



Research article

On an unsupervised method for parameter selection for the elastic net[†]

Zeljko Kereta* and Valeriya Naumova

Simula Research Laboratory, Martin Linges vei 25, Fornebu, Norway

[†] **This contribution is part of the Special Issue:** Mathematical aspects of machine learning

Guest Editors: Ernesto De Vito; Lorenzo Rosasco; Silvia Villa

Link: www.aimspress.com/mine/article/6026/special-articles

* **Correspondence:** Email: zeljko@simula.no.

Abstract: Despite recent advances in regularization theory, the issue of parameter selection still remains a challenge for most applications. In a recent work the framework of statistical learning was used to approximate the optimal Tikhonov regularization parameter from noisy data. In this work, we improve their results and extend the analysis to the elastic net regularization. Furthermore, we design a data-driven, automated algorithm for the computation of an approximate regularization parameter. Our analysis combines statistical learning theory with insights from regularization theory. We compare our approach with state-of-the-art parameter selection criteria and show that it has superior accuracy.

Keywords: parameter selection; elastic net regularization; data-driven regularization; iterative thresholding; sub-gaussian vectors; matrix concentration inequalities

1. Introduction

Inverse problems deal with the recovery of an unknown quantity of interest $\mathbf{x} \in \mathbb{R}^d$ from a corrupted observation $\mathbf{y} \in \mathbb{R}^m$. In most cases the relationship between \mathbf{x} and \mathbf{y} is linear, and can be approximately described by

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \sigma\mathbf{w}, \quad (1.1)$$

where $\mathbf{A} \in \mathbb{R}^{m \times d}$ is a known linear forward operator, \mathbf{w} is a zero-mean isotropic random vector, modeling the noise, and $\sigma > 0$ is the noise level. Inverse problems of this type are ubiquitous in image processing, compressed sensing and other scientific fields. In image processing applications they model tasks such as denoising, where \mathbf{A} is the identity; deblurring, where \mathbf{A} is a convolution operator; and inpainting, where \mathbf{A} is a masking operator.

The recovery of the original signal \mathbf{x} from the corrupted observation \mathbf{y} is an ill-posed inverse

problem. Thus, the theory of inverse problems suggests the use of suitable regularization techniques [16]. Specifically, in case of Gaussian noise, \mathbf{x} is approximated with the minimizer of a regularized functional

$$\operatorname{argmin}_{\mathbf{z} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + \lambda J(\mathbf{z}), \quad (1.2)$$

where $\|\cdot\|_2$ is the Euclidean norm modeling *data-fidelity*, J is a *penalty term* encoding an *a priori* knowledge on the nature of the true solution, and λ is a *regularization parameter* determining a trade-off between these two terms. Having the penalty term fixed, a central issue concerns the selection of λ . The optimal parameter λ_{opt} is the one that minimizes the discrepancy between the minimizer $\mathbf{z}^\lambda(\mathbf{y})$ of (1.2) and the exact solution \mathbf{x} , i.e.,

$$\lambda_{\text{opt}} = \operatorname{argmin}_{\lambda \in (0, +\infty)} \|\mathbf{z}^\lambda - \mathbf{x}\|_2. \quad (1.3)$$

If a dataset of pairs of clean signals and noisy observations were available, we could try to learn a suitable performance regression functional $\mathbf{y} \mapsto \lambda_{\text{opt}}$ that would allow to select a regularization parameter by knowing only the noisy observation. Unfortunately though, due to the curse of dimensionality accurate learning such a high-dimensional function requires a number of samples that scales exponentially with the ambient dimension [10, 26]. A common approach to mitigate these effects is to assume that the relevant data are supported on structures of substantially lower dimensionality. On a practical level, regularization parameters are instead typically selected heuristically or through cross-validation.

But, in many applications \mathbf{x} is unknown (thus we cannot use supervised methods) and there is no available information about the noise \mathbf{w} or the noise level σ . This is the case in classical regularization theory, where the clean image is unknown and hence λ_{opt} is approximated using prior knowledge or an estimate of the noise. We add that classical regularization theory is mostly concerned with data that belongs to a function space, and correspondingly most parameter selection methods focus on the recovery of the minimum least-squares norm solution.

Choosing a good approximation to λ_{opt} is a non-trivial, problem-dependent task that has motivated significant amounts of research over the last decades. However, there is still no framework that allows a fast and efficient parameter selection, particularly in a completely unsupervised setting. In this paper, we aim at (partially) closing this gap and provide a novel concept for automated parameter selection by recasting the problem in the framework of statistical learning theory. Specifically, inspired by recent and (to our knowledge) first results in this spirit [10] we propose a method for computing an accurate regularized solution $\mathbf{z}^\lambda(\mathbf{y})$ for the elastic net, with a nearly optimal regularization parameter λ , that uses a dimension reduction preprocessing step with the help of a dataset of corrupted data samples. We emphasize that the method is unsupervised and requires minimal human interference.

Existing parameter selection methods. Parameter selection rules used in regularization theory can be broadly classified as those based on the discrepancy principle [1, 25], generalized cross-validation (GCV) [19], balancing principle [23, 34], quasi-optimality [21, 31] and various estimations of the mean-squared error (MSE) (see [11, 27] and references therein). GCV is a particularly popular parameter rule for linear methods since it gives a closed form for the regularization parameter and does not require tuning of any additional parameters or the knowledge of the noise. In specialized cases GCV can

be extended to nonlinear problems [36], but the regularization parameter is no longer given in closed form nor through an implicit equation. Balancing principle is a stable method that has received a lot of attention in the inverse problems community and has also been studied in the framework of learning theory, but requires tuning of additional parameters. Quasi-optimality is one of the simplest parameter choice methods. It does not require any information about the problem but it is not as stable as the (hardened) balancing principle as it does not account for statistical outliers [3]. Discrepancy and MSE-based principles still remain the preferred methods for parameter selection for nonlinear estimators due to their simplicity and accuracy. We refer to a recent rather comprehensive comparative study on the existing approaches [3].

In order to select the regularization parameter most existing methods require the regularized solution $\mathbf{z}^\lambda(\mathbf{y})$ to be computed over a predefined grid of values of λ . The regularization parameters are then chosen according to some criteria, e.g., loss over a validation set. To find regularization parameters by an exhaustive search is a computationally expensive task, especially in the high-dimensional data scenario, with often no guarantees on the quality of approximation. Moreover, most criteria presuppose that some *a priori* information is available, such as an accurate estimate of the noise level (in e.g., discrepancy principle) or bounds on the noise error (in e.g., balancing principle) and require additional, method-specific parameters to be preselected.

Elastic net regularization. Elastic net regularization was proposed by Zou and Hastie [37], as

$$\mathbf{z}^\lambda(\mathbf{y}) = \underset{\mathbf{z} \in \mathbb{R}^d}{\operatorname{argmin}} \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + \lambda \left(\|\mathbf{z}\|_1 + \alpha \|\mathbf{z}\|_2^2 \right), \quad (1.4)$$

where $\alpha \geq 0$ is a hyperparameter controlling the trade-off between ℓ_1 and ℓ_2 penalty terms. Our motivation for considering the elastic net is that it produces sparse models comparable to the Lasso (and is thus well suited for problems with data on lower dimensional structures), while often achieving superior accuracy in real-world and simulated data. Moreover, the elastic net overcomes the main limitations of ℓ_1 minimization. Namely, it encourages the grouping effect, which is relevant for many real-life applications such as microarray classification and face detection (see [9] and references therein).

To solve (1.4) the authors in [37] rewrite the elastic net functional as Lasso regularization with augmented datum, use LARS [14] to reconstruct the entire solution path, and apply cross-validation to select the optimal regularization parameter. Later work [9] studies theoretical properties of (1.4) in the context of learning theory, analyzes the underlying functional and uses *iterative soft-thresholding* to compute the solution. For the parameter choice the authors provide an adaptive version of the balancing principle [23, 34]. The rule aims to balance approximation and sample errors, which have contrasting behavior with respect to the tuning parameter, but requires (potentially) many evaluations of $\mathbf{z}^\lambda(\mathbf{y})$. We will rework some of the arguments from [9] for the computation of \mathbf{z}^λ , while keeping our focus on an efficient approach for parameter learning. In [22] the authors propose an *active set* algorithm for solving (1.4). Addressing the problem in the framework of classical regularization theory, the authors consider the discrepancy principle [6, 25] for determining the parameter. This requires estimations of the solution \mathbf{z}^λ for many parameter values, and a pre-tuning of other, method-specific parameters. Moreover, it is assumed that the noise level is known, which is often not the case in practice.

The authors in [24] use a hybrid alternating method for tuning parameters λ and α for the model fitting problem $\mathbf{y}_i = \mathbf{a}_i^\top \mathbf{x}$, $i = 1, \dots, n$, where $\mathbf{y}_i \in \mathbb{R}$, $\mathbf{a}_i \in \mathbb{R}^p$ and $\mathbf{x} \in \mathbb{R}^p$. The first step is to update the

solution \mathbf{z}^λ , using coordinate descent, and then to update λ and α in one iteration. The main advantage is the efficiency, as one does not need to calculate \mathbf{z}^λ for multiple parameters at once, but rather on a much coarser parameter grid. The method is in spirit similar to LARS, but has better scalability. It requires that a non-convex problem is solved, and hence has inherent limitations. Moreover, it cannot be used in the setting of general inverse problems (1.1), where the design matrix A is fixed and each response \mathbf{y} is generated by a new clean signal \mathbf{x} . In summary, we are not aware of any parameter selection rule for the elastic net that allows to select λ without *a priori* assumptions and without extensive manual adjustments.

This work. We leverage the work [10] where the authors propose a data-driven method for determining the optimal parameter for Tikhonov regularization, under the assumption that a training set of independent observations $\mathbf{y}_1, \dots, \mathbf{y}_N$ is made available, each of them associated with an (unknown) signal $\mathbf{x}_1, \dots, \mathbf{x}_N$ through $\mathbf{y}_i = A\mathbf{x}_i + \sigma\mathbf{w}_i$. The starting point of the method is to find an empirical proxy $\widehat{\mathbf{x}}$ of the real solution \mathbf{x} by assuming that $\mathbf{x}_1, \dots, \mathbf{x}_N$ are distributed over a lower-dimensional linear subspace and then select the regularization parameter as

$$\widehat{\lambda}_{\text{opt}} = \underset{\lambda \in (0, +\infty)}{\operatorname{argmin}} \left\| \mathbf{z}_{\text{Tik}}^\lambda(\mathbf{y}) - \widehat{\mathbf{x}} \right\|_2, \quad (1.5)$$

where $\mathbf{z}_{\text{Tik}}^\lambda(\mathbf{y})$ is the minimizer of the Tikhonov functional $\min_{\mathbf{z} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{z}\|_2^2$. The analysis and techniques regarding $\widehat{\mathbf{x}}$ are independent from the chosen optimization scheme, whereas the selection of $\widehat{\lambda}_{\text{opt}}$ is defined by the regularization scheme. However, it is worth mentioning that if A^\dagger is not injective, $\widehat{\mathbf{x}}$ is not a good proxy of \mathbf{x} . For Tikhonov regularization this is not an issue as, without loss of generality, we can always assume that A is injective. Specifically, one can replace \mathbf{x} in (1.3) with $\mathbf{x}^\dagger = A^\dagger A\mathbf{x}$ and recall that $\mathbf{z}_{\text{Tik}}^\lambda$ belongs to $\ker^\perp(A)$ for all λ . Therefore, for wider applicability of the suggested framework, it is important to address the selection of $\widehat{\lambda}_{\text{opt}}$ for a wider class of regularizers and inverse problems.

In this paper, we extend the framework of [10] by providing the analysis for the elastic net regularization, and improving the theoretical results regarding the empirical estimators through stronger concentration bounds, see Lemma 2.2. Moreover, we develop an efficient, fully automated algorithm for computing a nearly optimal regularization parameter, that we call OptEN. The algorithm uses backtracking line search to minimize a given loss function: the discrepancy between the elastic net solution $\mathbf{z}^\lambda(\mathbf{y})$ and the empirical estimator $\widehat{\mathbf{x}}$, which is well suited if A is injective. In case of a non-injective A we develop an adjusted loss function that again ensures near optimal performance. Lastly, whereas in [10] the dimensionality of the subspace to which the real solution \mathbf{x} belongs to is assumed to be known, and which is needed to construct the empirical estimator $\widehat{\mathbf{x}}$, in this work we develop two strategies for its estimation. The first is based on the spectrum of the empirical covariance matrix of noisy observations, and is used when \mathbf{x} lies in a true lower dimensional linear subspace. The second strategy is used in cases where \mathbf{x} only approximately lies in a lower-dimensional subspace, as is the case *e.g.* in wavelet denoising. The algorithm is extensively tested on synthetic and real-world examples. The last point is the main practical contribution of our paper. We add that our goal is not to introduce a new regularization paradigm but rather to design a fast and unsupervised method for determining a near optimal regularization parameter for existing regularization methods. In summary, we analyze our problem in two settings:

- (i) *simplified case* $A = \text{Id}$: (corresponding to image denoising): We restate the lower level problem and show that in case of a bounded \mathbf{w} it admits a unique minimizer, which motivates the algorithm. Furthermore, we provide a bound on $|\lambda_{\text{opt}} - \widehat{\lambda}_{\text{opt}}|$ for independent Bernoulli random noise and discuss the number of samples needed for optimal learning, see Proposition 4.4. Though the latter model might be oversimplified, it captures the essence of the problem and our experiments confirm the results in more general settings.
- (ii) *general case*: for a general matrix A we provide an unsupervised, efficient, and accurate algorithm for the computation of an approximate optimal parameter. We study the performance of our algorithm, comparing it to state-of-the-art parameter choice methods on synthetic and image denoising problems. The obtained results show that our approach achieves superior accuracy.

1.1. Outline

In Section 2, we describe the assumptions on the linear inverse problem and the distribution of data that we study, and we define and prove bounds regarding empirical estimators. In Section 3 we discuss minimizers of the elastic net (1.4), and define loss functions that will be used for parameter selection. Section 4 provides the main theoretical results of the paper regarding loss functionals and their minimizers. In Section 5 we present an efficient and accurate algorithm for the computation of an approximate optimal parameter. We study the performance of our method through several numerical experiments on synthetic and imaging data in Section 6. Therein, we compare our method with state-of-the-art parameter selection criteria in terms of accuracy of the solution recovery, closeness to the optimal parameter, sparse recovery and computational time. For imaging tasks our focus is on wavelet-based denoising where we work on synthetic images and real-world brain MRIs. We conclude with a brief discussion about future directions in Section 7. The Appendix contains proofs of auxiliary results.

1.2. Notation

The Euclidean and the ℓ_1 -norms of $\mathbf{u} = (u_1, \dots, u_d)^\top$ are denoted by $\|\mathbf{u}\|_2$ and $\|\mathbf{u}\|_1$, respectively. The modulus function $|\cdot|$, the *sign* function $\text{sgn}(\cdot)$, and the *positive part* function $(\cdot)_+$ are defined component-wise for $i = 1, \dots, d$, by $|\mathbf{u}|_i = |u_i|$, $\text{sgn}(\mathbf{u})_i = \text{sgn}(u_i)$, and $((\mathbf{u})_+)_i = (u_i)_+$, where for any $u \in \mathbb{R}$

$$\text{sgn}(u) = \begin{cases} 1, & \text{if } u > 0, \\ 0, & \text{if } u = 0, \\ -1, & \text{if } u < 0, \end{cases} \quad \text{and} \quad (u)_+ = \max\{0, u\}.$$

The canonical basis of \mathbb{R}^d is denoted by $\{\mathbf{e}_i\}_{i=1, \dots, d}$. We denote the transpose of a matrix \mathbf{M} by \mathbf{M}^\top , the Moore-Penrose pseudo inverse by \mathbf{M}^\dagger , and the spectral norm by $\|\mathbf{M}\|_2$. Furthermore, $\text{range}(\mathbf{M})$ and $\text{ker}(\mathbf{M})$ are the range and the null space of \mathbf{M} , respectively. For a square-matrix \mathbf{M} , we use $\text{trace}(\mathbf{M})$ to denote its trace. The identity matrix is denoted by Id and we use $\mathbb{1}_{\mathcal{D}}$ for the indicator function of a set $\mathcal{D} \subset \mathbb{R}^d$. For any $\mathbf{v} \in \mathbb{R}^d$, $\mathbf{v} \otimes \mathbf{v}$ is the rank one operator acting on $\mathbf{w} \in \mathbb{R}^d$ as $(\mathbf{v}^\top \mathbf{w}) \mathbf{v}$.

A random vector $\boldsymbol{\xi}$ is called *sub-Gaussian* if

$$\|\boldsymbol{\xi}\|_{\psi_2} := \sup_{\|\mathbf{v}\|_2=1} \sup_{q \geq 1} q^{-\frac{1}{2}} \mathbb{E} \left[\left| \mathbf{v}^\top \boldsymbol{\xi} \right|^q \right]^{\frac{1}{q}} < +\infty.$$

The value $\|\xi\|_{\psi_2}$ is the sub-Gaussian norm of ξ , with which the space of sub-Gaussian vectors becomes a normed vector space [33]. The (non-centered) covariance of a random vector ξ is denoted as

$$\Sigma(\xi) := \text{Cov}(\xi) = \mathbb{E}[\xi \otimes \xi].$$

We write $a \lesssim b$ if there exists an absolute constant $C > 0$ such that $a \leq Cb$.

2. The empirical estimator for linear inverse problems

We consider the following stochastic linear inverse problem: given a deterministic matrix $A \in \mathbb{R}^{m \times d}$, we are interested in recovering a vector $\mathbf{x} \in \mathbb{R}^d$ from a noisy observation $\mathbf{y} \in \mathbb{R}^m$ obeying

$$\mathbf{y} = A\mathbf{x} + \sigma\mathbf{w}, \quad (2.1)$$

where

(A1) the unknown datum $\mathbf{x} \in \mathbb{R}^d$ is a sub-Gaussian vector, such that $\|\mathbf{x}\|_{\psi_2} = 1$;

(A2) there exists a sub-family $1 \leq i_1 < \dots < i_h \leq d$ of h indices, with $h \ll d$, such that

$$\mathcal{V} := \text{range}(\Sigma(\mathbf{x})) = \text{span}\{\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_h}\}$$

and $\ker(A) \cap \mathcal{V} = \{\mathbf{0}\}$;

(A3) the noise $\mathbf{w} \in \mathbb{R}^m$ is an independent sub-Gaussian vector, such that $\|\sigma\mathbf{w}\|_{\psi_2} \leq 1$, $\Sigma(\mathbf{w}) = \text{Id}$ and $\sigma > 0$ is the noise level.

Conditions (A1) and (A3) are standard assumptions on the distributions of the exact datum \mathbf{x} and the noise $\sigma\mathbf{w}$, ensuring that the tails have fast decay. Note also that from a theoretical standpoint, normalization conditions on \mathbf{x} and \mathbf{w} can always be satisfied by rescaling.

As discussed in the introduction, we consider the recovery of \mathbf{x} by means of regularized minimization, $\mathbf{z}^\lambda(\mathbf{y}) = \text{argmin}_{\mathbf{z} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + \lambda J(\mathbf{z})$, where $J(\mathbf{z})$ is the elastic net functional. More specifically, we are interested in the selection of the regularization parameter λ . To do so, we follow [10] and leverage assumption (A2) to construct an empirical estimator of \mathbf{x} , which will be used to select λ by minimizing the suitable notion of discrepancy between $\mathbf{z}^\lambda(\mathbf{y})$ and the constructed empirical estimator.

We begin by discussing the construction of the empirical estimator, from the given training data of noisy samples, and developing concentration bounds. First, it follows from the definition that \mathcal{V} is the smallest subspace such that $\mathbf{x} \in \mathcal{V}$, almost surely. Thus, by (A2), the exact datum \mathbf{x} almost surely has at most h non-zero entries, and since $\ker(A) \cap \mathcal{V} = \{\mathbf{0}\}$, it is a unique vector with that property. Define now $\mathcal{W} = \text{range}(\Sigma(A\mathbf{x}))$. The following simple result was shown for an injective A in [10]; here we extend it to the general case.

Lemma 2.1. *Under Assumption (A2) we have $\dim \mathcal{W} = h$ and $\mathcal{W} = A\mathcal{V}$.*

Proof. A direct computation gives

$$\Sigma(A\mathbf{x}) = \mathbb{E}[A\mathbf{x} \otimes A\mathbf{x}] = A\Sigma(\mathbf{x})A^\top = A\P\Sigma(\mathbf{x})(A\P)^\top,$$

where \mathbf{P} denotes the orthogonal projection onto \mathcal{V} . Assumption (A2) says that \mathbf{A} is injective on \mathcal{V} , and thus $\Sigma(\mathbf{A}\mathbf{x})$ and $\Sigma(\mathbf{x})$ both have rank h . Furthermore, $(\mathbf{A}\mathbf{P})^\top$ maps \mathbb{R}^m onto \mathcal{V} , so that

$$\text{range}(\Sigma(\mathbf{A}\mathbf{x})) = (\mathbf{A}\mathbf{P}\Sigma(\mathbf{x}))\mathcal{V} = \mathbf{A}\mathbf{P}\mathcal{V} = \mathcal{W},$$

where $\Sigma(\mathbf{x})\mathcal{V} = \mathcal{V}$, by the definition of \mathcal{V} and since $\Sigma(\mathbf{x})$ is symmetric. \square

2.1. Empirical estimators and concentration bounds

Lemma 2.1 suggests that \mathcal{V} could be directly recovered if \mathbf{A} were invertible and \mathcal{W} were known. In most practical situations though, neither of those assumptions is satisfied: we only have access to noisy observations and \mathbf{A} could not only be non-invertible, but also non-injective. We will address this issue by recasting the problem to a statistical learning framework, similar to [10]. Namely, suppose we are given observation samples $\mathbf{y}_1, \dots, \mathbf{y}_N$ such that $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i + \sigma\mathbf{w}_i$ for $i = 1, \dots, N$, where \mathbf{x}_i , \mathbf{w}_i and σ are unknown, and let

$$\widehat{\Sigma}(\mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \mathbf{y}_i \otimes \mathbf{y}_i$$

be the empirical covariance of the observations. Standard statistical theory suggests that $\widehat{\Sigma}(\mathbf{y})$ is a good approximation to $\Sigma(\mathbf{y})$ provided N is large enough. As a consequence, we will show that a vector space spanned by the first h eigenvectors of $\widehat{\Sigma}(\mathbf{y})$, denoted by $\widehat{\mathcal{W}}$, is a good estimator of \mathcal{W} .

To justify the above claims, observe first that since $\Sigma(\mathbf{w}) = \text{Id}$ holds by (A3), we have

$$\Sigma(\mathbf{y}) = \Sigma(\mathbf{A}\mathbf{x}) + \sigma^2 \text{Id}. \quad (2.2)$$

Therefore, $\Sigma(\mathbf{y})$ and $\Sigma(\mathbf{A}\mathbf{x})$ have the same eigenvectors and the spectrum of $\Sigma(\mathbf{y})$ is just a shift of the spectrum of $\Sigma(\mathbf{A}\mathbf{x})$ by σ^2 . Let $\lambda_1 \geq \dots \geq \lambda_h$ be the non-zero eigenvalues of $\Sigma(\mathbf{A}\mathbf{x})$, counting for multiplicity, and $\mu_1 \geq \dots \geq \mu_m$ and $\widehat{\mu}_1 \geq \dots \geq \widehat{\mu}_m$ be the eigenvalues of $\Sigma(\mathbf{y})$ and $\widehat{\Sigma}(\mathbf{y})$, respectively. From (2.2) it follows

$$\begin{cases} \mu_i = \lambda_i + \sigma^2, & \text{for } i = 1, \dots, h, \\ \mu_i = \sigma^2, & \text{for } i = h + 1, \dots, m. \end{cases} \quad (2.3)$$

Let Π be the (orthogonal) projection onto \mathcal{W} , which has rank h due to Lemma 2.1, and let $\widehat{\Pi}$ be the (orthogonal) projection onto $\widehat{\mathcal{W}}$. We now show the fundamental tool of our study: that $\widehat{\Pi}$ is an accurate and an unbiased approximation of Π for a sufficiently small noise level σ . We distinguish between bounded and unbounded \mathbf{y} and improve upon results in [10].

Lemma 2.2. *Assume that $\sigma^2 < \lambda_h$ and $\|\Sigma(\mathbf{y})\|_2 \geq 1$. Given $u > 0$, with probability greater than $1 - 2 \exp(-u)$*

$$\|\widehat{\Pi} - \Pi\|_2 \lesssim \frac{\lambda_1}{\lambda_h} \left(\sqrt{\frac{h + \sigma^2 m + u}{N}} + \frac{h + \sigma^2 m + u}{N} \right), \quad (2.4)$$

provided $N \gtrsim (h + \sigma^2 m + u)$. Furthermore, if \mathbf{y} is bounded, then with probability greater than $1 - \exp(-u)$

$$\|\widehat{\Pi} - \Pi\|_2 \lesssim \frac{\lambda_1}{\lambda_h} \left(\sqrt{\frac{\log(h + m) + u}{N}} + \frac{\log(h + m) + u}{N} \right), \quad (2.5)$$

provided $N \gtrsim (\log(2m) + u)$.

Proof. We will first show (2.4). Using Theorem 9.2.4 and Exercise 9.2.5 in [33], we have

$$\left\| \Sigma(\mathbf{y}) - \widehat{\Sigma}(\mathbf{y}) \right\|_2 \lesssim \|\Sigma(\mathbf{y})\|_2 \left(\sqrt{\frac{r+u}{N}} + \frac{r+u}{N} \right), \quad (2.6)$$

with probability greater than $1 - 2 \exp(-u)$, where $r = \text{trace}(\Sigma(\mathbf{y})) / \|\Sigma(\mathbf{y})\|_2$ is the *stable rank* of $\Sigma(\mathbf{y})^{1/2}$. Using $\text{trace}(\Sigma(\mathbf{y})) \leq \lambda_1 h + m\sigma^2$ and $\|\Sigma(\mathbf{y})\|_2 = \lambda_1 + \sigma^2 \leq 2\lambda_1$, we get

$$\left\| \Sigma(\mathbf{y}) - \widehat{\Sigma}(\mathbf{y}) \right\|_2 \lesssim \lambda_1 \left(\sqrt{\frac{h + \sigma^2 m + u}{N}} + \frac{h + \sigma^2 m + u}{N} \right), \quad (2.7)$$

where we used $\|\Sigma(\mathbf{y})\|_2 \geq 1$ to bound r . Let $\mu^* = \mu_h > \sigma^2$. By (2.3) it follows that Π is the projection onto the linear span of those eigenvectors of $\Sigma(\mathbf{y})$ whose corresponding eigenvalue is greater than or equal to μ^* . Using $\mu_h - \mu_{h+1} = \lambda_h$, by (2.7) we have

$$\epsilon := \left\| \Sigma(\mathbf{y}) - \widehat{\Sigma}(\mathbf{y}) \right\|_2 < \frac{\mu_h - \mu_{h+1}}{2} = \frac{\lambda_h}{2}, \quad (2.8)$$

provided $N \gtrsim (h + \sigma^2 m + u)$. Let now Π_{μ^*} be the projection onto the linear span of those eigenvectors of $\widehat{\Sigma}(\mathbf{y})$ whose corresponding eigenvalue is greater than or equal to μ^* . As a consequence of Theorem 7.3.1 in [5], there exists an eigenvalue $\widehat{\mu}^*$ of $\widehat{\Sigma}(\mathbf{y})$ such that

$$|\mu^* - \widehat{\mu}^*| \leq \epsilon, \text{ and } \dim \Pi_{\mu^*} = \dim \Pi \quad (2.9)$$

$$\widehat{\mu}_j \leq \mu_{h+1} + \epsilon = \sigma^2 + \epsilon, \quad \forall \widehat{\mu}_j < \widehat{\mu}^* \quad (2.10)$$

$$\left\| \Pi_{\mu^*} - \Pi \right\|_2 \leq \frac{1}{\lambda_h - \epsilon} \left\| (\text{Id} - \Pi_{\mu^*})(\widehat{\Sigma}(\mathbf{y}) - \Sigma(\mathbf{y}))\Pi \right\|_2. \quad (2.11)$$

By (2.9) it follows that $\widehat{\mu}^* = \widehat{\mu}_h$ so that $\Pi_{\mu^*} = \widehat{\Pi}$ and hence

$$\left\| \widehat{\Pi} - \Pi \right\|_2 \leq \frac{1}{\lambda_h - \epsilon} \left\| \widehat{\Sigma}(\mathbf{y})\Pi - \Sigma(\mathbf{y})\Pi \right\|_2 \leq \frac{2}{\lambda_h} \left\| \widehat{\Sigma}(\mathbf{y})\Pi - \Sigma(\mathbf{y})\Pi \right\|_2. \quad (2.12)$$

Since $\left\| \widehat{\Sigma}(\mathbf{y})\Pi - \Sigma(\mathbf{y})\Pi \right\|_2 \leq \left\| \widehat{\Sigma}(\mathbf{y}) - \Sigma(\mathbf{y}) \right\|_2$, the claim follows by (2.7).

Assume now that $\|\mathbf{y}\|_2 \leq \sqrt{L}$ holds almost surely and consider a family of independent $m \times h$ matrices

$$\mathbf{S}_i = (\mathbf{y}_i \otimes \mathbf{y}_i)\Pi - \Sigma(\mathbf{y})\Pi, \quad i = 1, \dots, N.$$

Since $\frac{1}{N} \sum_{i=1}^N \mathbf{S}_i = \widehat{\Sigma}(\mathbf{y})\Pi - \Sigma(\mathbf{y})\Pi$ we can apply the matrix Bernstein inequality for rectangular matrices (Theorem 6.1.1. in [32]). Thus, for $u > 0$ we have

$$\mathbb{P} \left(\left\| \widehat{\Sigma}(\mathbf{y})\Pi - \Sigma(\mathbf{y})\Pi \right\|_2 \geq u \right) \leq (m+h) \exp \left(\frac{-Ns^2}{M + 2Lu/3} \right),$$

where $M > 0$ is a matrix variance constant independent of m , h , and d , such that

$$\max \left\{ \mathbb{E} \left\| \mathbf{S}_i^\top \mathbf{S}_i \right\|_2, \mathbb{E} \left\| \mathbf{S}_i \mathbf{S}_i^\top \right\|_2 \right\} \leq M.$$

A direction computation gives $M \leq L \|\Sigma(\mathbf{y})\|_2$. It follows that

$$\left\| \widehat{\Sigma}(\mathbf{y})\Pi - \Sigma(\mathbf{y})\Pi \right\|_2 \lesssim \lambda_1 \left(\sqrt{\frac{\log(h+m)+u}{N}} + \frac{\log(h+m)+u}{N} \right),$$

holds with probability greater than $1 - \exp(-u)$ for every $u > 0$. Moreover, by analogous argumentation (2.8) holds provided $N \gtrsim (\log(2m) + u)$, see 7 for details. Thus, (2.5) follows by applying (2.12). \square

The previous result comes with a certain caveat. Namely, the proof implicitly assumes that either h or the spectral gap are known (which informs the choice of the approximate projector $\widehat{\Pi}$). In practice however, the desired eigenspace can only be detected if there is a spectral gap and if it corresponds to the eigenspace we want to recover, i.e., if $\lambda_h > \delta$, where

$$\delta = \max_{i=1,\dots,h-1} (\lambda_i - \lambda_{i+1}) = \max_{i=1,\dots,h-1} (\mu_i - \mu_{i+1}). \quad (2.13)$$

Proposition 2.3. *Assume $\lambda_h > \delta$, where δ is given by (2.13). Then the empirical covariance matrix has a spectral gap at the h -th eigenvalue, with probability greater than $1 - 2 \exp(-u)$, provided $\delta < \lambda_h$ and $N \gtrsim \frac{\lambda_1^2}{(\lambda_h - \delta)^2} (h + u)$.*

Proof. Assume $\|\Sigma(\mathbf{y}) - \widehat{\Sigma}(\mathbf{y})\|_2 < \epsilon$, for $\epsilon > 0$. Since $\max_{j=1,\dots,m} |\widehat{\mu}_j - \mu_j| \leq \|\Sigma(\mathbf{y}) - \widehat{\Sigma}(\mathbf{y})\|_2$ we get

$$|\widehat{\mu}_j - \widehat{\mu}_{j+1}| \leq 2\epsilon + |\mu_j - \mu_{j+1}|,$$

by adding and subtracting μ_j and μ_{j+1} inside the first term. Thus, if $j > h$ then $|\widehat{\mu}_j - \widehat{\mu}_{j+1}| \leq 2\epsilon$ by (2.3), and if $j < h$ then $|\widehat{\mu}_j - \widehat{\mu}_{j+1}| \leq 2\epsilon + \delta$, by the definition of δ in (2.13). For $j = h$ on the other hand we have $|\widehat{\mu}_h - \widehat{\mu}_{h+1}| > |\mu_h - \mu_{h+1}| - 2\epsilon$. In conclusion,

$$\operatorname{argmax}_{i=1,\dots,m-1} (\widehat{\mu}_i - \widehat{\mu}_{i+1}) = h$$

holds provided provided $\epsilon < \frac{\lambda_h - \delta}{4}$. Using (2.6) the claim follows. \square

It is clear that if $\delta > \lambda_h$ then the spectral gap of the empirical covariance matrix is achieved at $\operatorname{argmax}_{i=1,\dots,h-1} (\lambda_i - \lambda_{i+1})$, which is smaller than h . In practice though, the situation is not as pessimistic as this observation would suggest and we can rely on a wealth of *ad hoc* remedies. We devote more attention to this question in Section 6.1, and suggest alternative heuristics for estimating the intrinsic dimension h .

We are ready to define our empirical estimator of \mathbf{x} . Let $Q = AA^\dagger$ be the orthogonal projection onto $\operatorname{range}(A) = \ker^\perp(A^\top)$, and $P = A^\dagger A$ the orthogonal projection onto $\operatorname{range}(A^\top) = \ker^\perp(A)$. The *empirical estimator* of \mathbf{x} is defined as

$$\widehat{\mathbf{x}} = A^\dagger \widehat{\Pi} \mathbf{y}. \quad (2.14)$$

In the following, we use $\widehat{\mathbf{x}}$ to learn a nearly optimal regularization parameter for parameter selection for the elastic net, but it can in principle be used for a broader range of problems This is due to the fact that $\widehat{\mathbf{x}}$ is in principle independent of the choice of the optimization scheme. That is, it only leverages the structural assumption (A2).

Remark 2.4. One might want to consider $\widehat{\mathbf{x}}$ as an approximate solution by itself, and completely avoid regularization and thus the issue of parameter choice. Even though $\widehat{\mathbf{x}}$ can be a good estimator of \mathbf{x} in certain situations, it will however often perform poorly. Namely, when A is not injective, when the training set size N is small, or when the noise level is small, the empirical estimator $\widehat{\mathbf{x}}$ will perform worse than the regularized solution $\mathbf{z}^\lambda(\mathbf{y})$. This can be seen by the empirical evidence, cf. Table 2, where $\widehat{\mathbf{x}}$ is the worst performing method. In addition, $\widehat{\mathbf{x}}$ does not preserve the structure of the original signal, e.g., $\widehat{\mathbf{x}}$ will in general not be sparse for a sparse \mathbf{x} . Thus, regularization is indeed needed to ensure the desired structure and optimize the reconstruction performance. Lastly, we remind that we are interested in using an estimator $\widehat{\mathbf{x}}$ for which $|\lambda_{opt} - \widehat{\lambda}_{opt}|$ is small with high probability (i.e., the one that can be used to derive an accurate parameter selection) and we are not interested in directly controlling $\|\mathbf{x} - \widehat{\mathbf{x}}\|_2$, which is the goal in manifold learning [4].

Observe also that $\widehat{\boldsymbol{\eta}} = \mathbf{y} - \widehat{\Pi}\mathbf{y}$ the empirical estimator $\widehat{\mathbf{x}}$ satisfies the (empirical) inverse problem

$$A\widehat{\mathbf{x}} + Q\widehat{\boldsymbol{\eta}} = Q\mathbf{y}. \quad (2.15)$$

Thus, a direct consequence of (2.15) is that minimizers of the empirical and of the original problem coincide.

Lemma 2.5. Let $\widehat{\mathbf{z}}^\lambda(\mathbf{y}) = \operatorname{argmin}_{\mathbf{z} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{z} - Q\mathbf{y}\|_2^2 + \lambda J(\mathbf{z})$, where $Q = AA^\dagger$, for $A \in \mathbb{R}^{m \times d}$, and J is the (elastic net) penalty term. Then $\widehat{\mathbf{z}}^\lambda(\mathbf{y}) = \mathbf{z}^\lambda(\mathbf{y})$, where $\mathbf{z}^\lambda(\mathbf{y})$ is given by (1.4).

Proof. We compute $\|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 = \|\mathbf{A}\mathbf{z} - Q\mathbf{y} + (Q - \operatorname{Id})\mathbf{y}\|_2^2$. Since Q is an orthogonal projection onto $\operatorname{range}(A)$ it follows $(Q - \operatorname{Id})\mathbf{y} \in \operatorname{range}^\perp(A)$. Using Pythagoras' theorem we thus have $\|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 = \|\mathbf{A}\mathbf{z} - Q\mathbf{y}\|_2^2 + \|(Q - \operatorname{Id})\mathbf{y}\|_2^2$. Since the second term does not depend on \mathbf{z} we get

$$\operatorname{argmin}_{\mathbf{z} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + \lambda J(\mathbf{z}) = \operatorname{argmin}_{\mathbf{z} \in \mathbb{R}^d} \|\mathbf{A}\mathbf{z} - Q\mathbf{y}\|_2^2 + \lambda J(\mathbf{z}).$$

□

3. Loss functionals for parameter selection for the elastic net

In this section we discuss the elastic net minimization, and the loss functionals that will be used for the selection of the regularization parameter.

3.1. Elastic net minimization

From now on we focus on the parameter choice for the elastic net, where $J(\mathbf{z}) = \|\mathbf{z}\|_1 + \alpha \|\mathbf{z}\|_2^2$, so that

$$\mathbf{z}^\lambda(\mathbf{y}) = \operatorname{argmin}_{\mathbf{z} \in \mathbb{R}^m} \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + \lambda \left(\|\mathbf{z}\|_1 + \alpha \|\mathbf{z}\|_2^2 \right). \quad (3.1)$$

The term $\|\mathbf{z}\|_1$, enforces the sparsity of the solution, whereas $\|\mathbf{z}\|_2^2$ enforces smoothness and ensures that in case of highly correlated features we can correctly retrieve all the relevant ones. We first recall some basic facts about existence, uniqueness and sensitivity of elastic net solutions with respect to regularization parameters [22].

Lemma 3.1. *The elastic net functional is strictly convex and coercive. Moreover, for each $\lambda > 0$ the minimizer of (1.2) exists, is unique and the mapping $\lambda \mapsto \mathbf{z}^\lambda$ is continuous with respect to $\lambda > 0$.*

In the remainder of this paper we will recast (3.1) as

$$\mathbf{z}^t(\mathbf{y}) = \underset{\mathbf{z} \in \mathbb{R}^m}{\operatorname{argmin}} t \|\mathbf{A}\mathbf{z} - \mathbf{y}\|_2^2 + (1-t) (\|\mathbf{z}\|_1 + \alpha \|\mathbf{z}\|_2^2), \quad (3.2)$$

where $t \in [0, 1]$ and $\alpha > 0$ is a fixed parameter. For $t \in (0, 1)$ the solutions of (3.2) correspond to solutions of (3.1) for $\lambda = \frac{1-t}{t}$. On the other hand, for $t = 0$ we get $\mathbf{z}^0(\mathbf{y}) = \mathbf{0}$, and for $t = 1$ we define $\mathbf{z}^1(\mathbf{y}) := \mathbf{x}^\alpha$, where

$$\mathbf{x}^\alpha = \underset{\mathbf{z} \in \mathcal{N}}{\operatorname{argmin}} (\|\mathbf{z}\|_1 + \alpha \|\mathbf{z}\|_2^2), \text{ for } \mathcal{N} = \{\mathbf{z} \in \mathbb{R}^m \mid \mathbf{A}^\top \mathbf{A}\mathbf{z} = \mathbf{A}^\top \mathbf{y}\}. \quad (3.3)$$

This definition is driven by the following observations. First, the set $\mathcal{N} = \mathbf{A}^\dagger \mathbf{y} \oplus \ker(\mathbf{A})$ is non-empty (since \mathbf{A} has finite rank). Furthermore, it was shown in [9] and [22] that in case of the elastic net minimization

$$\lim_{t \rightarrow 1} \mathbf{z}^t = \mathbf{x}^\alpha. \quad (3.4)$$

In other words, \mathbf{x}^α plays the role of the Moore-Penrose solution in linear regularization schemes [16]. By Lemma 3.1 the minimizer of (3.2) always exist and is unique, the map $t \mapsto \mathbf{z}^t$ is continuous for $t \in (0, 1)$. Equation (3.4) implies that $t \mapsto \mathbf{z}^t$ is continuous at $t = 1$, and later in (3.7) we show that the continuity also holds at $t = 0$.

Solution via soft-thresholding. The elastic net does not admit a closed form solution in case of a general forward matrix \mathbf{A} . In Zou and Hastie [37] the elastic net problem is recast as a Lasso problem with augmented data, which can then be solved by many different algorithms (e.g., the LARS method [14]). Alternative algorithms compute the elastic net minimizer directly, and are generally either of the *active set* [22] or the *iterative soft-thresholding-type* [9]. Here we adhere to iterative soft-thresholding, and rework the arguments in [9] to show that the solution to (3.2) can be obtained through fixed point iterations for all $t \in [0, 1]$. To begin, define the soft-thresholding function by

$$\mathcal{S}_\tau(\mathbf{u}) = \operatorname{sgn}(\mathbf{u}) \left(\left| \mathbf{u} \right| - \frac{\tau}{2} \right)_+, \quad (3.5)$$

and the corresponding soft-thresholding operator $\mathcal{S}_\tau(\mathbf{u})$, acting component-wise on vectors $\mathbf{u} \in \mathbb{R}^m$. The next lemma is a direct reworking of the arguments in [9, Corollary 1] and states that (3.2) is a fixed point of a contractive map.

Lemma 3.2. *The solution to (3.2), for $\mathbf{A} \in \mathbb{R}^{m \times d}$, $\mathbf{y} \in \mathbb{R}^m$ and $t \in (0, 1)$, satisfies $\mathbf{z} = \mathcal{T}_t(\mathbf{z})$, where the map $\mathcal{T}_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a contraction and is defined by*

$$\mathbf{z}^t = \mathcal{T}_t(\mathbf{z}) = \frac{1}{\tau t + (1-t)\alpha} \mathcal{S}_{1-t} (t(\theta \mathbf{I}d - \mathbf{A}^\top \mathbf{A}) \mathbf{z} + t \mathbf{A}^\top \mathbf{y}), \quad (3.6)$$

with the Lipschitz constant

$$\frac{t(\sigma_M^2 - \sigma_m^2)}{t(\sigma_M^2 + \sigma_m^2) + 2\alpha(1-t)} < 1,$$

where $\theta = \frac{\sigma_m^2 + \sigma_M^2}{2}$, and σ_m and σ_M are the smallest and the largest singular values of the matrix \mathbf{A} , respectively.

For $t = 0$, the solution is $\mathbf{z}^0 = \mathbf{0}$, which is consistent with (3.6). Furthermore, by (3.5) and (3.6), we get

$$\mathbf{z}^t = 0 \quad \text{if} \quad 0 \leq t \leq \frac{1}{1 + 2\|\mathbf{A}^\top \mathbf{y}\|_2}. \quad (3.7)$$

Our definition of the solution at $t = 1$ in (3.3) also satisfies $\mathbf{z} = \mathcal{T}_t(\mathbf{z})$ since \mathcal{S}_{1-t} is identity and thus $\mathbf{z}^{t=1}(\mathbf{y}) = \mathbf{A}^\dagger \mathbf{y}$, though \mathcal{T}_1 is not a contraction. In summary, the solutions are consistent with Lemma 3.1, as expected.

Closed form solution. In the case of orthogonal design [37], i.e., $\mathbf{A}^\top \mathbf{A} = \text{Id}$, the solution of (3.1) is given by

$$\mathbf{z}^\lambda = \frac{(|\mathbf{A}^\top \mathbf{y}| - \lambda/2)_+}{1 + \alpha\lambda} \text{sgn}(\mathbf{A}^\top \mathbf{y}), \quad (3.8)$$

where the product is understood component-wise. Plugging $\lambda = \frac{1-t}{t}$ into (3.8) we have

$$\mathbf{z}^t(\mathbf{y}) = \frac{(t(1 + 2|\mathbf{A}^\top \mathbf{y}|) - 1)_+}{2(t(1 - \alpha) + \alpha)} \text{sgn}(\mathbf{A}^\top \mathbf{y}). \quad (3.9)$$

3.2. Quadratic loss functionals

To select the regularization parameters we go back to first principles and consider quadratic loss functionals.

Definition 3.3. Functions $R, \widehat{R}: [0, 1] \rightarrow \mathbb{R}$, defined by

$$R(t) = \|\mathbf{z}^t(\mathbf{y}) - \mathbf{x}\|_2^2, \quad \widehat{R}(t) = \|\mathbf{z}^t(\mathbf{y}) - \widehat{\mathbf{x}}\|_2^2, \quad (3.10)$$

are called the true and the empirical quadratic loss, respectively. Furthermore, define

$$t_{\text{opt}} = \underset{t \in [0,1]}{\text{argmin}} R(t), \quad \widehat{t}_{\text{opt}} = \underset{t \in [0,1]}{\text{argmin}} \widehat{R}(t). \quad (3.11)$$

In view of Lemma 3.1 and the discussion in Section 3.1, the benefits of recasting λ to $[0, 1]$ are clear: R and \widehat{R} are both continuous, defined on a bounded interval, and, hence, achieve a minimum. Thus, our aim is to minimize \widehat{R} while ensuring $|t_{\text{opt}} - \widehat{t}_{\text{opt}}|$ is small.

Let us discuss some difficulties associated with elastic net minimization which need to be addressed. On one hand, a closed form solution of (3.2) is available only when $\mathbf{A}^\top \mathbf{A} = \text{Id}$ and is otherwise only approximated. Furthermore, as we will see below, loss functionals R and \widehat{R} are globally neither differentiable nor convex, but rather only piecewise. These two issues suggest that the analysis of their minimizers in full detail is challenging. Therefore, in the following we split the analysis into a simplified case for $\mathbf{A}^\top \mathbf{A} = \text{Id}$ where we can provide guarantees, and the general case where we provide an efficient algorithm. Furthermore, we need to amend the empirical loss function \widehat{R} in the case when \mathbf{A} is non-injective. This is due to the fact that in case of non-linear methods $R(t)$ cannot be reliably estimated outside the kernel of \mathbf{A} , see [15] and Figure 1. We follow the idea of SURE-based methods [18], which provide an unbiased estimate of $R(t)$ by projecting the regularized solution onto $\ker^\perp(\mathbf{A})$. Namely, we define *projected and modified loss functions* $\widehat{R}_p, \widehat{R}_m: [0, 1] \rightarrow \mathbb{R}$ by

$$\widehat{R}_p(t) = \|\mathbf{P}\mathbf{z}^t(\mathbf{y}) - \widehat{\mathbf{x}}\|_2^2, \quad \text{and} \quad \widehat{R}_m(t) = \|\mathbf{A}\mathbf{z}^t(\mathbf{y}) - \widehat{\Pi}\mathbf{y}\|_2^2, \quad (3.12)$$

where $P = A^\dagger A$ is the orthogonal projection onto $\ker^\perp(A)$. Define also

$$\hat{t}_p = \operatorname{argmin}_{t \in [0,1]} \widehat{R}_p(t), \text{ and } \hat{t}_m = \operatorname{argmin}_{t \in [0,1]} \widehat{R}_m(t). \quad (3.13)$$

Note that to define $\widehat{R}_m(t)$ we used the fact that $\widehat{\mathbf{x}} = A^\dagger \widehat{\Pi} \mathbf{y}$, and thus compared to \widehat{R}_p , we avoid the computation of the Moore-Penrose inverse A^\dagger , which might be either costly to compute or indeed numerically unstable if A is poorly conditioned. As we will show in Section 6, and can see in the right-most panel in Figure 1, using \widehat{R}_p and \widehat{R}_m instead of \widehat{R} when A is non-injective dramatically improves the performance. Note that projecting onto $\ker^\perp(A)$ makes the loss functional smoother (dampening the gradients).




Figure 1. Empirical and true losses for $m = 500$, $d = 60$, $h = 5$, $N = 50$, $\alpha = 10^{-3}$, $\sigma = 0.08$, and zero mean isotropic Gaussians \mathbf{x} and \mathbf{w} . In the left panel A is injective and \widehat{t}_{opt} is a good proxy for t_{opt} . In the middle panel $\text{rank}(A) = 40$ and we see that \widehat{t}_{opt} does not approximate t_{opt} well. On the other hand, the right panel shows that in case of a non-injective matrix, \widehat{R}_m and \widehat{R}_p improve the performance.

In the next section we analyze the minimizers t_{opt} and \widehat{t}_{opt} , and develop bounds on $|t_{\text{opt}} - \widehat{t}_{\text{opt}}|$ in simplified settings. These results motivate us to design an algorithm, which we call OptEN, that minimizes the corresponding empirical loss function ($\widehat{R}(t)$, $\widehat{R}_p(t)$, or $\widehat{R}_m(t)$). This is discussed in Section 5.

4. Parameter error

Since the elastic net solution is in general not available in closed form, a rigorous study of the parameter error is challenging in full generality. Therefore, we restrict our attention to simplified cases, though we emphasize that our approach in practice performs well on significantly broader model assumptions, which we will show in Section 6. In case of orthogonal design $A^\top A = \text{Id}$ we can, without loss of generality, assume $A = \text{Id}$ (otherwise redefine \mathbf{y} as $A^\top \mathbf{y}$). Let now $\mathbf{y} = \mathbf{x} + \sigma \mathbf{w}$ and assume $|y_1| \geq \dots \geq |y_m|$. Plugging (3.9) into (3.10) we get

$$R(t) = \sum_{i=1}^m \left(\frac{(t(1 + 2|y_i|) - 1)_+}{2(t(1 - \alpha) + \alpha)} \operatorname{sgn}(y_i) - x_i \right)^2.$$

Define $\mathbf{b}_i = 1 + 2|y_i|$, for $i = 1, \dots, m$. Loss function $R(t)$ is continuous on $[0, 1]$, and differentiable on intervals

$$\mathcal{I}_0 = [0, \mathbf{b}_1^{-1}), \quad \mathcal{I}_m = (\mathbf{b}_m^{-1}, 1], \text{ and } \mathcal{I}_k = (\mathbf{b}_k^{-1}, \mathbf{b}_{k+1}^{-1}), \text{ for } k = 1, \dots, m - 1.$$

Considering one interval at a time a direct computation yields the minimizer of each $R|_{\mathcal{I}_k}$. Let $k = 1, \dots, m-1$ and define $a_i = \text{sgn}(y_i)(1 + 2\alpha |y_i|)$, $c_i = \text{sgn}(y_i) + 2x_i(\alpha - 1) + 2y_i$, and $d_i = \text{sgn}(y_i) + 2\alpha x_i$, and consider the expression $t_k = \frac{\sum_{i=1}^k a_i d_i}{\sum_{i=1}^k a_i c_i}$. The minimizer $t^{*,k}$ of $R|_{\mathcal{I}_k}$ is then given by

$$t^{*,k} = \begin{cases} \mathbf{b}_k, & \text{for } t_k < \mathbf{b}_k \\ t_k, & \text{for } \mathbf{b}_k \leq t_k \leq \mathbf{b}_{k+1} \\ \mathbf{b}_{k+1}, & \text{for } t_k > \mathbf{b}_{k+1} \end{cases} \quad (4.1)$$

Therefore, t_k is the minimizer if it is inside the interval \mathcal{I}_k , and otherwise the minimizer is the corresponding edge of the interval. Further details regarding the derivation are in Section 7 in the Appendix. An analogous expression holds for $k = m$, whereas $R(t)$ is constant on \mathcal{I}_0 , as argued in (3.7). Therefore, the minimizer of $R(t)$ is $t_{\text{opt}} = \text{argmin}_{k=0, \dots, m} R(t^{*,k})$. The empirical loss function $\widehat{R}(t)$ is also continuous on $[0, 1]$ and is piecewise differentiable on the same set of intervals since they depend only on \mathbf{y} . Consequently, minimizers $\widehat{t}^{*,k}$ of $\widehat{R}(t)$ are also of the form (4.1), where we only ought to replace x_i by \widehat{x}_i .

Notice that unless further assumptions are made, minimizers t_{opt} and \widehat{t}_{opt} are not given explicitly: we still need to evaluate $R(t)$ and $\widehat{R}(t)$ at $m+1$ locations, and it is not clear that there are no local minima or that the minimizer is unique. We will now show that in case of bounded noise there is indeed only one minimum and that it concentrates near $t = 1.0$ for moderate noise levels. This analysis will also give a theoretical intuition that will drive our algorithm. Furthermore, we will show that in a simplified case of Bernoulli noise we get explicit bounds on the parameter error.

4.1. Bounded noise

Consider now the case of bounded noise such that there is a gap between the noise and the signal. We show that there exists a unique minimizer and there are no local minima. For simplicity of computation, we let $\alpha = 1$, though the results hold for all $\alpha > 0$. Let $\mathbf{y} = \mathbf{x} + \sigma \mathbf{w}$ and assume $\mathbf{x} = (x_1, \dots, x_h, 0, \dots, 0)^\top$ where $|x_i| > 2\sigma |w_j|$ for all $i = 1, \dots, h$ and $j = 1, \dots, m$. Without loss of generality, we assume that \mathbf{y} is ordered so that

$$|x_i + \sigma w_i| > |x_j + \sigma w_j| \text{ for } 1 \leq i < j \leq h \text{ and } |w_i| \geq |w_j| \text{ for } 1 \leq i < j \leq m.$$

The loss functional $R(t) = \|\mathbf{z}^t - \mathbf{x}\|_2^2$ is thus piecewise differentiable on intervals \mathcal{I}_k , where $\mathbf{b}_i = 1 + 2|x_i + \sigma w_i|$ for $i = 1, \dots, h$, and $\mathbf{b}_i = 1 + 2\sigma |w_i|$ for $i = h+1, \dots, m$. Also, we have $\mathbf{b}_i \geq \mathbf{b}_j$ for $i \leq j$. We will show that $R(t)$ can be monotonically increasing only for t larger than some $\vartheta_j > \mathbf{b}_{h+1}^{-1}$. Thus, for all t smaller than ϑ_j , $R(t)$ is a monotonically decreasing function. Let thus $t \in \mathcal{I}_j$ for $j < h$. The function $R(t)$ is continuously differentiable in \mathcal{I}_j , so it is sufficient to show that $R'(t)$ is positive. A direct computation, as in Section 4 and in particular in 7 in the Appendix, gives

$$R'(t) \geq 0 \text{ if } t \geq \frac{\sum_{i=1}^j \mathbf{b}_i (1 + 2 \text{sgn}(y_i) x_i)}{\sum_{i=1}^j \mathbf{b}_i^2} =: \vartheta_j. \quad (4.2)$$

It suffices to show $\vartheta_j \geq \mathbf{b}_{j+1}^{-1}$. Since $\text{sgn}(y_i) = \text{sgn}(x_i)$ for $i \leq j < h$, we have

$$\mathbf{b}_{j+1} (1 + 2|x_i|) - \mathbf{b}_i > 4|x_i| |x_{j+1} + \sigma w_{j+1}| > 0. \quad (4.3)$$

Therefore, $\mathbf{b}_i (1 + 2 \operatorname{sgn}(y_i)x_i) \geq \mathbf{b}_{j+1}^{-1} \mathbf{b}_i^2$, and the claim follows. Extending the same analysis to $t \in \mathcal{I}_h$ we ought to show $\mathbf{b}_{h+1} (1 + 2|x_i|) - \mathbf{b}_i > 0$ for $1 \leq i \leq h$. Computing gives

$$\mathbf{b}_{h+1} (1 + 2|x_i|) - \mathbf{b}_i > 2\sigma |\mathbf{w}_{h+1}| (1 + 2|x_i|) > 0.$$

Hence, $t_{\text{opt}} > \mathbf{b}_{h+1}^{-1}$, as desired. We will now show that $R(t)$ admits only one minimizer. Assume there exists t^* such that $t^* \in \mathcal{I}_{j^*}$ for some $j^* > h$ and $R'(t^*) = 0$. This means

$$t^* = \frac{\sum_{i=1}^h \mathbf{b}_i (1 + 2 \operatorname{sgn}(y_i)x_i) + \sum_{i=h+1}^{j^*} \mathbf{b}_i}{\sum_{i=1}^{j^*} \mathbf{b}_i^2} = \vartheta_{j^*}, \text{ and } \mathbf{b}_{j^*}^{-1} < \vartheta_{j^*} < \mathbf{b}_{j^*+1}^{-1}.$$

We proceed by induction showing that $R(t)$ is increasing for all $j > j^*$. For $t \in \mathcal{I}_j$ with $j > h$, it follows

$$R'(t) \leq 0 \text{ if } t \leq \frac{\sum_{i=1}^h \mathbf{b}_i (1 + 2 \operatorname{sgn}(y_i)x_i) + \sum_{i=h+1}^j \mathbf{b}_i}{\sum_{i=1}^j \mathbf{b}_i^2} =: \vartheta_j.$$

Let us show $\vartheta_j < \mathbf{b}_j^{-1}$ for $j = j^* + 1$. We have

$$\vartheta_j = \frac{\sum_{i=1}^h \mathbf{b}_i (1 + 2 \operatorname{sgn}(y_i)x_i) + \sum_{i=h+1}^j \mathbf{b}_i}{\sum_{i=1}^j \mathbf{b}_i^2} = \frac{\vartheta_{j^*} \sum_{i=1}^{j^*} \mathbf{b}_i^2 + \mathbf{b}_j}{\sum_{i=1}^j \mathbf{b}_i^2} \leq \mathbf{b}_j^{-1},$$

where we used the fact $\vartheta_{j^*} \leq \mathbf{b}_j^{-1} = \mathbf{b}_{j^*+1}^{-1}$, and $\mathbf{b}_j > \mathbf{b}_{j^*}$. The rest of the proof then follows by mathematical induction. Analogous computation yields the same type of a result for the empirical loss function \widehat{R} .

Lemma 4.1. *Let $\mathbf{y} = \mathbf{x} + \sigma \mathbf{w}$, assume $\mathbf{x} = (x_1, \dots, x_h, 0, \dots, 0)^\top$ where $|x_i| > 2\sigma |\mathbf{w}_j|$ for all $i = 1, \dots, h$ and $j = 1, \dots, m$, and let $\alpha > 0$. Loss function $R(t)$ is then either monotonically decreasing on the entire interval $[0, 1]$, or it is decreasing until some interval \mathcal{I}_{j^*} , for $j^* > h + 1$ where it achieves a (unique) minimum, and it is monotonically increasing on all the subsequent intervals. The same holds for \widehat{R} .*

4.2. Bernoulli noise

Lemma 4.1 states that $R(t)$ and $\widehat{R}(t)$ achieve a unique minimum in $[0, 1]$, and that they are monotonically decreasing before, and monotonically increasing after this minimum. Furthermore, the minimizer is bigger than $(1 + 2\sigma |\mathbf{w}_{h+1}|)^{-1}$, which means that for moderate noise levels, it will be close to 1. The issue is that minimizers ϑ_{j^*} and $\widehat{\vartheta}_{\widehat{j}^*}$ do not need to lie in the same interval, i.e., $j^* \neq \widehat{j}^*$, and thus they cannot be directly compared. Instead, we consider a simplified model that encodes the main features of the problem. In particular, let

$$\mathbf{y} = \mathbf{x} + \sigma \mathbf{w}, \text{ where } \mathbb{P}(\mathbf{w}_i = \pm 1) = \frac{1}{2}, \quad (4.4)$$

and assume $\mathbf{x} = (x_1, \dots, x_h, 0, \dots, 0)^\top$, and $|x_i| \geq 2\sigma$, for $i = 1, \dots, h$. As before, without loss of generality we can assume $|y_1| \geq |y_2| \geq \dots \geq |y_m|$. It then follows $\mathbf{b}_i = 2|x_i \pm \sigma| + 1$, for $1 \leq i \leq h$, and $\mathbf{b}_i = 2\sigma + 1$ otherwise. Moreover, $\mathbf{b}_j > \mathbf{b}_{h+1}$, for all $j = 1, \dots, h$. In the following, we will for the sake of simplicity consider the case $\alpha = 1$. The details regarding the general case, $\alpha \neq 1$, are in the Appendix. First, as in Section 4.1 we know that $t_{\text{opt}} \geq \mathbf{b}_{h+1}^{-1}$. We can now explicitly compute the minimizer of $R(t)$.

Lemma 4.2. Let \mathbf{x} satisfy (4.4) and assume $\mathbf{x} = (x_1, \dots, x_h, 0, \dots, 0)^\top$, and $|x_i| \geq 2\sigma$, for $i = 1, \dots, h$. True loss functional $R(t)$ is minimized for $t_{opt} = \min\{t^*, 1\} \in [\mathbf{b}_{h+1}^{-1}, 1]$, where

$$t^* = \frac{\sum_{i=1}^h \mathbf{b}_i (1 + 2 \operatorname{sgn}(y_i) x_i) + \mathbf{b}_{h+1}(m-h)}{\sum_{i=1}^h \mathbf{b}_i^2 + (m-h)\mathbf{b}_{h+1}^2}. \quad (4.5)$$

Proof. Considering (4.2) for $t \in (\mathbf{b}_{h+1}^{-1}, 1)$ we get

$$2R'(t) = \left(\sum_{i=1}^h \mathbf{b}_i^2 + (m-h)\mathbf{b}_{h+1}^2 \right) t - \sum_{i=1}^h \mathbf{b}_i (1 + 2 \operatorname{sgn}(y_i) x_i) - \mathbf{b}_{h+1}(m-h). \quad (4.6)$$

The root of (4.6) is exactly (4.5). Arguing as in (4.3) we have $t^* > \frac{1}{\mathbf{b}_{h+1}}$. Restricting to $[0, 1]$ the claim follows. \square

Remark 4.3. The minimizer given by Lemma 4.2 will be in $[0, 1]$ provided $\sum_{i=1}^h |y_i| \leq (m-h)\sigma$ and $h \leq m/2$.

For the empirical loss function it is in general not true that $\widehat{x}_i = 0$ for $i > h$, nor is $\frac{y_i - \widehat{x}_i}{\sigma}$ a Bernoulli random variable. However, \mathbf{y} and \mathbf{b}_i 's remain the same, and an entirely analogous computation gives

$$\widehat{t}^* = t^* + \frac{\sum_{i=1}^m \mathbf{b}_i \operatorname{sgn}(y_i) (x_i - \widehat{x}_i)}{\sum_{\mathbf{b}_i > 1/t^*} \mathbf{b}_i^2}. \quad (4.7)$$

We can now bound the approximation error for the optimal regularization parameter.

Theorem 4.4. Assume that the conditions in Lemma 4.2 hold, and that $t_{opt} < 1$ and $\sigma \frac{h}{m} < 1$. Given $u > 0$, with probability of at least $1 - 2 \exp(-u)$ we have

$$|t_{opt} - \widehat{t}_{opt}| \leq \frac{\lambda_1}{\lambda_h} \left(\sqrt{\frac{h + \sigma^2 m + u}{N}} + \frac{h + \sigma^2 m + u}{N} \right) + \sigma \sqrt{\frac{h}{m}}, \quad (4.8)$$

provided $N \gtrsim (h + \sigma^2 m + u)$. Assume now \mathbf{y} is bounded. With probability greater than $1 - 3 \exp(-u)$ we then have

$$|t_{opt} - \widehat{t}_{opt}| \leq \frac{\lambda_1}{\lambda_h} \left(\sqrt{\frac{\log(h+m) + u}{N}} + \frac{\log(h+m) + u}{N} \right) + \sigma \sqrt{\frac{h}{m}}, \quad (4.9)$$

provided $N \gtrsim (\log 2m + u)$.

Proof. By assumptions we have $t_{opt} = t^*$ by Lemma 4.2. We can now rewrite (4.5) and (4.7) as

$$t^* = \frac{\|\mathbf{b}\|_1 + 2 \langle (\operatorname{sgn} \mathbf{y} \cdot \mathbf{b}), \mathbf{x} \rangle}{\|\mathbf{b}\|_2^2}, \quad \widehat{t}^* = \frac{\|\mathbf{b}\|_1 + 2 \langle (\operatorname{sgn} \mathbf{y} \cdot \mathbf{b}), \widehat{\mathbf{x}} \rangle}{\|\mathbf{b}\|_2^2}.$$

Therefore, using (4.7) we have

$$t^* - \widehat{t}^* = \frac{\langle \mathbf{v}, \mathbf{x} - \widehat{\mathbf{x}} \rangle}{\|\mathbf{b}\|_2^2},$$

with \mathbf{v} defined by $v_j = 2 \operatorname{sgn}(y_j) \mathbf{b}_j$ for $j = 1, \dots, h$, and $v_j = \operatorname{sgn}(w_j)$ for $j > h$. Thus, $\|\mathbf{v}\|_2 \leq \|\mathbf{b}\|_2$ and we have

$$|t^* - \widehat{t}^*| \leq 2 \frac{\|\mathbf{x} - \widehat{\mathbf{x}}\|_2}{\|\mathbf{b}\|_2} \leq 2 \left(\frac{\|\mathbf{y}\|_2}{\|\mathbf{b}\|_2} \|\Pi - \widehat{\Pi}\|_2 + \sigma \frac{\|\Pi \mathbf{w}\|_2}{\|\mathbf{b}\|_2} \right).$$

Using $\|\mathbf{b}\|_2^2 = m + 4\|\mathbf{y}\|_1 + 4\|\mathbf{y}\|_2^2$ and $\|\Pi \mathbf{w}\|_2 = h$, and provided $N \gtrsim (h + u + \sigma^2 m)$, by Lemma 2.2 we have

$$|t^* - \widehat{t}^*| \leq C \frac{\lambda_1}{\lambda_h} \left(\sqrt{\frac{h + \sigma^2 m + u}{N}} + \frac{h + \sigma^2 m + u}{N} \right) + \sigma \sqrt{\frac{h}{m}},$$

with probability greater than $1 - 2 \exp(-u)$, and $C > 0$ is a constant. Since $|t^* - \widehat{t}^*| \geq |t^* - \widehat{t}_{\text{opt}}|$ the claim follows. The proof for a bounded \mathbf{y} is entirely analogous. \square

Remark 4.5. *Theorem 4.4 is valid not only for $\alpha = 1$ but for all α . Namely, in Appendix 7 we show that for any $\alpha > 0$ the following holds*

$$|t_{\text{opt}} - \widehat{t}_{\text{opt}}| \lesssim \frac{\lambda_1}{\lambda_h} \left(\sqrt{\frac{h + u + \sigma^2 m}{N}} + \frac{h + u + \sigma^2 m}{N} \right) + \sigma \sqrt{\frac{h}{m}}.$$

5. OptEN algorithm

Driven by insights in Section 4, we are ready to present an efficient heuristic algorithm for learning the *Optimal* regularization parameter for the *Elastic Net* (OptEN). The algorithm is based on the minimization of a given loss function (\widehat{R} , \widehat{R}_m or \widehat{R}_p). In Section 4 we showed that in a simplified, yet instructive, setting that the optimal parameter tends to be in the vicinity of $t = 1$, depending on the noise level and the signal-to-noise gap. This is supported by experimental evidence in more general situations such as for non-injective \mathbf{A} , as we will see in Section 6. Moreover, the loss function is monotonically decreasing as we get away from $t = 1$. These observations drive our algorithm which assumes that the minimizer lies in a valley not too far from $t = 1$, see Figure 1. Therefore, we will perform a *line search* on the graph of a given loss function, starting from $t = 1$.

Line search methods follow iterations $t_{k+1} = t_k + s_k \mathbf{p}_k$, where \mathbf{p}_k is the *search direction* and s_k the *step size*:

- **Search direction.** We select \mathbf{p}_k by estimating $\widehat{R}'(t_k)$ with central differences, $\widehat{R}'(t) \sim \Delta_\epsilon R(t) := \frac{R(t+\epsilon) - R(t-\epsilon)}{2\epsilon}$ where $\epsilon > 0$ is small enough. For $t = 1$ we instead use $\widetilde{\Delta}_\epsilon R(1) := \frac{R(1) - R(1-\epsilon)}{\epsilon}$. Then set $\mathbf{p}_k = -\Delta_\epsilon R(t)$.
- **Step size.** We estimate s_k with the backtracking line search (consult [2] for an overview of line search methods).

Our approach is presented in Algorithm 1, while an extensive numerical study is provided in the next section.

Algorithm 1 OptEN algorithm for approximating the optimal elastic net regularization parameter using backtracking line search

Input: $\mathbf{y}_1, \dots, \mathbf{y}_N, \mathbf{y} \in \mathbb{R}^m$, $\mathbf{A} \in \mathbb{R}^{m \times d}$;

Compute $\widehat{\mathbf{x}}$ according to (2.14).

Set a loss function $\rightarrow \widehat{R}$ or \widehat{R}_p or \widehat{R}_m . In the rest of the algorithm we will refer to it as \tilde{R} ;

Set $\epsilon > 0$, $\text{tol} > 0$, $\text{tol2} > 0$, $0 < \alpha < 1$, and $c_1, \beta, \gamma > 0$;

Set $k \leftarrow 0$, $t_0 = 1$

Compute $r_1 = \tilde{R}(1)$, $\tilde{r}_1 = \tilde{R}(1 - \epsilon)$, $\mathbf{p}_0 = (r_1 - \tilde{r}_1)/\epsilon$, and $r_2 = \varphi(\gamma_0)$;

repeat

$\tilde{t} = t_k + \alpha \mathbf{p}_k$;

$\varphi_0 = r_1$, $\varphi'_0 = -\mathbf{p}_k^2$, $\varphi_1 = R(\tilde{t})$;

if $\varphi_1 - \varphi_0 < c_1 \varphi'_0 \alpha$ **then**

$s_k = \alpha$;

else

$s_k = -\frac{1}{2} \frac{\varphi'_0 \alpha^2}{\varphi_1 - \varphi_0 - \varphi'_0 \alpha}$;

end if

if $|s_k| < \text{tol2}$ or $|s_{k-1}/s_k| > \gamma$ **then**

$s_k = s_{k-1} \cdot \beta$;

end if

Set $t_{k+1} = t_k + s_k \mathbf{p}_k$;

Compute $r_1 = \tilde{R}(t_{k+1})$, $\mathbf{p}_{k+1} = (\tilde{R}(t_{k+1} + \epsilon) - \tilde{R}(t_{k+1} - \epsilon))/(2\epsilon)$;

$k \leftarrow k + 1$;

until $|\mathbf{p}_k| < \text{tol}$ or $k < \text{max_iter}$;

Output: Approximate regularization parameter $\hat{t} := t_k$.

6. Experimental results

We now study the performance of our approach and show its adaptivity to different scenarios by conducting experiments on synthetic and imaging data. In the first set of experiments we perform a thorough comparison of our method with state-of-the-art parameter selection rules by exploring their behavior with respect to noise level and other notions. The second set of experiments deals with image denoising where we use wavelet-based thresholding with the elastic net. We consider two data-sets: natural images and a real-world brain MRI data. Note that we do not aim to compare our method with state-of-the-art denoising methods, but rather only with state-of-the-art methods regarding the selection of the regularization parameter for the elastic net. We start with a discussion of methods that can be used for the automatic detection of the sparsity level h and show that when a sufficient amount of training points is given, we can reliably estimate h .

6.1. Estimating the sparsity level

In real applications the sparsity level of a vector is either not available or is only an approximate notion, i.e., the desired vector is sparse only when we threshold its entries. Such regimes require h to be estimated, which in our case means looking at the spectrum of the corresponding covariance matrix.

This question belongs to the class of low-rank matrix recovery problems since what we are trying to recover is the geometry (i.e., projection onto the range) of the noiseless, lower rank matrix $\Sigma(\mathbf{A}\mathbf{x})$, using only the covariance matrix of noisy observations $\widehat{\Sigma}(\mathbf{y})$, which is of full rank. Thus, estimating h boils down to thresholding singular values of the empirical covariance matrix according to some spectral criterion that exploits the underlying structure.

For a positive definite matrix with singular values $\lambda_1 \geq \dots \geq \lambda_m > 0$ commonly used spectral criteria are

- (i) The spectral gap $\operatorname{argmax}_k |\lambda_k - \lambda_{k+1}|$;
- (ii) The relative gap $\operatorname{argmax}_k \left(1 - \frac{\lambda_k}{\lambda_{k+1}}\right)$;
- (iii) The cumulative spectral energy $\sum_{i=1}^k \lambda_i / \sum_{i=1}^m \lambda_i$;
- (iv) The relative cumulative spectral energy $\sum_{i=1}^k \lambda_i / \sum_{i=1}^{k+1} \lambda_i$.

For the latter two criteria one sets a threshold, say 0.95, and selects \tilde{h} as the first k for which the corresponding spectral energy reaches that threshold.

We study the behavior of these four criteria on three different types of forward matrices \mathbf{A} : random Gaussian, random circulant Rademacher, and random Toeplitz Gaussian matrices. These matrices were chosen because they have different spectral behavior and commonly appear in inverse problems. In each case \mathbf{A} is a 100×100 real matrix, normalized so that $\|\mathbf{A}\|_2 = 1$, and we take $N = 100$ samples \mathbf{x}_i , sampled according to $\mathbf{x}_i = \boldsymbol{\xi}_i + \mathbf{v}_i$, where $\xi_i \sim \mathcal{N}(0, 1)$ and $v_i = 4 \operatorname{sgn}(\xi_i)$ for $1 \leq i \leq 20$, and $\xi_i = v_i = 0$ otherwise, $\mathbf{w}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{100})$ and $\sigma = 0.3$. We compute $\widehat{\Sigma}(\mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \mathbf{y}_i \otimes \mathbf{y}_i$ for $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i + \sigma\mathbf{w}_i$.

In Figure 2, we show the application of the aforementioned spectral criteria to $\widehat{\Sigma}(\mathbf{y})$. It is clear from the results that all four methods would fail if used without taking further information into account. For example, the spectral gap criterion (in the upper right panel) would dictate the selection of $h = 3$, but a more careful look at the plot suggests that the behavior of the spectral gap changes dramatically around $h = 20$, which corresponds to the true h . Such *ad hoc* solutions are sensible and often improve the performance but can be hard to quantify, especially on real data.

The last spectral criteria, $1 - \lambda_k/\lambda_{k+1}$, is perhaps the most promising, but is also subject to demands on N , as shown in the bottom row of Figure 2. Namely, if N is not large enough then $1 - \lambda_k/\lambda_{k+1}$ has a heavy tail on the spectrum of $\widehat{\Sigma}(\mathbf{y})$ and would thus suggest a large h , as in the bottom left corner of the Figure. Instead, in Section 6.2 we look for the relative gap within the first $m/2$ singular vectors. We note that in the case of the first three spectral criteria the situation does not change as the number of samples increases. On the other hand, for the last criterion it does: heavy tails flatten back to zero for all three choices of random matrices, see the plot in the bottom right corner of Figure 2.

In Section 6.3 we will apply our algorithm to wavelet denoising where there is no natural choice of h since wavelet coefficients of images are not truly sparse. We will instead consider two scenarios; when we are given an oracle h (i.e., the h giving the highest PSNR), and when h has to be estimated from data.

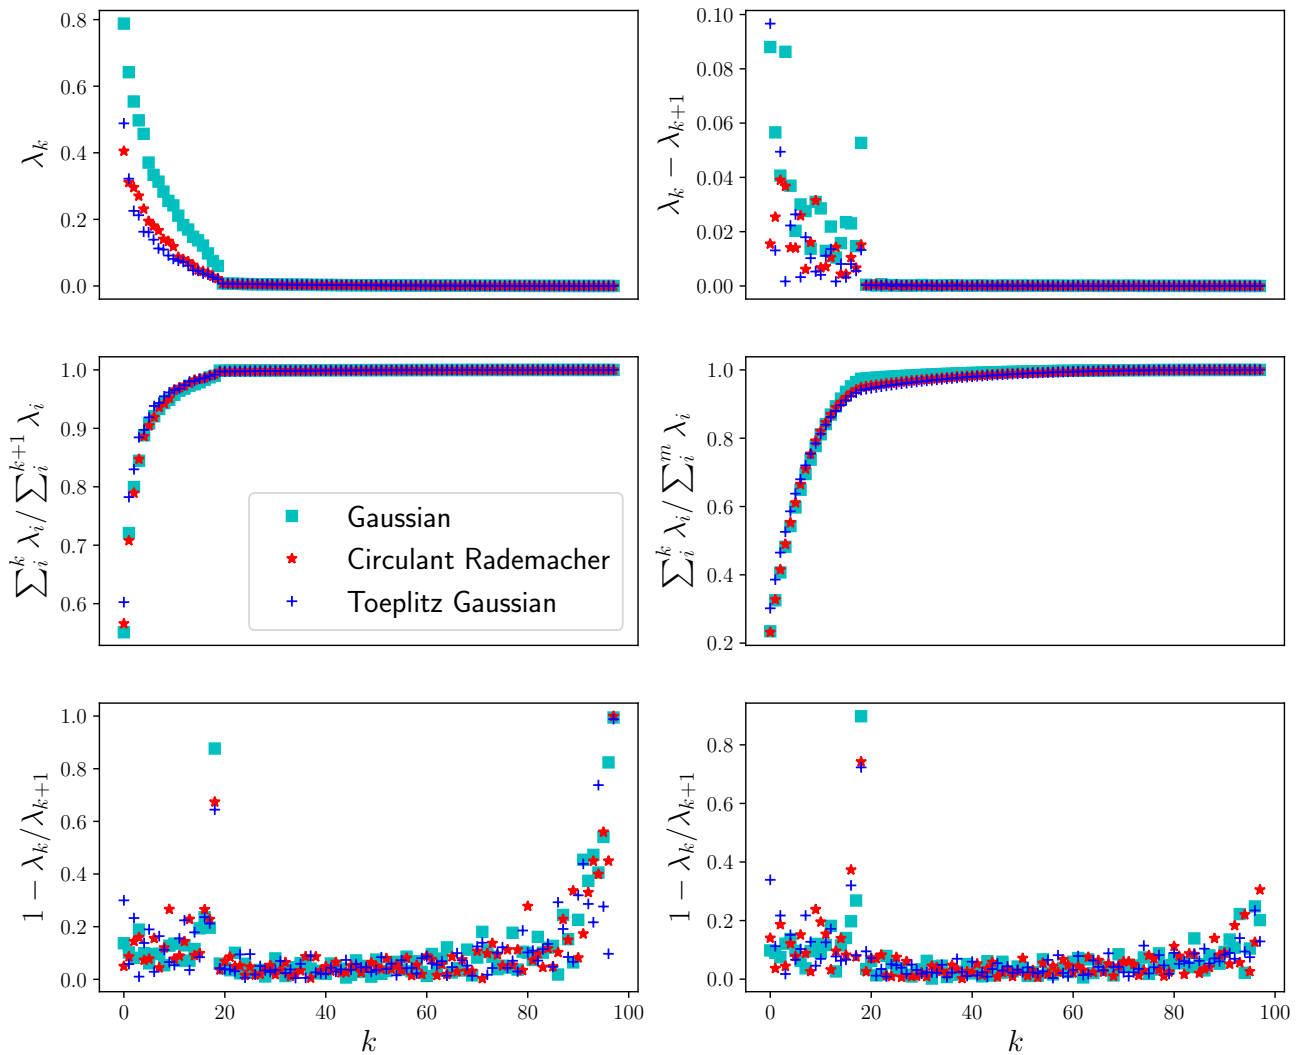


Figure 2. Plot in the upper left corner shows the spectrum of three different types of matrices. The remaining plots (not including the one in the bottom right corner) consider different notions of the spectral gap for $N = 100$. The plot in the bottom right considers the last criteria, $1 - \frac{\lambda_k}{\lambda_{k+1}}$ but for $N = 150$ samples, showing that the behavior for large k changes dramatically, compared to the plot in the bottom right for $N = 100$.

6.2. Synthetic examples

Experimental setting. We consider the inverse problem of the type $\mathbf{y} = \mathbf{A}\mathbf{x} + \sigma\mathbf{w}$, where the data are generated according to $\mathbf{x} = \boldsymbol{\xi} + \mathbf{v}$, where

- (D1) $\mathbf{A} \in \mathbb{R}^{m \times d}$ is a random Gaussian matrix such that $\|\mathbf{A}\|_2 = 1$,
- (D2) $\xi_i \sim \mathcal{N}(0, 1)$, and $v_i = 4 \operatorname{sgn}(\xi_i)$ for $1 \leq i \leq h$; $\xi_i = v_i = 0$ otherwise,
- (D3) $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \operatorname{Id}_m)$.

The rationale behind distributional choices in (D2) is twofold. First, having \mathbf{v} be a non-constant vector ensures that $\mathcal{V} = \operatorname{range}(\boldsymbol{\Sigma}(\mathbf{x}))$ is truly and fully h -dimensional (*i.e.* if \mathbf{v} were constant there

would be one dominant singular vector). Second, forcing $|v_i| = 4$ for $i = 1, \dots, h$ ensures that the data are not concentrated around the origin (which happens e.g., if $v_i = 0$ for all i in (D2)) and that there is a gap between the original signal and the noise. The gap can be measured by $\text{SparseSNR} := \frac{\max_{1 \leq i \leq m} |\sigma w_i|}{\min_{1 \leq i \leq h} |x_i|}$. We add that the conclusions and the results of this section, and of Section 6.1, stay the same in case of the more usual distribution assumptions, i.e., for $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \text{Id}_h)$ where Id_h is a diagonal matrix with exactly h entries set to 1 and the rest to 0, and noise levels as in Section 6.3.

Comparison. We compare our algorithm with the following parameter selection methods: the discrepancy principle [25], monotone error rule [29], quasi optimality [30], L-curve [20], (Monte-Carlo) balancing principle [23] and its elastic net counterpart [9], (Monte-Carlo) generalized cross-validation [18] and nonlinear cross validation [17]. In the remainder of this paper we refer to other methods by their acronyms, and to our method as *OptEN*. The first five methods are commonly used in inverse problems (a detailed account and an experimental study can be found in [3]), whereas Monte-Carlo and nonlinear cross-validation are adaptations of generalized cross-validation for non-linear regularization methods.

Before presenting the results, we provide a concise description of considered methods. Most of the methods require some additional information about the problem, predominantly the noise level σ , to be either known or estimated, which affects their performance. We provide the true noise level whenever a given method requires it and furthermore, we perform judicious testing and tuning of all other quantities, taking into account recommendations from relevant literature. We consider a regularization parameter sequence $t_n = \frac{1}{1+\nu_0 q^n}$, where $n \in \{0, 1, \dots, N_{\max}\}$, and $\nu_0 > 0$, $q > 0$ and $N_{\max} \in \mathbb{N}$ are preselected*. For each n we denote the corresponding elastic net solution as $\mathbf{z}_n := \mathbf{z}^{t_n}$.

Discrepancy Principle [DP]

Discrepancy principle is one of the oldest parameter choice rules which selects a solution so that the norm of the residual is at the noise level. Thus, the regularization parameter is chosen by the first $n \in \mathbb{N}$ such that

$$\|\mathbf{A}\mathbf{z}_n - \mathbf{y}\|_2 \leq \tau\sigma\sqrt{m}, \quad (6.1)$$

where we fix $\tau = 1$.

Monotone Error Rule [ME]

This rule is based on the observation that the monotone decrease of the error $\|\mathbf{z}_n - \mathbf{x}\|_2$ can only be guaranteed for large values of the regularization parameter. Therefore, the best parameter t_{n^*} is chosen as the first t -value for which one can ensure that the error is monotonically decreasing. The parameter is then chosen by the smallest n such that

$$\frac{\langle \mathbf{A}\mathbf{z}_n - \mathbf{y}, \mathbf{A}^{-\top}(\mathbf{z}_n - \mathbf{z}_{n+1}) \rangle}{\|\mathbf{A}^{-\top}(\mathbf{z}_n - \mathbf{z}_{n+1})\|_2} \leq \tau\sigma\sqrt{m}. \quad (6.2)$$

We fix $\tau = 1$ for our experiments. The left hand side of (6.2) is replaced with (6.1) whenever the denominator is 0.

*This is an adaptation of the parameter sequence from [3] that reflects our reparametrization from λ to t , as in (3.2)

Quasi-Optimality Criterion [QO]

Quasi-optimality is a parameter rule that does not need the noise level, and thus has enjoyed reasonable success in practice, especially for Tikhonov regularization and truncated singular value decomposition. The regularization parameter is chosen according to

$$n_{\star} = \operatorname{argmin}_{n \leq N_{\max}} \|\mathbf{z}_n - \mathbf{z}_{n+1}\|_2. \quad (6.3)$$

L-curve method [LC]

The criterion is based on the fact that the log – log plot of $(\|\mathbf{A}\mathbf{z}_n - \mathbf{y}\|_2, \|\mathbf{z}_n\|_2)$ often has a distinct L-shape. As the points on the vertical part correspond to under-smoothed solutions, and those on the horizontal part correspond to over-smoothed solutions, the optimal parameter is chosen at the elbow of that *L-curve*. There exist several versions of the method; here we use the following criterion

$$n_{\star} = \operatorname{argmin}_{n \leq N_{\max}} \{\|\mathbf{A}\mathbf{z}_n - \mathbf{y}\|_2 \|\mathbf{z}_n\|_2\}. \quad (6.4)$$

(Monte-Carlo) Balancing Principle [BP]

The principle aims to balance two error contributions, approximation and sampling errors, which have an opposite behavior with respect to the tuning parameter. More precisely, we select the parameter by

$$n^* = \operatorname{argmin}_n \{t_n \|\mathbf{z}_n - \mathbf{z}_k\|_2 \leq 4\kappa\sigma\rho(k), k = n, \dots, N_{\max}\},$$

where $\kappa > 0$ is a tuning parameter. More computationally friendly, yet equally accurate, versions of the balancing principle are also available [3]. As our main focus on the accuracy of the parameter choice, we will use the original and more computationally heavy version of the balancing principle.

The value of $\sigma\rho(k)$ is in general unknown but it can be estimated in case of white noise. Following [3], we calculate $\rho(k)^2 \approx \operatorname{mean}\{\|\mathbf{A}_n^{-1}\xi_i\|_2^2\}$, where $\xi_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$, $1 \leq i \leq L$ (we use $L = 4$), and \mathbf{A}_n^{-1} is the map that assigns \mathbf{y} to \mathbf{z}_n .

Elastic Nets Balancing Principle [ENBP]

In [9] the authors propose a reformulation of the balancing principle for elastic net,

$$n^* = \operatorname{argmin}_n \{t_n \|\mathbf{z}_k - \mathbf{z}_{k+1}\|_2 \leq \frac{4C}{\sqrt{d}\alpha\nu_0 q^{k+1}}, k = N_{\max} - 1, \dots, n\},$$

The method stops the first time two solutions are sufficiently far apart. The constant C needs to be selected, and in our experience this task requires a delicate touch.

(Monte-Carlo) Generalized Cross-Validation [GCV]

The rule stems from the ordinary cross-validation, which considers all the *leave-one-out* regularized solutions and chooses the parameter that minimizes the average of the squared prediction errors. Specifically, GCV selects n according to

$$n_{\star} = \operatorname{argmin}_{n \leq N_{\max}} \frac{m^{-1} \|\mathbf{A}\mathbf{z}_n - \mathbf{y}\|_2^2}{\left(m^{-1} \operatorname{tr}(\mathbf{I}_m - \mathbf{A}\mathbf{A}_n^{-1})\right)^2}, \quad (6.5)$$

where A_n^{-1} is the map such that $\mathbf{z}_n = A_n^{-1}(\mathbf{y})$. In the case of the elastic net, the map A_n^{-1} is not linear and, thus, we cannot assign a meaning to its trace. Instead, we follow the ideas of [18] and estimate the trace stochastically using only one data sample.

Nonlinear Generalized Cross-Validation [NGCV]

In [17] the authors reconfigure GCV for non-linear shrinkage methods, and n is selected according to

$$n_* = \operatorname{argmin}_{n \leq N_{\max}} \frac{\|\mathbf{A}\mathbf{z}_n - \mathbf{y}\|_2^2}{m^{-1}(1 - ds/m)^2},$$

where $s = \frac{\|\mathbf{z}_n\|_y}{\|\mathbf{z}^*\|_y}$ with $\|\cdot\|_y := \|\cdot\|_1 + \alpha \|\cdot\|_2^2$.

Comparison/Error. For each method we compute: the (normalized) error in approximating the optimal regularization parameter, the (normalized) error, false discovery proportion (FDP), true positive proportion (TPP), and the computational time. TPP and FDP are measures that quantify the notions of true and false discovery of relevant features in sparsity based regression tasks [28]. FDP is the ratio between false discoveries and the total number of discoveries,

$$\text{FDP}(t) = \frac{\#\left[j : z_j^t \neq 0 \text{ and } x_j = 0\right]}{\max\left(\#\left[j : z_j^t \neq 0\right], 1\right)}.$$

TPP on the other hand is the ratio between true (i.e., correct) discoveries in the reconstruction and true discoveries in the original signal,

$$\text{TPP}(t) = \frac{\#\left[j : z_j^t \neq 0 \text{ and } x_j \neq 0\right]}{h}.$$

Thus, to recover the structure of the original sparse data we want FDP close to 0 and TPP close to 1. It is known that there is often an explicit (and sometimes even quantifiable) trade-off between FDP and TPP, in the sense that the support overestimation is an (undesirable) side-effect of full support recovery. In other words, a consequence of true support discovery is often a non-trivial false support discovery [28]. When computing FDP and TPP we will rather than demand for an entry to be exactly zero, instead threshold the values (with 0.5 being the threshold).

Testing setup. To compute the true optimal parameter t_{opt} , we run a dense grid search on $[0, 1]$ using the true expected loss $\|\mathbf{z}^t - \mathbf{x}\|_2^2$. As suggested in [3] we use $\tau = 1$ and provide the true noise level σ for discrepancy principle and monotone error rule; balancing principle uses $\kappa = 1/4$ and true σ ; elastic net balancing principle uses $C = 1/2500$. The parameter grid for DP, ME, BP, QO, LC, BP, GCV, and NGCV is defined by $v_0 = 1$, $q = 0.95$, and $N_{\max} = 100$ (thus, $t_0 = 0.5$ and $t_{N_{\max}} = 0.9941143171$), whereas for ENBP, we use $t_0 = 0.05$, $q = 1.05$, and $N_{\max} = 100$. The tests are conducted for $m \in \{500, 900\}$, $d \in \{100, 200\}$ and $h \in \{10, 20, 30\}$, where all combinations of $\alpha \in \{10^{-5}, 10^{-3}, 10^{-2}, 10^{-1}\}$ and $\sigma \in \{0.05, 0.1, 0.2, 0.3\}$ are considered. To compute the empirical estimator $\widehat{\mathbf{x}}$, we generate $N = 50$ independent random samples of the training data (\mathbf{x}, \mathbf{y}) .

Results in Table 1 are averaged over 100 independent runs for $\alpha = 10^{-3}$, $m = 500$, $d = 100$, $h = 10$, where $\mathcal{V} = \text{span}\{\mathbf{e}_1, \dots, \mathbf{e}_h\}$, and $\sigma = 0.3$, which corresponds to $\text{SparseSNR} \approx 0.17$. The first row in the table, $x_{t_{opt}}$, describes the elastic net minimizer for which the *true* optimal regularization parameter is provided.

Table 1. Comparison of errors for regularization parameter selection methods, with an injective matrix $\mathbf{A} \in \mathbb{R}^{500 \times 100}$ and $h = 10$, $\alpha = 10^{-3}$, $\sigma = 0.3$. The values are averages over 100 independent runs. We write in bold the method that achieved the best result.

| Method | $\frac{ t_{opt} - \hat{t} }{t_{opt}}$ | $\frac{\ \mathbf{x} - \mathbf{z}^t\ _2}{\ \mathbf{x}\ _2}$ | FDP(\hat{t}) | TPP(\hat{t}) | computational time [s] |
|------------------------|---------------------------------------|--|------------------|------------------|------------------------|
| $x_{t_{opt}}$ | 0 | 0.0984 | 0.1583 | 1.000 | 0 |
| $\widehat{\mathbf{x}}$ | N/A | 0.1004 | 0.000 | 1.000 | N/A |
| OptEN | 0.0254 | 0.0994 | 0.160 | 1.000 | 3.716 |
| DP | 0.1196 | 0.1118 | 0.087 | 1.000 | 0.7421 |
| ME | 0.2493 | 0.1376 | 0.010 | 1.000 | 0.1757 |
| QO | 0.4543 | 0.1843 | 0.716 | 1.000 | 7.193 |
| LC | 0.1414 | 0.1174 | 0.188 | 1.000 | 1.564 |
| BP | 0.0877 | 0.1057 | 0.095 | 1.000 | 36.77 |
| ENBP | 0.2728 | 0.1450 | 0.007 | 1.000 | 4.355 |
| GCV | 0.4548 | 0.1844 | 0.716 | 1.000 | 15.91 |
| NGCV | 0.3597 | 0.1577 | 0.638 | 1.000 | 7.306 |

Discussion. OptEN always returns the value which is the closest to the optimal regularization parameter, and its results are in general comparable to the ones provided by the minimizer with the optimal parameter. However, one can observe that other methods, e.g., discrepancy principle, provide a better balance between FDP and TPP (returning solutions that are more sparse), though at a cost of a larger approximation error. Balancing principle also provides very good results, but it is slow unless an effort is made to improve its computational time. Moreover, we observed that the performance of all methods that require the noise level σ to be known deteriorates if we do not provide the exact value of the noise level, but only its rough estimate.

The overall results are mostly consistent over all experimental scenarios we looked at, with a couple of exceptions. As expected, FDP and estimation errors deteriorate not only for larger σ but also for larger α , though the ranking of the methods and the patterns of behavior remain the same. This is due to the fact that as α increases the elastic net penalty sacrifices sparsity for smoothness. The empirical estimator $\widehat{\mathbf{x}}$ is a very accurate estimator of the original signal, and it sometimes outperforms even the elastic net solution that uses the optimal parameter. However, as it has been observed in [10] for Tikhonov regularization and confirmed in Figure 3, the performance of the empirical estimator worsens in the small noise regime.

Comparison with empirical estimator: effects of σ and N_{train}

We study the behavior of the relative estimation error with respect to σ and the number of training samples. We compare OptEN with the empirical estimator $\widehat{\mathbf{x}}$, DP, NGCV, and BP. We use $m = 500$, $d = 100$, $h = 10$ and (D1)-(D3) with σ ranging from 0.1 to 0.5 in the first experiment, whereas we vary

the number of training samples N from 20 to 60 in the second experiments as depicted in Figure 3. Our method again outperforms other considered parameter selection rules. However, the empirical estimator performs slightly better than OptEN for larger noise levels (it is also better than the optimal elastic net solution), and it performs worse for lower noise levels. This is essentially due to the fact that $\widehat{\mathbf{x}}$ is never truly sparse, but has a lower (thresholded) FDP. Namely, as a projection onto an h -dimensional space the non-zero entries of $\widehat{\mathbf{x}}$ are very small, whereas the non-zero entries obtained with elastic net are larger and their size depending on the noise level. The parameter α plays a similar role; for small α OptEN beats $\widehat{\mathbf{x}}$, and for larger α the situation is reversed.




Figure 3. In the left panel is the behavior of OptEN, the empirical estimator, discrepancy principle, balancing principle, and nonlinear GCV with respect to σ is shown. In the right panel the behavior of our method for different values of σ as the number of samples N increases is shown. Dashed lines represent the error achieved by taking the true optimal parameter for the corresponding σ .

Non-injective matrices

We now conduct experiments with non-injective matrices. The setting is as in Table 1, where now $A \in \mathbb{R}^{500 \times 100}$ with $\text{rank}(A) = 40$. As mentioned in Section 3.2, we test our method by minimizing the projected loss functional \widehat{R}_p and the modified error functional \widehat{R}_m . The results can be found in Table 2. Our method (using both the projected and modified functionals) again outperforms standard parameter selection rules in terms of the precision accuracy, and loses out to some methods when it comes to FDP and TPP. We also observe that the performance of the empirical estimator deteriorates and $\widehat{\mathbf{x}}$ indeed should not be used as the solution itself but some additional regularization is required.

Table 2. Comparison of errors for regularization parameter selection methods, with a matrix $A \in \mathbb{R}^{500 \times 100}$, $\text{rank}(A) = 40$, and $h = 10$, $\alpha = 10^{-3}$, $\sigma = 0.3$. The values are averages over 100 independent runs. We write in bold the method that achieved the best result.

| Method | $\frac{ t_{opt} - \hat{t} }{t_{opt}}$ | $\frac{\ \mathbf{x} - \mathbf{z}^*\ _2}{\ \mathbf{x}\ _2}$ | $FDP(\hat{t})$ | $TPP(\hat{t})$ | computational time [s] |
|--|---------------------------------------|--|----------------|----------------|------------------------|
| using t_{opt} | 0 | 0.5704 | 0.4918 | 0.9450 | 0 |
| empirical estimator $\hat{\mathbf{x}}$ | N/A | 0.7978 | 0.8047 | 1.000 | N/A |
| projected OptEN | 0.0718 | 0.6033 | 0.507 | 0.930 | 8.16 |
| modified OptEN | 0.0763 | 0.6046 | 0.497 | 0.926 | 15.57 |
| DP | 0.1316 | 0.6343 | 0.528 | 0.926 | 1.93 |
| ME | 0.3203 | 0.7234 | 0.290 | 0.765 | 0.15 |
| QO | 0.3167 | 0.7857 | 0.819 | 0.997 | 7.04 |
| LC | 0.3389 | 0.7426 | 0.304 | 0.749 | 7.00 |
| GCV | 0.3172 | 0.7865 | 0.819 | 0.997 | 14.04 |
| NGCV | 0.3172 | 0.7865 | 0.819 | 0.997 | 7.06 |
| BP | 0.2636 | 0.7179 | 0.757 | 0.974 | 35.01 |
| ENBP | 0.2133 | 0.6549 | 0.362 | 0.847 | 3.69 |

6.3. Image denoising

The task of image denoising is to find an estimate \mathbf{Z} of an unknown image \mathbf{X} from a noisy measurement \mathbf{Y} , where $\mathbf{Y} = \mathbf{X} + \sigma\boldsymbol{\Xi}$, and $\boldsymbol{\Xi}$ denotes isotropic white noise. The goal is to improve the image quality by removing noise while preserving important image features such as edges and homogeneous regions.

There are a large number of methods addressing image denoising, starting from ‘classical’ wavelet thresholding [12, 13] and non-linear filters, to stochastic and variational methods [7, 8]. Since the primary goal of this paper is to evaluate how the proposed approach performs as a parameter selection method for the elastic net, here we only compare our method with other state-of-the-art parameter selection methods for the elastic net, and do not compare the elastic net with image denoising methods in general. In particular, we compare OptEN with the discrepancy principle and the balancing principle (i.e., the top performers from previous experiments). In all cases the results show that OptEN has superior performance and selects nearly optimal parameters, see Table 3.

Wavelet-based denoising. We denoise the noisy image \mathbf{Y} by minimizing

$$(1 - t) \|\mathbf{Z} - \mathbf{Y}\|_2^2 + t(\|\mathbb{W}\mathbf{Z}\|_1 + \alpha \|\mathbf{Z}\|_2^2), \text{ for } \mathbf{Z} \in [0, 1]^p, \quad (6.6)$$

where \mathbb{W} is the wavelet transform using the family of db4 wavelets, $\alpha = 10^{-3}$, and p is the number of pixels in the image. The wavelet transform sparsifies natural images, and we thus select the empirical estimator in the wavelet domain. Moreover, in the limit with respect to the number of samples $N \rightarrow \infty$, the empirical projection $\widehat{\Pi}\mathbf{Y}$ is for a given h equivalent to a hard thresholding of $\mathbb{W}\mathbf{Y}$ that preserves its h largest wavelet coefficients. Thus, for image denoising we do not use samples \mathbf{Y}_i but instead only threshold $\mathbb{W}\mathbf{Y}$ for a well chosen h . Here h cannot be chosen by searching for a gap in $\mathbb{W}\mathbf{Y}$, since it most often does not exist. Instead, we say that the *true* h is the one that minimizes the MSE of

the reconstructed image. In our first set of experiments the empirical estimator \widehat{X} is chosen by hard thresholding Y , where h is optimal.

Table 3. Results on wavelet denoising of noisy images with different noise levels using the elastic net minimization. Each column defines the method used to select the regularization parameter. In bold is the method that achieved the best result. Columns titled *noisy* correspond to PSNR and SSIM values of the initial noisy image. Columns titled *using λ_{opt}* correspond to the best values achievable for the selected the elastic net functional, where we find the optimal parameter by a grid search on the true loss functional. We write in bold the method that achieved the best result.

| | PSNR | | | | | SSIM | | | | |
|------------------|-----------------|-------|--------------|-------|-------|-----------------|-------|--------------|-------|-------|
| | using t_{opt} | noisy | OptEN | DP | BP | using t_{opt} | noisy | OptEN | DP | BP |
| space shuttle | | | | | | | | | | |
| $\sigma = 0.05$ | 32.33 | 26.27 | 32.32 | 30.26 | 30.63 | 0.929 | 0.687 | 0.929 | 0.850 | 0.862 |
| $\sigma = 0.075$ | 30.26 | 22.87 | 30.25 | 27.32 | 27.69 | 0.908 | 0.516 | 0.908 | 0.748 | 0.765 |
| $\sigma = 0.1$ | 28.77 | 20.50 | 28.76 | 25.32 | 25.60 | 0.892 | 0.394 | 0.891 | 0.660 | 0.676 |
| cherries | | | | | | | | | | |
| $\sigma = 0.05$ | 35.80 | 26.08 | 35.79 | 30.52 | 31.08 | 0.973 | 0.647 | 0.972 | 0.837 | 0.855 |
| $\sigma = 0.075$ | 33.59 | 22.63 | 33.59 | 27.40 | 27.77 | 0.964 | 0.463 | 0.964 | 0.720 | 0.737 |
| $\sigma = 0.1$ | 31.94 | 20.24 | 31.82 | 25.26 | 25.54 | 0.958 | 0.339 | 0.942 | 0.617 | 0.633 |
| cat | | | | | | | | | | |
| $\sigma = 0.05$ | 29.34 | 26.02 | 29.33 | 28.72 | 28.93 | 0.890 | 0.767 | 0.890 | 0.861 | 0.868 |
| $\sigma = 0.075$ | 27.06 | 22.51 | 27.03 | 25.85 | 26.04 | 0.825 | 0.612 | 0.823 | 0.758 | 0.767 |
| $\sigma = 0.1$ | 25.69 | 20.04 | 25.24 | 23.88 | 24.03 | 0.771 | 0.486 | 0.741 | 0.664 | 0.671 |
| mud flow | | | | | | | | | | |
| $\sigma = 0.05$ | 28.38 | 26.02 | 28.37 | 28.20 | 28.30 | 0.870 | 0.777 | 0.869 | 0.856 | 0.860 |
| $\sigma = 0.075$ | 26.07 | 22.50 | 26.06 | 25.46 | 25.60 | 0.795 | 0.629 | 0.795 | 0.755 | 0.762 |
| $\sigma = 0.1$ | 24.71 | 20.01 | 24.46 | 23.58 | 23.70 | 0.735 | 0.506 | 0.717 | 0.662 | 0.668 |
| IHC | | | | | | | | | | |
| $\sigma = 0.05$ | 29.33 | 26.02 | 29.33 | 28.73 | 28.92 | 0.890 | 0.767 | 0.889 | 0.861 | 0.868 |
| $\sigma = 0.075$ | 27.06 | 22.51 | 27.04 | 25.86 | 26.04 | 0.825 | 0.612 | 0.823 | 0.759 | 0.767 |
| $\sigma = 0.1$ | 25.70 | 20.04 | 25.39 | 23.85 | 24.04 | 0.772 | 0.486 | 0.748 | 0.662 | 0.671 |

6.3.1. Denoising with an oracle h

Data and learning setup. We consider five grayscale images: space shuttle, cherries, cat, mud flow, and IHC, each of size 512×512 pixels. For BP and DP the regularization parameter is selected from a sequence of parameter values $t_n = \frac{1}{1+\nu_0 q^n}$ with $\nu_0 = 1$, $q = 0.95$, and $N_{\max} = 100$, same as before. Moreover, we fix $\tau = 1$, $\kappa = 1/4$ for BP and DP, and provide them with the true noise level. For OptEN the empirical estimator is computed with an oracle h , i.e., the one returning the lowest MSE.

Comparison/Error. We use two performance metrics: peak signal-to-noise ratio (PSNR) and the structural similarity index (SSIM) between the original image \mathbf{X} and the recovered version \mathbf{Z} . PSNR is a standard pixel-based performance metric, defined through the MSE by

$$\text{PSNR}(\mathbf{X}, \mathbf{Z}) = 10 \log_{10} \left(\frac{\max_{1 \leq i \leq p} X_i - \min_{1 \leq i \leq p} X_i}{\text{MSE}(\mathbf{X}, \mathbf{Z})} \right), \quad \text{MSE}(\mathbf{X}, \mathbf{Z}) = \frac{1}{p} \|\mathbf{X} - \mathbf{Z}\|_2^2.$$

MSE and PSNR are ubiquitous in image and signal analysis due to their simplicity and suitability for optimization tasks, but are also infamous for their inability to capture features salient for human perception of image quality and fidelity [35]. SSIM on the other hand, is a structure-based performance metric that tries to address this issue by using easy-to-compute structural statistics to estimate image similarity. It is defined through

$$\text{SSIM}(\mathbf{X}, \mathbf{Z}) = \left(\frac{2\bar{X}\bar{Z} + C_1}{\bar{X}^2 + \bar{Z}^2 + C_1} \right) \left(\frac{2\text{std}(\mathbf{X})\text{std}(\mathbf{Z}) + C_2}{\text{std}(\mathbf{X})^2 + \text{std}(\mathbf{Z})^2 + C_2} \right),$$

where \bar{X} , \bar{Z} are the means, and $\text{std}(\mathbf{X})$, $\text{std}(\mathbf{Z})$ are the standard deviations of pixels of corresponding images \mathbf{X} and \mathbf{Z} , and C_1, C_2 are positive constants[†].

Table 3 provides the PSNR and SSIM values generated by all algorithms on the considered images, while Figure 4 shows the result of denoising on a 128×128 detail of each image for $\sigma = 0.075$. We can see that our method achieves the highest PSNR on all images and that this effect is more pronounced for larger noise values.

6.3.2. Denoising with a heuristically chosen h

In this set of experiments we study the performance of our method in a situation where the optimal h is not known *a priori*. We run experiments on a real-world dataset of brain images[‡], which consists of *in vivo* MRIs of 13 patients with brain tumor, taken pre-surgery. For each patient we took an MRI slice, isolated the area around the brain and then added additional isotropic white noise with $\sigma \in \{0.05, 0.075, 0.1\}$. We then select h by a heuristically driven procedure. Namely, for each image \mathbf{Y} we set an initial $h_0 \in \mathbb{N}$ and determine \hat{t}_0 by performing Algorithm 1, where $\widehat{\mathbf{X}}_{h_0}$ is constructed by taking h_0 largest coefficients of $\mathbb{W}\mathbf{Y}$. We then set $h_1 = h_0 + h_{\text{step}}$, repeat the procedure, and continue iteratively for h_k . The iterations are stopped once the corresponding \hat{t}_k start to decrease or become discontinuous (since heuristically this corresponds to a decrease in the PSNR of the corresponding elastic-net regularized solution). h_0 and h_{step} are chosen according to the size of the image. The behavior of this criterion can be seen in Figure 5, and it shows that if h is too large the empirical estimator is virtually the same as \mathbf{Y} . In other words, the minimizer of $\|\mathbf{Z}' - \widehat{\mathbf{X}}_h\|_2$ is $t = 1$ (i.e., $\lambda = 0$), which we observe in Figure 5.

The resulting reconstruction for $\sigma = 0.1$ can be seen in Figure 6 on four images. The effects of denoising are visually not as striking as the results in Section 6.3.1. We attribute this to the fact that PSNR gains with the best possible choice of parameter using elastic net are quite small, namely, PSNR of the noisy image improves only by around 5 – 7%, when taking the optimal parameter (see Table 4).

[†]We take $C_1 = 0.01$, $C_2 = 0.03$ by the convention of python's `skimage` package

[‡]Obtained from http://nist.mni.mcgill.ca/?page_id=672, therein referred to as *group 2*

Other parameter selection rules (discrepancy principle and balancing principle) did not improve the PSNR and are thus not presented.




Figure 4. Comparing the effects of denoising for different parameter selection rules.





Figure 5. A visual representation of our heuristic criteria for selecting h . Here we choose $h = 4400$. The flat line at the end corresponds to $\hat{t}_k = 1$.

Table 4. PSNR values for the results in Figure 6.

| using t_{opt} | noisy | OptEN |
|------------------------|--------|--------|
| 28.608 | 27.297 | 28.432 |
| 28.329 | 26.792 | 28.079 |
| 28.221 | 26.735 | 27.935 |
| 28.501 | 26.906 | 28.240 |

**Figure 6.** Denoising results on a brain image data set. The PSNR values are in Table 4.

7. Conclusions and future work

In this paper, we presented an approach for the estimation of the optimal regularization parameter for the elastic net. Theoretical guarantees are available only in simplified scenarios but we used insights gained therein to steer and create an efficient algorithm. The algorithm exhibits excellent prediction

accuracy, including in cases when there are no theoretical guarantees. Comparison with state-of-the-art methods show a clear superiority of our method, under the studied testing scenarios. Moreover, whereas other studied methods require adjusting a number of additional parameters in order to achieve satisfactory results, our method is entirely autonomous given a sufficient number of training samples.

We aim to use the ideas presented in this paper in further studies. Namely, we will study the behavior of the solution with respect to the other hyperparameter, α , and consider other optimization schemes, predominantly focusing on imaging applications. We will also work on developing an optimization scheme for a joint minimization of both the regularization functional and the loss functional for the regularization parameter.

Acknowledgments

V. Naumova and Z. Kereta acknowledge the support from RCN-funded FunDaHD project No 251149/O70.

Conflict of interest

The authors declare no conflict of interest.

References

1. S. W. Anzengruber, R. Ramlau, Morozov's discrepancy principle for Tikhonov-type functionals with nonlinear operators, *Inverse Probl.*, **26** (2010), 025001.
2. A. Astolfi, *Optimization: An introduction*, 2006.
3. F. Bauer, M. A. Lukas, Comparing parameter choice methods for regularization of ill-posed problems, *Math. Comput. Simulat.*, **81** (2011), 1795–1841.
4. M. Belkin, P. Niyogi, V. Sindhvani, Manifold regularization: A geometric framework for learning from labeled and unlabeled examples, *J. Mach. Learn. Res.*, **7** (2006), 2399–2434.
5. R. Bhatia, *Matrix analysis*, New York: Springer-Verlag, 1997.
6. T. Bonesky, Morozov's discrepancy principle and Tikhonov-type functionals, *Inverse Probl.*, **25** (2009), 015015.
7. A. Chambolle, R. DeVore, N. Lee, B. Lucier, Nonlinear wavelet image processing: variational problems, compression, and noise removal through wavelet shrinkage, *IEEE Trans. Image Process.*, **7** (1998), 319–335.
8. A. Chambolle, P. L. Lions, Image recovery via total variation minimization and related problems, *Numer. Math.*, **76** (1997), 167–188.
9. C. De Mol, E. De Vito, L. Rosasco, Elastic-net regularisation in learning theory, *J. Complexity*, **25** (2009), 201–230.
10. E. De Vito, M. Fornasier, V. Naumova, A machine learning approach to optimal Tikhonov regularization I: affine manifolds, *Anal. Appl.*, 2021, in press.

11. C.-A. Deledalle, S. Vaiter, J. Fadili, G. Peyré, Stein Unbiased Gradient estimator of the Risk (SUGAR) for multiple parameter selection, *SIAM J. Imaging Sci.*, **7** (2014), 2448–2487.
12. D. L. Donoho, De-noising by soft-thresholding, *IEEE Trans. Inform. Theory*, **41** (1995), 613–627.
13. D. L. Donoho, I. Johnstone, Ideal spatial adaptation via wavelet shrinkage, *Biometrika*, **81** (1994), 425–455.
14. B. Efron, T. Hastie, I. Johnstone, R. Tibshirani, Least angle regression, *The Annals of Statistics*, **32** (2004), 407–499.
15. Y. C. Eldar, Generalized SURE for exponential families: applications to regularization, *IEEE Trans. Signal Process.*, **57** (2009), 471–481.
16. H. W. Engl, M. Hanke, A. Neubauer, *Regularization of inverse problems*, Dordrecht: Kluwer Academic Publishers, 1996.
17. W. J. Fu, Nonlinear GCV and quasi-GCV for shrinkage models, *J. Stat. Plan. Infer.*, **131** (2005), 333–347.
18. R. Giryes, M. Elad, Y. C. Eldar, The projected GSURE for automatic parameter tuning in iterative shrinkage methods, *Appl. Comput. Harmonic Anal.*, **30** (2011), 407–422.
19. G. H. Golub, M. Heath, G. Wahba, Generalized cross-validation as a method for choosing a good ridge parameter, *Technometrics*, **21** (1979), 215–223.
20. P. C. Hansen, Analysis of discrete ill-posed problems by means of the L-curve, *SIAM Rev.*, **34** (1992), 561–580.
21. B. Hofmann, *Regularization for applied inverse and ill-posed problems: a numerical approach*, Springer-Verlag, 2013.
22. B. Jin, D. Lorenz, S. Schiffler, Elastic-net regularisation: error estimates and active set methods, *Inverse Probl.*, **25** (2009), 115022.
23. O. Lepskii, On a problem of adaptive estimation in Gaussian white noise, *Theory Probab. Appl.*, **35** (1991), 454–466.
24. A. Lorbert, P. Ramadge, Descent methods for tuning parameter refinement, In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 2010, 469–476.
25. V. A. Morozov, *Methods for solving incorrectly posed problems*, Springer Science & Business Media, 2012.
26. E. Novak, H. Woźniakowski, Optimal order of convergence and (in)tractability of multivariate approximation of smooth functions, *Constr. Approx.*, **30** (2009), 457–473.
27. C. M. Stein, Estimation of the mean of a multivariate normal distribution, *Ann. Statist.*, **9** (1981), 1135–1151.
28. W. Su, M. Bogdan, E. Candés, False discoveries occur early on the Lasso path, *Ann. Statist.*, **45** (2017), 2133–2150.
29. U. Tautenhahn, U. Hämarik, The use of monotonicity for choosing the regularization parameter in ill-posed problems, *Inverse Probl.*, **15** (1999), 1487–1505.
30. A. Tikhonov, V. Glasko, Use of the best rate of adaptive estimation in some inverse problems, *USSR Computational Mathematics and Mathematical Physics*, **5** (1965), 93–107.

31. A. N. Tikhonov, V. Y. Arsenin, *Solutions of ill-posed problems*, Washington, D. C.: V. H. Winston & Sons, New York: John Wiley & Sons, 1977.
32. J. A. Tropp, *An introduction to matrix concentration inequalities*, Now Foundations and Trends, 2015.
33. R. Vershynin, *High-dimensional probability: An introduction with applications in data science*, Cambridge University Press, 2018.
34. E. De Vito, S. Pereverzyev, L. Rosasco, Adaptive kernel methods using the balancing principle, *Found. Comput. Math.*, **10** (2010), 455–479.
35. Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, Image quality assessment: from error visibility to structural similarity, *IEEE Trans. Image Process.*, **13** (2004), 600–612.
36. S. N. Wood, Modelling and smoothing parameter estimation with multiple quadratic penalties, *J. Roy. Stat. Soc. B*, **62** (2000), 413–428.
37. H. Zou, T. Hastie, Regularisation and variable selection via the elastic net, *J. Roy. Stat. Soc. B*, **67** (2005), 301–320.

Appendix - Supplementary proofs

Proofs for Section 4

We will here add more details regarding per-interval minimisers of the loss function $R(t)$, defined in (3.10). First, we note the expression

$$R(t) = \sum_{i=1}^m \left(\frac{(t(1 + 2|y_i|) - 1)_+}{2(t(1 - \alpha) + \alpha)} \operatorname{sgn}(y_i) - x_i \right)^2 =: \sum_{i=1}^m r_i(t),$$

established in 4. Define $\mathbf{b}_i = 1 + 2|y_i|$, for $i = 1, \dots, m$. Loss function $R(t)$ is continuous on $[0, 1]$, and differentiable on intervals

$$\mathcal{I}_0 = [0, \mathbf{b}_1^{-1}), \quad \mathcal{I}_m = (\mathbf{b}_m^{-1}, 1], \quad \text{and } \mathcal{I}_k = (\mathbf{b}_k^{-1}, \mathbf{b}_{k+1}^{-1}), \quad \text{for } k = 1, \dots, m - 1.$$

On \mathcal{I}_0 the loss function is constant, $R(t) = \sum_{i=1}^m x_i^2$. Define now $a_i = \operatorname{sgn}(y_i)(1 + 2\alpha|y_i|)$, $c_i = \operatorname{sgn}(y_i) + 2x_i(\alpha - 1) + 2y_i$, $d_i = \operatorname{sgn}(y_i) + 2\alpha x_i$. If $1 < k < m$ then $r_i(t) = x_i^2$ for $i > k$, and for $i \leq k$ we have $r_i(t) = \frac{1}{4} \left(\frac{tc_i - d_i}{t(1 - \alpha) + \alpha} \right)^2$, and thus

$$r'_i(t) = -\frac{1}{2}(d_i(1 - \alpha) + c_i\alpha) \frac{tc_i - d_i}{(\alpha + t(1 - \alpha))^3} = -\frac{1}{2}a_i \frac{tc_i - d_i}{(\alpha + t(1 - \alpha))^3}.$$

Therefore, $R'_{\mathcal{I}_k}(t) = -\frac{1}{2(\alpha + t(1 - \alpha))^3} \sum_{i=1}^k a_i(tc_i - d_i)$. Equating the above expression with zero and restricting the solution to the interval \mathcal{I}_k , yields the desired expression.

Proofs for Theorem 2.2

We will here add an analogue of Eq (2.8) for the case of bounded \mathbf{y} . Assume $\|\mathbf{y}\|_2 \leq \sqrt{L}$ holds almost surely and consider a random matrix $\mathbf{R} = \mathbf{y} \otimes \mathbf{y}$. Then $\mathbb{E}\mathbf{R} = \Sigma(\mathbf{y})$, and $\|\mathbf{R}\|_2 \leq L$. Furthermore, $\mathbf{R}^\top = \mathbf{R}$ and

$$m_2(\mathbf{R}) = \max \left\{ \|\mathbb{E}[\mathbf{R}\mathbf{R}^\top]\|_2, \|\mathbb{E}[\mathbf{R}^\top\mathbf{R}]\|_2 \right\} = \|\mathbb{E}\mathbf{R}^\top\mathbf{R}\|_2 \leq L \|\Sigma(\mathbf{y})\|_2.$$

Let now $\mathbf{y}_i \sim \mathbf{y}$ and define a family of independent $m \times m$ matrices

$$\mathbf{R}_i = \mathbf{y}_i \otimes \mathbf{y}_i, \quad i = 1, \dots, N,$$

so that $\mathbf{R}_i \sim \mathbf{R}$. The empirical covariance $\widehat{\Sigma}(\mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \mathbf{R}_i$ is then the matrix sampling estimator and by Corollary 6.2.1 from [32] we have that for all $s \geq 0$

$$\mathbb{P}\left(\left\|\widehat{\Sigma}(\mathbf{y}) - \Sigma(\mathbf{y})\right\|_2 \geq s\right) \leq 2m \exp\left(-\frac{Ns^2/2}{m_2(\mathbf{R}) + 2Ls/3}\right).$$

Writing now $2m \exp\left(-\frac{Ns^2/2}{m_2(\mathbf{R}) + 2Ls/3}\right) = \exp(-u)$, we have a quadratic equation for s , whose solution is

$$s = \frac{4L\varepsilon + \sqrt{(4L\varepsilon)^2 + 72Nm_2(\mathbf{R})\varepsilon}}{6N}$$

for $\varepsilon := u + 4 \log(2m)$. It then follows

$$\begin{aligned} s &\leq \frac{4L\varepsilon + 3\sqrt{2Nm_2(\mathbf{R})\varepsilon}}{3N} \leq \frac{\max(4L, 3\sqrt{2m_2(\mathbf{R})})}{3} \frac{\varepsilon + \sqrt{N\varepsilon}}{N} \\ &\leq C \left(\frac{u + \log(2m)}{N} + \sqrt{\frac{u + \log(2m)}{N}} \right) \end{aligned}$$

for $C = \frac{\max(4L, 3\sqrt{2m_2(\mathbf{R})})}{3}$. Plugging it all together we have that with probability at least $1 - \exp(-u)$

$$\left\|\widehat{\Sigma}(\mathbf{y}) - \Sigma(\mathbf{y})\right\|_2 \lesssim \frac{u + \log(2m)}{N} + \sqrt{\frac{u + \log(2m)}{N}}.$$

Thus, provided $N \gtrsim u + \log(2m)$ we have

$$\left\|\widehat{\Sigma}(\mathbf{y}) - \Sigma(\mathbf{y})\right\|_2 \leq \lambda_h/2.$$

Computations for $\alpha \neq 1$ in Section 4.2

Let $\mathbf{y} = \mathbf{x} + \sigma\mathbf{w}$, where $\mathbb{P}(w_i = \pm 1) = \frac{1}{2}$, and assume $\mathbf{x} = (x_1, \dots, x_h, 0, \dots, 0)^\top$, and $|x_i| \geq 2\sigma$, for $i = 1, \dots, h$. In the following we will use $\mathbf{v}_{1:k}$ to denote a vector in \mathbb{R}^k that consists of the first k entries of a vector $\mathbf{v} \in \mathbb{R}^m$, and denote $\boldsymbol{\eta} = \sigma\mathbf{w}$. In Section 4 we showed that the minimum of $R(t) = \|\mathbf{z}^t(\mathbf{y}) - \mathbf{x}\|_2^2$ and $\widehat{R}(t) = \|\mathbf{z}^t(\mathbf{y}) - \widehat{\mathbf{x}}\|_2^2$ in each sub-interval \mathcal{I}_k , for $k = 1, \dots, m$ of $[0, 1]$ is of the form

$$t^{*,k} = \frac{\sum_{i=1}^k a_i d_i}{\sum_{i=1}^k a_i c_i}, \quad \widehat{t}^{*,k} = \frac{\sum_{i=1}^k a_i \widehat{d}_i}{\sum_{i=1}^k a_i \widehat{c}_i}$$

for

$$\begin{aligned} a_i &= \mathbf{s}_i(1 + 2\alpha |y_i|), \quad c_i = \mathbf{s}_i + 2x_i(-1 + \alpha) + 2y_i, \quad \widehat{c}_i = \mathbf{s}_i + 2\widehat{x}_i(-1 + \alpha) + 2y_i, \\ d_i &= \mathbf{s}_i + 2\alpha x_i, \quad \widehat{d}_i = \mathbf{s}_i + 2\alpha \widehat{x}_i, \end{aligned}$$

where $\mathbf{s}_i = \text{sgn}(y_i)$. Denoting $\mathbf{err} = 2(\widehat{\mathbf{x}} - \mathbf{x})$, we write

$$\widehat{d}_i - d_i = 2\alpha(\widehat{x}_i - x_i) = \alpha \text{err}_i, \text{ and } \widehat{c}_i - c_i = 2(\alpha - 1)(\widehat{x}_i - x_i) = (\alpha - 1)\text{err}_i.$$

We now have

$$\begin{aligned} \widehat{t}^{*,k} - \widehat{t}^{*,k} &= \frac{\sum_{i=1}^k a_i d_i \sum_{i=1}^k a_i \widehat{c}_i - \sum_{i=1}^k a_i \widehat{d}_i \sum_{i=1}^k a_i c_i}{\sum_{i=1}^k a_i c_i \sum_{i=1}^k a_i \widehat{c}_i} \\ &= \frac{\sum_{i=1}^k a_i^2 (d_i \widehat{c}_i - \widehat{d}_i c_i) - \sum_{1 \leq i < j \leq k} a_i a_j (d_i \widehat{c}_j + d_j \widehat{c}_i - \widehat{d}_i c_j - \widehat{d}_j c_i)}{\sum_{i=1}^k a_i c_i \sum_{i=1}^k a_i \widehat{c}_i}. \end{aligned}$$

Writing down each of the terms in the numerator we get

$$d_i \widehat{c}_i - \widehat{d}_i c_i = -2\alpha \text{err}_i (y_i - x_i) - \text{err}_i d_i = -\text{err}_i a_i,$$

where we use the fact that $c_i = d_i + 2(y_i - x_i)$. We also get

$$d_i \widehat{c}_j + d_j \widehat{c}_i - \widehat{d}_i c_j - \widehat{d}_j c_i = \text{err}_j ((\alpha - 1)d_i - \alpha c_i) + \text{err}_i ((\alpha - 1)d_j - \alpha c_j) = -(\text{err}_j a_i + \text{err}_i a_j).$$

Thus,

$$\begin{aligned} \sum_{i=1}^k a_i d_i \sum_{i=1}^k a_i \widehat{c}_i - \sum_{i=1}^k a_i \widehat{d}_i \sum_{i=1}^k a_i c_i &= -\left(\sum_{i=1}^k a_i^3 \text{err}_i + \sum_{1 \leq i < j \leq k} a_i a_j (a_i \text{err}_j + a_j \text{err}_i) \right) \\ &= -\sum_{i=1}^k \text{err}_i \left(a_i^3 + a_i \sum_{i \neq j} a_j^2 \right) = -\|\mathbf{a}_{1:k}\|_2^2 \sum_{i=1}^k a_i \text{err}_i \end{aligned}$$

Turning our attention to the denominator we have

$$\begin{aligned} \sum_{i=1}^k a_i c_i \sum_{i=1}^k a_i \widehat{c}_i &= \sum_{i=1}^k a_i^2 c_i \widehat{c}_i + \sum_{1 \leq i < j \leq k} a_i a_j (c_i \widehat{c}_j + \widehat{c}_i c_j) \\ &= \left(\sum_{i=1}^k a_i c_i \right)^2 + (\alpha - 1) \left(\sum_{i=1}^k a_i c_i \right) \sum_{j=1}^k a_j \text{err}_j, \end{aligned}$$

due to

$$c_i \widehat{c}_j + \widehat{c}_i c_j = 2c_i c_j + (\alpha - 1)(\text{err}_j c_i + \text{err}_i c_j), \text{ and } c_i \widehat{c}_i = c_i^2 + (\alpha - 1)c_i \text{err}_i.$$

Putting it all together and rewriting we have

$$\widehat{t}^{*,k} - \widehat{t}^{*,k} = \frac{-\|\mathbf{a}_{1:k}\|_2^2 \langle \mathbf{a}_{1:k}, \mathbf{err}_{1:k} \rangle}{\left(\|\mathbf{a}_{1:k}\|_2^2 - 2(\alpha - 1) \langle \mathbf{a}_{1:k}, \boldsymbol{\eta}_{1:k} \rangle \right) \left(\|\mathbf{a}_{1:k}\|_2^2 - 2(\alpha - 1) \langle \mathbf{a}_{1:k}, (\mathbf{y} - \widehat{\mathbf{x}})_{1:k} \rangle \right)}$$

and recall $\boldsymbol{\eta}_{1:k} = (\mathbf{y} - \mathbf{x})_{1:k}$. Taking now $k = m$ we have by Cauchy-Schwartz inequality

$$|t_{\text{opt}} - \widehat{t}_{\text{opt}}| \leq \frac{\|\mathbf{a}\|_2^3}{\left| \left(\|\mathbf{a}\|_2^2 - 2(\alpha - 1) \langle \mathbf{a}, \boldsymbol{\eta} \rangle \right) \left(\|\mathbf{a}\|_2^2 - 2(\alpha - 1) \langle \mathbf{a}, (\mathbf{y} - \widehat{\mathbf{x}}) \rangle \right) \right|} \|\mathbf{x} - \widehat{\mathbf{x}}\|_2.$$

The term $\|\mathbf{x} - \widehat{\mathbf{x}}\|_2$ can be bounded as in Section 4. To bound the first factor we compute

$$\left| \|\mathbf{a}\|_2^2 - 2(\alpha - 1)\langle \mathbf{a}, \boldsymbol{\eta} \rangle \right| = \|\mathbf{a}\|_2^2 \left| 1 - 2(\alpha - 1) \frac{\langle \frac{\mathbf{a}}{\|\mathbf{a}\|_2}, \boldsymbol{\eta} \rangle}{\|\mathbf{a}\|_2} \right|.$$

Provided[§] $\left| 2(\alpha - 1) \frac{\langle \frac{\mathbf{a}}{\|\mathbf{a}\|_2}, \boldsymbol{\eta} \rangle}{\|\mathbf{a}\|_2} \right| \leq \frac{\sqrt{2}}{2} \|\mathbf{a}\|_2$ we have

$$\left| 1 - 2(\alpha - 1) \frac{\langle \frac{\mathbf{a}}{\|\mathbf{a}\|_2}, \boldsymbol{\eta} \rangle}{\|\mathbf{a}\|_2} \right|^{-1} \leq 2 \left| 1 + 2(\alpha - 1) \frac{\langle \frac{\mathbf{a}}{\|\mathbf{a}\|_2}, \boldsymbol{\eta} \rangle}{\|\mathbf{a}\|_2} \right| \leq 2 \left(1 + 2|\alpha - 1| \frac{\|\boldsymbol{\eta}\|_2}{\|\mathbf{a}\|_2} \right),$$

and

$$\left| 1 - 2(\alpha - 1) \frac{\langle \frac{\mathbf{a}}{\|\mathbf{a}\|_2}, \mathbf{y} - \widehat{\mathbf{x}} \rangle}{\|\mathbf{a}\|_2} \right|^{-1} \leq 2 \left(1 + 2|\alpha - 1| \frac{\|\mathbf{y} - \widehat{\mathbf{x}}\|_2}{\|\mathbf{a}\|_2} \right) \leq 2 \left(1 + 2|\alpha - 1| \frac{\|\mathbf{y}\|_2}{\|\mathbf{a}\|_2} \right).$$

where in the last line we used $\mathbf{y} - \widehat{\mathbf{x}} = (\text{Id} - \widehat{\Pi})\mathbf{y}$ and the fact $\|\text{Id} - \text{P}\|_2 = \|\text{P}\|_2$ for non-trivial (neither null nor identity) orthogonal projections P. For $\alpha \geq 1$ we have

$$2|\alpha - 1| \frac{\|\mathbf{y}\|_2}{\|\mathbf{a}\|_2} \leq 1, \text{ and } 2|\alpha - 1| \frac{\|\boldsymbol{\eta}\|_2}{\|\mathbf{a}\|_2} \leq 1,$$

using $\|\mathbf{a}\|_2 \geq 2|\alpha| \|\mathbf{y}\|_2$, and the signal-to-noise gap in the last inequality. On the other hand, for $0 < \alpha < 1$ we have $\|\mathbf{y}\| = \|\mathbf{x} + \boldsymbol{\eta}\| \lesssim \sqrt{h} + \sigma \sqrt{m}$, with high probability, giving

$$\frac{\|\mathbf{a}\|_2^3}{\left| \left(\|\mathbf{a}\|_2^2 - 2(\alpha - 1)\langle \mathbf{a}, \boldsymbol{\eta} \rangle \right) \left(\|\mathbf{a}\|_2^2 - 2(\alpha - 1)\langle \mathbf{a}, \mathbf{y} - \widehat{\mathbf{x}} \rangle \right) \right|} \lesssim \frac{1}{\sqrt{m}}.$$

In conclusion, for $\alpha > 0$ we have

$$|t_{\text{opt}} - \widehat{t}_{\text{opt}}| \lesssim \|\Pi - \widehat{\Pi}\|_2 + \sigma \sqrt{\frac{h}{m}},$$

as desired.



AIMS Press

©2022 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>)

[§]This holds for example if $(\alpha - 1)^2 \sigma^2 m \lesssim m + 4\alpha \|\mathbf{y}\|_1 + 4\alpha^2 \|\mathbf{y}\|_2^2$