

5-2008

Properties and Refinements of The Fused Lasso

Alessandro Rinaldo

Carnegie Mellon University, arinaldo@stat.cmu.edu

Follow this and additional works at: <http://repository.cmu.edu/statistics>

 Part of the [Statistics and Probability Commons](#)

This Technical Report is brought to you for free and open access by the Dietrich College of Humanities and Social Sciences at Research Showcase @ CMU. It has been accepted for inclusion in Department of Statistics by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Properties and Refinements of The Fused Lasso

Alessandro Rinaldo*
Department of Statistics
Carnegie Mellon University
Pittsburgh, PA 15213-3890 USA

Abstract

We consider estimating an unknown signal, which is both blocky and sparse, corrupted by additive noise. We study three interrelated least squares procedures and their asymptotic properties. The first procedure is the fused lasso, put forward by [Friedman et al. \(2007\)](#), which we modify into a different estimator, called the fused adaptive lasso, with better properties. The other two estimators we discuss solve least squares problems on sieves, one constraining the maximal ℓ_1 norm and the maximal total variation seminorm, the other restricting the number of blocks and the number of nonzero coordinates of the signal. We derive conditions for the recovery of the true block partition and the true sparsity patterns by the fused lasso and the fused adaptive lasso, and convergence rates for the sieve estimators, explicitly in terms of the constraining parameters.

1 Introduction

We consider the non-parametric regression model

$$y_i = \mu_i^0 + \epsilon_i, \quad i = 1, \dots, n,$$

where $\mu^0 \in \mathbb{R}^n$ is the unknown vector of mean values to be estimated using the observations y , and the errors ϵ_i are assumed to be independent with either Gaussian or sub-Gaussian distributions and bounded variances. We are concerned with the more specialized settings where μ^0 can be both *sparse*, with a possibly very large number of zero entries, and *blocky*, namely the number of coordinates where μ^0 changes its values can be much smaller than n . [Figure 1](#) shows an instance of data generated by corrupting with additive noise a blocky and sparse signal (see [Section 2.4](#) for details about this example). Formally, we assume that there exists a partition $\{\mathcal{B}_1^0, \dots, \mathcal{B}_{J_0}^0\}$ of $\{1, \dots, n\}$ into sets of consecutive indexes, from now on a *block partition*, and a vector $\nu^0 \in \mathbb{R}^{J_0}$, which may be sparse, such that the true mean vector can be written as

$$\mu^0 = \sum_{j=1}^{J_0} \nu_j^0 1_{\mathcal{B}_j^0}, \tag{1}$$

where $1_{\mathcal{B}}$ is the indicator function of the set $\mathcal{B} \subseteq \{1, \dots, n\}$, i.e. the n -dimensional vector whose i -th coordinate is 1 if $i \in \mathcal{B}$ and 0 otherwise. The partition $\{\mathcal{B}_1^0, \dots, \mathcal{B}_{J_0}^0\}$, its size J_0 , the vector ν^0 of block values and its zero coordinates are all unknown, and our goal is to produce estimates of those or related quantities that are accurate when n is large enough.

In particular, we investigate the asymptotic properties of three different but interrelated methods for to the recovery of the unknown mean vector μ^0 under the assumption [\(1\)](#).

*supported in part by NSF grant DMS-0631589 and a grant from the Pennsylvania Department of Health through the Commonwealth Universal Research Enhancement Program.




Figure 1: Signal (solid line) plus noise for the example described in Section 2.4.

The first methodology we study, which is the central focus of this work, is the *fused lasso* procedure of Friedman et al. (2007). The *fused lasso* is the penalized least squares estimator

$$\hat{\mu}^{FL} = \operatorname{argmin}_{\mu \in \mathbb{R}^n} \left\{ \sum_{i=1}^n (y_i - \mu_i)^2 + 2\lambda_{1,n} \|\mu\|_1 + 2\lambda_{2,n} \|\mu\|_{\text{TV}} \right\}, \quad (2)$$

where $\|\mu\|_1 \equiv \sum_{i=1}^n |\mu_i|$ is the ℓ_1^n norm and $\|\mu\|_{\text{TV}} \equiv \sum_{i=2}^n |\mu_i - \mu_{i-1}|$ the total variation seminorm of μ , respectively, and $(\lambda_{1,n}, \lambda_{2,n})$ are positive tuning parameters to be chosen appropriately. The solution to the convex program (2) can be computed in a fast and efficient way using the algorithm developed in Friedman et al. (2007), where the properties of the fused lasso solution are considered from the optimization theory standpoint. Our analysis will lead us to develop a modified version of the fused lasso, which we call the fused adaptive lasso, that has improved properties. Figure 2 shows an example of an fused adaptive lasso fit to the data displayed in Figure 1.

In our second approach, we turn to a different convex optimization program, namely

$$\begin{aligned} & \operatorname{argmin}_{\mu \in \mathbb{R}^n} \sum_{i=1}^n (y_i - \mu_i)^2 \\ \text{s.t. } & \|\mu\|_1 \leq L_n, \quad \|\mu\|_{\text{TV}} \leq T_n, \end{aligned} \quad (3)$$

for some nonnegative constants L_n and T_n . Notice that, in this alternative formulation, which is akin to the least squares method on sieves, a solution different from y is obtained provided $\|y\|_1 > L_n$ or $\|y\|_{\text{TV}} > T_n$. The link with the fused lasso estimator is clear: the objective function in the fused lasso problem (2) is the Lagrangian function of (3), and, in fact, the two problems are equivalent from the point of view convexity theory.

Our third and final method for the recovery a sparse and blocky signal is also related to sieve least square procedures, and is more naturally tailored to the model assumption (1). Specifically, we study the solution to the highly non-convex optimization problem

$$\begin{aligned} & \operatorname{argmin}_{\mu \in \mathbb{R}^n} \sum_{i=1}^n (y_i - \mu_i)^2 \\ \text{s.t. } & |\{i: \mu_i \neq 0\}| \leq S_n, \quad 1 + |\{i: \mu_i - \mu_{i-1} \neq 0, 2 \leq i \leq n\}| \leq J_n, \end{aligned} \quad (4)$$

where S_n and J_n are nonnegative constants. Although lack of convexity makes this problem computationally difficult when n is large, the theoretical relevance of this third formulation stems from the fact that (3) is, effectively, a convex relaxation of (4).

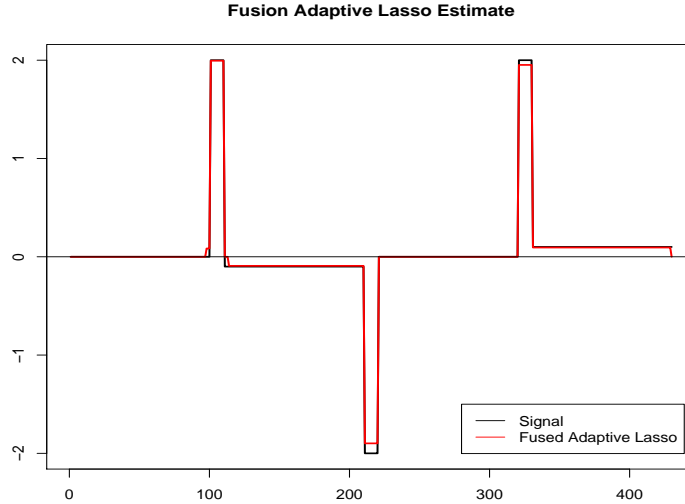


Figure 2: A fusion adaptive lasso estimate for the example from Section 2.4, using the most biased fusion estimator shown in Figure 3 the oracle threshold for the lasso penalty, as described in Section 2.3.

Our approach to the study of the estimators defined by (2), (3) and (4) is asymptotic, as we allow the block representation for the unobserved signal μ^0 to change with n in such a way that the problem of the recovery of a noisy signal under the model (1) may become increasingly difficult. Despite being quite closely related as optimization problems, from an inferential perspective the three procedures under investigation each sheds some light on different and, in some way, complementary aspects of this problem. In essence, our results provide conditions for the (sequences) of regularization parameters $\lambda_{1,n}$, $\lambda_{2,n}$, L_n , J_n and S_n to guarantee various degrees of recovery of μ^0 .

The idea of using the total variation seminorm in penalized least squares problem has been exploited and studied in many applications, for example in signal processing, parametric and nonparametric regression, image denoising. From the algorithmic viewpoint, this idea was originally brought up by Rudin et al. (1992); for more recent developments, see e.g., Dobson and Vogel (1997) and Caselles et al. (2007). See also DeVore (1998). The original motivation for this article was the recent work by Friedman et al. (2007), who devise efficient coordinate-wise descent algorithms for a variety of convex problems. In particular, they propose a novel approach based on penalized least squares problem using simultaneously the total variation and the ℓ_1 penalties, which favors solutions that are both blocky and sparse. In the classical nonparametric framework of statistical functional estimation, two important contributions in the development and analysis of total variation-based methods come from Mammen and van de Geer (1997) and Davies and Kovac (2001). Specifically, Mammen and van de Geer (1997) consider least squares splines with adaptively chosen knots and derive, among other things, consistency rates for both one and two-dimensional problems. Using a different approach, Davies and Kovac (2001) propose a very simple and effective procedure, the taut-string algorithm, to consistently estimate at an almost optimal rate the number and location of local maxima for an unknown function. Both methods impose virtually no assumptions on the degree of smoothness of the true underlying function. More recently, Boysen et al. (2008) study jump-penalized least square regression problems, where the underlying function is assumed to be a linear combination of a finite number of indicator functions of intervals in $[0, 1]$, and derive consistency rates under different metrics on functional spaces.

Our work differs from the contributions based on a nonparametric function estimation framework in various aspects, some of which are closely related to the methodology and scope of Friedman et al. (2007). First, we are only interested to the recovery of mean vectors under the model assumption (1), and do not necessarily view them as n values of function on $[0, 1]$. Nonetheless, we remark there is a simple reformulation

of our problem as nonparametric functional estimation one. In fact, suppose we observe n datapoints of the form

$$y_i = \frac{1}{n} \int_{\frac{i-1}{n}}^{\frac{i}{n}} \mu^0(t) dt + \epsilon_i, \quad i = 1, \dots, n,$$

from an unknown function $\mu^0: [0, 1] \rightarrow \mathbb{R}$. Setting $\mu_i^0 = \frac{1}{n} \int_{\frac{i-1}{n}}^{\frac{i}{n}} \mu^0(t) dt$ would return our original model (see also [Boysen et al., 2008](#), for a similar model). Furthermore, we are concerned with the simultaneous recovery of both the block partition and of the sparsity patten of μ^0 . Overall, our analysis yields conditions for consistency of the block partition and block sparsity estimates by model (2) and its variant described in Section 2.3, and explicit rates of consistency of both sieve solutions (3) and (4).

The article is organized as follows. In Section 2 we study the fused lasso estimator. After deriving in Section 2.1 an explicit formula for the fused lasso solution, we establish conditions under which both the fused lasso procedure are sparsistent, in the sense of a being weakly consistent estimator of the true partitions and of the set of nonzero coordinates of μ^0 . In Section 2.3 we propose a simple modification of the fused lasso, which we call the fused adaptive lasso, that achieves sparsistence under milder conditions and also allows to derive an oracle inequality for the empirical risk. Finally, in Section 3 we derive consistency rates for the estimators defined in (3) and (4), which depends explicitly on the parameters L_n and T_n , and of S and J , respectively. The proofs are relegated to the Appendix.

We conclude this introductory section by fixing the notation that we will be using throughout the article. For a vector $\mu \in \mathbb{R}^n$, we let $\mathcal{S}(\mu) = \{i: \mu_i \neq 0\}$ denote its support and $\mathcal{J}(\mu) = \{i: \mu_i = \mu_{i-1} \neq 0, i \geq 2\}$ the set of coordinates where μ changes its value. Furthermore, notice that we can always write

$$\mu = \sum_{j=1}^J \nu_j \mathbf{1}_{\mathcal{B}_j},$$

from some (possibly trivial) block partition $\{\mathcal{B}_1, \dots, \mathcal{B}_J\}$, with $1 \leq J \leq n$, and some vector $\nu \in \mathbb{R}^J$. Then, we will write $\mathcal{J}(\mu) = \{j: \nu_j \neq 0\}$ for the sets of non-zero blocks of μ . On a final note, although all the quantities defined so far may change with n , for ease of readability, we do not always make this dependence explicit in our notation.

2 Properties and refinements of the fused lasso estimator

The crucial feature of the fused lasso solution 2 that makes it ideal for the present problem is of being simultaneously blocky, because of the total variation penalty $\|\cdot\|_{\text{TV}}$, and sparse, because of the ℓ_1 penalty $\|\cdot\|_2$. The central goal of this section is to characterize the asymptotic behavior of the regularization parameters $\lambda_{1,n}$ and $\lambda_{2,n}$ so that, as $n \rightarrow \infty$, the blockiness and sparsity pattern of the the fused lasso estimates match the ones of the unknown signal μ^0 , with overwhelming probability. We first consider the fused lasso estimator as originally proposed in [Friedman et al. \(2007\)](#) and then a simple variant, the fused adaptive lasso, which has better asymptotic properties. For this modified version, we also derive an oracle inequality. We will make a simplifying assumption on the errors:

- (E) The errors ϵ_i , $1 \leq i \leq n$ are identically distributed centered Gaussian variables with variance σ_n^2 such that $\sigma_n \rightarrow 0$.

In the typical scenario we have in mind, $\sigma_n = \frac{\sigma}{\sqrt{n}}$. Assumption (E) is by no means necessary and it can be easily relaxed to the case of sub-Gaussian centered, as we point out later.

2.1 The fused lasso solution

Below, we provide a explicit formula for the fused lasso solution that offers some insight on its properties and suggests possible improvements. By inspecting (2), as both penalty functions $\|\cdot\|_1$ and $\|\cdot\|_{\text{TV}}$ are convex

and the objective function is strictly convex, $\widehat{\mu}^{FL}$ is uniquely determined as the solution to the subgradient equation

$$\widehat{\mu}^{FL} = y - \lambda_{1,n}s_1 - \lambda_{2,n}s_2, \quad (5)$$

where $s_1 \in \partial\|\widehat{\mu}^{FL}\|_1$ and $s_2 \in \partial\|\widehat{\mu}^{FL}\|_{TV}$. For a vector $x \in \mathbb{R}^n$, the subgradient $\partial\|x\|_1$ is a subset of \mathbb{R}^n consisting of vectors s such that $s_i = \text{sgn}(x_i)$, where, with some abuse of notation, we will denote with $\text{sgn}(\cdot)$ the (possibly set-valued) function on \mathbb{R} given by

$$\text{sgn}(x) = \begin{cases} 1 & \text{if } x > 0 \\ -1 & \text{if } x < 0 \\ z & \text{if } x = 0, \end{cases}$$

where z is any number in $[-1, 1]$. The subgradient $\partial\|x\|_{TV}$ has slightly more elaborated form, which is given in Lemma 6.1 in the Appendix.

An explicit expression for $\widehat{\mu}^{FL}$ can be obtained in terms of the *fusion* estimator

$$\widehat{\mu}^F = \underset{\mu \in \mathbb{R}^n}{\text{argmin}} \left\{ \sum_{i=1}^n (y_i - \mu_i)^2 + 2\lambda_2 \|\mu\|_{TV} \right\}. \quad (6)$$

(Notice that, by the same arguments above, $\widehat{\mu}^F$ is unique.) This fusion estimator solves a regularized least squares problem with a penalty on the total variation of the signal and works by fusing together adjacent coordinates that have a similar values to produce a blocky estimate of the form (1). As remarked in the introduction, Mammen and van de Geer (1997) and Davies and Kovac (2001) establish asymptotic properties of fusion-type estimators of slightly different nature than in this paper in a nonparametric functional estimation framework.

For a given solution $\widehat{\mu}^F$ to (6), there exists a block partition $\{\widehat{\mathcal{B}}_1, \dots, \widehat{\mathcal{B}}_{\widehat{J}}\}$ and a unique vector $\widehat{\nu} \in \mathbb{R}^{\widehat{J}}$ such that

$$\widehat{\mu}^F = \sum_{j=1}^{\widehat{J}} \widehat{\nu}_j 1_{\widehat{\mathcal{B}}_j}. \quad (7)$$

We take note that both the number \widehat{J} and the elements of the partition $\{\widehat{\mathcal{B}}_1, \dots, \widehat{\mathcal{B}}_{\widehat{J}}\}$ are random quantities, and that, by construction, no two consecutive entries of $\widehat{\nu}$ are identical. Using (7), the individual entries of the vector $\widehat{\nu}$ can be obtained explicitly, as shown next.

Lemma 2.1. *Let $\widehat{\nu} \in \mathbb{R}^{\widehat{J}}$ satisfy (7) and $\widehat{b}_j = |\widehat{\mathcal{B}}_j|$, for $1 \leq j \leq \widehat{J}$. Then,*

$$\widehat{\nu}_j = \frac{1}{\widehat{b}_j} \sum_{i \in \widehat{\mathcal{B}}_j} y_i + \widehat{c}_j,$$

where

$$\widehat{c}_1 = \begin{cases} -\frac{\lambda_{2,n}}{\widehat{b}_j} & \text{if } \widehat{\nu}_2 - \widehat{\nu}_1 > 0 \\ \frac{\lambda_{2,n}}{\widehat{b}_j} & \text{if } \widehat{\nu}_2 - \widehat{\nu}_1 < 0, \end{cases} \quad (8)$$

$$\widehat{c}_{\widehat{J}} = \begin{cases} \frac{\lambda_{2,n}}{\widehat{b}_j} & \text{if } \widehat{\nu}_{\widehat{J}} - \widehat{\nu}_{\widehat{J}-1} > 0 \\ -\frac{\lambda_{2,n}}{\widehat{b}_j} & \text{if } \widehat{\nu}_{\widehat{J}} - \widehat{\nu}_{\widehat{J}-1} < 0. \end{cases} \quad (9)$$

and, for $1 < j < \widehat{J}$,

$$\widehat{c}_j = \begin{cases} \frac{2\lambda_{2,n}}{\widehat{b}_j} & \text{if } \widehat{\nu}_{j+1} - \widehat{\nu}_j > 0, \widehat{\nu}_j - \widehat{\nu}_{j-1} < 0 \\ -\frac{2\lambda_{2,n}}{\widehat{b}_j} & \text{if } \widehat{\nu}_{j+1} - \widehat{\nu}_j < 0, \widehat{\nu}_j - \widehat{\nu}_{j-1} > 0 \\ 0 & \text{if } (\widehat{\nu}_j - \widehat{\nu}_{j-1})(\widehat{\nu}_{j+1} - \widehat{\nu}_j) = 1. \end{cases} \quad (10)$$

By Proposition 1 in [Friedman et al. \(2007\)](#), the fused lasso estimator is obtained by soft-thresholding of the individual coordinates of $\hat{\mu}^F$, so that we immediately obtain the next result.

Corollary 2.2. *The fused lasso estimator $\hat{\mu}^{FL}$ is*

$$\hat{\mu}_i^{FL} = \begin{cases} \hat{\mu}_i^F - \lambda_{1,n} & \hat{\mu}_i^F \geq \lambda_1 \\ 0 & |\hat{\mu}_i^F| < \lambda_{1,n} \\ \hat{\mu}_i^F + \lambda_{1,n} & \hat{\mu}_i^F \leq -\lambda_1, \end{cases} \quad i = 1, \dots, n, \quad (11)$$

where $\hat{\mu}^F$ is the fusion estimator.

Remarks

1. As is apparent from Lemma 2.1, the individual blocks found by the fusion solution $\hat{\mu}^F$ are each biased by a term whose magnitude depends directly on the regularization parameter $\lambda_{2,n}$ and, inversely, on the size of the estimated block itself. That is, the larger the estimated blocks the smaller the effect of the bias. This term is simply a vertical shift which is positive if the block is a local local maximum, negative if it is a local minimum, and is absent otherwise. See Figure 3. It is worth pointing out that we observed the same type of behavior for the solution obtained using the taut-string algorithm of [Davies and Kovac \(2001\)](#), where the magnitude of the vertical shift is controlled by the size of the tube around the integrated process.
2. The regularization parameter $\lambda_{1,n}$, which modulates the size of the sparsity penalty, also causes some bias effect due to the soft-thresholding. However, unlike the bias determined by the total variation penalty, this second type of bias is of the same magnitude for all the non-zero coordinates, as can be seen directly from (11). An easy fix, which is considered in Section 2.3, would be to adaptively penalize the estimated blocks differently, depending on their sizes, with larger blocks penalized less.

2.2 Sparsistency for the fused lasso

In this section we provide conditions under which the block partition $\{\mathcal{B}_1^0, \dots, \mathcal{B}_{J_0}^0\}$ and the block sparsity pattern $\mathcal{JS}(\mu^0)$ of μ^0 can be estimated consistently (see equation 1). We break down our analysis into two parts, dealing separately with the fusion estimator $\hat{\mu}^F$ first, which can be used to recover $\{\mathcal{B}_1^0, \dots, \mathcal{B}_{J_0}^0\}$, and then with the fused lasso solution $\hat{\mu}^{FL}$, from which the set $\mathcal{JS}(\mu^0)$ can be estimated. In Section 2.3, we show how this second task can be accomplished more effectively by a modified version of the fused lasso estimator.

2.2.1 Recovery of true blocks by fusion only

We first derive sufficient conditions for the fusion estimator to recover correctly the block partition of μ^0 . Let $\mathcal{J}_0 = \mathcal{J}(\mu^0)$ be the set of jumps of μ^0 and $J_0 = |\mathcal{J}_0| + 1$ the cardinality of the associated block partition. Similarly, let $\hat{\mathcal{J}} = \mathcal{J}(\hat{\mu}^F)$ be the set of jumps for the fusion estimate, given in (7).

Theorem 2.3. *Assume (E) and (1). If, for some $\delta > 0$,*

1. $\frac{\lambda_{2,n}}{\sigma_n} \rightarrow \infty$ and $\frac{\lambda_{2,n}}{\sigma_n \sqrt{\log(n-J_0)}} > \frac{1}{2\sqrt{2}}(1 + \delta)$,
2. $\frac{\alpha_n}{\sigma_n} \rightarrow \infty$, $\frac{\alpha_n}{\sigma_n \sqrt{\log J_0}} > \sqrt{16}(1 + \delta)$ and $\lambda_{2,n} < \frac{\alpha_n}{8}$,

where $\alpha_n = \min_{i \in \mathcal{J}_0} |\mu_i^0 - \mu_{i-1}^0|$, then

$$\lim_n \mathbb{P} \left(\{\hat{\mathcal{J}} = \mathcal{J}_0\} \cap \{\text{sgn}(\hat{\mu}_i^F - \hat{\mu}_{i-1}^F) = \text{sgn}(\mu_i^0 - \mu_{i-1}^0), \forall i \in \mathcal{J}_0\} \right) = 1 \quad (12)$$

Remark.

1. In the proof of Theorem 2.3, instead of Slepian's inequality, one could use Markov's inequality and well-known bounds on the supremum of centered subgaussian vectors (see, e.g., Lemma 2.3 in Massart, 2007) to derive slightly stronger sufficient conditions for (12), which however hold for the larger class of subgaussian errors. We give these conditions without a proof:

$$(a) \lim_n \frac{\sigma_n \sqrt{2 \log |\mathcal{J}(\mu_0)| + 2\lambda_{2,n}}}{\alpha_n} = 0,$$

$$(b) \lim_n \frac{\lambda_{2,n}}{\sigma_n \sqrt{\log |\mathcal{J}_0^c|}} = \infty.$$

Furthermore, the errors need not be identically distributed. In fact, the proof of the Theorem holds almost unchanged if, for example, one only assumes that the individual variances are of order $O(1/\sqrt{n})$.

2. Equation (12) actually implies not only that \mathcal{J}_0 can be consistently estimated, but also that the true signs of the jumps can be recovered with overwhelming probability, a feature known in the lasso literature as sign consistency (see, e.g. Wainwright, 2006; Zhao and Yu, 2006). In the present settings, sign consistency of the fusion estimate implies the following, nice feature of $\hat{\mu}^F$:

Corollary 2.4. *The fusion estimator $\hat{\mu}^F$ can consistently recover the local maxima and local minima of μ^0 .*

3. The magnitude α_n of the smallest jumps of μ^0 is a fundamental quantity, whose asymptotic behavior determines whether recovery of the true blocks obtains. In particular, if α_n vanishes at a rate faster than $1/\sigma_n$, then no recovery is possible. In a way, this guarantees some form of asymptotic distinguishability that prevents adjacent blocks from looking too similar.

2.2.2 Recovery of true blocks and true non-zero coordinates by the fused lasso

Let $\mathcal{JS}_0 = \mathcal{JS}(\mu^0)$ be set of non-zero blocks of μ^0 and $K_0 = |\mathcal{JS}_0|$ its cardinality. Let $\widehat{\mathcal{JS}} = \mathcal{JS}(\hat{\mu}^{FL})$ be the equivalent quantity defined using the fused lasso estimate $\hat{\mu}^{FL}$. Consider the event

$$\mathcal{R}_{1,n} = \left\{ \mathcal{JS}_0 = \widehat{\mathcal{JS}} \right\} \cap \left\{ \text{sgn}(\hat{\nu}_j) = \text{sgn}(\nu_j^0), \forall j \in \mathcal{JS}_0 \right\}$$

that soft-thresholding $\hat{\mu}^F$ with penalty $\lambda_{1,n}$ will return the nonzero blocks of μ^0 .

Theorem 2.5. *If the conditions of Theorem 2.3 are satisfied and, for some $\delta > 0$,*

1. $\frac{\lambda_{1,n} \sqrt{b_{\min}^0}}{\sigma_n} \rightarrow \infty$ and $\frac{\lambda_{1,n} \sqrt{b_{\min}^0}}{\sigma_n \sqrt{\log(J_0 - K_0)}} > 2\sqrt{2}(1 + \delta)$;
2. $\frac{2\lambda_{2,n}}{b_{\min}^0} < \frac{\lambda_{1,n}}{2}$, for all n large enough;
3. $\frac{\rho_n \sqrt{b_{\min}^0}}{\sigma_n} \rightarrow \infty$, $\frac{\rho_n \sqrt{b_{\min}^0}}{\sigma_n \sqrt{\log K_0}} > \sqrt{18}(1 + \delta)$ and $\lambda_{1,n} < \frac{\rho_n}{3}$ for all n large enough;
4. $\frac{2\lambda_{2,n}}{b_{\min}^0} < \frac{\rho_n}{3}$, for all n large enough,

where $\rho_n = \min_{j \in \mathcal{K}_0} |\nu_j^0|$ and $b_{\min}^0 = \min_{1 \leq j \leq J_0} \{b_j^0\}$, then,

$$\lim_n \mathbb{P}(\mathcal{R}_{1,n}) = 1.$$

Remarks.

1. As it was the case for Theorem 2.3, the assumption of Gaussian errors is not essential and can be relaxed, and, in fact, Remark 1. above still applies.

2. The previous result implies that the fused lasso is not only consistent, but, in fact, sign consistent, so that the signs of the non-zero blocks are estimated correctly.
3. The magnitude ρ_n of the smallest non-zero block value cannot decrease to zero too fast, otherwise the sparsity pattern cannot be fully recovered, just as we pointed out in Remark 3. above for the fusion solution.
4. The conditions of Theorem 2.5 appear to be quite cumbersome, mainly for two reasons. First, the regularization parameters $\lambda_{1,n}$ and $\lambda_{2,n}$ interact with each other. As a result, it appears necessary to impose assumption 2. in order to guarantee that the two different bias terms they each determine will not disrupt the recovery process. Secondly, it seems necessary to keep track of the size b_{\min}^0 of the minimal block. This additional bookkeeping is due to the fact that the sparsity penalty is enforced globally, in the sense that all coordinates are penalized in equal amount, thus ignoring the fact that longer blocks require less regularization (see Remark 1. after Lemma 2.1).

2.3 The fused adaptive lasso: sparsistency and an oracle inequality

Motivated by the stringent nature of the conditions of Theorem 2.5, below we propose a refinement of the fused lasso estimator, which we call the *fused adaptive lasso*. Overall, this slightly different estimator enjoys better asymptotic properties than the fused lasso, at no additional complexity cost.

The fused adaptive lasso is obtained with the following two-step procedure:

1. Fusion step: compute the fusion solution $\hat{\mu}^F$ using the fusion regularization parameter $\lambda_{2,n}$, as in (6), and the corresponding block-partition $(\hat{\mathcal{B}}_1, \dots, \hat{\mathcal{B}}_{\hat{J}})$ (see equation 7). Obtain

$$\hat{\mu}^{AF} = \sum_{j=1}^{\hat{J}} \bar{y}_j \mathbf{1}_{\hat{\mathcal{B}}_j}, \quad (13)$$

where

$$\bar{y}_j = \frac{1}{b_j} \sum_{i \in \hat{\mathcal{B}}_j} y_i, \quad 1 \leq j \leq \hat{J}.$$

2. Adaptive lasso step: compute the fused adaptive lasso solution

$$\hat{\mu}^{FAL} = \operatorname{argmin}_{\mu \in \mathbb{R}^n} \|\hat{\mu}^{AF} - \mu\|_2^2 + \sum_{i=1}^n \tilde{\lambda}_i |\mu_i|, \quad (14)$$

where the n -dimensional random vector $\tilde{\lambda}$ of penalties is

$$\lambda = \lambda_1 \sum_{j=1}^{\hat{J}} \frac{1}{\sqrt{b_j}} \mathbf{1}_{\hat{\mathcal{B}}_j}, \quad (15)$$

with $\lambda_{1,n}$ as the ℓ_1 regularization parameter.

Remarks

1. The fused adaptive lasso differs from the fused lasso in two fundamental aspects. First, as easily seen from Equation (13), the bias term in the fusion solution due to the terms c_j , which depends on the regularization parameter $\lambda_{2,n}$, is absent (see Lemma 2.1). Equivalently, the fusion estimator is only used to estimate the block partition of μ^0 , and, provided this estimated block partition is correct, the block values are estimated unbiasedly with the sample averages. Using the fusion procedure as an estimator of the block partition rather than μ^0 has the other advantage of decoupling the estimation

from the model selection problem, thus freeing, to some extent, the user from the task of carefully choosing an optimal penalty $\lambda_{2,n}$. In fact, recovery of the true partition can be obtained even if the problem is overpenalized and, therefore, the resulting estimator $\widehat{\mu}^F$ is highly biased.

Secondly, the penalty terms used for thresholding individual blocks are rescaled by the squared root of the length of the estimated blocks. The rationale for using this rescaling is very simple. In fact, suppose that, for some j_1, j_2 , $\widehat{b}_{j_1} \gg \widehat{b}_{j_2}$. Since the variance of the j -th block average \bar{y}_j is $\frac{\sigma^2}{b_j}$, \bar{y}_{j_1} has a much smaller standard error than \bar{y}_{j_2} and, therefore, should be penalized less heavily. The adequate reduction in the sparsity penalty of \bar{y}_{j_1} versus \bar{y}_{j_2} is precisely the difference in their standard errors, hence the choice of rescaling by square root of the block lengths. The advantage of adaptively thresholding the block values in this manner is that the procedure will be more effective at identifying longer non-zero blocks whose values are quite close to 0.

In Section 2.4 we explain both these improvements concretely with a numerical example.

2. In step 2. the vector $\widehat{\mu}$ is straightforward to compute via soft-thresholding of the individual coordinates of $\widehat{\mu}^{AF}$ with coordinate-dependent thresholds:

$$\widehat{\mu}_i^{FAL} = \begin{cases} \widehat{\mu}_i^{AF} - \lambda_i & \widehat{\mu}_i^{AF} \geq \lambda_i \\ 0 & |\widehat{\mu}_i^{AF}| < \lambda_i \\ \widehat{\mu}_i^{AF} + \lambda_i & \widehat{\mu}_i^{AF} \leq -\lambda_i, \end{cases} \quad 1 \leq i \leq n.$$

3. Instead of the soft-thresholded block estimate of step 2., whose coordinates are biased by the threshold penalties, one may consider instead the corresponding estimate $\widetilde{\mu}$ based on the hard-threshold, where

$$\widetilde{\mu}_i = \widehat{\mu}_i^{AF} \mathbf{1}\{|\widehat{\mu}_i^{AF}| \geq \lambda_i\}, \quad 1 \leq i \leq n.$$

One of the asymptotic advantages of the fused adaptive lasso versus the ordinary fused lasso is that block recovery obtains under milder conditions than Theorem 2.5, without the need to consider the fusion penalty parameter $\lambda_{2,n}$ and the length of the minimal block. In some sense, the fused adaptive lasso can adapt more flexibly to the block sparsity than the fused lasso.

Proposition 2.6. *Assume that the conditions of Theorem 2.3 are satisfied. Then*

$$\lim_n \mathbb{P}\{\mathcal{R}_{1,n}\} = 1$$

if, for some $\delta > 0$,

1. $\frac{\lambda_{1,n}}{\sigma_n} \rightarrow \infty$ and $\frac{\lambda_{1,n}}{\sigma_n \sqrt{\log(J_0 - K_0)}} > \sqrt{2}(1 + \delta)$;
2. $\frac{\rho_n}{\sigma_n} \rightarrow \infty$, $\frac{\rho_n}{\sigma_n \sqrt{\log K_0}} > 2\sqrt{2}(1 + \delta)$ and $\lambda_{1,n} < \frac{\rho_n}{2}$ for all n large enough,

where $\rho_n = \min_{j \in \mathcal{K}_0} |\nu_j^0|$.

A second advantage of the fused adaptive lasso stems from the oracle property derived below. Consider the ideal situation where we have available an oracle letting us know the K^0 sets $\mathcal{B}_{j_k}^0$, $k = 1, \dots, K^0$, of the true block partition of μ^0 for which $|\nu_{j_k}^0| > \sigma_n / \sqrt{b_{j_k}^0}$. Notice that, from this information, one can recover the true partition. The oracle estimate $\widehat{\mu}^O$ is the vector with coordinates

$$\widehat{\mu}_i^O = \begin{cases} \frac{1}{b_{j_k}^0} \sum_{z \in \mathcal{B}_{j_k}^0} y_z & \text{if } i \in \mathcal{B}_{j_k} \\ 0 & \text{otherwise.} \end{cases}$$

This procedure amounts to setting to 0 the estimates for the coordinates belonging to blocks whose true mean value is smaller than $\sigma_n/\sqrt{b_j^0}$. The corresponding ideal risk is

$$\mathbb{E}\|\widehat{\mu}^O - \mu^0\|_2^2 = \sum_i \sum_{j_k} 1\{i \in \mathcal{B}_{j_k}^0\} \min \left\{ \frac{\sigma_n^2}{b_{j_k}^0}, (\nu_j^0)^2 \right\} = K_0 \sigma_n^2 + \sum_{j \notin \mathcal{JS}_0} b_j^0 (\nu_j^0)^2. \quad (16)$$

Note, in particular, that

$$\mathbb{E}\|\widehat{\mu}^O - \mu^0\|_2^2 \leq \sum_i \min \{ \sigma_n^2, \mu_i^2 \},$$

with equality if and only if $b_j^0 = 1$ for all j , where the expression on the right hand side is the ideal risk for the oracle estimator based on thresholding of individual coordinates, rather than of blocks. Therefore, if μ^0 has a block structure, as is assumed here, this different oracle will be able to achieve a smaller ideal risk.

Before stating our oracle result, we need some additional notation. Recall that any $\mu \in \mathbb{R}^n$ can always be written as

$$\mu = \sum_{j=1}^J \nu_j 1_{\mathcal{B}_j}, \quad (17)$$

for some (possibly trivial) block partition $(\mathcal{B}_1, \dots, \mathcal{B}_J)$ of $\{1, \dots, n\}$, with $J \leq n$. Let μ^1 and μ^2 be vectors in \mathbb{R}^n with block partitions $\{\mathcal{B}_1^1, \dots, \mathcal{B}_{J_1}^1\}$ and $\{\mathcal{B}_1^2, \dots, \mathcal{B}_{J_2}^2\}$, respectively, where $J_1, J_2 \leq n$. Then, they satisfy (17), for some vectors $\nu^1 \in \mathbb{R}^{J_1}$ and $\nu^2 \in \mathbb{R}^{J_2}$, respectively. Let $\{\mathcal{L}_1, \dots, \mathcal{L}_m\}$ be the partition of $\{1, \dots, n\}$ obtained as the refinement of the block partitions of μ^1 and μ^2 , i.e. for every $l = 1, \dots, m$, $\mathcal{L}_l = \mathcal{B}_{j_1}^1 \cap \mathcal{B}_{j_2}^2$, for some j_1 and j_2 . We define the quantity

$$\mathcal{JS}(\mu^1; \mu^2) = \{l: \mathcal{L}_l = \mathcal{B}_{j_1}^1 \cap \mathcal{B}_{j_2}^2, \nu_{j_1}^1 \neq 0\}.$$

Theorem 2.7. *Assume that μ^0 satisfies (1) and that*

$$\alpha_n = o\left(\sqrt{\frac{\log n}{n}}\right). \quad (18)$$

Let $\sigma_n^2 = \frac{\sigma^2}{n}$, $\lambda_{2,n} = \sqrt{\sigma_n^2 \log n}$ and $\lambda_{1,n} = 2\sqrt{\sigma_n^2 \log \widehat{J}}$, where \widehat{J} is obtained by solving the fusion problem (6) in the first step of the adaptive fused-lasso procedure. For any vector $\mu \in \mathbb{R}^n$, set

$$V(\mu) = 32|\mathcal{JS}(\mu; \mu^0)|\sigma_n^2 \log J_0.$$

Then, for any $\delta \in [0, 1)$,

$$\lim_n \mathbb{P} \left\{ \|\widehat{\mu}^{FAL} - \mu_0\|_2^2 \leq \frac{2 + \delta}{2 - \delta} \inf_{\mu \in \mathbb{R}^n} \{V(\mu) + \|\mu - \mu^0\|_2^2\} \right\} = 1. \quad (19)$$

Remarks

1. The assumption in equation (18) is crucial in our proof, as it guarantees that recovery of the true block partition of μ^0 by fusion, which is necessary for mimicking the oracle solution $\widehat{\mu}^O$, is feasible.
2. The proof of theorem 2.7 shows that $V(\mu)$ is minimized by vectors such that

$$|\mathcal{JS}(\mu; \mu^0)| = |\mathcal{JS}(\mu^0)| = K^0,$$

i.e. vectors whose block partition matches the the true block partition. Therefore equation (19) shows that the adaptive fused-lasso achieves the same oracle rates granted by ideal risk (16), up a term that is logarithmic in J_0 .

3. If it is further assumed that $\|\mu^0\|_\infty < C$ uniformly in n , for some constant C , the result (19) can be strengthened to

$$\mathbb{E}\|\widehat{\mu} - \mu_0\|_2^2 \leq \frac{2 + \delta}{2 - \delta} \inf_{\mu \in \mathbb{R}^n} \{V(\mu) + \|\mu - \mu^0\|_2^2\} + o(1).$$

2.4 An illustrative Example

We discuss a stylized numerical example for the purpose of clearly illustrating the two crucial advantages of the fused adaptive lasso, namely the use of the fusion penalty only for recovering the true block partition, and the block-dependent rescaling of the lasso penalty. See Remark 1. before Proposition 2.6 for details.




Figure 3: Different fusion estimates for the data described in section 2.4. The dashed line correspond to the true mean vector, while the three lines correspond to the fusion estimates with different regularization parameters.

We simulate one sample according to the model

$$y_i = \mu_i^0 + \epsilon_i,$$

where

$$\mu_i^0 = \begin{cases} 0 & 1 \leq i \leq 100, \\ 2 & 101 \leq i \leq 110, \\ -0.1 & 111 \leq i \leq 210, \\ -2 & 211 \leq i \leq 220, \\ 0 & 221 \leq i \leq 320, \\ 2 & 321 \leq i \leq 330, \\ 0.1 & 331 \leq i \leq 430, \end{cases}$$

and the errors are independent Gaussian variables with mean zero and standard deviation $\sigma = 0.2$. Figure 1 shows the data along with the true signal. Notice that some of the coordinates of μ^0 are in absolute value less than σ , a fact that, as we will see, if μ^0 were not blocky, would make the recovery of those coordinates infeasible. Figure 3 portrays the simulated data and three fusion estimates $\hat{\mu}^F$, each of them solving (6) for different values of $\lambda_{2,n}$. The dashed line corresponds to the true mean vector μ^0 . The excessive amount of penalization is apparent from the large bias in all these estimates. Nonetheless, the block partition each of these estimates produce match, in fact, very closely the true block partition.

Figure 4 shows the modified fusion estimate $\hat{\mu}^{AF}$ given in (13) using the fusion estimate from Figure 3 with the largest amount of bias, along with the true mean vector μ^0 , displayed as a dashed line. Because the block partition was estimated correctly, the estimate $\hat{\mu}^{AF}$ is almost indistinguishable from the true vector μ^0 . For this particular dataset, the adaptive lasso step would set to zero correctly the first and fifth block,




Figure 4: The modified fusion estimate $\hat{\mu}^{AF}$ of equation (13), using the fusion estimate from Figure 3 with the lowest total variation. The dashed gray line, which is almost indistinguishable from the estimate is the true signal μ^0 . The vertical lines enclose the third and seventh blocks, whose value is in magnitude half the standard deviation of the errors.

but not the third and seventh blocks, which in Figure 4 are enclosed by black vertical lines. In fact, although the true value of those blocks is in magnitude half the standard deviation of the errors, σ , the standard error for both the block estimates is roughly $\sigma/10$. This is taken into account in the adaptive lasso step, but not in the lasso step, where even the ideal soft threshold, i.e. σ , would be too high, thus incorrectly setting to zero both these blocks.

Finally, we simulated 1000 datasets according to the model described here and compute the empirical mean squared errors for the fused adaptive lasso estimates, using for the penalty terms the values indicated in Theorem 2.7. Figure 5 shows the histogram of the empirical mean squared errors, with the vertical line representing the true mean squared error $\frac{1}{n}\mathbb{E}\|y - \mu^0\|^2$, namely σ^2 . Notice how the empirical mean squared errors are larger than the true value, the usual price paid for adaptivity. We remark that there are other ways of choosing the penalty parameters, for example using cross validation.

3 Sieve Methods

In this section we study the rates of convergence for the sieve least squares solutions (3) and (4). For convenience, consistency is measured with respect to the normalized Euclidian norm $\|x\|_n = \frac{1}{\sqrt{n}}\sqrt{\sum_{i=1}^n x_i^2}$. Accordingly, we change our assumption on the errors as follows:

- (E') The errors $(\epsilon_1, \dots, \epsilon_n)$ are independent sub-gaussian variables with variances bounded by σ^2 , uniformly in n .

Notice that the results and settings of previous sections can be adapted in a straightforward way to the present framework.

We first study the estimator given in (3). To that end, consider the class of vectors

$$\mathcal{C}_{\text{TV}}(T_n) = \{\mu \in \mathbb{R}^n : \|\mu\|_{\text{TV}} \leq T_n, \|\mu\|_{\infty} \leq C\},$$




Figure 5: Distributions of the empirical mean squared errors from 1000 simulations from the model described on Section 2.4 using the fused adaptive lasso with penalty parameters chosen according to Theorem 2.7. The vertical line represents σ^2 .

where C is a finite constant that does not depend on n , and the ℓ_1 ball of radius L_n

$$\mathcal{C}_{\ell_1}(L_n) = \{\mu \in \mathbb{R}^n : \|\mu\|_1 \leq L_n\},$$

with both numbers T_n and L_n being allowed to grow unboundedly with n . Then, we can rewrite (3) as

$$\hat{\mu}^{TL} = \operatorname{argmin}_{\mu \in \mathcal{C}_{\text{TV}}(T_n) \cap \mathcal{C}_{\ell_1}(L_n)} \|y - \mu\|_2^2.$$

Below we derive the consistency rate for $\hat{\mu}^{TL}$ in terms of the sequences T_n and L_n by dealing separately with the two sieves.

Theorem 3.1. *Assume (E') and $\mu_n^0 \in \mathcal{C}_{\ell_1}(L_n) \cap \mathcal{C}_{\text{TV}}(T_n)$. Let*

$$\hat{\mu}^T = \inf_{\mu \in \mathcal{C}_{\text{TV}}(T_n)} \|y - \mu\|_2^2. \quad (20)$$

and

$$\hat{\mu}^L = \operatorname{argmin}_{\mu \in \mathcal{C}_{\ell_1}(L_n)} \|y - \mu\|_2.$$

Then,

$$\|\hat{\mu}^F - \mu^0\|_n = O_P\left(T_n^{1/3} n^{-1/3}\right),$$

so that $\hat{\mu}^F$ is consistent provided that $T_n = o(n)$, and

$$\|\hat{\mu}^L - \mu^0\|_n = O_P\left(\sqrt{\frac{L_n(\log n)^{3/2}}{n}}\right), \quad (21)$$

so that $\hat{\mu}^L$ is consistent provided that

$$L_n = o\left(\frac{n}{(\log n)^{3/2}}\right).$$

As a result,

$$\|\hat{\mu}^{TL} - \mu^0\|_n = O_P\left(\frac{L_n(\log n)^{3/2}}{n} \wedge \left(\frac{T_n}{n}\right)^{1/3}\right). \quad (22)$$

Remarks.

1. It appears that the requirement for the vectors in $\mathcal{C}_{\text{TV}}(T_n)$ to be uniformly bounded cannot be relaxed without affecting negatively the rate of consistency or without introducing additional assumptions. See for example Theorem 9.2 in [van de Geer \(2000\)](#).
2. The rate of consistency for $\hat{\mu}^F$ should be compared with the analogous rate derived in Theorem 9 of [Mammen and van de Geer \(1997\)](#) for the penalized version of the least squares problem (20).
3. The rate given in (21) is not the sharpest possible. In fact, an application of Theorem 5 of [Donoho and Johnstone \(1994\)](#) yields for $\hat{\mu}^L$ the improved minimax rate

$$\sqrt{\frac{L_n}{n}} (\log n)^{1/4}$$

for the case of *iid* Gaussian errors, from which we can infer a maximal rate of growth $L_n = o\left(\frac{n}{\sqrt{\log n}}\right)$.

4. We make no claims that the rate given (22), which is just the minimum of the rates for two separate sieve least squares problems, is sharp. Better rates may be obtained from sharper estimate of the metric entropy of the set $\mathcal{C}_{\ell_1}(L_n) \cap \mathcal{C}_{\text{TV}}(T_n)$.
5. **On the relationship between L_n and T_n .**

The total variation and ℓ_1 constraints are not independent of each other. One can easily verify that

$$T_n^{\max} \equiv \max_{x \in \mathcal{C}_{\ell_1}(L_n)} \|x\|_{\text{TV}} = 2L_n.$$

On the other hand, every vector $x \in \mathbb{R}^n$ such that $\|x\|_{\text{TV}} = T_n$ can be written as

$$x = m + t,$$

where $\|t\|_{\text{TV}} = T_n$, $m = 1_n \bar{x}_n$, with $\bar{x}_n = \frac{1}{n} \sum_i x_i$, and $\frac{1}{n} \sum_i m_i t_i = 0$. Notice that m can be estimated at the rate $\frac{1}{\sqrt{n}}$ so the convergence rates for $\hat{\mu}^T$ depends on how well t can be estimated. Next, notice that

$$L_n^{\max} \equiv \max_{x \in \mathcal{C}_{\text{TV}}(T_n), x=m+t} \|t\|_1 = \frac{T_n}{2} \frac{n}{n-1},$$

where $m + t$ is the decomposition of x discussed above. Therefore, over the set $\mathcal{C}_{\text{TV}}(T_n) \cap \mathcal{C}_{\ell_1}(L_n)$, we obtain the relationship

$$T_n^{\max} \sim 2L_n^{\max}. \quad (23)$$

Our final result concerns the estimator resulting from the non-convex sieve least squares problem (4). Define the set

$$\mathcal{C}(S_n, J_n) = \{\mu \in \mathbb{R}^n : |\mathcal{S}_n(\mu)| \leq S_n\} \cap \{\mu \in \mathbb{R}^n : |\mathcal{J}_n(\mu)| + 1 \leq J_n\},$$

consisting of vectors in \mathbb{R}^n that have at most S_n non-zero coordinates and take on at most J_n different values. We further impose the following, fairly weak assumption, which does not preclude the coordinates of μ^0 from becoming increasingly large in magnitude:

- (R) The set $\mathcal{C}(S_n, J_n)$ is contained in a S_n -dimensional cube centered at the origin with volume R_n such that

$$\log R_n = o(n).$$

Theorem 3.2. *Assume (E') and (R) and let $\hat{\mu}^{SJ} = \operatorname{argmin}_{\mu \in \mathcal{C}(S_n, J_n)} \|y - \mu\|_2^2$.*

1. If $S_n = o\left(\frac{n}{\log n}\right)$, then

$$\|\hat{\mu}^{SJ} - \mu^0\|_n = O_P\left(\sqrt{\frac{J_n}{n}}\right). \quad (24)$$

2. When $S_n = n$, (24) still holds, provided $J_n = o\left(\frac{n}{\log n}\right)$.

Remarks.

1. Assumption (R) was introduced for convenience to simplify the proof and it is possible that it may be relaxed.
2. The rate on S_n is in accordance with the persistence rate derived in Greenshtein (2006, Theorem 1) for related least squares regression problems on sieves.
3. If J_0 is bounded, uniformly in n , the consistency rate we obtain is parametric. See Boysen et al. (2008) for a similar result.

4 Discussion and future directions

In this work we tackle the task of estimating a blocky and sparse signal using three different methodologies, whose asymptotic properties we investigate. We first study the fused lasso estimator proposed in Friedman et al. (2007) and derive conditions under which it recovers with overwhelming probability and for n large enough the block partition and thdifferent one, with better asymptotic guarantees. We also study consistency rates of sieve least square problems under two types of constraints, one on the maximal radiuses of the ℓ_1 and $\|\cdot\|_{\text{TV}}$ balls, and the other on the maximal number of blocks and non-zero coordinates. Overall, these results complement each other in providing different types of asymptotic information for the for the task at hand and complement other analyses already existing in the statistical literature.

There are a number of generalizations of the results derived in this work. We mention only the ones that seem the most natural to us. The first extension involves considering a corrupted version of a signal $\mu^0 \in \mathbb{R}^n \times \mathbb{R}^n$, corresponding to the denoising problem of $n \times n$ images, for which total variation methods have proved quite effective. Another interesting direction would be to assume a known slowly-varying variance function, for example with given Lipschitz constant, and incorporate this information directly into the penalty functions for the fused adaptive lasso. We conjecture that it is possible to generalize our techniques and results to characterize these more complex settings. Finally, we believe it would be of value to investigate the possibility of building confidence balls and, in particular, confidence bands for the entire signal or for some of its local maxima or minima based on the estimators we consider here.

5 Acknowledgments

I am extremely thankful to Larry Wasserman for bringing this problem to my attention, and for many long discussions and indispensable advice.

6 Appendix

Lemma 6.1. *Let $\|\cdot\|_{\text{TV}}: \mathbb{R}^k \rightarrow \mathbb{R}$ be the fused penalty $\|x\|_{\text{TV}} = \sum_{i=2}^k |x_i - x_{i-1}|$. Then $\|\cdot\|_{\text{TV}}$ is convex and, for any $x \in \mathbb{R}^k$, the subdifferential $\partial\|x\|_{\text{TV}}$ is the set of all vectors $s \in \mathbb{R}^k$ such that*

$$s_i = \begin{cases} -w_2 & \text{if } i = 1 \\ w_i - w_{i+1} & \text{if } 1 < i < k \\ w_k & \text{if } i = k \end{cases} \tag{25}$$

where $w_i = \text{sgn}(x_i - x_{i-1})$, for $2 \leq i \leq k$.

Proof of Lemma 6.1. Let L be a $(k-1) \times k$ matrix with entries $L_{i,i} = -1$ and $L_{i,i+1} = 1$ for $1 \leq i \leq (k-1)$ and 0 otherwise. Then, for any $x \in \mathbb{R}^k$, $\|x\|_{\text{TV}} = \|Lx\|_1$. Convexity of $\|\cdot\|_{\text{TV}}$ follows from the fact that

it is the composition of a linear functional by the ℓ_1 norm, which is convex. Next, by the definition of the subdifferential of the ℓ_1 norm, for any $y \in \mathbb{R}^k$

$$\|Ly\|_1 \geq \|Lx\|_1 + \langle L(y-x), w \rangle \quad (26)$$

holds if and only if $w \in \mathcal{W}_x \subset \mathbb{R}^{k-1}$, where \mathcal{W}_x is the set of all vectors w such that $w_i = \text{sgn}((Lx)_i)$. Equation (26) is equivalent to

$$\|y\|_{\text{TV}} \geq \|x\|_{\text{TV}} + \langle y-x, s \rangle,$$

for each k -dimensional vector s such that $s = L^\top w$ for some $w \in \mathcal{W}_x$. This set is described by equation (25) and is, therefore, $\partial\|x\|_{\text{TV}}$. \blacksquare

Proof of Lemma 2.1. From the subgradient condition (5) with $\lambda_{1,n} = 0$, we obtain

$$\widehat{v}_j = \frac{1}{\widehat{b}_j} \sum_{i \in \widehat{\mathcal{B}}_j} y_i = \frac{1}{\widehat{b}_j} \sum_{i \in \widehat{\mathcal{B}}_j} y_i - \frac{\lambda_{2,n}}{\widehat{b}_j} \sum_{i \in \widehat{\mathcal{B}}_j} s_i.$$

Using (25), a simple telescoping argument leads to

$$\sum_{i \in \mathcal{B}_j} s_i = w_{i_j} - w_{i_{j+1}} = \begin{cases} 2 & \text{if } (\widehat{v}_{j+1} - \widehat{v}_j) > 0, (\widehat{v}_j - \widehat{v}_{j-1}) < 0 \\ -2 & \text{if } (\widehat{v}_{j+1} - \widehat{v}_j) < 0, (\widehat{v}_j - \widehat{v}_{j-1}) > 0 \\ 0 & \text{if } (\widehat{v}_j - \widehat{v}_{j-1})(\widehat{v}_{j+1} - \widehat{v}_j) = 1, \end{cases}$$

where $i_j = \min\{i: i \in \widehat{\mathcal{B}}_j\}$. This gives (10). It remains to consider the cases $j = 1$ and $j = \widehat{J}$. If $j = 1$, $\sum_{i \in \mathcal{B}_1} s_i = -w_{i_2}$, and if $j = \widehat{J}$, $\sum_{i \in \mathcal{B}_J} s_i = w_{i_J}$, from which (8) and (9) follow, respectively. \blacksquare

Proof of Theorem 2.3. Let

$$\mathcal{R}_{\lambda_{2,n}} = \{\widehat{\mathcal{J}} = \mathcal{J}_0\} \cap \{\text{sgn}(\widehat{\mu}_i^F - \widehat{\mu}_{i-1}^F) = \text{sgn}(\mu_i^0 - \mu_{i-1}^0), \forall i \in \mathcal{J}_0\} \quad (27)$$

and, for $2 \leq i \leq n$, let $d_i^0 = \mu_i^0 - \mu_{i-1}^0$, $\widehat{d}_i = \widehat{\mu}_i^F - \widehat{\mu}_{i-1}^F$ and $d_i^\epsilon = \epsilon_i - \epsilon_{i-1}$. Using the subgradient conditions (25), the event $\mathcal{R}_{\lambda_{2,n}}$ occurs if and only if, for all $i \notin \mathcal{J}_0$,

$$d_i^\epsilon = \lambda_{2,n} \left(2\text{sgn}(d_i^0) - \text{sgn}(\widehat{d}_{i-1}) - \text{sgn}(\widehat{d}_{i+1}) \right),$$

and, for all $i \in \mathcal{J}_0$,

$$\widehat{d}_i = d_i^0 + d_i^\epsilon - \lambda_{2,n} \left(2\text{sgn}(d_i^0) - \text{sgn}(\widehat{d}_{i-1}) - \text{sgn}(\widehat{d}_{i+1}) \right),$$

where in both equations for $x = 0$, $\text{sgn}(x)$ is the set $[-1, 1]$. As a result, the event $\mathcal{R}_{\lambda_{2,n}}$ occurs in probability if both

$$\max_{i \notin \mathcal{J}_0} |d_i^\epsilon| < \lambda_{2,n} |2\text{sgn}(d_i^0) - \text{sgn}(\widehat{d}_{i-1}) - \text{sgn}(\widehat{d}_{i+1})| < 4\lambda_{2,n}, \quad (28)$$

and

$$\min_{i \in \mathcal{J}_0} \left| d_i^0 + d_i^\epsilon - \lambda_{2,n} \left(2\text{sgn}(d_i^0) - \text{sgn}(\widehat{d}_{i-1}) - \text{sgn}(\widehat{d}_{i+1}) \right) \right| > 0, \quad (29)$$

hold with probability tending to 1 and $n \rightarrow \infty$.

We first consider Equation (28). Notice that, for each $2 \leq i \neq j \leq n$, $\mathbb{E}d_i^\epsilon = 0$, $\text{Var}d_i^\epsilon = 2\sigma_n^2$ and

$$\text{Cov}(d_i^\epsilon, d_j^\epsilon) = \begin{cases} -\sigma_n^2 & \text{if } |i-j| = 1 \\ 0 & \text{otherwise} . \end{cases}$$

For $2 \leq i \leq n$, let $d_i^* \sim N(0, 2\sigma_n^2)$ be independent, so that

$$\begin{cases} \mathbb{E}(d_i^\epsilon d_j^\epsilon) \leq \mathbb{E}(d_i^* d_j^*) & \text{for all } 2 \leq i \neq j \leq n \\ \mathbb{E}(d_i^\epsilon)^2 = \mathbb{E}(d_i^*)^2 & \text{for all } 2 \leq i \leq n. \end{cases}$$

Then, by Slepian's inequality (see, e.g. [Ledoux and Talagrand, 1991](#)),

$$\mathbb{P} \left\{ \max_{i \in \mathcal{J}_0^c} |d_i^\epsilon| \geq 4\lambda_{2,n} \right\} \leq \mathbb{P} \left\{ \max_{i \in \mathcal{J}_0^c} |d_i^*| \geq 4\lambda_{2,n} \right\}.$$

By the Chernoff's bound for standard Gaussian variables, followed by the union bound,

$$\mathbb{P} \left\{ \max_{i \in \mathcal{J}_0^c} |d_i^*| \geq 4\lambda_{2,n} \right\} \leq 2 \exp \left\{ -8 \frac{\lambda_{2,n}^2}{\sigma_n^2} + \log |\mathcal{J}_0^c| \right\},$$

which vanishes if condition 1. is satisfied. In order to verify [\(29\)](#), it is sufficient to show that, with probability tending to 1 as $n \rightarrow \infty$,

$$\max_{i \in \mathcal{J}_0} \left| d_i^\epsilon - \lambda_{2,n} \left(2\text{sgn}(d_i^0) - \text{sgn}(\widehat{d}_{i-1}) - \text{sgn}(\widehat{d}_{i+1}) \right) \right| \leq \alpha_n,$$

where $\alpha_n = \min_{i \in \mathcal{J}_0} |d_i^0|$. Using the triangle inequality and the fact

$$\max_i \left| 2\text{sgn}(d_i^0) - \text{sgn}(\widehat{d}_{i-1}) - \text{sgn}(\widehat{d}_{i+1}) \right| \leq 4,$$

equation [\(29\)](#) is, in turn, implied by

$$\max_{i \in \mathcal{J}_0} |d_i^\epsilon| + 4\lambda_{2,n} \leq \alpha_n.$$

Using Slepian inequality again and standard Gaussian tail bounds,

$$\mathbb{P} \left\{ \max_{i \in \mathcal{J}_0} |d_i^\epsilon| \geq \frac{\alpha_n}{2} \right\} \leq 2 \exp \left\{ -\frac{\alpha_n^2}{16\sigma_n^2} + \log |\mathcal{J}_0| \right\},$$

so that [\(29\)](#) holds in probability if condition 2. holds. ■

Proof of Theorem 2.5. It is enough to show that the event

$$\mathcal{R}_{\lambda_{1,n}} \cap \mathcal{R}_{\lambda_{2,n}}$$

occurs in probability for $n \rightarrow \infty$. Because the conditions of [Theorem 2.3](#) are assumed, $\lim_n \mathbb{P} \{ \mathcal{R}_{\lambda_{2,n}} \} = 1$ which implies that we can restrict our analysis to the set $\mathcal{R}_{\lambda_{2,n}}$, where $\widehat{\mathcal{J}} = \mathcal{J}_0$ and $\widehat{\mathcal{B}}_j = \mathcal{B}_j^0$, for $1 \leq j \leq J_0$. Next, from [Corollary 2.2](#), it is immediate to verify that the fused-lasso solution is

$$\widehat{\mu}^{FL} = \sum_{j=1}^{\widehat{\mathcal{J}}} 1_{\widehat{\mathcal{B}}_j} \widehat{\nu}_j^T$$

where $\widehat{\nu}_j^T = \text{sgn}(\widehat{\nu}_j)(\widehat{\nu}_j = \lambda_{1,n})_+$ is the soft-thresholded version of $\widehat{\nu}_j$. Therefore, in order to verify the claim, one needs to consider the simpler lasso problem applied to the vector $\widehat{\nu}$. Inspecting the sub-gradient condition for this problem, and by arguments similar to the ones used above, it follows that $\lim_n \mathbb{P}(\mathcal{R}_{\lambda_{1,n}}) = 1$ obtains provided both

$$\max_{j \in \mathcal{K}_0^c} \left| \frac{1}{b_j^0} \sum_{i \in \mathcal{B}_j^0} \epsilon_i + c_j \right| < \lambda_{1,n} \tag{30}$$

and

$$\max_{j \in \mathcal{K}_0} \left| \frac{1}{b_j^0} \sum_{i \in \mathcal{B}_j^0} \epsilon_i + c_j - \lambda_{1,n} \right| < \rho_n \tag{31}$$

hold with probability tending to 1 as $n \rightarrow \infty$, where the quantities c_j are given in Lemma 2.1. Letting $X_j = \frac{1}{b_j^0} \sum_{i \in \mathcal{B}_j^0} \epsilon_i$, notice that $X_j \sim N(0, \frac{\sigma_n^2}{b_j^0})$ and that (X_1, \dots, X_{J_0}) are independent. Then, a combination of the Chernoff's and the union bounds yields

$$\mathbb{P} \left\{ \max_{j \in \mathcal{J}S_0^c} \left| \frac{1}{b_j^0} \sum_{i \in \mathcal{B}_j^0} \epsilon_i \right| \geq \frac{\lambda_{1,n}}{2} \right\} \leq \sum_{j \in \mathcal{J}S_0^c} \exp \left\{ -\frac{\lambda_{1,n}^2 b_j^0}{8\sigma_n^2} \right\} \leq \exp \left\{ -\frac{\lambda_{1,n}^2 b_{\min}^0}{8\sigma_n^2} + \log |\mathcal{J}S_0^c| \right\},$$

and

$$\mathbb{P} \left\{ \max_{j \in \mathcal{J}S_0} \left| \frac{1}{b_j^0} \sum_{i \in \mathcal{B}_j^0} \epsilon_i \right| \geq \frac{\rho_n}{3} \right\} \leq \sum_{j \in \mathcal{J}S_0} \exp \left\{ -\frac{\rho_n^2 b_j^0}{18\sigma_n^2} \right\} \leq \exp \left\{ -\frac{\rho_n^2 b_{\min}^0}{18\sigma_n^2} + \log |\mathcal{J}S_0| \right\},$$

which give large deviations bounds for the error sums in (30) and (31). Conditions 1. and 3. guarantees that the above probabilities vanish for $n \rightarrow \infty$. Thus, with the additional conditions 2. and 4., the inequalities (30) and (31) are verified in probability. \blacksquare

Proof of Proposition 2.6. The proof is virtually identical to the proof of Theorem 2.5, the main differences stemming from the facts that the bias terms $c_j = 0$ for all $1 \leq j \leq J_0$, and

$$\frac{1}{\sqrt{b_j^0}} \sum_{i \in \mathcal{B}_j^0} \epsilon_i \sim N(0, \sigma_n^2).$$

We omit the details. \blacksquare

Proof of Theorem 2.7. Let $\hat{\mu}_F$ be the fusion estimate using the penalty $\lambda_{2,n}$. Then, because of assumption (18), and with the specific choice of $\lambda_{2,n}$ and σ_n^2 given in the statement, it can be verified that the conditions of Theorem 2.3 are met. Thus, the event

$$\mathcal{F} = \{\hat{\mathcal{J}} = \mathcal{J}^0\} \cap \{\hat{\mathcal{B}}_j = \mathcal{B}_j^0, 1 \leq j \leq \mathcal{J}^0\}$$

has probability arbitrarily close to 1, for all n large enough. On this event \mathcal{F} , we next investigate the adaptive fused-lasso $\hat{\mu}$. Because $\hat{\mu}$ is the minimizer of (14), for any $\mu \in \mathbb{R}^n$,

$$\|\hat{\mu}^{AF} - \hat{\mu}\|_2^2 + 2 \sum_i \lambda_i |\hat{\mu}_i| \leq \|\hat{\mu}^{AF} - \mu\|_2^2 + 2 \sum_i \lambda_i |\mu_i|,$$

where $\hat{\mu}^{AF}$ and λ are given in (13) and (15) respectively. Adding and subtracting μ_0 inside both terms $\|\hat{\mu}^{AF} - \hat{\mu}\|_2^2$ and $\|\hat{\mu}^{AF} - \mu\|_2^2$ yields

$$\|\hat{\mu} - \mu_0\|_2^2 \leq \|\mu - \mu_0\|_2^2 + 2 \sum_i \lambda_i (|\mu_i| - |\hat{\mu}_i|) + 2\langle \epsilon^*, \hat{\mu} - \mu \rangle, \quad (32)$$

where, on \mathcal{F} , $\epsilon^* = \hat{\mu}^{AF} - \mu^0 = \sum_{j=1}^{J_0} X_j \mathbf{1}_{\mathcal{B}_j^0}$, with $X_j \sim N(0, \frac{\sigma_n^2}{b_j^0})$ and (X_1, \dots, X_{J_0}) independent. Next, consider the sub-event $\mathcal{A} \subseteq \mathcal{F}$ given by

$$\mathcal{A} = \{|\epsilon_i^*| \leq \lambda_i, \text{ for each } i = 1, \dots, n\} = \left\{ |X_j| \leq \lambda_{1,n} / \sqrt{b_j^0}, \text{ for each } j = 1, \dots, J_0 \right\}.$$

Then,

$$\mathbb{P}(\mathcal{A}) = \mathbb{P} \left\{ \max_j |\zeta_j| \leq \lambda_{1,n} \right\},$$

where $(\zeta_1, \dots, \zeta_{J_0})$ are i.i.d $N(0, \sigma_n^2)$. Notice that because of the choice of $\lambda_{1,n}$, $\lim_n \mathbb{P}\mathcal{A} = 1$ by standard large deviation bounds for Gaussians (see also the proof of Theorem 2.3). Next, on \mathcal{A} , we have

$$2\langle \epsilon^*, \hat{\mu} - \mu \rangle \leq 2 \sum_{i \in \mathcal{S}(\mu)} \lambda_i |\hat{\mu}_i - \mu_i| + 2 \sum_{i \notin \mathcal{S}(\mu)} \lambda_i |\hat{\mu}_i|. \quad (33)$$

The decomposition

$$2 \sum_i \lambda_i (|\mu_i| - |\widehat{\mu}_i|) = 2 \sum_{i \in \mathcal{S}(\mu)} \lambda_i |\mu_i| - 2 \sum_{i \in \mathcal{S}(\mu)} \lambda_i |\widehat{\mu}_i| - 2 \sum_{i \notin \mathcal{S}(\mu)} \lambda_i |\widehat{\mu}_i|,$$

along with equation (33) and the triangle inequality yields, on \mathcal{A} ,

$$2 \sum_i \lambda_i (|\mu_i| - |\widehat{\mu}_i|) + 2 \langle \epsilon^*, \widehat{\mu} - \mu \rangle \leq 4 \sum_{i \in \mathcal{S}(\mu)} \lambda_i |\widehat{\mu}_i - \mu_i|.$$

The previous display and equation (32) lead to the inequality

$$\|\widehat{\mu} - \mu_0\|_2^2 \leq \|\mu - \mu_0\|_2^2 + 4 \sum_{i \in \mathcal{S}(\mu)} \lambda_i |\widehat{\mu}_i - \mu_i|, \quad (34)$$

valid on \mathcal{A} . Next, it is easy to see that

$$\sum_{i \in \mathcal{S}(\mu)} \lambda_i^2 = \sum_{j \in \mathcal{JS}(\mu)} b_j \lambda_i^2 \leq \lambda_{1,n}^2 \sum_{l \in \mathcal{JS}(\mu; \mu^0)} 1 = \lambda_{1,n}^2 |\mathcal{JS}(\mu; \mu^0)|,$$

and, in particular,

$$\sum_{i \in \mathcal{S}(\mu)} \lambda_i^2 = \lambda_1^2 |\mathcal{JS}(\mu^0)|$$

if and only if $\mathcal{JS}(\mu) = \mathcal{JS}(\mu^0)$.

Therefore, by the Cauchy-Swartz inequality, the second term on the right hand side of (34) can be bounded on \mathcal{A} as follows:

$$4 \sum_{i \in \mathcal{S}(\mu)} \lambda_i |\widehat{\mu}_i - \mu_i| \leq 4 \lambda_{1,n} \sqrt{|\mathcal{JS}(\mu; \mu^0)|} \|\widehat{\mu} - \mu\|_2.$$

Then, using the triangle inequality, (34) becomes

$$\|\widehat{\mu} - \mu_0\|_2^2 \leq \|\mu - \mu_0\|_2^2 + 4 \lambda_{1,n} \sqrt{|\mathcal{JS}(\mu; \mu^0)|} (\|\widehat{\mu} - \mu^0\|_2 + \|\mu_0 - \mu\|_2).$$

On \mathcal{A} , the same arguments used in the second part of the proof of Lemma 3.7 in [van de Geer \(2007\)](#) establish the inequality in the claim. Since $\lim_n \mathbb{P}(\mathcal{A}) = 1$, the first result follows. \blacksquare

Proof of Theorem 3.1. Let $N(\delta, \mathcal{F}_n, \|\cdot\|_n)$ denote the δ -covering number of the set $\mathcal{F}_n \subset \mathbb{R}^n$ with respect to the norm $\|\cdot\|_n$ and notice that, for any $C > 0$,

$$N(\delta, C\mathcal{F}_n, \|\cdot\|_n) = N\left(\frac{\delta}{C}, \mathcal{F}_n, \|\cdot\|_n\right).$$

Furthermore, observe that $\mathcal{C}_{\text{TV}}(T_n) = T_n \mathcal{C}(1)$. By a theorem of [Birman and Solomjak \(1967\)](#) (see, e.g., [Lorentz et al., 1996](#), Theorem 6.1), the δ -metric entropy of $\mathcal{C}_{\text{TV}}(T_n)$ with respect to the $L^2(\mathbb{P}_n)$ norm is

$$C \frac{T_n}{\delta},$$

for some constant C independent of n . Letting $\Psi(\delta) = \int_0^\delta \sqrt{C \frac{T_n}{s}} = \sqrt{T_n C \delta}$, the solution to

$$\sqrt{n} \delta_n^2 \gtrsim \Psi(\delta_n)$$

gives

$$\delta_n \gtrsim \frac{T_n^{1/3}}{n^{1/3}},$$

where the symbol \lesssim indicates inequality up to a universal constant. The result now follows from Theorem 3.4.1 of [van der Vaart and Wellner \(1996\)](#) (see also the discussion on pages 331-332 of the same reference). In order to establish (21), we use Lemma 4.3 in [Loubes and van de Geer \(2002\)](#) to get that the metric entropy of $\mathcal{C}_{\ell_1}(L_n)$ is

$$H(\delta, \mathcal{C}_{\ell_1}(L_n), \|\cdot\|_n) \leq C \frac{L_n^2}{n\delta^2} \left(\log n + \log \frac{L_n}{\sqrt{n}\delta} \right),$$

for some constant C independent of n . Notice that the entropy integral of $\sqrt{H(\delta, \mathcal{C}_{\ell_1}(L_n), \|\cdot\|_n)}$ diverges on any neighborhood of 0. By Theorem 9.1 in [van de Geer \(2000\)](#), the rate of consistency δ_n for $\hat{\mu}^L$ with respect to the norm $\|\cdot\|_n$ is given by the solution to

$$\sqrt{n}\delta_n^2 \gtrsim \Psi(\delta_n) \tag{35}$$

where

$$\Psi(\delta_n) \geq \int_{A\delta_n^2}^{\delta_n} \sqrt{H(x, \mathcal{C}_{\ell_1}(L_n))} dx,$$

with A a constant independent of n . Equation (35) is satisfied for a sequence δ_n satisfying

$$\sqrt{n}\delta_n^2 \gtrsim \frac{L_n \sqrt{\log n}}{\sqrt{n}} \log 1/\delta_n,$$

which gives the rate (21). ■

Proof of Theorem 3.2. Let $H(\delta_n, \mathcal{C}(S_n, J_n), \|\cdot\|_n)$ denote the metric entropy of $\mathcal{C}(S_n, J_n)$ with respect to the norm $\|\cdot\|_n$. By Lemma 6.2 and assumption (C2), for $\delta_n < 1$, the equation

$$\sqrt{n}\delta_n^2 \gtrsim \int_0^{\delta_n} \sqrt{\log H(x, \mathcal{C}(S_n, J_n), \|\cdot\|_n)} dx$$

leads to

$$\delta_n \gtrsim \sqrt{\frac{S_n}{n}} \log \sqrt{\frac{1}{\delta_n}} + o(1),$$

because $S_n = o\left(\frac{n}{\log n}\right)$ and $j_n \leq s_n$. The sequence $\delta_n = \sqrt{\frac{J_n}{n}}$ satisfies the conditions of Theorem 3.4.1 of [van der Vaart and Wellner \(1996\)](#), thus proving (24). The second claim in the theorem is proved similarly, where now the left hand side of (36) in Lemma 6.2 is bounded by $C_{1,n}$ only. ■

Lemma 6.2. *For the distance induced by the norm $\|x\|_n = \frac{1}{\sqrt{n}} \sqrt{\sum_{i=1}^n x_i^2}$, the metric entropy of $\mathcal{C}(S_n, J_n)$ satisfies*

$$H(\delta, \mathcal{C}(S_n, J_n), \|\cdot\|_n) \leq C_{1,n} + C_{2,n}, \tag{36}$$

where

$$C_{1,n} = \frac{J_n}{S_n} \log R_n + J_n \left(\log \frac{\sqrt{n}}{\delta} + \frac{1}{2} \log S_n \right) + S_n \log(S_n + J_n - 1),$$

and

$$C_{2,n} = \log S_n + S_n \log n.$$

Proof of Lemma 6.2. For fixed $\delta > 0$, we will construct an δ -grid of $\mathcal{C}(S_n, J_n)$ based on the Euclidian distance. For every choice of S_n non-zero entries of μ , we regard μ as a vector in \mathbb{R}^{S_n} which is block-wise constant with J_n blocks. Then, there exist J_n positive integer numbers d_1, \dots, d_{J_n} such that $\sum_l d_l = S_n$ and one can think of μ as the concatenation of J_n vectors μ_1, \dots, μ_{J_n} each having constant entries, where $\mu_l \in \mathbb{R}^{d_l}$, $l = 1, \dots, J_n$. Each μ_l can be any point along the main diagonal of the d_l -dimensional cube center at 0 with edge length R_n^{1/S_n} and volume $R_n^{d_l/S_n}$. The length of the main diagonal of each such cube

is $R_n^{1/S_n} \sqrt{d_l}$. Therefore, for any specific choice of S_n non-zero coordinates, the slice in the corresponding S_n -dimensional cube centered at 0 and with edge length R_n^{1/S_n} consisting of the set of vectors in \mathcal{B}_n with discontinuity profile (d_1, \dots, d_{J_n}) is the set

$$\mathcal{R}_n = \prod_{l=1}^{J_n} \ell \left(R_n^{1/S_n}, d_l \right),$$

where $\ell(R, d_l)$ denotes the closed line segment in \mathbb{R}^{S_n} between the points $\pi_{d_l}(\mathbf{1}R)$ and $\pi_{d_l}(-\mathbf{1}R)$, where $\mathbf{1}$ is the S_n -dimensional vector with coordinates all equal to 1 and π_{d_l} the function from \mathbb{R}^{S_n} onto \mathbb{R}^{S_n} given by $\pi_{d_l}(x) = y$ with $y_i = 0$ for $i \leq \sum_{j=1}^l d_j - 1$ or $i \geq \sum_{j=1}^{l+1} d_j$ and $y_i = x_i$ otherwise. Notice that the length of each $\ell \left(R_n^{1/S_n}, d_l \right)$ is precisely $R_n^{1/S_n} \sqrt{d_l}$. If $J_n = S_n$, \mathcal{R}_n is the S_n -dimensional cube centered at 0 with volume R_n , while if $J_n < S_n$ the set \mathcal{R}_n is a hyper-rectangle (not fully dimensional) which can be embedded as a hyper-rectangle in \mathbb{R}^{J_n} centered at 0 and with edge lengths equal to the lengths of $\ell \left(R_n^{1/S_n}, d_l \right)$, for $l = 1, \dots, J_n$. As a result, it is immediate to see that the volume of \mathcal{R}_n can be calculated as

$$\prod_l R_n^{1/S_n} \sqrt{d_l} = R_n^{J_n/S_n} \prod_l \sqrt{d_l}.$$

Next, partition each of the J_n perpendicular sides of \mathcal{R}_n into intervals of length $\delta \sqrt{\frac{d_l}{S_n}}$, $l = 1, \dots, J_n$. This gives a partition of \mathcal{R}_n into smaller hyper-rectangle of edge lengths $\delta \sqrt{\frac{d_l}{S_n}}$, for $l = 1, \dots, J_n$. Every point in \mathcal{R}_n is within Euclidian distance δ from the center of one of the small hyper-rectangles, so that the centers of those smaller hyper-rectangles form an δ -grid for \mathcal{R}_n . By a volume comparison, the cardinality of such a grid is

$$\frac{R_n^{J_n/S_n} \prod_l \sqrt{d_l}}{\prod_l \delta \sqrt{\frac{d_l}{S_n}}} = \left(R_n^{1/S_n} \frac{\sqrt{S_n}}{\delta} \right)^{J_n}.$$

For fixed S_n , the number of distinct block patterns with cardinality at most J_n is equal to the the number of non-negative solutions to $d_1 + d_2 + \dots + d_{J_n} = S_n$, which is

$$\binom{S_n + J_n - 1}{J_n} \leq (S_n + J_n - 1)^{J_n},$$

(see, e.g. [Stanley, 2000](#)). Thus, the logarithm of cardinality of this δ -grid is

$$\frac{J_n}{S_n} \log R_n + J_n \left(\log \frac{1}{\delta} + \frac{1}{2} \log S_n \right) + J_n \log(S_n + J_n - 1). \quad (37)$$

Next, the number of subsets of $\{1, \dots, n\}$ of size at most S_n is

$$\sum_{i=1}^{S_n} \binom{n}{i} \leq S_n n^{S_n}.$$

Thus, the logarithm of the cardinality for an δ grid over \mathcal{B}_n is bounded by (37) plus the quantity

$$\log S_n + J_n \log n.$$

The result for the $\|\cdot\|_n$ norms now follows by replacing δ with δ/\sqrt{n} in equation (37). ■

References

- Piece-wise polynomial approximations of functions in the class W_p^α , *Mathematics of the USSR Sbornik*, 73, 295–317.
- Boysen, L., Kempe, A., Liebscher, V., Munk, A. and Wittich, Ol. (2008). Consistencies and Rates of Convergence of Jump-Penalized Least Squares Estimators, preprint.
- Caselles, V., Chambolle, A. and Novaga, M. (2007). The discontinuity set of solutions of the TV denoising problem and some extensions, preprint.
- Davies, P.L. and Kovac, A. (2001). Local extremes, runs, strings and multiresolution, *Annals of Statistics*, 29(1), 1–65.
- DeVore, R. A. (1998). Nonlinear approximation, *Acta Numerica* 7, 51–150.
- Dobson, D.C. and Vogel, C.R. (1997). Convergence of an Iterative Method for Total Variation Denoising, *SIAM Journal on Numerical Analysis*, 34(5), 1779–1791.
- Donoho, D.L. and Johnstone, I.M. (1994). Minimax Risk over ℓ_p -balls for ℓ_q -error. *Probability Theory and Related Fields*, 99, 277–303.
- Friedman, J. Hastie, T., Hoglinh, H. and Tibshirani, R. (2007). Pathways Coordinate Optimization, *Annals of Applied Statistics*, 1(2), 302–332.
- Greenshtein, E. (2006). Best subset selection, persistence in high-dimensional statistical learning and optimization under l_1 constraint, *The Annals of Statistics*, 34(5), 2367–2386.
- Ledoux, M. and Talagrand, M. (1991). *Probability in Banach spaces*. Springer (1991).
- Lorentz, G.G., Golitschek, M.V. and Makovoz, Y. (1996). *Constructive Approximation: Advanced Problems*, Springer.
- Loubes, J.M. and van de Geer, (2002). S. Adaptive estimation in regression, using soft thresholding type penalties, *Statistica Neerlandica*, 56(4), 454–479.
- Mammen, E. and van de Geer, S. (1997). Locally Adaptive Regression Splines, *Annals of Statistics*, 25(1), 387–413.
- Massart, P. (2007). *Concentration Inequalities and Model Selection*, Lecture Notes in Mathematics, 1896, Springer.
- Rudin, L.I, Osher, S. and Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms, *Physica D*, 60, 259–268.
- Stanley, R. P. (2000). *Enumerative Combinatorics*, Cambridge University Press.
- van de Geer, S. (2007). *Oracle Inequalities and Regularization*, Lectures on Empirical Processes, European Mathematical Society.
- van de Geer, S. (2000). *Empirical Processes in M-Estimation*, Cambridge University Press.
- van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence of Empirical Processes*, Springer.
- Wainwright, M. J. (2006). Sharp thresholds for high-dimensional and noisy recovery of sparsity, Technical Report 708, Department of Statistics, UC Berkeley.
- Zhao, P. and Yu, B (2006). On model selection consistency of Lasso. *Journal of Machine Learning Research*, 7, 2541-2563.