive” serialization is used, where the component updates are performed in sequence. To avoid zeros on the y -axis (log scale), we added a small value to all the suboptimality values (dotted line).*

A.19 Simulated homogeneously-smooth data

Figure 7 shows the results of a homogeneous simulation, as in Section 5.1 and Figure 4, except that for the base component trends we used sinusoids of equal (and spatially-constant) frequency:

$$g_{0j}(x_j) = \sin(10\pi x_j), \quad j = 1, \dots, 10,$$

and we defined the component functions as $f_{0j} = a_j g_{0j} - b_j$, $j = 1, \dots, d$, where a_j, b_j were chosen to standardize f_{0j} (give it zero empirical mean and unit empirical norm), for $j = 1, \dots, d$.




Figure 7: Results from a simulation setup identical to that described in Section 5.1, i.e., identical to that used to produce Figure 4, except with homogeneous smoothness in the underlying component functions.

References

- Aleksandr Ivanovich Aptekarev, Sergey Denisov, and Dmitry Nikolaevich Tulyakov. On a problem by Steklov. *Journal of American Mathematical Society*, 29(4):1117–1165, 2016.
- Peter Bartlett, Olivier Bousquet, and Shahar Mendelson. Local Rademacher complexities. *Annals of Statistics*, 33(4):1497–1537, 2005.
- Heinz Bauschke and Jonathan Borwein. On projection algorithms for solving convex feasibility problems. *SIAM Review*, 38(3):367–426, 1996.
- M. Birman and M. Solomyak. Piecewise-polynomial approximations of functions of the classes W_p^α . *Mathematics of the USSR-Sbornik*, 73(115):331–335, 1967. In Russian.
- Steve Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternative direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- Leo Breiman and Jerome Friedman. Estimating optimal transformations for multiple regression and correlation. *Journal of the American Statistical Association*, 80(391):614–619, 1985.
- Andrews Buja, Trevor Hastie, and Robert Tibshirani. Linear smoothers and additive models. *Annals of Statistics*, 17(2):453–510, 1989.
- Carl de Boor. *A Practical Guide to Splines*. Springer, 1978.
- David L. Donoho and Iain M. Johnstone. Minimax estimation via wavelet shrinkage. *Annals of Statistics*, 26(8):879–921, 1998.
- Richard M. Dudley. The sizes of compact subsets of Hilbert space and continuity of Gaussian processes. *Journal of Functional Analysis*, 1(3):290–330, 1967.
- Bradley Efron. How biased is the apparent error rate of a prediction rule? *Journal of the American Statistical Association*, 81(394):461–470, 1986.

- Ludwig Fahrmeir and Stefan Lang. Bayesian inference for generalized additive mixed models based on Markov random field priors. *Journal of the Royal Statistical Society: Series C*, 50(2):201–220, 2001.
- Jerome Friedman and Werner Stuetzle. Projection pursuit regression. *Journal of the American Statistical Association*, 76(376):817–823, 1981.
- Chong Gu and Grace Wahba. Minimizing GCV/GML scores with multiple smoothing parameters via the Newton method. *SIAM Journal of Scientific and Statistical Computing*, 12(2):383–398, 1991.
- Asad Haris, Noah Simon, and Ali Shojaie. Generalized sparse additive models. 2018.
- Trevor Hastie. Non-parametric logistic regression. Technical Report, Stanford University, 1983.
- Trevor Hastie and Robert Tibshirani. *Generalized additive models*. Chapman & Hall, 1990.
- Daniel Hsu, Sham Kakde, and Tong Zhang. A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability*, 17(52):1–6, 2012.
- Seung-Jean Kim, Kwangmoo Koh, Stephen Boyd, and Dimitry Gorinevsky. ℓ_1 trend filtering. *SIAM Review*, 51(2):339–360, 2009.
- Young-Ju Kim and Chong Gu. Smoothing spline Gaussian regression: more scalable computation via efficient approximation. *Journal of the Royal Statistical Society: Series B*, 66(2):337–356, 2004.
- Andrey N. Kolmogorov and Vladimir M. Tikhomirov. ϵ -entropy and ϵ -capacity of sets in function spaces. *Uspekhi Matematicheskikh Nauk*, 14(2):3–86, 1959.
- Yi Lin and Hao Helen Zhang. Component selection and smoothing in multivariate nonparametric regression. *Annals of Statistics*, 34(5):2272–2297, 2006.
- Yin Lou, Jacob Bien, Rich Caruana, and Johannes Gehrke. Sparse partially linear additive models. *Journal of Computational and Graphical Statistics*, 25(4):1126–1140, 2016.
- Enno Mammen. Nonparametric regression under qualitative smoothness assumptions. *Annals of Statistics*, 19(2):741–759, 1991.
- Enno Mammen and Sara van de Geer. Locally adaptive regression splines. *Annals of Statistics*, 25(1):387–413, 1997.
- Lukas Meier, Sara van de Geer, and Peter Bühlmann. High-dimensional additive modeling. *Annals of Statistics*, 37(6):3779–3821, 2009.
- Masashi Okamoto. Distinctness of the eigenvalues of a quadratic form in a multivariate sample. *Annals of Statistics*, 1(4):763–765, 1973.
- Ashley Petersen, Daniela Witten, and Noah Simon. Fused lasso additive model. *Journal of Computational and Graphical Statistics*, 25(4):1005–1025, 2016.
- Aaditya Ramdas and Ryan J. Tibshirani. Fast and flexible ADMM algorithms for trend filtering. *Journal of Computational and Graphical Statistics*, 25(3):839–858, 2016.
- Garvesh Raskutti, Martin J. Wainwright, and Bin Yu. Minimax-optimal rates for sparse additive models over kernel classes via convex programming. *Journal of Machine Learning Research*, 13:389–427, 2012.
- Pradeep Ravikumar, John Lafferty, Han Liu, and Larry Wasserman. Sparse additive models. *Journal of the Royal Statistical Society: Series B*, 71(5):1009–1030, 2009.

- Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
- Havard Rue, Sara Martino, and Nicolas Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B*, 71(2):319–392, 2009.
- David Ruppert, Matt P. Wand, and Raymond J. Carroll. *Semiparametric Regression*. Cambridge University Press, 2003.
- Sylvain Sardy and Paul Tseng. AMlet, RAMlet and GAMlet: Automatic nonlinear fitting of additive model, robust and generalized with wavelets. *Journal of Computational and Graphical Statistics*, 13(2):283–309, 2004.
- Gabriel Steidl, Stephan Didas, and Julia Neumann. Splines in higher order TV regularization. *International Journal of Computer Vision*, 70(3):214–255, 2006.
- Charles Stone. Additive regression and other nonparametric models. *Annals of Statistics*, 13(2):689–705, 1985.
- Zhiqiang Tan and Cun-Hui Zhang. Penalized estimation in additive regression with high-dimensional data. arXiv: 1704.07229, 2017.
- Robert Tibshirani. Non-parametric estimation of relative risk. Technical Report, Stanford University, 1983.
- Robert Tibshirani, Michael Saunders, Saharon Rosset, Ji Zhu, and Keith Knight. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B*, 67(1):91–108, 2005.
- Ryan J. Tibshirani. The lasso problem and uniqueness. *Electronic Journal of Statistics*, 7:1456–1490, 2013.
- Ryan J. Tibshirani. Adaptive piecewise polynomial estimation via trend filtering. *Annals of Statistics*, 42(1):285–323, 2014.
- Ryan J. Tibshirani. Degrees of freedom and model search. *Statistica Sinica*, 25(3):1265–1296, 2015.
- Ryan J. Tibshirani. Dykstra’s algorithm, ADMM, and coordinate descent: Connections, insights, and extensions. *Advances in Neural Information Processing Systems*, 30, 2017.
- Ryan J. Tibshirani and Jonathan Taylor. The solution path of the generalized lasso. *Annals of Statistics*, 39(3):1335–1371, 2011.
- Ryan J. Tibshirani and Jonathan Taylor. Degrees of freedom in lasso problems. *Annals of Statistics*, 40(2):1198–1232, 2012.
- Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.
- Paul Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications*, 109(3):475–494, 2001.
- Sara van de Geer. Estimating a regression function. *Annals of Statistics*, 18(2):907–924, 1990.
- Sara van de Geer. *Empirical Processes in M-Estimation*. Cambridge University Press, 2000.

- Sara van de Geer. On the uniform convergence of empirical norms and inner products, with application to causal inference. *Electronic Journal of Statistics*, 8:543–574, 2014.
- Eeke van der Burg and Jan de Leeuw. Non-linear canonical correlation. *British Journal of Mathematical and Statistical Psychological*, 36(1):54–80, 1983.
- Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic View*. Cambridge University Press, 2017. To appear.
- Yu-Xiang Wang, Alexander Smola, and Ryan J. Tibshirani. The falling factorial basis and its statistical applications. *International Conference on Machine Learning*, 31, 2014.
- Simon N. Wood. Modelling and smoothing parameter estimation with multiple quadratic penalties. *Journal of the Royal Statistical Society: Series B*, 62(2):413–428, 2000.
- Simon N. Wood. Stable and efficient multiple smoothing parameter estimation for generalized additive models. *Journal of the American Statistical Association*, 99(467):673–686, 2004.
- Simon N. Wood. Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B*, 73(1):3–36, 2011.
- Simon N. Wood. *Generalized Additive Models: An Introduction with R*. Chapman & Hall/CRC Press, 2017.
- Simon N. Wood, Yannig Goude, and Simon Shaw. Generalized additive models for large data sets. *Journal of the Royal Statistical Society: Series C*, 64(1):139–155, 2015.
- Simon N. Wood, Natalya Pya, and Benjamin Säfken. Smoothing parameter and model selection for general smooth models. *Journal of the American Statistical Association*, 111(516):1548–1575, 2016.
- Yuhong Yang and Andrew Barron. Information-theoretic determination of minimax rates of convergence. *Annals of Statistics*, 27(5):1564–1599, 1999.
- Shuanglin Zhang and Man-Yu Wong. Wavelet threshold estimation for additive regression models. *Annals of Statistics*, 31(1):152–173, 2003.