

Total Variation Classes Beyond 1d: Minimax Rates, and the Limitations of Linear Smoothers

Veeranjaneyulu Sadhanala*

Yu-Xiang Wang*

Ryan J. Tibshirani

Carnegie Mellon University
 (*These authors contributed equally)

Abstract

We consider the problem of estimating a function defined over n locations on a d -dimensional grid (having all side lengths equal to $n^{1/d}$). When the function is constrained to have discrete total variation bounded by C_n , we derive the minimax optimal (squared) ℓ_2 estimation error rate, parametrized by n and C_n . Total variation denoising, also known as the fused lasso, is seen to be rate optimal. Several simpler estimators exist, such as Laplacian smoothing and Laplacian eigenmaps. A natural question is: can these simpler estimators perform just as well? We prove that these estimators, and more broadly all estimators given by linear transformations of the input data, are suboptimal over the class of functions with bounded variation. This extends fundamental findings of [Donoho and Johnstone \[1998\]](#) on 1-dimensional total variation spaces to higher dimensions. The implication is that the computationally simpler methods cannot be used for such sophisticated denoising tasks, without sacrificing statistical accuracy. We also derive minimax rates for discrete Sobolev spaces over d -dimensional grids, which are, in some sense, smaller than the total variation function spaces. Indeed, these are small enough spaces that linear estimators can be optimal—and a few well-known ones are, such as Laplacian smoothing and Laplacian eigenmaps, as we show. Lastly, we investigate the problem of adaptivity of the total variation denoiser to these smaller Sobolev function spaces.

1 Introduction

Let $G = (V, E)$ be a d -dimensional grid graph, i.e., a lattice graph, with equal side lengths. Label the nodes as $V = \{1, \dots, n\}$, and edges as $E = \{e_1, \dots, e_m\}$. Consider data $y = (y_1, \dots, y_n) \in \mathbb{R}^n$ observed over the nodes, from a model

$$y_i \sim N(\theta_{0,i}, \sigma^2), \quad \text{i.i.d., for } i = 1, \dots, n, \quad (1)$$

where $\theta_0 = (\theta_{0,1}, \dots, \theta_{0,n}) \in \mathbb{R}^n$ is an unknown mean parameter to be estimated, and $\sigma^2 > 0$ is the marginal noise variance. It is assumed that θ_0 displays some kind of regularity over the grid G , e.g., $\theta_0 \in \mathcal{T}_d(C_n)$ for some $C_n > 0$, where

$$\mathcal{T}_d(C_n) = \{\theta : \|D\theta\|_1 \leq C_n\}, \quad (2)$$

and $D \in \mathbb{R}^{m \times n}$ is the edge incidence matrix of G . This has ℓ th row $D_\ell = (0, \dots, -1, \dots, 1, \dots, 0)$, with a -1 in the i th location, and 1 in the j th location, provided that the ℓ th edge is $e_\ell = (i, j)$ with $i < j$. Equivalently, $L = D^T D$ is the graph Laplacian matrix of G , and thus

$$\|D\theta\|_1 = \sum_{(i,j) \in E} |\theta_i - \theta_j|, \quad \text{and} \quad \|D\theta\|_2^2 = \theta^T L \theta = \sum_{(i,j) \in E} (\theta_i - \theta_j)^2.$$

We will refer to the class in (2) as a *discrete total variation (TV) class*, and to the quantity $\|D\theta_0\|_1$ as the discrete total variation of θ_0 , though for simplicity we will often drop the word “discrete”.

The problem of estimating θ_0 given a total variation bound as in (2) is of great importance in both nonparametric statistics and signal processing, and has many applications, e.g., changepoint detection for 1d grids, and image denoising for 2d and 3d grids. There has been much methodological and computational work devoted to this problem, resulting in practically efficient estimators in dimensions 1, 2, 3, and beyond. However, theoretical performance, and in particular optimality, is only really well-understood in the 1-dimensional setting. This paper seeks to change that, and offers theory in d -dimensions that parallel more classical results known in the 1-dimensional case.

Estimators under consideration. Central role to our work is the *total variation (TV) denoising* or *fused lasso* estimator (e.g., Rudin et al. [1992], Vogel and Oman [1996], Tibshirani et al. [2005], Chambolle and Darbon [2009], Hoefling [2010], Tibshirani and Taylor [2011], Sharpnack et al. [2012], Barbero and Sra [2014]), defined by the convex optimization problem

$$\hat{\theta}^{\text{TV}} = \underset{\theta \in \mathbb{R}^n}{\operatorname{argmin}} \|y - \theta\|_2^2 + \lambda \|D\theta\|_1, \quad (3)$$

where $\lambda \geq 0$ is a tuning parameter. Another pair of methods that we study carefully are *Laplacian smoothing* and *Laplacian eigenmaps*, which are most commonly seen in the context of clustering, dimensionality reduction, and semi-supervised learning, but are also useful tools for estimation in a regression setting like ours (e.g., Belkin and Niyogi [2002, 2003], Smola and Kondor [2003], Zhu et al. [2003], Belkin and Niyogi [2004], Zhou et al. [2005], Belkin et al. [2005], Belkin and Niyogi [2005], Ando and Zhang [2006], Sharpnack and Singh [2010]). The Laplacian smoothing estimator is given by

$$\hat{\theta}^{\text{LS}} = \underset{\theta \in \mathbb{R}^n}{\operatorname{argmin}} \|y - \theta\|_2^2 + \lambda \|D\theta\|_2^2, \quad \text{i.e.,} \quad \hat{\theta}^{\text{LS}} = (I + \lambda L)^{-1}y, \quad (4)$$

for a tuning parameter $\lambda \geq 0$, where in the second expression we have written $\hat{\theta}^{\text{LS}}$ in closed-form (this is possible since it is the minimizer of a convex quadratic). For Laplacian eigenmaps, we must introduce the eigendecomposition of the graph Laplacian, $L = V\Sigma V^T$, where $\Sigma = \operatorname{diag}(\rho_1, \dots, \rho_n)$ with $0 = \rho_1 < \rho_2 \leq \dots \leq \rho_n$, and where $V = [V_1, V_2, \dots, V_n] \in \mathbb{R}^{n \times n}$ has orthonormal columns. The Laplacian eigenmaps estimator is

$$\hat{\theta}^{\text{LE}} = V_{[k]}V_{[k]}^T y, \quad \text{where} \quad V_{[k]} = [V_1, V_2, \dots, V_k] \in \mathbb{R}^{n \times k}, \quad (5)$$

where now $k \in \{1, \dots, n\}$ acts as a tuning parameter.

Laplacian smoothing and Laplacian eigenmaps are appealing because they are (relatively) simple: they are just linear transformations of the data y . Indeed, as we are considering G to be a grid, both estimators in (4), (5) can be computed very quickly, in nearly $O(n)$ time, since the columns of V here are discrete cosine transform (DCT) basis vectors when $d = 1$, or Kronecker products thereof, when $d \geq 2$ (e.g., Conte and de Boor [1980], Godunov and Ryabenkii [1987], Kunsch [1994], Ng et al. [1999], Wang et al. [2008]). The TV denoising estimator in (3), on the other hand, cannot be expressed in closed-form, and is much more difficult to compute, especially when $d \geq 2$, though several advances have been made over the years (see the references above, and in particular Barbero and Sra [2014] for an efficient operator-splitting algorithm and nice literature survey). Importantly, these computational difficulties are often worth it: TV denoising often practically outperforms ℓ_2 -regularized estimators like Laplacian smoothing (and also Laplacian eigenmaps) in image denoising tasks, as it is able to better preserve sharp edges and object boundaries (this is now widely accepted, early references are, e.g., Acar and Vogel [1994], Dobson and Santosa [1996], Chambolle and Lions [1997]). See Figure 1 for an example, using the often-studied ‘‘cameraman’’ image.

In the 1d setting, classical theory from nonparametric statistics draws a clear distinction between the performance of TV denoising and estimators like Laplacian smoothing and Laplacian eigenmaps. Perhaps surprisingly, this theory has not yet been fully developed in dimensions $d \geq 2$. Arguably, the comparison

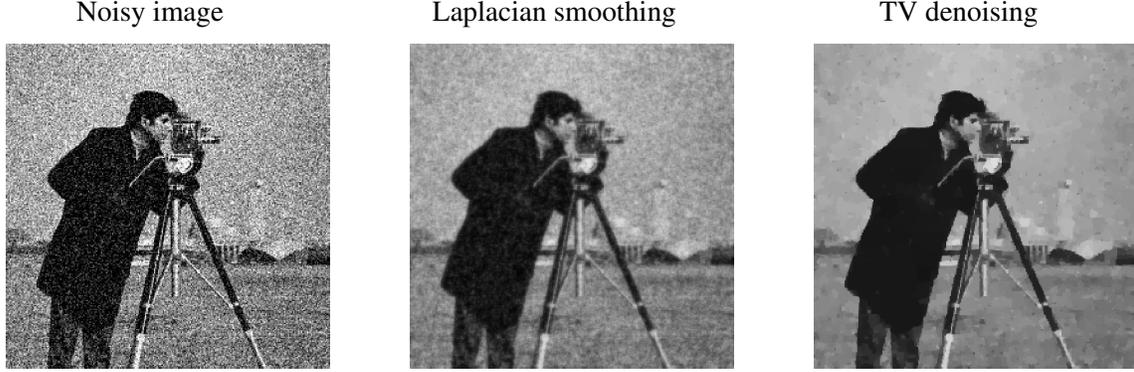


Figure 1: Comparison of Laplacian smoothing and TV denoising for the common “cameraman” image. TV denoising provides a more visually appealing result, and also achieves about a 35% reduction in MSE compared to Laplacian smoothing (MSE being measured to the original image). Both methods were tuned optimally.

between TV denoising and Laplacian smoothing and Laplacian eigenmaps is even more interesting in higher dimensions, because the computational gap between the methods is even larger (the former method being much more expensive, say in 2d and 3d, than the latter two). Shortly, we review the 1d theory, and what is known in d -dimensions, for $d \geq 2$. First, we introduce notation.

Notation. For deterministic (nonrandom) sequences a_n, b_n we write $a_n = O(b_n)$ to denote that a_n/b_n is upper bounded for all n large enough, and $a_n \asymp b_n$ to denote that both $a_n = O(b_n)$ and $a_n^{-1} = O(b_n^{-1})$. Also, for random sequences A_n, B_n , we write $A_n = O_{\mathbb{P}}(B_n)$ to denote that A_n/B_n is bounded in probability. We abbreviate $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. For an estimator $\hat{\theta}$ of the parameter θ_0 in (1), we define its mean squared error (MSE) to be

$$\text{MSE}(\hat{\theta}, \theta_0) = \frac{1}{n} \|\hat{\theta} - \theta_0\|_2^2.$$

The risk of $\hat{\theta}$ is the expectation of its MSE, and for a set $\mathcal{K} \subseteq \mathbb{R}^n$, we define the minimax risk and minimax linear risk to be

$$R(\mathcal{K}) = \inf_{\hat{\theta}} \sup_{\theta_0 \in \mathcal{K}} \mathbb{E}[\text{MSE}(\hat{\theta}, \theta_0)] \quad \text{and} \quad R_L(\mathcal{K}) = \inf_{\hat{\theta} \text{ linear}} \sup_{\theta_0 \in \mathcal{K}} \mathbb{E}[\text{MSE}(\hat{\theta}, \theta_0)],$$

respectively, where the infimum on in the first expression is over all estimators $\hat{\theta}$, and in the second expression over all *linear estimators* $\hat{\theta}$, meaning that $\hat{\theta} = Sy$ for a matrix $S \in \mathbb{R}^{n \times n}$. We will also refer to linear estimators as *linear smoothers*. Note that both Laplacian smoothing in (4) and Laplacian eigenmaps in (5) are linear smoothers, but TV denoising in (3) is not. Lastly, in somewhat of an abuse of nomenclature, we will often call the parameter θ_0 in (1) a function, and a set of possible values for θ_0 as in (2) a function space; this comes from thinking of the components of θ_0 as the evaluations of an underlying function over n locations on the grid. This embedding has no formal importance, but it is convenient notationally, and matches the notation in nonparametric statistics.

Review: TV denoising in 1d. The classical nonparametric statistics literature [Donoho et al., 1990, Donoho and Johnstone, 1998, Mammen and van de Geer, 1997] provides a more or less complete story for estimation under total variation constraints in 1d. See also Tibshirani [2014] for a translation of these results to a setting more consistent (notationally) to that in the current paper. Assume that $d = 1$ and $C_n = C > 0$, a constant (not growing with n). The results in Donoho and Johnstone [1998] imply that

$$R(\mathcal{T}_1(C)) \asymp n^{-2/3}. \tag{6}$$

Further, [Mammen and van de Geer \[1997\]](#) showed that the TV denoiser $\hat{\theta}^{\text{TV}}$ in (3), with $\lambda \asymp n^{1/3}$, satisfies

$$\text{MSE}(\hat{\theta}^{\text{TV}}, \theta_0) = O_{\mathbb{P}}(n^{-2/3}), \quad (7)$$

for all $\theta_0 \in \mathcal{T}_1(C)$, and is thus minimax rate optimal over $\mathcal{T}_1(C)$. (In assessing rates here and throughout, we do not distinguish between convergence in expectation versus convergence in probability.) Wavelet denoising, under various choices of wavelet bases, also achieves the minimax rate. However, many simpler estimators do not. To be more precise, it is shown in [Donoho and Johnstone \[1998\]](#) that

$$R_L(\mathcal{T}_1(C)) \asymp n^{-1/2}. \quad (8)$$

Therefore, a substantial number of commonly used nonparametric estimators—such as running mean estimators, smoothing splines, kernel smoothing, Laplacian smoothing, and Laplacian eigenmaps, which are all linear smoothers—have a major deficiency when it comes to estimating functions of bounded variation. Roughly speaking, they will require many more samples to estimate θ_0 within the same degree of accuracy as an optimal method like TV or wavelet denoising (on the order of $\epsilon^{-1/2}$ times more samples to achieve an MSE of ϵ). Further theory and empirical examples (e.g., [Donoho and Johnstone \[1994a, 1998\]](#), [Tibshirani \[2014\]](#)) offer the following perspective: linear smoothers cannot cope with functions in $T(C)$ that have spatially inhomogeneous smoothness, i.e., that vary smoothly at some locations and vary wildly at others. Linear smoothers can only produce estimates that are smooth throughout, or wiggly throughout, but not a mix of the two. They can hence perform well over smaller, more homogeneous function classes like Sobolev or Holder classes, but not larger ones like total variation classes (or more generally, Besov and Triebel classes), and for these, one must use more sophisticated, nonlinear techniques. A motivating question: does such a gap persist in higher dimensions, between optimal nonlinear and linear estimators, and if so, how big is it?

Review: TV denoising in multiple dimensions. Recently, [Wang et al. \[2016\]](#) established rates for TV denoising over various graph models, including grids, and [Hutter and Rigollet \[2016\]](#) made improvements, particularly in the case of d -dimensional grids with $d \geq 2$. We can combine Propositions 4 and 6 of [Hutter and Rigollet \[2016\]](#) with Theorem 3 of [Wang et al. \[2016\]](#) to give the following result: if $d \geq 2$, and C_n is an arbitrary sequence (potentially unbounded with n), then the TV denoiser $\hat{\theta}^{\text{TV}}$ in (3) satisfies, over all $\theta_0 \in \mathcal{T}_d(C_n)$,

$$\text{MSE}(\hat{\theta}^{\text{TV}}, \theta_0) = O_{\mathbb{P}}\left(\frac{C_n \log n}{n}\right) \text{ for } d = 2, \text{ and } \text{MSE}(\hat{\theta}^{\text{TV}}, \theta_0) = O_{\mathbb{P}}\left(\frac{C_n \sqrt{\log n}}{n}\right) \text{ for } d \geq 3, \quad (9)$$

with $\lambda \asymp \log n$ for $d = 2$, and $\lambda \asymp \sqrt{\log n}$ for $d \geq 3$. Note that, at first glance, this is a very different result from the 1d case. We expand on this next.

2 Summary of results

A gap in multiple dimensions. For estimation of θ_0 in (1) when $d \geq 2$, consider, e.g., the simplest possible linear smoother: the mean estimator, $\hat{\theta}^{\text{mean}} = \bar{y}\mathbf{1}$ (where $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^n$, the vector of all 1s). Lemma 4, given below, implies that over $\theta_0 \in \mathcal{T}_d(C_n)$, the MSE of the mean estimator is bounded in probability by $C_n^2 \log n/n$ for $d = 2$, and C_n^2/n for $d \geq 3$. Compare this to (9). When $C_n = C > 0$ is a constant, i.e., when the TV of θ_0 is assumed to be bounded (which is assumed for the 1d results in (6), (7), (8)), this means that the TV denoiser and the mean estimator converge to θ_0 *at the same rate*, basically (ignoring log terms), the “parametric rate” of $1/n$, for estimating a finite-dimensional parameter! That TV denoising and such a trivial linear smoother perform comparably over 2d and 3d grids could not be farther from the story in 1d, where TV denoising is separated by an unbridgeable gap from *all* linear smoothers, as shown in (6), (7), (8).

Our results in Section 3 clarify this conundrum, and can be summarized by three points.

- We argue in Section 3.1 that there is a proper “canonical” scaling for the TV class defined in (2). E.g., when $d = 1$, this yields $C_n \asymp 1$, a constant, but when $d = 2$, this yields $C_n \asymp \sqrt{n}$, and C_n also diverges with n for all $d \geq 3$. Sticking with $d = 2$ as an interesting example, we see that under such a scaling, the MSE rates achieved by TV denoising and the mean estimator respectively, are drastically different; ignoring log terms, these are

$$\frac{C_n}{n} \asymp \frac{1}{\sqrt{n}} \quad \text{and} \quad \frac{C_n^2}{n} \asymp 1, \quad (10)$$

respectively. Hence, TV denoising has an MSE rate of $1/\sqrt{n}$, in a setting where the mean estimator has a *constant* rate, i.e., a setting where it is not even known to be consistent.

- We show in Section 3.3 that our choice to study the mean estimator here is not somehow “unlucky” (it is not a particularly bad linear smoother, nor is the upper bound on its MSE loose): the minimax linear risk over $\mathcal{T}_d(C_n)$ is on the order C_n^2/n , for all $d \geq 2$. Thus, even the best linear smoothers have the same poor performance as the mean over $\mathcal{T}_d(C_n)$.
- We show in Section 3.2 that the TV estimator is (essentially) minimax optimal over $\mathcal{T}_d(C_n)$, as the minimax risk over this class scales as C_n/n (ignoring log terms).

To summarize, these results reveal a significant gap between linear smoothers and optimal estimators like TV denoising, for estimation over $\mathcal{T}_d(C_n)$ in d dimensions, with $d \geq 2$, as long as C_n scales appropriately. Roughly speaking, the TV classes encompass a challenging setting for estimation because they are very broad, containing a wide array of functions—both globally smooth functions, said to have homogeneous smoothness, and functions with vastly different levels of smoothness at different grid locations, said to have heterogeneous smoothness. Linear smoothers cannot handle heterogeneous smoothness, and only nonlinear methods can enjoy good estimation properties over the entirety of $\mathcal{T}_d(C_n)$. To reiterate, a telling example is $d = 2$ with the canonical scaling $C_n \asymp \sqrt{n}$, where we see that TV denoising achieves the optimal $1/\sqrt{n}$ rate (up to log factors), meanwhile, the best linear smoothers have max risk that is constant over $\mathcal{T}_2(\sqrt{n})$. See Figure 2 for an illustration.

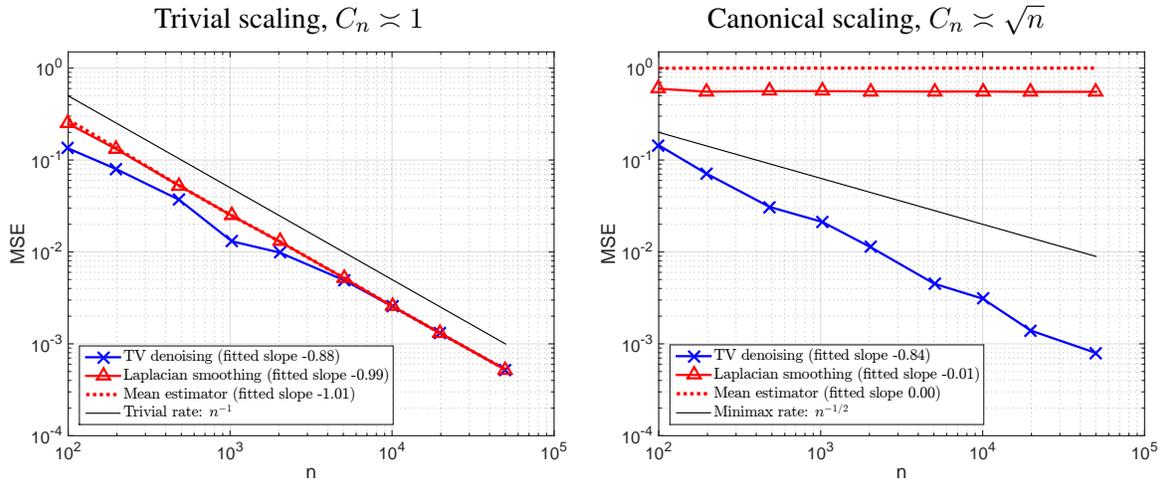


Figure 2: MSE curves for estimation over a 2d grid, under two very different scalings of C_n : constant and \sqrt{n} . The parameter θ_0 was a “one-hot” signal, with all but one component equal to 0. For each n , the results were averaged over 5 repetitions, and Laplacian smoothing and TV denoising were tuned for optimal average MSE.

Minimax rates over smaller function spaces, and adaptivity. Sections 4 and 5 are focused on different function spaces, discrete Sobolev spaces, which are ℓ_2 analogs of discrete TV spaces as we have defined them in (2). Under the canonical scaling of Section 3.1, Sobolev spaces are contained in TV spaces, and the former can be roughly thought of as containing functions of more homogeneous smoothness. The story now is more optimistic for linear smoothers, and the following is a summary.

- In Section 4, we derive minimax rates for Sobolev spaces, and prove that linear smoothers—in particular, Laplacian smoothing and Laplacian eigenmaps—are optimal over these spaces.
- In Section 5, we discuss an interesting phenomenon, a phase transition of sorts, at $d = 3$ dimensions. When $d = 1$ or 2 , the minimax rates for a TV space and its inscribed Sobolev space match; when $d \geq 3$, they do not, and the inscribed Sobolev space has a faster minimax rate. Aside from being an interesting statement about the TV and Sobolev function spaces in high dimensions, this raises an important question of adaptivity over the smaller Sobolev function spaces. As the minimax rates match for $d = 1$ and 2 , any method optimal over TV spaces in these dimensions, such as TV denoising, is automatically optimal over the inscribed Sobolev spaces. But the question remains open for $d \geq 3$ —does, e.g., TV denoising adapt to the faster minimax rate over Sobolev spaces? We present empirical evidence to suggest that this may be true, and leave a formal study to future work.

Other considerations and extensions. There are many problems related to the one that we study in this paper. Clearly, minimax rates for the TV and Sobolev classes over general graphs, not just d -dimensional grids, are of interest. Our minimax lower bounds for TV classes actually apply to generic graphs with bounded max degree, though it is unclear whether to what extent they are sharp beyond grids; a detailed study will be left to future work. Another related topic is that of higher-order smoothness classes, i.e., classes containing functions whose *derivatives* are (say) of bounded variation. The natural extension of TV denoising here is called *trend filtering*, defined via the regularization of discrete higher-order derivatives. In the 1d setting, minimax rates, the optimality of trend filtering, and the suboptimality of linear smoothers is already well-understood [Tibshirani, 2014]. Trend filtering has been defined and studied to some extent on general graphs [Wang et al., 2016], but no notions of optimality have been investigated beyond 1d. This will also be left to future work. Lastly, it is worth mentioning that there are other estimators (i.e., other than the ones we study in detail) that attain or nearly attain minimax rates over various classes we consider in this paper. E.g., wavelet denoising is known to be optimal over TV classes in 1d [Donoho and Johnstone, 1998]; and comparing recent upper bounds from Needell and Ward [2013], Hutter and Rigollet [2016] with the lower bounds in this work, we see that wavelet denoising is also nearly minimax in 2d (ignoring log terms).

3 Analysis over TV classes

3.1 Canonical scaling for TV and Sobolev classes

We start by establishing what we call a “canonical” scaling for the radius C_n of the TV ball $\mathcal{T}_d(C_n)$ in (2), as well as the radius C'_n of the Sobolev ball $\mathcal{S}_d(C'_n)$, defined as

$$\mathcal{S}_d(C'_n) = \{\theta : \|D\theta\|_2 \leq C'_n\}. \quad (11)$$

Proper scalings for C_n, C'_n will be critical for properly interpreting our new results in d dimensions, in a way that is comparable to known results for $d = 1$ (which are usually stated in terms of the 1d scalings $C_n \asymp 1, C'_n \asymp 1/\sqrt{n}$). To study (2), (11), it helps to introduce a third function space,

$$\mathcal{H}_d(1) = \left\{ \theta : \theta_i = f(i_1/\ell \dots, i_d/\ell), i = 1, \dots, n, \text{ for some } f \in \mathcal{H}_d^{\text{cont}}(1) \right\}. \quad (12)$$

Above, we have mapped each location i on the grid to a multi-index $(i_1, \dots, i_d) \in \{1, \dots, \ell\}^d$, where $\ell = n^{1/d}$, and $\mathcal{H}_d^{\text{cont}}(1)$ denotes the (usual) continuous Holder space on $[0, 1]^d$, i.e., functions that are 1-Lipschitz with respect to the ℓ_∞ norm. We seek an embedding that is analogous to the embedding of continuous Holder, Sobolev, and total variation spaces in 1d functional analysis, namely,

$$\mathcal{H}_d(1) \subseteq \mathcal{S}_d(C'_n) \subseteq \mathcal{T}_d(C_n). \quad (13)$$

Our first lemma provides a choice of C_n, C'_n that makes the above true. Its proof, as with all proofs in this paper, can be found in the appendix.

Lemma 1. *For $d \geq 1$, the embedding in (13) holds with choices $C_n \asymp n^{1-1/d}$ and $C'_n \asymp n^{1/2-1/d}$. Such choices are called the canonical scaling for the function classes in (2), (11).*

As a sanity check, both the (usual) continuous Holder and Sobolev function spaces in d dimensions are known to have minimax risks that scale as $n^{-2/(2+d)}$, in a standard nonparametric regression setup (e.g., Györfi et al. [2002]). Under the canonical scaling $C'_n \asymp n^{1/2-1/d}$, our results in Section 4 show that the discrete Sobolev class $\mathcal{S}_d(n^{1/2-1/d})$ also admits a minimax rate of $n^{-2/(2+d)}$.

3.2 Minimax rates over TV classes

The following is a lower bound for the minimax risk of the TV class $\mathcal{T}_d(C_n)$ in (2).

Theorem 2. *Assume $n \geq 2$, and denote $d_{\max} = 2d$. Then, for constants $c > 0$, $\rho_1 \in (2.34, 2.35)$,*

$$R(\mathcal{T}_d(C_n)) \geq c \cdot \begin{cases} \frac{\sigma C_n \sqrt{1 + \log(\sigma d_{\max} n / C_n)}}{d_{\max} n} & \text{if } C_n \in [\sigma d_{\max} \sqrt{\log n}, \sigma d_{\max} n / \sqrt{\rho_1}] \\ C_n^2 / (d_{\max}^2 n) \vee \sigma^2 / n & \text{if } C_n < \sigma d_{\max} \sqrt{\log n} \\ \sigma^2 / \rho_1 & \text{if } C_n > \sigma d_{\max} n / \sqrt{\rho_1} \end{cases}. \quad (14)$$

The proof uses a simplifying reduction of the TV class, via $\mathcal{T}_d(C_n) \supseteq B_1(C_n / d_{\max})$, the latter set denoting the ℓ_1 ball of radius C_n / d_{\max} in \mathbb{R}^n . It then invokes a sharp characterization of the minimax risk in normal means problems over ℓ_p balls due to Birge and Massart [2001]. Several remarks are in order.

Remark 1. The first line on the right-hand side in (14) often provides the most useful lower bound. To see this, recall that under the canonical scaling for TV classes, we have $C_n = n^{1-1/d}$. For all $d \geq 2$, this certainly implies $C_n \in [\sigma d_{\max} \sqrt{\log n}, \sigma d_{\max} n / \sqrt{\rho_1}]$, for large n .

Remark 2. Even though its construction is very simple, the lower bound on the minimax risk in (14) is sharp or nearly sharp in many interesting cases. Assume that $C_n \in [\sigma d_{\max} \sqrt{\log n}, \sigma d_{\max} n / \sqrt{\rho_1}]$. The lower bound rate is $C_n \sqrt{\log(n/C_n)} / n$. When $d = 2$, we see that this is very close to the upper bound rate of $C_n \log n / n$ achieved by the TV denoiser, as stated in (9). These two differ by at most a $\log n$ factor (achieved when $C_n \asymp n$). When $d \geq 3$, we see that the lower bound rate is even closer to the upper bound rate of $C_n \sqrt{\log n} / n$ achieved by the TV denoiser, as in (9). These two now differ by at most a $\sqrt{\log n}$ factor (again achieved when $C_n \asymp n$). We hence conclude that the TV denoiser is essentially minimax optimal in all dimensions $d \geq 2$.

Remark 3. When $d = 1$, and (say) $C_n \asymp 1$, the lower bound rate of $\sqrt{\log n} / n$ given by Theorem 2 is not sharp; we know from Donoho and Johnstone [1998] (recall (6)) that the minimax rate over $\mathcal{T}_1(1)$ is $n^{-2/3}$. The result in the theorem (and also Theorem 3) in fact holds more generally, beyond grids: for an arbitrary graph G , its edge incidence matrix D , and $\mathcal{T}_d(C_n)$ as defined in (2), the result holds for d_{\max} equal to the max degree of G . It is unclear to what extent this is sharp, for different graph models.

3.3 Minimax linear rates over TV classes

We now turn to a lower bound on the minimax linear risk of the TV class $\mathcal{T}_d(C_n)$ in (2).

Theorem 3. *Recall the notation $d_{\max} = 2d$. Then*

$$R_L(\mathcal{T}_d(C_n)) \geq \frac{\sigma^2 C_n^2}{C_n^2 + \sigma^2 d_{\max}^2 n} \vee \frac{\sigma^2}{n} \geq \frac{1}{2} \left(\frac{C_n^2}{d_{\max}^2 n} \wedge \sigma^2 \right) \vee \frac{\sigma^2}{n}. \quad (15)$$

The proof relies on an elegant meta-theorem on minimax rates from [Donoho et al. \[1990\]](#), which uses the concept of a “quadratically convex” set, whose minimax linear risk is the same as that of its hardest rectangular subproblem. An alternative proof can be given entirely from first principles. See the appendix.

Remark 4. When C_n^2 grows with n , but not too fast (scales as \sqrt{n} , at most), the lower bound rate in (15) will be C_n^2/n . Compared to the C_n/n minimax rate from [Theorem 2](#) (ignoring log terms), we see a clear gap between optimal nonlinear and linear estimators. In fact, under the canonical scaling $C_n \asymp n^{1-1/d}$, for any $d \geq 2$, this gap is seemingly huge: the lower bound for the minimax linear rate will be a constant, whereas the minimax rate from [Theorem 2](#) (ignoring log terms) will be $n^{-1/d}$.

We now show that the lower bound in [Theorem 3](#) is essentially tight, and remarkably, it is certified by analyzing two trivial linear estimators: the mean estimator and the identity estimator.

Lemma 4. *Let M_n denote the largest column norm of D^\dagger . For the mean estimator $\hat{\theta}^{\text{mean}} = \bar{y}\mathbf{1}$,*

$$\sup_{\theta_0 \in \mathcal{T}_d(C_n)} \mathbb{E}[\text{MSE}(\hat{\theta}^{\text{mean}}, \theta_0)] \leq \frac{\sigma^2 + C_n^2 M_n^2}{n},$$

From [Proposition 4](#) in [Hutter and Rigollet \[2016\]](#), we have $M_n = O(\sqrt{\log n})$ when $d = 2$ and $M_n = O(1)$ when $d \geq 3$.

The risk of the identity estimator $\hat{\theta}^{\text{id}} = y$ is clearly σ^2 . Combining this logic with [Lemma 4](#) gives the upper bound $R_L(\mathcal{T}_d(C_n)) \leq (\sigma^2 + C_n^2 M_n^2)/n \wedge \sigma^2$. Comparing this with the lower bound described in [Remark 4](#), we see that the two rates basically match, modulo the M_n^2 factor in the upper bound, which only provides an extra $\log n$ factor when $d = 2$. The takeaway message: in the sense of max risk, the best linear smoother does not perform much better than the trivial estimators.

Additional empirical experiments, similar to those shown in [Figure 2](#), are given in the appendix.

4 Analysis over Sobolev classes

Our first result here is a lower bound on the minimax risk of the Sobolev class $\mathcal{S}_d(C'_n)$ in (11).

Theorem 5. *For a universal constant $c > 0$,*

$$R(\mathcal{S}_d(C'_n)) \geq \frac{c}{n} \left((n\sigma^2)^{\frac{2}{d+2}} (C'_n)^{\frac{2d}{d+2}} \wedge n\sigma^2 \wedge n^{2/d} (C'_n)^2 \right) + \frac{\sigma^2}{n}.$$

Elegant tools for minimax analysis from [Donoho et al. \[1990\]](#), which leverage the fact that the ellipsoid $\mathcal{S}_d(C'_n)$ is orthosymmetric and quadratically convex (after a rotation), are used to prove the result.

The next theorem gives upper bounds, certifying that the above lower bound is tight, and showing that Laplacian eigenmaps and Laplacian smoothing, both linear smoothers, are optimal over $\mathcal{S}_d(C'_n)$ for all d and for $d = 1, 2$, or 3 respectively.

Function class	Dimension 1	Dimension 2	Dimension $d \geq 3$
TV ball $\mathcal{T}_d(n^{1-1/d})$	$n^{-2/3}$	$n^{-1/2} \sqrt{\log n}$	$n^{-1/d} \sqrt{\log n}$
Sobolev ball $\mathcal{S}_d(n^{1/2-1/d})$	$n^{-2/3}$	$n^{-1/2}$	$n^{-\frac{2}{2+d}}$

Table 1: Summary of rates for canonically-scaled TV and Sobolev spaces.

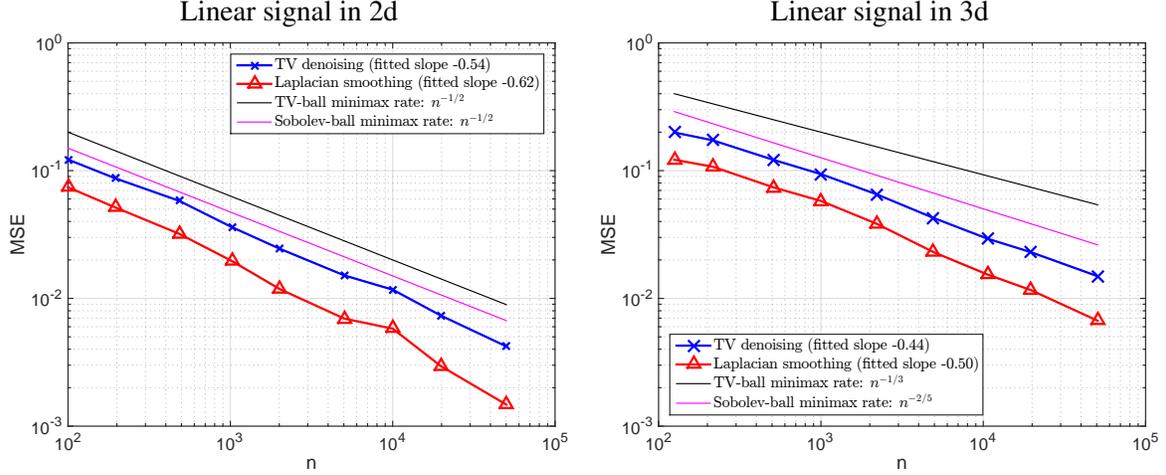


Figure 3: MSE curves for estimating a “linear” signal, a very smooth signal, over 2d and 3d grids. For each n , the results were averaged over 5 repetitions, and Laplacian smoothing and TV denoising were tuned for best average MSE performance. The signal was set to satisfy $\|D\theta_0\|_2 \asymp n^{1/2-1/d}$, matching the canonical scaling.

Theorem 6. For Laplacian eigenmaps, $\hat{\theta}^{\text{LE}}$ in (5), with $k \asymp ((n(C'_n)^d)^{2/(d+2)} \vee 1) \wedge n$, we have

$$\sup_{\theta_0 \in \mathcal{S}_d(C'_n)} \mathbb{E}[\text{MSE}(\hat{\theta}^{\text{LE}}, \theta_0)] \leq \frac{c}{n} \left((n\sigma^2)^{\frac{2}{d+2}} (C'_n)^{\frac{2d}{d+2}} \wedge n\sigma^2 \wedge n^{2/d} (C'_n)^2 \right) + \frac{c\sigma^2}{n},$$

for a universal constant $c > 0$, and n large enough. When $d = 1, 2$, or 3 , the same bound holds for Laplacian smoothing $\hat{\theta}^{\text{LS}}$ in (5), with $\lambda \asymp (n/(C'_n)^2)^{2/(d+2)}$ (and a possibly different constant c).

Remark 5. As shown in the proof, Laplacian smoothing is nearly minimax rate optimal over $\mathcal{S}_d(C'_n)$ when $d = 4$, just incurring an extra log factor. It is unclear to us whether this method is still (nearly) optimal when $d \geq 5$; based on insights from our proof technique, we conjecture that it is not.

5 A phase transition, and adaptivity

The TV and Sobolev classes in (2) and (11), respectively, display a curious relationship. We reflect on Theorems 2 and 5, under the canonical scaling with $C_n \asymp n^{1-1/d}$ and $C'_n \asymp n^{1/2-1/d}$, that, recall, guarantees $\mathcal{S}_d(C'_n) \subseteq \mathcal{T}_d(C_n)$. (This is done for concreteness; statements could also be made outside of this case, subject to an appropriate relationship with $C_n/C'_n \asymp \sqrt{n}$.) When $d = 1$, both the TV and Sobolev classes have a minimax rate of $n^{-2/3}$ (this TV result is actually due to Donoho and Johnstone [1998], as stated in (6), not Theorem 2). When $d = 2$, both the TV and Sobolev classes again have the same minimax rate of $n^{-1/2}$, the caveat being that the rate for TV class has an extra $\sqrt{\log n}$ factor. But for all $d \geq 3$, the rates for the canonical TV and Sobolev classes differ, and the smaller Sobolev spaces have faster rates than their inscribing TV spaces. This may be viewed as a phase transition at $d = 3$; see Table 1.

We may paraphrase to say that 2d is just like 1d, in that expanding the Sobolev ball into a larger TV ball does not hurt the minimax rate, and methods like TV denoising are automatically *adaptive*, i.e., optimal over

both the bigger and smaller classes. However, as soon as we enter the 3d world, it is no longer clear whether TV denoising can adapt to the smaller, inscribed Sobolev ball, whose minimax rate is faster, $n^{-2/5}$ versus $n^{-1/3}$ (ignoring log factors). Theoretically, this is an interesting open problem that we do not approach in this paper and leave to future work.

We do, however, investigate the matter empirically: see Figure 3, where we run Laplacian smoothing and TV denoising on a highly smooth “linear” signal θ_0 . This is constructed so that each component θ_i is proportional to $i_1 + i_2 + \dots + i_d$ (using the multi-index notation (i_1, \dots, i_d) of (12) for grid location i), and the Sobolev norm is $\|D\theta_0\|_2 \asymp n^{1/2-1/d}$. Arguably, these are among the “hardest” types of functions for TV denoising to handle. The left panel, in 2d, is a case in which we know that TV denoising attains the minimax rate; the right panel, in 3d, is a case in which we do not, though empirically, TV denoising surely seems to be doing better than the slower minimax rate of $n^{-1/3}$ (ignoring log terms) that is associated with the larger TV ball.

Even if TV denoising is shown to be minimax optimal over the inscribed Sobolev balls when $d \geq 3$, note that this does not necessarily mean that we should scrap Laplacian smoothing in favor of TV denoising, in all problems. Laplacian smoothing is the unique Bayes estimator in a normal means model under a certain Markov random field prior (e.g., [Sharpnack and Singh \[2010\]](#)); statistical decision theory therefore tells that it is *admissible*, i.e., no other estimator—TV denoising included—can uniformly dominate it.

6 Discussion

We conclude with a quote from Albert Einstein: “Everything should be made as simple as possible, but no simpler”. In characterizing the minimax rates for TV classes, defined over d -dimensional grids, we have shown that simple methods like Laplacian smoothing and Laplacian eigenmaps—or even in fact, all linear estimators—must be passed up in favor of more sophisticated, nonlinear estimators, like TV denoising, if one wants to attain the optimal max risk. Such a result was previously known when $d = 1$; our work has extended it to all dimensions $d \geq 2$. We also characterized the minimax rates over discrete Sobolev classes, revealing an interesting phase transition where the optimal rates over TV and Sobolev spaces, suitably scaled, match when $d = 1$ and 2 but diverge for $d \geq 3$. It is an open question as to whether an estimator like TV denoising can be optimal over both spaces, for all d .

A Proofs

We present proofs of all results, according to the order in which they appear in the paper.

A.1 Proof of Lemma 1 (canonical scaling)

Suppose that $\theta \in \mathcal{H}_d(1)$ that is a discretization of a 1-Lipshitz function f , i.e., $\theta_i = f(i_1/\ell, \dots, i_d/\ell)$, $i = 1, \dots, n$. We first we compute and bound its squared Sobolev norm

$$\begin{aligned} \|D\theta\|_2^2 &= \sum_{(i,j) \in E} (\theta_i - \theta_j)^2 = \sum_{(i,j) \in E} (f(i_1/\ell, \dots, i_d/\ell) - f(j_1/\ell, \dots, j_d/\ell))^2 \\ &\leq \sum_{(i,j) \in E} \|(i_1/\ell, \dots, i_d/\ell) - (j_1/\ell, \dots, j_d/\ell)\|_\infty^2 \\ &= m/\ell^2, \end{aligned}$$

where, recall, we denote by $m = |E|$ the number of edges in the grid. In the second line we used the 1-Lipschitz property of f , and in the third we used that multi-indices corresponding to adjacent locations

on the grid are exactly 1 apart, in ℓ_∞ distance. Thus we see that setting $C'_n = \sqrt{m}/\ell$ gives the desired containment $\mathcal{S}_d(C'_n) \supseteq \mathcal{H}_d(1)$. It is always true that $m \asymp n$ for a d -dimensional grid (though the constant may depend on d), so that $\sqrt{m}/\ell \asymp n^{1/2-1/d}$. This completes the proof for the Sobolev class scaling.

As for TV class scaling, the result follows from the simple fact that $\|x\|_1 \leq \sqrt{m}\|x\|_2$ for any $x \in \mathbb{R}^m$, so that we may take $C_n = \sqrt{m}C'_n = n^{1-1/d}$. \square

A.2 Proof of Theorem 2 (minimax rates over TV classes)

Here and henceforth, we use the notation $B_p(r) = \{x : \|x\|_p \leq r\}$ for the ℓ_p ball of radius r , where $p, r > 0$ (and the ambient dimension will be determined based on the context).

We begin with a very simple lemma, that embeds an ℓ_1 ball inside the TV ball $\mathcal{T}_d(C_n)$.

Lemma 7. *Let G be a graph with maximum degree d_{\max} , and let $D \in \mathbb{R}^{m \times n}$ be its incidence matrix. Then for any $r > 0$, it holds that $B_1(r/d_{\max}) \subseteq T_d(r)$.*

Proof. Write D_i for the i th column of D . The proof follows from the observation that, for any θ ,

$$\|D\theta\|_1 = \left\| \sum_{i=1}^n D_i \theta_i \right\|_1 \leq \sum_{i=1}^n \|D_i\|_1 |\theta_i| \leq \left(\max_{i=1, \dots, n} \|D_i\|_1 \right) \|\theta\|_1 = d_{\max} \|\theta\|_1.$$

\square

To prove Theorem 2, we will rely on a result from Birge and Massart [2001], which gives a lower bound for the risk in a normal means problem, over ℓ_p balls. Another related, earlier result is that of Donoho and Johnstone [1994b]; however, the Birge and Massart result places no restrictions on the radius of the ball in question, whereas the Donoho and Johnstone result does. Translated into our notation, the Birge and Massart result is as follows.

Lemma 8 (Proposition 5 of Birge and Massart [2001]). *Assume i.i.d. observations $y_i \sim N(\theta_{0,i}, \sigma^2)$, $i = 1, \dots, n$, and $n \geq 2$. Then the minimax risk over the ℓ_p ball $B_p(r_n)$, where $0 < p < 2$, satisfies*

$$n \cdot R(B_p(r_n)) \geq c \cdot \begin{cases} \sigma^{2-p} r_n^p \left[1 + \log \left(\frac{\sigma^p n}{r_n^p} \right) \right]^{1-p/2} & \text{if } \sigma \sqrt{\log n} \leq r_n \leq \sigma n^{1/p} / \sqrt{\rho_p} \\ r_n^2 & \text{if } r_n < \sigma \sqrt{\log n} \\ \sigma^2 n / \rho_p & \text{if } r_n > \sigma n^{1/p} / \sqrt{\rho_p} \end{cases}.$$

Here $c > 0$ is a universal constant, and $\rho_p > 1.76$ is the unique solution of $\rho_p \log \rho_p = 2/p$.

Finally, applying Lemma 8 to $B_1(C_n/d_{\max})$ almost gives the lower bound as stated in Theorem 2. However, note that the minimax risk in question is trivially lower bounded by σ^2/n , because

$$\begin{aligned} \inf_{\hat{\theta}} \sup_{\theta_0 \in \mathcal{T}_d(C_n)} \frac{1}{n} \mathbb{E} \|\hat{\theta} - \theta_0\|_2^2 &\geq \inf_{\hat{\theta}} \sup_{\theta_0: \theta_{0,1} = \dots = \theta_{0,n}} \frac{1}{n} \sum_{i=1}^n \mathbb{E} (\hat{\theta}_i - \theta_{0,1})^2 \\ &= \inf_{\hat{\theta}_1} \sup_{\theta_{0,1}} \mathbb{E} (\hat{\theta}_1 - \theta_{0,1})^2 \\ &= \frac{\sigma^2}{n}. \end{aligned}$$

In the second to last line, the problem is to estimate a 1-dimensional mean parameter $\theta_{0,1}$, given the observations $y_i \sim N(\theta_{0,1}, \sigma^2)$, i.i.d., for $i = 1, \dots, n$; this has a well-known minimax risk of σ^2/n . What this means for our TV problem: to derive a lower bound for the minimax rate over $\mathcal{T}_d(C_n)$, we may take the maximum of the result of applying Lemma 8 to $B_1(C_n/d_{\max})$ and σ^2/n . One can see that the term σ^2/n only plays a role for small C_n , i.e., it effects the case when $C_n < \sigma d_{\max} \sqrt{\log n}$, where the lower bound becomes $C_n^2/(d_{\max}^2 n) \vee \sigma^2/n$. \square

A.3 Proof of Theorem 3 (minimax linear rates over TV classes)

First we recall a few definitions, from [Donoho et al. \[1990\]](#). Given a set $A \subseteq \mathbb{R}^k$, its *quadratically convex hull* $\text{qconv}(A)$ is defined as

$$\begin{aligned} \text{qconv}(A) &= \{(x_1, \dots, x_k) : (x_1^2, \dots, x_k^2) \in \text{conv}(A_+^2)\}, \quad \text{where} \\ A_+^2 &= \{(a_1^2, \dots, a_k^2) : a \in A, a_i \geq 0, i = 1, \dots, k\}. \end{aligned}$$

(Here $\text{conv}(B)$ denotes the convex hull of a set B .) Furthermore, the set A is called *quadratically convex* provided that $\text{qconv}(A) = A$. Also, A is called *orthosymmetric* provided that $(a_1, \dots, a_k) \in A$ implies $(\sigma_1 a_1, \dots, \sigma_k a_k) \in A$, for any choice of signs $\sigma_1, \dots, \sigma_k \in \{-1, 1\}$.

Now we proceed with the proof. Following from equation (7.2) of [Donoho et al. \[1990\]](#),

$$\text{qconv}(B_1(C_n/d_{\max})) = B_2(C_n/d_{\max}).$$

Theorem 11 of [Donoho et al. \[1990\]](#) states that, for orthosymmetric, compact sets, such as $B_1(C_n/d_{\max})$, the minimax linear risk equals that of its quadratically convex hull. Moreover, Theorem 7 of [Donoho et al. \[1990\]](#) tells us that for sets that are orthosymmetric, compact, convex, and quadratically convex, such as $B_2(C_n/d_{\max})$, the minimax linear risk is the same as the minimax linear risk over the worst rectangular subproblem. We consider $B_\infty(C_n/(d_{\max}\sqrt{n}))$, and abbreviate $r_n = C_n/(d_{\max}\sqrt{n})$. It is fruitful to study rectangles because the problem separates across dimensions, as in

$$\begin{aligned} \inf_{\hat{\theta} \text{ linear}} \sup_{\theta_0 \in B_\infty(r_n)} \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n (\hat{\theta}_i - \theta_{0,i})^2 \right] &= \frac{1}{n} \sum_{i=1}^n \left[\inf_{\hat{\theta}_i \text{ linear}} \sup_{|\theta_{0,i}| \leq r_n} \mathbb{E} (\hat{\theta}_i - \theta_{0,i})^2 \right] \\ &= \inf_{\hat{\theta}_1 \text{ linear}} \sup_{|\theta_{0,1}| \leq r_n} \mathbb{E} (\hat{\theta}_1 - \theta_{0,1})^2. \end{aligned}$$

Thus it suffices to compute the minimax linear risk over the 1d class $\{\theta_{0,1} : |\theta_{0,1}| \leq r_n\}$. It is easily shown (e.g., see Section 2 of [Donoho et al. \[1990\]](#)) that this is $r_n^2 \sigma^2 / (r_n^2 + \sigma^2)$, and so this is precisely the minimax linear risk for $B_2(C_n/d_{\max})$, and for $B_1(C_n/d_{\max})$.

To get the first lower bound as stated in the theorem, we simply take a maximum of $r_n^2 \sigma^2 / (r_n^2 + \sigma^2)$ and σ^2/n , as the latter is the minimax risk for estimating a 1-dimensional mean parameter given n observations in a normal model with variance σ^2 , recall the end of the proof of Theorem 2. To get the second, we use the fact that $2ab/(a+b) \geq \min\{a, b\}$. This completes the proof. \square

A.4 Alternative proof of Theorem 3

Here, we reprove Theorem 3 using elementary arguments. We write $y = \theta_0 + \epsilon$, for $\epsilon \sim N(0, \sigma^2 I)$. Given an arbitrary linear estimator, $\hat{\theta} = Sy$ for a matrix $S \in \mathbb{R}^{n \times n}$, observe that

$$\begin{aligned} \mathbb{E}[\text{MSE}(\hat{\theta}, \theta_0)] &= \frac{1}{n} \mathbb{E} \|\hat{\theta} - \theta_0\|_2^2 = \frac{1}{n} \mathbb{E} \|S(\theta_0 + \epsilon) - \theta_0\|_2^2 \\ &= \frac{1}{n} \mathbb{E} \|S\epsilon\|_2^2 + \frac{1}{n} \|(S - I)\theta_0\|_2^2 \\ &= \frac{\sigma^2}{n} \|S\|_F^2 + \frac{1}{n} \|(S - I)\theta_0\|_2^2, \end{aligned} \tag{16}$$

which we may view as the variance and (squared) bias terms, respectively. Now denote by e_i the i th standard basis vector, and consider

$$\begin{aligned}
\frac{\sigma^2}{n} \|S\|_F^2 + \left(\sup_{\theta_0: \|D\theta_0\|_1 \leq C_n} \frac{1}{n} \|(S - I)\theta_0\|_2^2 \right) &\geq \frac{\sigma^2}{n} \|S\|_F^2 + \frac{C_n^2}{d_{\max}^2 n} \left(\max_{i=1, \dots, n} \|(I - S)e_i\|_2^2 \right) \\
&\geq \frac{\sigma^2}{n} \|S\|_F^2 + \frac{C_n^2}{d_{\max}^2 n^2} \sum_{i=1}^n \|(I - S)e_i\|_2^2 \\
&= \frac{\sigma^2}{n} \|S\|_F^2 + \frac{C_n^2}{d_{\max}^2 n^2} \|(I - S)\|_F^2 \\
&\geq \frac{\sigma^2}{n} \sum_{i=1}^n S_{ii}^2 + \frac{C_n^2}{d_{\max}^2 n^2} \sum_{i=1}^n (1 - S_{ii})^2 \\
&= \frac{1}{n} \sum_{i=1}^n \left(\sigma^2 S_{ii}^2 + \frac{C_n^2}{d_{\max}^2 n} (1 - S_{ii})^2 \right).
\end{aligned}$$

Here S_{ii} , $i = 1, \dots, n$ denote the diagonal entries of S . To bound each term in the sum, we apply the simple inequality $ax^2 + b(1 - x)^2 \geq ab/(a + b)$ for all x (since a short calculation shows that the quadratic in x here is minimized at $x = b/(a + b)$). We may continue on lower bounding the last displayed expression, giving

$$\frac{\sigma^2}{n} \|S\|_F^2 + \left(\sup_{\theta_0: \|D\theta_0\|_1 \leq C_n} \frac{1}{n} \|(S - I)\theta_0\|_2^2 \right) \geq \frac{\sigma^2 C_n^2}{C_n^2 + \sigma^2 d_{\max}^2 n}.$$

Lastly, we may take the maximum of this with σ^2/n in order to derive a final lower bound, as argued in the proof of Theorem 3. \square

A.5 Proof of Lemma 4 (mean estimator over TV classes)

For this estimator, the smoother matrix is $S = \mathbf{1}\mathbf{1}^T/n$ and so $\|S\|_F^2 = 1$. From (16), we have

$$\mathbb{E}[\text{MSE}(\hat{\theta}^{\text{mean}}, \theta_0)] = \frac{\sigma^2}{n} + \frac{1}{n} \|\theta_0 - \bar{\theta}_0 \mathbf{1}\|_2^2,$$

where $\bar{\theta}_0 = (1/n) \sum_{i=1}^n \theta_{0,i}$. Now

$$\begin{aligned}
\sup_{\theta_0: \|D\theta_0\|_1 \leq C_n} \frac{1}{n} \|\theta_0 - \bar{\theta}_0 \mathbf{1}\|_2^2 &= \sup_{x \in \text{row}(D): \|Dx\|_1 \leq C_n} \frac{1}{n} \|x\|_2^2 \\
&= \sup_{z \in \text{col}(D): \|z\|_1 \leq C_n} \frac{1}{n} \|D^\dagger z\|_2^2 \\
&\leq \sup_{z: \|z\|_1 \leq C_n} \frac{1}{n} \|D^\dagger z\|_2^2 \\
&= \frac{C_n^2}{n} \max_{i=1, \dots, n} \|D_i^\dagger\|_2^2 \\
&\leq \frac{C_n^2 M_n^2}{n},
\end{aligned}$$

which establishes the desired bound. \square

A.6 Proof of Theorem 5 (minimax rates over Sobolev classes)

Recall that we denote by $L = V\Sigma V^T$ the eigendecomposition of the graph Laplacian $L = D^T D$, where $\Sigma = \text{diag}(\rho_1, \dots, \rho_n)$ with $0 = \rho_1 < \rho_2 \leq \dots \leq \rho_n$, and where $V \in \mathbb{R}^{n \times n}$ has orthonormal columns. Also denote by $D = U\Sigma^{1/2}V^T$ the singular value decomposition of the edge incidence matrix D , where $U \in \mathbb{R}^{m \times n}$ has orthonormal columns.¹ First notice that

$$\|D\theta_0\|_2 = \|U\Sigma^{1/2}V^T\theta_0\|_2 = \|\Sigma^{1/2}V^T\theta_0\|_2.$$

This suggests that a rotation by V^T will further simplify the minimax risk over $\mathcal{S}_d(C'_n)$, i.e.,

$$\begin{aligned} \inf_{\hat{\theta}} \sup_{\theta_0: \|\Sigma^{1/2}V^T\theta_0\|_2 \leq C'_n} \frac{1}{n} \mathbb{E} \|\hat{\theta} - \theta_0\|_2^2 &= \inf_{\hat{\theta}} \sup_{\theta_0: \|\Sigma^{1/2}V^T\theta_0\|_2 \leq C'_n} \frac{1}{n} \mathbb{E} \|V^T\hat{\theta} - V^T\theta_0\|_2^2 \\ &= \inf_{\hat{\gamma}} \sup_{\gamma_0: \|\Sigma^{1/2}\gamma_0\|_2 \leq C'_n} \frac{1}{n} \mathbb{E} \|\hat{\gamma} - \gamma_0\|_2^2, \end{aligned} \quad (17)$$

where we have rotated and now consider the new parameter $\gamma_0 = V^T\theta_0$, constrained to lie in

$$\mathcal{E}_d(C'_n) = \left\{ \gamma : \sum_{i=2}^n \rho_i \gamma_i^2 \leq (C'_n)^2 \right\}.$$

To be clear, in the rotated setting (17) we observe a vector $y' = V^T y \sim N(\gamma_0, \sigma^2 I)$, and the goal is to estimate the mean parameter γ_0 . Since there are no constraints along the first dimension, we can separate out the MSE in (17) into that incurred on the first component, and all other components. Decomposing $\gamma_0 = (\alpha_0, \beta_0) \in \mathbb{R}^{1 \times (n-1)}$, with similar notation for an estimator $\hat{\gamma}$,

$$\begin{aligned} \inf_{\hat{\gamma}} \sup_{\gamma_0 \in \mathcal{E}_d(C'_n)} \frac{1}{n} \mathbb{E} \|\hat{\gamma} - \gamma_0\|_2^2 &= \inf_{\hat{\alpha}} \sup_{\alpha_0} \frac{1}{n} \mathbb{E} (\hat{\alpha} - \alpha_0)^2 + \inf_{\hat{\beta}} \sup_{\beta_0 \in P_{-1}(\mathcal{E}_d(C'_n))} \frac{1}{n} \mathbb{E} \|\hat{\beta} - \beta_0\|_2^2 \\ &= \frac{\sigma^2}{n} + \inf_{\hat{\beta}} \sup_{\beta_0 \in P_{-1}(\mathcal{E}_d(C'_n))} \frac{1}{n} \mathbb{E} \|\hat{\beta} - \beta_0\|_2^2, \end{aligned} \quad (18)$$

where P_{-1} projects onto all coordinate axes but the 1st, i.e., $P_{-1}(x) = (0, x_2, \dots, x_n)$, and in the second line we have used the fact that the minimax risk for estimating a 1-dimensional parameter α_0 given an observation $z \sim N(\alpha_0, \sigma^2)$ is simply σ^2 .

Let us lower bound the second term in (18), i.e., $R(P_{-1}(\mathcal{E}_d(C'_n)))$. The ellipsoid $P_{-1}(\mathcal{E}_d(C'_n))$ is orthosymmetric, compact, convex, and quadratically convex, hence Theorem 7 in Donoho et al. [1990] tells us that its minimax linear risk is the minimax linear risk of its hardest rectangular subproblem. Further, Lemma 6 in Donoho et al. [1990] then tells us the minimax linear risk of its hardest rectangular subproblem is, up to a constant factor, the same as the minimax (nonlinear) risk of the full problem. More precisely, Lemma 6 and Theorem 7 from Donoho et al. [1990] imply

$$R(P_{-1}(\mathcal{E}_d(C'_n))) \geq \frac{4}{5} R_L(P_{-1}(\mathcal{E}_d(C'_n))) = \sup_{H \subseteq P_{-1}(\mathcal{E}_d(C'_n))} R_L(H), \quad (19)$$

where the supremum above is taken over all rectangular subproblems, i.e., all rectangles H contained in $P_{-1}(\mathcal{E}_d(C'_n))$.

¹When $d = 1$, we have $m = n - 1$ edges, and so it is not possible for U to have orthonormal columns; however, we can just take its first column to be all 0s, and take the rest as the eigenbasis for \mathbb{R}^{n-1} , and all the arguments given here will go through.

To study rectangular subproblems, it helps to reintroduce the multi-index notation for a location i on the d -dimensional grid, writing this as $(i_1, \dots, i_d) \in \{1, \dots, \ell\}^d$, where $\ell = n^{1/d}$. For a parameter $2 \leq \tau \leq \ell$, we consider rectangular subsets of the form²

$$H(\tau) = \{\beta \in \mathbb{R}^{n-1} : |\beta_i| \leq t_i(\tau), i = 2, \dots, n\}, \quad \text{where}$$

$$t_i(\tau) = \begin{cases} C'_n / (\sum_{j_1, \dots, j_d \leq \tau} \rho_{j_1, \dots, j_d})^{1/2} & \text{if } i_1, \dots, i_d \leq \tau \\ 0 & \text{otherwise} \end{cases}, \quad \text{for } i = 2, \dots, n.$$

It is not hard to check that $H(\tau) \subseteq \{\beta \in \mathbb{R}^{n-1} : \sum_{i=2}^n \rho_i \beta_i^2 \leq (C'_n)^2\} = P_{-1}(\mathcal{E}_d(C'_n))$. Then, from (19),

$$\begin{aligned} R(P_{-1}(\mathcal{E}_d(C'_n))) &\geq \sup_{\tau} R_L(H(\tau)) = \sup_{\tau} \frac{1}{n} \sum_{i=1}^n \frac{t_i(\tau)^2 \sigma^2}{t_i(\tau)^2 + \sigma^2} \\ &= \sup_{\tau} \frac{1}{n} \frac{(\tau^d - 1) \sigma^2 (C'_n)^2}{(C'_n)^2 + \sum_{j_1, \dots, j_d \leq \tau} \rho_{j_1, \dots, j_d}}. \end{aligned}$$

The first equality is due to the fact that the minimax risk for rectangles decouples across dimensions, and the 1d minimax linear risk is straightforward to compute for an interval, as argued in the proof Theorem 3; the second equality simply comes from a short calculation following the definition of $t_i(\tau)$, $i = 2, \dots, n$. Applying Lemma 9, on the eigenvalues of the graph Laplacian matrix L for a d -dimensional grid, we have that for a constant $c > 0$,

$$\frac{(\tau^d - 1) \sigma^2 (C'_n)^2}{(C'_n)^2 + \sum_{j_1, \dots, j_d \leq \tau} \rho_{j_1, \dots, j_d}} \geq \frac{(\tau^d - 1) \sigma^2 (C'_n)^2}{(C'_n)^2 + c \sigma^2 \tau^{d+2} / \ell^2} \geq \frac{1}{2} \frac{\sigma^2 (C'_n)^2}{(C'_n)^2 \tau^{-d} + c \sigma^2 \tau^2 / \ell^2}.$$

We can choose τ to maximize the expression on the right above, given by

$$\tau^* = \left(\frac{\ell^2 (C'_n)^2}{c \sigma^2} \right)^{\frac{1}{d+2}}.$$

When $2 \leq \tau^* \leq \ell$, this provides us with the lower bound on the minimax risk

$$R(P_{-1}(\mathcal{E}_d(C'_n))) \geq R_L(H(\tau^*)) \geq \frac{1}{2n} \frac{\tau^d \sigma^2 (C'_n)^2}{2(c\sigma^2)^{\frac{d}{d+2}} (C'_n)^{\frac{4}{d+2}} \ell^{-\frac{2d}{d+2}}} = \frac{c_1}{n} (n\sigma^2)^{\frac{2}{d+2}} (C'_n)^{\frac{2d}{d+2}}, \quad (20)$$

for a constant $c_1 > 0$. When $\tau^* < 2$, we can use $\tau = 2$ as lower bound on the minimax risk,

$$R(P_{-1}(\mathcal{E}_d(C'_n))) \geq R_L(H(2)) \geq \frac{1}{2n} \frac{\sigma^2 \ell^2 (C'_n)^2}{\ell^2 (C'_n)^2 2^{-d} + c \sigma^2 2^2} \geq \frac{c_2}{n} \ell^2 (C'_n)^2, \quad (21)$$

for a constant $c_2 > 0$, where in the last inequality, we used the fact that $\ell^2 (C'_n)^2 \leq c \sigma^2 2^{d+2}$ (just a constant) since we are in the case $\tau^* < 2$. Finally, when $\tau^* > \ell$, we can use $\tau = \ell$ as a lower bound on the minimax risk,

$$R(P_{-1}(\mathcal{E}_d(C'_n))) \geq R_L(H(\ell)) \geq \frac{1}{2n} \frac{\sigma^2 (C'_n)^2}{\ell^{-d} (C'_n)^2 + c \sigma^2} \geq c_3 \sigma^2, \quad (22)$$

for a constant $c_3 > 0$, where in the last inequality, we used that $c \sigma^2 \leq \ell^{-d} (C'_n)^2$ as we are in the case $\tau^* > \ell$. Taking a minimum of the lower bounds in (20), (21), (22), as a way to navigate the cases, gives us a final lower bound on $R(P_{-1}(\mathcal{E}_d(C'_n)))$, and completes the proof.

²Here, albeit unconventional, it helps to index $\beta \in H(\tau) \subseteq \mathbb{R}^{n-1}$ according to components $i = 2, \dots, n$, rather than $i = 1, \dots, n-1$. This is so that we may keep the index variable i to be in correspondence with positions on the grid.

A.7 Proof of Theorem 6 (Laplacian eigenmaps and Laplacian smoothing over Sobolev classes)

We will prove the results for Laplacian eigenmaps and Laplacian separately.

Laplacian eigenmaps. The smoother matrix for this estimator is $S_k = V_{[k]}V_{[k]}^T$, for a tuning parameter $k = 1, \dots, n$. From (16),

$$\mathbb{E}[\text{MSE}(\hat{\theta}^{\text{LE}}, \theta_0)] = \frac{\sigma^2}{n}k + \frac{1}{n}\|(I - S_k)\theta_0\|_2^2.$$

Now we write $k = \tau^d$, and analyze the max risk of the second term,

$$\begin{aligned} \sup_{\theta_0: \|D\theta_0\|_2 \leq C'_n} \frac{1}{n}\|(I - S_k)\theta_0\|_2^2 &= \sup_{z: \|z\|_2 \leq C'_n} \frac{1}{n}\|(I - S_k)D^\dagger z\|_2^2 \\ &= \frac{(C'_n)^2}{n}\sigma_{\max}^2((I - S_k)D^\dagger) \\ &\leq \frac{(C'_n)^2}{n} \frac{1}{4\sin^2(\pi\tau/(2\ell))} \\ &\leq \frac{(C'_n)^2}{n} \frac{4\ell^2}{\pi^2\tau^2}. \end{aligned}$$

Here we denote by $\sigma_{\max}(A)$ the maximum singular value of a matrix A . The last inequality above used the simple lower bound $\sin(x) \geq x/2$ for $x \in [0, \pi/2]$. The earlier inequality used that

$$(I - S_k)D^\dagger = (I - V_{[k]}V_{[k]}^T)V^T(\Sigma^\dagger)^{1/2}U^T = [0, \dots, 0, V_{k+1}, \dots, V_n](\Sigma^\dagger)^{1/2}U^T,$$

where we have kept the same notation for the singular value decomposition of D as in the proof of Theorem 5. Therefore $\sigma_{\max}^2((I - S_k)D^\dagger)$ is the reciprocal of the $(k + 1)$ st smallest eigenvalue ρ_{k+1} of the graph Laplacian L . For any subset A of the eigenvalues Λ of the Laplacian, with $|A| = k$, $\rho_{k+1} \geq \min \Lambda \setminus A$. So for a d -dimensional grid,

$$\begin{aligned} \rho_{k+1} &\geq \min \Lambda \setminus \{\rho_{i_1, \dots, i_d} : 1 \leq i_j \leq \tau \text{ for all } 1 \leq j \leq d\} \\ &= 4\sin^2(\pi\tau/(2\ell)) \end{aligned}$$

where recall $\ell = n^{1/d}$, as explained by (23), in the proof of Lemma 9.

Hence, we have established

$$\sup_{\theta_0: \|D\theta_0\|_2 \leq C'_n} \mathbb{E}[\text{MSE}(\hat{\theta}^{\text{LE}}, \theta_0)] \leq \frac{\sigma^2}{n} + \frac{\sigma^2}{n}\tau^d + \frac{(C'_n)^2}{n} \frac{4\ell^2}{\pi^2\tau^2}.$$

Choosing τ to balance the two terms on the right-hand side above results in $\tau^* = (2\ell C'_n/(\pi\sigma))^{\frac{2}{d+2}}$. Plugging in this choice of τ , while utilizing the bounds $1 \leq \tau \leq \ell$, very similar to the arguments given at the end of the proof of Theorem 5, gives the result for Laplacian eigenmaps.

Laplacian smoothing. The smoother matrix for this estimator is $S_\lambda = (I + \lambda L)^{-1}$, for a tuning parameter $\lambda \geq 0$. From (16),

$$\mathbb{E}[\text{MSE}(\hat{\theta}^{\text{LS}}, \theta_0)] = \frac{\sigma^2}{n} \sum_{i=1}^n \frac{1}{(1 + \lambda\rho_i)^2} + \frac{1}{n}\|(I - S_\lambda)\theta_0\|_2^2.$$

When $d = 1, 2$, or 3 , the first term upper is bounded by $c_1\sigma^2/n + c_2\sigma^2/\lambda^{d/2}$, for some constants $c_1, c_2 > 0$, by Lemma 10. As for the second term,

$$\begin{aligned}
\sup_{\theta_0: \|D\theta_0\|_2 \leq C'_n} \frac{1}{n} \|(I - S_\lambda)\theta_0\|_2^2 &= \sup_{z: \|z\|_2 \leq C'_n} \|(I - S_\lambda)D^\dagger z\|_2^2 \\
&= \frac{(C'_n)^2}{n} \sigma_{\max}^2((I - S_\lambda)D^\dagger) \\
&= \frac{(C'_n)^2}{n} \max_{i=2, \dots, n} \left(1 - \frac{1}{1 + \lambda\rho_i}\right)^2 \frac{1}{\rho_i} \\
&= \frac{(C'_n)^2}{n} \lambda \max_{i=2, \dots, n} \frac{\lambda\rho_i}{(1 + \lambda\rho_i)^2} \\
&\leq \frac{(C'_n)^2 \lambda}{4n}.
\end{aligned}$$

In the third equality we have used the fact the eigenvectors of $I - S_\lambda$ are the left singular vectors of D^\dagger , and in the last inequality we have used the simple upper bound $f(x) = x/(1+x)^2 \leq 1/4$ for $x \geq 0$ (this function being maximized at $x = 1$).

Therefore, from what we have shown,

$$\sup_{\theta_0: \|D\theta_0\|_2 \leq C'_n} \mathbb{E}[\text{MSE}(\hat{\theta}^{\text{LS}}, \theta_0)] \leq \frac{c_1\sigma^2}{n} + \frac{c_2\sigma^2}{\lambda^{d/2}} + \frac{(C'_n)^2 \lambda}{4n}.$$

Choosing λ to balance the two terms on the right-hand side above gives $\lambda^* = c(n/(C'_n)^2)^{2/(d+2)}$, for a constant $c > 0$. Plugging in this choice, and using upper bounds from the trivial cases $\lambda = 0$ and $\lambda = \infty$ when C'_n is very small or very large, respectively, gives the result for Laplacian smoothing. \square

Remark 6. When $d = 4$, Lemma 10 gives a slightly worse upper bound on $\sum_{i=1}^n 1/(1 + \lambda\rho_i)^2$, with an “extra” term $(nc_2/\lambda^{d/2}) \log(1 + c_3\lambda)$, for constants $c_2, c_3 > 0$. It is not hard to show, by tracing through the same arguments as given above that we can use this to establish an upper bound on the max risk of

$$\sup_{\theta_0 \in \mathcal{S}_d(C'_n)} \mathbb{E}[\text{MSE}(\hat{\theta}^{\text{LE}}, \theta_0)] \leq \frac{c}{n} \left((n\sigma^2)^{\frac{2}{d+2}} (C'_n)^{\frac{2d}{d+2}} \log(n/(C'_n)^2) \wedge n\sigma^2 \wedge n^{2/d} (C'_n)^2 \right) + \frac{c\sigma^2}{n},$$

only slightly worse than the minimax optimal rate, by a log factor.

When $d \geq 5$, our analysis provides a much worse bound for the max risk of Laplacian smoothing, as the integral denoted $I(d)$ in the proof of Lemma 10 grows very large when $d \geq 5$. We conjecture that this not due to slack in our proof technique, but rather, to the Laplacian smoothing estimator itself, since all inequalities the proof can be fairly tight.

B Utility lemmas used in the proofs of Theorems 5 and 6

This section contains some calculations on the partial sums of eigenvalues of the Laplacian matrix L , for d -dimensions grids. These are useful for the proofs of both Theorem 5 and Theorem 6.

Lemma 9. *Let $L \in \mathbb{R}^{n \times n}$ denote the graph Laplacian matrix of a d -dimensional grid graph, and ρ_{i_1, \dots, i_d} , $(i_1, \dots, i_d) \in \{1, \dots, \ell\}^d$ be its eigenvalues, where $\ell = n^{1/d}$. Then there exists a constant $c > 0$ (dependent on d) such that, for any $1 \leq \tau \leq \ell$,*

$$\sum_{(i_1, \dots, i_d) \in \{1, \dots, \tau\}^d} \rho_{i_1, \dots, i_d} \leq c \frac{\tau^{d+2}}{\ell^2}.$$

Proof. The eigenvalues of L can be written explicitly as

$$\rho_i = 4 \sin^2 \left(\frac{\pi(i_1 - 1)}{2\ell} \right) + \dots + 4 \sin^2 \left(\frac{\pi(i_d - 1)}{2\ell} \right), \quad (i_1, \dots, i_d) \in \{1, \dots, \ell\}^d. \quad (23)$$

This follows from known facts about the eigenvalues for the Laplacian matrix of a 1d grid, and the fact that the Laplacian matrix for higher-dimensional grids can be expressed in terms of a Kronecker sum of the Laplacian matrix of an appropriate 1d grid (e.g., [Conte and de Boor \[1980\]](#), [Kunsch \[1994\]](#), [Ng et al. \[1999\]](#), [Wang et al. \[2008, 2016\]](#), [Hutter and Rigollet \[2016\]](#)). We now use the fact that $\sin(x) \leq x$ for all $x \geq 0$, which gives us the upper bound

$$\begin{aligned} \sum_{(i_1, \dots, i_d) \in \{1, \dots, \tau\}^d} \rho_{i_1, \dots, i_d} &\leq \frac{\pi^2}{\ell^2} \sum_{(i_1, \dots, i_d) \in \{1, \dots, \tau\}^d} \left((i_1 - 1)^2 + \dots + (i_d - 1)^2 \right) \\ &\leq \frac{\pi^2 d}{\ell^2} \tau^{d-1} \sum_{i=1}^{\tau} (i - 1)^2 \\ &\leq \frac{\pi^2 d}{\ell^2} \tau^{d-1} \tau^3 \\ &= \frac{\pi^2 d}{\ell^2} \tau^{d+2}, \end{aligned}$$

as desired. \square

Lemma 10. *Let $L \in \mathbb{R}^{n \times n}$ denote the graph Laplacian matrix of a d -dimensional grid graph, and ρ_i , $i = 1, \dots, n$ be its eigenvalues. Let $\lambda \geq 0$ be arbitrary. For $d = 1, 2$, or 3 , there are constants $c_1, c_2 > 0$ such that*

$$\sum_{i=1}^n \frac{1}{(1 + \lambda \rho_i)^2} \leq c_1 + c_2 \frac{n}{\lambda^{d/2}}.$$

For $d = 4$, there are constants $c_1, c_2, c_3 > 0$ such that

$$\sum_{i=1}^n \frac{1}{(1 + \lambda \rho_i)^2} \leq c_1 + c_2 \frac{n}{\lambda^{d/2}} \left(1 + \log(1 + c_3 \lambda) \right).$$

Proof. We will use the explicit form of the eigenvalues as given in the proof of Lemma 9. In the expressions below, we use $c > 0$ to denote a constant whose value may change from line to line. Using the inequality $\sin x \geq x/2$ for $x \in [0, \pi/2]$,

$$\begin{aligned} \sum_{i=1}^n \frac{1}{(1 + \lambda \rho_i)^2} &\leq \sum_{(i_1, \dots, i_d) \in \{1, \dots, \ell\}^d} \frac{1}{\left(1 + \lambda \frac{\pi^2}{4\ell^2} \sum_{j=1}^d (i_j - 1)^2 \right)^2} \\ &\leq 1 + \int_{[0, \ell]^d} \frac{1}{\left(1 + \lambda \frac{\pi^2}{4} \sum_{j=1}^d x_j^2 / \ell^2 \right)^2} dx \\ &= 1 + c \int_0^{\ell\sqrt{d}} \frac{1}{\left(1 + \lambda \frac{\pi^2}{4} r^2 / \ell^2 \right)^2} r^{d-1} dr \\ &= 1 + c \frac{n}{\lambda^{d/2}} \underbrace{\int_0^{\frac{\pi}{2}\sqrt{\lambda d}} \frac{u^{d-1}}{(1 + u^2)^2} du}_{I(d)}. \end{aligned}$$

In the second inequality, we used the fact that the right-endpoint Riemann sum is always an underestimate for the integral of a function that is monotone nonincreasing in each coordinate. In the third, we made a change to spherical coordinates, and suppressed all of the angular variables, as they contribute at most a constant factor. It remains to compute $I(d)$, which can be done by symbolic integration:

$$\begin{aligned}
 I(1) &= \frac{\pi\sqrt{d}}{4(1 + \frac{\pi^2}{4}\lambda d)} + \frac{1}{2} \tan^{-1} \left(\frac{\pi}{2} \sqrt{\lambda d} \right) \leq \frac{1}{4} + \frac{\pi}{4}, \\
 I(2) &= \frac{1}{2} - \frac{1}{2(1 + \frac{\pi^2}{4}\lambda d)} \leq \frac{1}{2}, \\
 I(3) &= \frac{1}{2} \tan^{-1} \left(\frac{\pi}{2} \sqrt{\lambda d} \right) \leq \frac{\pi}{4}, \quad \text{and} \\
 I(4) &= \frac{1}{2} \log \left(1 + \frac{\pi^2}{4}\lambda d \right) + \frac{1}{2(1 + \frac{\pi^2}{4}\lambda d)} - \frac{1}{2} \leq \frac{1}{2} \log \left(1 + \frac{\pi^2}{4}\lambda d \right) + \frac{1}{2}.
 \end{aligned}$$

This completes the proof. □

C Additional experiments comparing TV denoising and Laplacian smoothing for piecewise constant functions

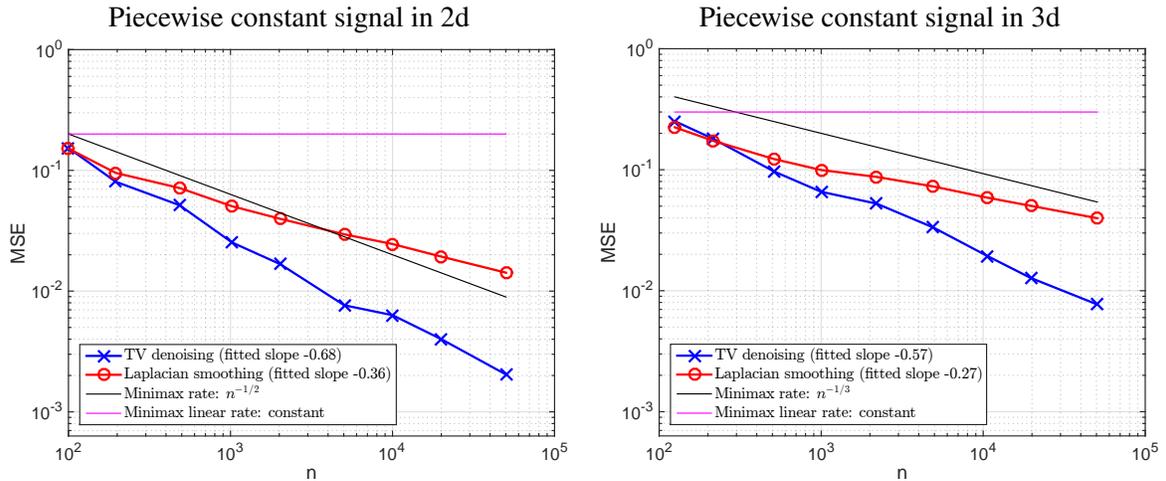


Figure 4: MSE curves for estimating a “piecewise constant” signal, having a single elevated region, over 2d and 3d grids. For each n , the results were averaged over 5 repetitions, and the Laplacian smoothing and TV denoising estimators were tuned for best average MSE performance. We set θ_0 to satisfy $\|D\theta_0\|_1 \asymp n^{1-1/d}$, matching the canonical scaling. Note that all estimators achieve better performance than that dictated by their minimax rates.

References

- Robert Acar and Curtis R. Vogel. Analysis of total variation penalty methods. *Inverse Problems*, 10: 1217–1229, 1994.
- Rie Ando and Tong Zhang. Learning on graph with Laplacian regularization. *Advances in Neural Information Processing Systems*, 9, 2006.

- Alvero Barbero and Suvrit Sra. Modular proximal optimization for multidimensional total-variation regularization. arXiv: 1411.0589, 2014.
- Mikhail Belkin and Partha Niyogi. Using manifold structure for partially labelled classification. *Advances in Neural Information Processing Systems*, 15, 2002.
- Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.
- Mikhail Belkin and Partha Niyogi. Semi-supervised learning on Riemannian manifolds. *Machine Learning*, 56(1–3):209–239, 2004.
- Mikhail Belkin and Partha Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. *Conference on Learning Theory (COLT-05)*, 18, 2005.
- Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani. On manifold regularization. *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 8, 2005.
- Lucien Birge and Pascal Massart. Gaussian model selection. *Journal of the European Mathematical Society*, 3(3):203–268, 2001.
- Antonin Chambolle and Jerome Darbon. On total variation minimization and surface evolution using parametric maximum flows. *International Journal of Computer Vision*, 84:288–307, 2009.
- Antonin Chambolle and Pierre-Louis Lions. Image recovery via total variation minimization and related problems. *Numerische Mathematik*, 76(2):167–188, 1997.
- Samuel Conte and Carl de Boor. *Elementary Numerical Analysis: An Algorithmic Approach*. McGraw-Hill, New York, 1980. International Series in Pure and Applied Mathematics.
- David Dobson and Fadil Santosa. Recovery of blocky images from noisy and blurred data. *SIAM Journal on Applied Mathematics*, 56(4):1181–1198, 1996.
- David Donoho and Iain Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455, 1994a.
- David Donoho and Iain Johnstone. Minimax risk over ℓ_p -balls for ℓ_q -error. *Probability Theory and Related Fields*, 99(2):277–303, 1994b.
- David Donoho and Iain Johnstone. Minimax estimation via wavelet shrinkage. *Annals of Statistics*, 26(8): 879–921, 1998.
- David Donoho, Richard Liu, and Brenda MacGibbon. Minimax risk over hyperrectangles, and implications. *Annals of Statistics*, 18(3):1416–1437, 1990.
- S. Godunov and V. Ryabenkii. *Difference Schemes: An Introduction to the Underlying Theory*. Elsevier, Amsterdam, 1987. Number 19 in Studies in Mathematics and Its Applications.
- Laszlo Györfi, Michael Kohler, Adam Krzyżak, and Harro Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer, New York, 2002.
- Holger Hoefling. A path algorithm for the fused lasso signal approximator. *Journal of Computational and Graphical Statistics*, 19(4):984–1006, 2010.

- Jan-Christian Hutter and Philippe Rigollet. Optimal rates for total variation denoising. In *Conference on Learning Theory (COLT-16)*, 2016. to appear.
- Hans Kunsch. Robust priors for smoothing and image restoration. *Annals of the Institute of Statistical Mathematics*, 46(1):1–19, 1994.
- Enno Mammen and Sara van de Geer. Locally adaptive regression splines. *Annals of Statistics*, 25(1): 387–413, 1997.
- Deanna Needell and Rachel Ward. Stable image reconstruction using total variation minimization. *SIAM Journal on Imaging Sciences*, 6(2):1035–1058, 2013.
- Michael Ng, Raymond Chan, and Wun-Cheung Tang. A fast algorithm for deblurring models with Neumann boundary conditions. *SIAM Journal on Scientific Computing*, 21(3):851–866, 1999.
- Leonid Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60:259–268, 1992.
- James Sharpnack and Aarti Singh. Identifying graph-structured activation patterns in networks. *Advances in Neural Information Processing Systems*, 13, 2010.
- James Sharpnack, Alessandro Rinaldo, and Aarti Singh. Sparsistency of the edge lasso over graphs. *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 15:1028–1036, 2012.
- Alexander Smola and Risi Kondor. Kernels and regularization on graphs. *Proceedings of the Annual Conference on Learning Theory*, 16, 2003.
- 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. Adaptive piecewise polynomial estimation via trend filtering. *Annals of Statistics*, 42(1): 285–323, 2014.
- Ryan J. Tibshirani and Jonathan Taylor. The solution path of the generalized lasso. *Annals of Statistics*, 39 (3):1335–1371, 2011.
- Curtis R. Vogel and M. E. Oman. Iterative methods for total variation denoising. *SIAM Journal on Scientific Computing*, 17(1):227–238, 1996.
- Yilun Wang, Junfeng Yang, Wotao Yin, and Yin Zhang. A new alternating minimization algorithm for total variation image reconstruction. *SIAM Journal on Imaging Sciences*, 1(3):248–272, 2008.
- Yu-Xiang Wang, James Sharpnack, Alex Smola, and Ryan J. Tibshirani. Trend filtering on graphs. *Journal of Machine Learning Research*, 2016. To appear.
- Dengyong Zhou, Jiayuan Huang, and Bernhard Scholkopf. Learning from labeled and unlabeled data on a directed graph. *Proceedings of the International Conference on Machine Learning*, 22, 2005.
- Xiaojin Zhu, Zoubin Ghahramani, and John Lafferty. Semi-supervised learning using Gaussian fields and harmonic functions. *International Conference on Machine Learning (ICML-03)*, 20, 2003.