
The Well-Tempered Lasso

Yuanzhi Li¹ Yoram Singer^{1 2}

Abstract

We study the complexity of the entire regularization path for least squares regression with 1-norm penalty, known as the Lasso. Every regression parameter in the Lasso changes linearly as a function of the regularization value. The number of changes is regarded as the Lasso’s complexity. Experimental results using exact path following exhibit polynomial complexity of the Lasso in the problem size. Alas, the path complexity of the Lasso on artificially designed regression problems is *exponential*. We use *smoothed analysis* as a mechanism for bridging the gap between worst case settings and the de facto low complexity. Our analysis assumes that the observed data has a *tiny* amount of intrinsic noise. We then prove that the Lasso’s complexity is polynomial in the problem size.

1. Introduction

In high dimensional learning problems, a sparse solution is often desired as it has better generalization and interpretation. In order to promote sparse solutions, a regularization term that penalizes for the 1-norm of the vector of parameters is often augmented to an empirical loss term. In regression problems, the empirical loss amounts to sum of the squared differences between the linear predictions and true targets. The task of linear regression with 1-norm penalty is known as Lasso (Tibshirani, 1996). To obtain meaningful solutions for the Lasso, it is required to pick a good value of the ℓ_1 regularizer. To automatically choose the best regularization value, algorithms that calculate all possible solutions were developed (Efron et al., 2004; Osborne et al., 2000; Tibshirani & Taylor, 2012).

These algorithms find the solution set for all possible regularization values, commonly referred to as the entire reg-

ularization path. The algorithms typically built upon the property that the Lasso regularization path is piecewise linear in the constituents of solution vector. As a result, their running times are also governed by the total number of linear segments. While experiments with real datasets suggest that the number of linear segments is in practice linear in the dimension of the problem (Rosset & Zhu, 2007), worst case settings (Mairal & Yu, 2012) can yield to exponentially many linear segments. The construction of exponentially complex regression problems of (Mairal & Yu, 2012) stands in stark contrast to the aforementioned methods. Provably polynomial path complexity has so far derived in vastly more restricted settings, such as the one described in (Dubiner & Singer, 2011).

We bridge the gap between the de facto complexity of the regularization path in real problems and the worst case analysis of the number of linear segments. We show that under fairly general models, the complexity of the entire regularization path is guaranteed to be polynomial in the dimension of the problem. As an important observation, settings which attain the worst case complexity of the regularization path often exhibit fragile algebraic structure. In contrast, natural datasets often comes with noise, which renders those highly frail structures improbable. This approach is called the *smoothed analysis*, introduced by the seminal paper of (Spielman & Teng, 2009).

The core of smoothed analysis is the assumption that the input data is subjected to a small *intrinsic* noise. Such noise may come from uncertainty in the physical measurements when collecting the data, irrational decisions in human feedback, or simply the rounding errors in the computation process used for obtaining the data. In this model, we let $\mathbf{X} \in \mathbb{R}^{n \times d}$ denote the data matrix where n is the number of observations and d is the dimension (number of free parameters). Smoothed analysis assumptions implies that \mathbf{X} is the sum of \mathbf{X}^h , an unknown fixed matrix, and \mathbf{G} , which consists of i.i.d. random samples from a normal distribution with a zero mean and low variance, $\mathbf{X} = \mathbf{X}^h + \mathbf{G}$. In this view, the data is neither *completely random* nor *completely arbitrary*. The *smoothed complexity* of the problem is measured as the expected complexity taken over the random choices of \mathbf{G} . Using this framework, it was proved that the smoothed running time of the simplex, k -means, and the Perceptron algorithm (Spielman & Teng, 2001; Arthur et al.,

¹Department of Computer Science, Princeton University ²Googol Brain. Correspondence to: Yuanzhi Li <yuanzhil@cs.princeton.edu>.

2009; Blum & Dunagan, 2002) is in fact polynomial while the worst case complexity of these problems is exponential.

We use the above smoothed analysis model to show that on “typical” instances, the total number of linear segments of the Lasso’s *exact* regularization path is polynomial in the problem size with high probability. Informally speaking and omitting technical details, our main result can be stated as follows.

Let $\mathbf{X} \in \mathbb{R}^{n \times d}$ be a data matrix of the form $\mathbf{X} = \mathbf{X}^h + \mathbf{G}$ for any fixed matrix \mathbf{X}^h and a random Gaussian matrix \mathbf{G} , $G_{ij} \sim \mathcal{N}(0, \sigma^2)$. Then, for an arbitrary vector of targets $y \in \mathbb{R}^n$, with high probability, the total number of linear segments of the Lasso’s *exact* regularization path for (\mathbf{X}, y) is *polynomial* in n , d , and $\frac{1}{\sigma}$.

Our result is conceptually different than the one presented in (Mairal & Yu, 2012). Mairal and Yu showed that there exists an *approximate* regularization path with a small number of linear segments. However, the analysis, while being novel and inspiring, does not shed light on why, in practice, the *exact* number of linear segments is small as the *approximated* path is unlikely to coincide with the *exact* path. Our analysis covers uncharted terrain and different aspects than the approximated path algorithms in (Mairal & Yu, 2012; Giesen et al., 2010). On one hand, we show that when the input data is “typical”, namely comes from a “naturally smooth” distribution, then with high probability, the total number of linear segments, of the *exact* Lasso path, is already polynomially small. This part of our analysis provides theoretical backing to the empirical findings reported in (Rosset & Zhu, 2007). On the other hand, when the input matrix is atypical and induces a high-complexity path, then we can also obtain a low-complexity approximate regularization path by adding a small amount of random noise to the data and then solve the Lasso’s regularization path on the perturbed instance *exactly*. We also verify our analysis experimentally in section 9. We show that even a *tiny* amount of perturbation to high-complexity data matrices, results in a dramatic drop in the number of linear segments.

The technique used in this paper is morally different from the smoothed analysis obtained for simplex (Spielman & Teng, 2001), k-means (Arthur et al., 2009), and the Perceptron (Blum & Dunagan, 2002), as there is no concrete algorithm involved. We develop a new framework which shows that when the total number of linear segments is excessively high, then we can *tightly* couple the solutions of the original, smoothed problems to another set of solutions in a manner that does not depend on \mathbf{G} . We then use the randomness of \mathbf{G} to show that such couplings are unlikely to exist, thus high complexity solutions are rare. We believe that our framework can be extended to other problems such as the regularization path of support vector machines.

2. Preliminaries

We use uppercase boldface letters, e.g. \mathbf{X} , to denote matrices and lowercase letters x, w to denote vectors, variables, and scalars. We use \mathbf{X}_i to denote the i ’th column of \mathbf{X} . Given a set \mathcal{S} , we denote by $\mathbf{X}_{\mathcal{S}} \in \mathbb{R}^{d \times |\mathcal{S}|}$ the sub-matrix of \mathbf{X} whose columns are \mathbf{X}_i for $i \in \mathcal{S}$. Analogously, $\mathbf{X}_{\bar{\mathcal{S}}}$ denotes the sub-matrix of \mathbf{X} with columns \mathbf{X}_i for $i \notin \mathcal{S}$. For a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ with $n \geq d$, we use the term *smallest (largest)* singular value of \mathbf{X} to denote the smallest (largest) right singular value of \mathbf{X} . We define the *generalized sign* of a scalar b as follows,

$$\text{sign}(b) = \begin{cases} +1 & b > 0 \\ -1 & b < 0 \\ 0 & b = 0 \end{cases}.$$

Let y be a vector in \mathbb{R}^n and let $\mathbf{X} = [\mathbf{X}_1, \dots, \mathbf{X}_d]$ be a matrix in $\mathbb{R}^{n \times d}$. The Lasso is the following regression problem,

$$w[\lambda] = \arg \min_{w \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{X}w - y\|_2^2 + \lambda \|w\|_1. \quad (1)$$

Here, $\lambda > 0$ is the regularization value. The value of λ influences the sparsity level of the solution $w[\lambda]$. The larger λ is the sparser the solution is. When \mathbf{X} is of full column rank, the solution to (1) is unique. We therefore denote it by $w[\lambda]$. We use

$$\mathcal{P} = \{w[\lambda] \mid \lambda > 0\}$$

to denote the set of all possible solution vectors. This set is also referred to as the *entire regularization path*.

To establish our main result we need a few technical lemmas. The first Lemma from (Mairal & Yu, 2012) provides optimality conditions for $w[\lambda]$.

Lemma 1. *Let $\lambda > 0$, the $w[\lambda]$ is the optimal solution iff it satisfies the following conditions,*

1. *There exists a vector $u[\lambda]$ s.t.*

$$\mathbf{X}^\top (\mathbf{X}w[\lambda] - y) = u[\lambda].$$

2. *Each coordinate of $u[\lambda]$ satisfies,*

$$u_i[\lambda] = \begin{cases} -\lambda \text{sign}(w_i[\lambda]) & |w_i[\lambda]| > 0 \\ \in [-\lambda, \lambda] & o.w. \end{cases}.$$

Let us denote the sign vector as $\text{sign}(w[\lambda])$, which is obtained by applying the generalized sign function $\text{sign}(\cdot)$ element-wise to $w[\lambda]$. The result of (Mairal & Yu, 2012) shows that \mathcal{P} is piecewise linear and unique in the following sense.

Lemma 2. *Suppose \mathbf{X} is of full column rank, then $\mathcal{P} = \{w[\lambda] \mid \lambda > 0\}$ is unique, well-defined, and $w[\lambda]$ is piecewise linear. Moreover, for any $\lambda_1, \lambda_2 > 0$, if the sign vectors at λ_1 and λ_2 are equal, $\text{sign}(w[\lambda_1]) = \text{sign}(w[\lambda_2])$, then $w[\lambda_1]$ and $w[\lambda_2]$ are in the same linear segment.*

We use $|\mathcal{P}|$ to denote the total number of linear segments in \mathcal{P} . We denote by $\alpha > 0$ the smallest singular value of \mathbf{X} . Without loss of generality, as we can rescale \mathbf{X} and y accordingly, we assume that $\|y\|_2 = 1$. To obtain our main result, we introduce the following smoothness assumption on the data matrix \mathbf{X} .

Assumption 3 (Smoothness). \mathbf{X} is generated according to,

$$\mathbf{X} = \mathbf{X}^h + \mathbf{G},$$

where $\mathbf{X}^h \in \mathbb{R}^{n \times d}$ ($n \geq d$) is a fixed unknown matrix with $\|\mathbf{X}^h\|_2 \leq 1$. Each entry of \mathbf{G} is an i.i.d. sample from the normal distribution with 0 mean and variance of $\frac{\sigma^2}{n}$.

We use $\mathcal{N}(0, \sigma^2/n)$ instead of $\mathcal{N}(0, \sigma^2)$ for the noise distribution \mathbf{G} . This choice implies that when σ is a constant the spectral norm of \mathbf{G} is also a constant in expectation, $\mathbb{E}[\|\mathbf{G}\|_2] = O(1)$, see for instance (Rudelson & Vershynin, 2010). Therefore, the signal-to-noise ratio satisfies,

$$\mathbb{E}[\|\mathbf{X}\|_2 \|\mathbf{G}\|_2^{-1}] = O(1).$$

It is convenient to think of σ as an arbitrary small constant. The analysis presented in the sequel employs a fixed constant c that does not depend on the problem size. We use $f(\cdot) \stackrel{c}{>} \text{poly}(\cdot)$ (analogously, $f(\cdot) \stackrel{c}{<} \text{poly}(\cdot)$) to denote the fact that the function f is everywhere greater (smaller) than a polynomial function up to a multiplicative constant. Equipped with the above conventions and the smoothness assumption, the following lemma, due to (Sankar et al., 2006), characterizes the extremal singular values of \mathbf{X} .

Lemma 4. Let $\delta > 0$. With probability of at least $1 - \delta$, the smallest, denoted α , and largest, denoted β , right singular values of \mathbf{X} satisfy,

$$\alpha \stackrel{c}{>} \frac{\delta\sigma}{d} \text{ and } \beta \stackrel{c}{<} 1 + \sigma \log(1/\delta).$$

The bound on β lets assume henceforth that β is $O(1)$ for any reasonable choices of σ and δ . In our analysis We describe explicitly dependencies on $\|\mathbf{X}\|$ for clarification and states the main results with $\beta = O(1)$.

The main result of the paper is states in the following theorem.

Theorem 5 (Lasso's Smoothed Complexity). Suppose assumption 3 holds for arbitrary $n, d \in \mathbb{Z}$ with $n \geq d$ and $\sigma \in (0, 1]$. Then, with a probability of at least $1 - \delta$ (over the random selection of \mathbf{G}), the complexity of the Lasso satisfies,

$$|\mathcal{P}| \stackrel{c}{<} n^{1.1} \left(\frac{d}{\delta\sigma} \right)^6.$$

3. Main Lemmas

To prove the main theorem, we introduce several properties of $\mathbf{X}, y, w[\lambda]$ and $u[\lambda]$ that are critical in the analysis of $|\mathcal{P}|$. We then use the smoothness assumption to bound these properties.

Definition 1 (Lipschitzness). Let $w_i[\lambda]$ and $u_i[\lambda]$ be the value of the i 'th coordinate of $w[\lambda]$ and $u[\lambda]$ respectively for $i \in [d]$. The coordinate-wise Lipschitz parameters of w and u are defined as,

$$\mathcal{L}_w = \max_{i \in [d]} \sup_{\lambda > 0} \left| \partial_\lambda w_i[\lambda] \right|, \quad \mathcal{L}_u = \max_{i \in [d]} \sup_{\lambda > 0} \left| \partial_\lambda u_i[\lambda] \right|.$$

By definition, \mathcal{L}_w and \mathcal{L}_u characterize how much each coordinate of $w[\lambda]$ and $u[\lambda]$ can change as we vary the value of λ . We later use the smoothness assumption to show that \mathcal{L}_w and \mathcal{L}_u are polynomially small. This implies that $w[\lambda]$ and $u[\lambda]$ would not change too fast with λ . However, Lipschitzness by itself does not give us a bound on $|\mathcal{P}|$ since $w[\lambda]$ can still oscillate around zero and induce an excessively large number of linear segments. Therefore, we also need the following property which defines the restricted distance between the column space of \mathbf{X} and y .

Definition 2 (Subspace distance). For any $s, \delta > 0$, let γ_s denote the largest value such that,

$$\Pr \left[\exists v \in \mathbb{R}^{d-s} \text{ s.t. } \|\mathbf{X}_{\bar{S}} v - y\|_2 \leq \gamma_s \right] \leq \delta,$$

for all $S \subset [d]$ of size s .

This definition quantifies the distance of y to a subspace spanned by $s \leq d$ columns of \mathbf{X} . Since $y \in \mathbb{R}^n$, $n \geq d$, and \mathbf{X} is smooth, it can be shown that y cannot be too close to the subspace spanned by $\mathbf{X}_{\bar{S}}$. That is, γ_s is inversely proportional to a polynomial in $n, d, 1/\sigma$. We interchangeably use in the following the original matrix \mathbf{X} with $v \in \mathbb{R}^d$ s.t. $v_i = 0$ for $i \in S$ and $\mathbf{X}_{\bar{S}}$ with $v \in \mathbb{R}^{d-s}$. Using the above properties, we prove the following theorem.

Theorem 6 (Exact Smooth Complexity). Let \mathbf{X} satisfy Assumption 3. Then, for all $s \in [d]$ and $\delta > 0$, with probability of at least $1 - \delta$ the complexity of the Lasso satisfies,

$$|\mathcal{P}| \stackrel{c}{<} 3^s \left(\frac{\sqrt{snd} \left(\frac{\mathcal{L}_w}{\alpha^2} + \mathcal{L}_u \right)}{\delta^2 \sigma \gamma_s} \right)^{\frac{s}{s-1}}.$$

The following lemma characterizes the (smoothed) values of $\mathcal{L}_w, \mathcal{L}_u$, and γ_s .

Lemma 7. Let \mathbf{X} satisfy Assumption 3. Then with probability of at least $1 - \delta$ the following properties hold,

$$\mathcal{L}_w, \mathcal{L}_u \stackrel{c}{<} \frac{\sqrt{d}}{\alpha^2}, \quad \gamma_s \stackrel{c}{>} \frac{\sigma}{\sqrt{dn}(d/\delta)^{2/s}}.$$

Applying the bounds on \mathcal{L}_w , \mathcal{L}_u , γ_s , and α to Theorem 6 while letting s be a sufficiently large constant, we can directly prove Theorem 5. In Section 5, we use the value of α to bound \mathcal{L}_w and \mathcal{L}_u . In Section 6, we employ the smoothness of \mathbf{X} to bound γ_s . Finally, in Section 7 we prove Theorem 6.

4. Proof sketch

Since $\|\mathbf{X}\|_2 = O(1)$, there exists a constant $\lambda_{\max} = \Omega(1)$ such that for $\lambda \geq \lambda_{\max}$, $w[\lambda]$ is the zero vector. Thus, we can divide $\lambda \in [0, \lambda_{\max}]$ into λ_{\max}/ν intervals, each of size ν for some (inversely polynomial) small ν . We then show that within every interval the total number of linear segments exceeds a fixed polynomial number with exponentially small probability. The total number of linear segments follows by taking a union bound over all intervals. We need to specifically address the following two questions in the analysis.

What if $|\mathcal{P}|$ within an intervals is excessively large?

We will show that when there are N linear segments in an interval, then there must be at least $\log_3(N)$ many coordinates of $w[\lambda]$, which we denote as the set \mathcal{S} , that change their sign. Since ν is small and $w[\lambda]$ is a Lipschitz function in λ , we know that those coordinates of $w[\lambda]$ must be close to zero. Therefore, we can show that $w[\lambda]$ is close to the optimal solution, $v[\lambda]$, of the Lasso problem when the entries of coordinates in \mathcal{S} are constrained to be *exactly* zero,

$$v[\lambda] = \arg \min_{w \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{X}w - y\|_2^2 + \lambda \|w\|_1 \text{ s.t. } \forall i \in \mathcal{S} : w_i = 0 \quad (2)$$

What if $w[\lambda]$ oscillates excessively around $v[\lambda]$?

From the optimality condition of $u[\lambda]$ and the smoothness of $u[\lambda]$, we also know that the coordinates in \mathcal{S} of $u[\lambda]$ must be close to either $-\lambda$ or λ . Thus, $u_{\mathcal{S}}[\lambda] \stackrel{\text{def}}{=} \mathbf{X}_{\mathcal{S}}^{\top}(\mathbf{X}w[\lambda] - y)$ is close to a vector on the scaled hypercube $\{-\lambda, \lambda\}^{|\mathcal{S}|}$. On the other hand, if $w[\lambda]$ is close to $v[\lambda]$, we know that $\mathbf{X}_{\mathcal{S}}^{\top}(\mathbf{X}v[\lambda] - y)$ is also close to a vector in $\{-\lambda, \lambda\}^{|\mathcal{S}|}$. However, $v[\lambda]$ does not depend on $\mathbf{X}_{\mathcal{S}}$ by construction. Therefore, the residual $\mathbf{X}v[\lambda] - y$ also does not depend on $\mathbf{X}_{\mathcal{S}}$. Thus, using the randomness etched in $\mathbf{X}_{\mathcal{S}}$ we can now show that $\mathbf{X}_{\mathcal{S}}^{\top}(\mathbf{X}v[\lambda] - y)$ is close to a vector in $\{-\lambda, \lambda\}^{|\mathcal{S}|}$ with probability which is exponentially small in the size of \mathcal{S} . Therefore, we know that w.h.p. the total number of linear segments in this interval is unlikely to be large.

5. Bounding \mathcal{L}_w and \mathcal{L}_u

Recall that we denote the smallest singular value of \mathbf{X} by α . We first show the following lemma regarding pertur-

bations of strongly convex functions. Also recall that a second-order smooth function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is α^2 -strongly convex if $\nabla^2 f(x) \geq \alpha^2$ for all $x \in \mathbb{R}^d$ and is \mathcal{L} -Lipschitz if $\|\nabla f(x)\|_2 \leq \mathcal{L}$ for all $x \in \mathbb{R}^d$.

Lemma 8 (Perturbation of strongly convex functions I). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a non-negative, α^2 -strongly convex function. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be a \mathcal{L} -Lipschitz non-negative convex function. For any $\lambda \geq 0$, let $z[\lambda]$ be the minimizer of $f(z) + \lambda g(z)$, then we have,*

$$\left\| \frac{dz[\lambda]}{d\lambda} \right\|_2 \leq \frac{\mathcal{L}}{\alpha^2}.$$

Proof of Lemma 8. For any $\tau \geq 0$ and $\lambda \geq 0$, let us abbreviate $z = z[\lambda]$ and denote $\varepsilon = z[\lambda + \tau] - z$. From α^2 -strong convexity of f at z and the optimality of z at λ , we know that

$$\frac{1}{2} \alpha^2 \|\varepsilon\|_2^2 + f(z) + \lambda g(z) \leq f(z + \varepsilon) + \lambda g(z + \varepsilon). \quad (3)$$

Moreover, using the optimality of $z + \varepsilon$ at $\lambda + \tau$, we know that,

$$\begin{aligned} \frac{1}{2} \alpha^2 \|\varepsilon\|_2^2 + f(z + \varepsilon) + (\lambda + \tau)g(z + \varepsilon) \\ \leq f(z) + (\lambda + \tau)g(z). \end{aligned} \quad (4)$$

Summing Eqs. (3) and (4) and rearranging terms yields,

$$\alpha^2 \|\varepsilon\|_2^2 \leq \tau(g(z) - g(z + \varepsilon)) \leq \tau \|\varepsilon\|_2 \mathcal{L},$$

where the last inequality is due to the Lipschitzness assumption on g . Therefore, we get that $\|\varepsilon\|_2/\tau \leq \mathcal{L}/\alpha^2$. Letting $\tau \rightarrow 0^+$ completes the proof. \square

Using Lemma 8 with $f(w) = \frac{1}{2} \|\mathbf{X}w - y\|_2^2$ and $g(w) = \|w\|_1$, we obtain Lipschitz properties for $w[\lambda]$ and $u[\lambda]$. Since we assume that the minimum singular value of \mathbf{X} is α then $f(w)$ is α^2 -strongly convex. In addition, the norm of $\nabla g(w)$ is clearly at most \sqrt{d} . To simplify notation, when $w[\lambda]$ is not differentiable at a point λ , we define $dw[\lambda]/d\lambda = 0$. Due to Lipschitzness and strong convexity all vectors in the subgradient set $\partial_{\lambda} w[\lambda]$ include this particular choice for a subgradient. In summary we get the following corollary.

Corollary 9 (Lipschitzness of w). *For any $\lambda \geq 0$ it holds that,*

$$\left\| \frac{dw[\lambda]}{d\lambda} \right\|_2 \leq \frac{\sqrt{d}}{\alpha^2}.$$

Since by definition, $u[\lambda] = \mathbf{X}_i^{\top}(\mathbf{X}w[\lambda] - y)$, we obtain a similar corollary for u .

Corollary 10 (Lipschitzness of u). *For any $\lambda \geq 0$ it holds that,*

$$\left\| \frac{du[\lambda]}{d\lambda} \right\|_2 \leq \frac{\|\mathbf{X}\|_2^2 \sqrt{d}}{\alpha^2}.$$

Recall that $\|\mathbf{X}\|_2 = O(1)$, thus, from the above corollaries we get

$$\mathcal{L}_w, \mathcal{L}_u \stackrel{c}{\leq} \sqrt{d}/\alpha^2. \quad (5)$$

We also use in the next section the following Lemma.

Lemma 11 (Perturbation of strongly convex functions II). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be an α^2 -strongly convex function and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ an \mathcal{L} -Lipschitz convex function. Let z_1 and z_2 be the minimizers of $f(z)$ and $f(z) + g(z)$, respectively, then*

$$\|z_1 - z_2\|_2 \leq \frac{\mathcal{L}}{\alpha^2}.$$

Proof of Lemma 11. Let $\varepsilon = z_2 - z_1$. From strong convexity of f and optimality of z_1 , we know that

$$\frac{1}{2}\alpha^2\|\varepsilon\|_2^2 + f(z_1) \leq f(z_2). \quad (6)$$

Due to the optimality of z_2 for $f(z) + g(z)$ we get,

$$\frac{1}{2}\alpha^2\|\varepsilon\|_2^2 + f(z_2) + g(z_2) \leq f(z_1) + g(z_1). \quad (7)$$

Summing Eq. (6) and (7) and rearranging terms gives,

$$\alpha^2\|\varepsilon\|_2^2 \leq g(z_1) - g(z_2) \leq \mathcal{L}\|\varepsilon\|_2.$$

We therefore get that $\|\varepsilon\|_2 \leq \frac{\mathcal{L}}{\alpha^2}$. \square

6. Bounding γ_s

In this section we prove that Assumption 3 also yields a bound on γ_s .

Lemma 12. *Let $y \in \mathbb{S}^{n-1}$ be an arbitrary unit vector. Then for any $s \in \{1, \dots, d\}$ and a set $\mathcal{S} \subset [d]$ of size s , the following holds,*

$$\Pr \left[\exists v \in \mathbb{R}^{d-s} \text{ s.t. } \|\mathbf{X}_{\mathcal{S}}v - y\|_2 \stackrel{c}{\leq} \frac{\sigma\delta^{\frac{1}{s}}}{\sqrt{dn}} \right] \leq \delta.$$

For brevity, we denote the event above by \mathcal{D}_γ where $\gamma = \Omega\left(\frac{\sigma\delta^{\frac{1}{s}}}{\sqrt{dn}}\right)$.

Proof. The proof relies on the following simple fact about Gaussian random variables. For any unit vector u , scalar $\tau > 0$, $g \in \mathbb{R}$, and $i \in [d]$ the following holds,

$$\Pr \left[|\langle u, \mathbf{X}_i \rangle - g| \leq \tau \right] \leq \frac{e\tau\sqrt{n}}{\sigma}, \quad (8)$$

where \mathbf{X}_i is a (column) vector with elements distributed i.i.d according to $\mathcal{N}(0, \sigma^2/n)$.

Let us assume that there exists $v \in \mathbb{R}^{d-s}$ such that $\|\mathbf{X}_{\mathcal{S}}v - y\|_2 \leq \gamma$. From Lemma 4, we know that with probability of at least $1 - \delta$, $\|\mathbf{X}\|_2 \stackrel{c}{\leq} 1$. Since we assume that $\|y\|_2 = 1$, it holds that

$$\begin{aligned} \gamma &\geq \|\mathbf{X}_{\mathcal{S}}v - y\|_2 \geq \|y\|_2 - \|\mathbf{X}v\|_2 \geq 1 - \|v\|_2 \\ &\Rightarrow \|v\|_2 \geq 1 - \gamma. \end{aligned}$$

Since γ is smaller than $1/2$, there must exist a coordinate i for which $|v_i| \geq \frac{1}{2\sqrt{d}}$.

Now, let us fix \mathbf{G}_j to its observed values for all $j \neq i$. In addition, let us denote by \mathbf{U} the subspace spanned by

$$\{(\mathbf{X}^h)_i\} \cup \{\mathbf{X}_j \mid \forall j \neq i, j \in \mathcal{S}\} \cup \{y\}.$$

We know that \mathbf{U} is of dimension $d - s + 1$. Consider an orthonormal basis for $\{u_1, \dots, u_{s-1}\} \in \mathbb{R}^n$ for the subspace $\text{span}(\{\mathbf{X}_i\}) - \mathbf{U}$. Suppose there exists v such that $\|\mathbf{X}_{\mathcal{S}}v - y\|_2 \leq \gamma$. Then, for every $j \in [s-1]$, multiplying by u_j yields

$$|\langle u_j, \mathbf{X}_i \rangle v_i| \leq \gamma \Rightarrow |\langle u_j, \mathbf{X}_i \rangle| \leq \frac{\gamma}{|v_i|} \leq 2\sqrt{d}\gamma.$$

Note that any pair $j \neq j'$, the inner products $\langle u_j, \mathbf{X}_i \rangle$ and $\langle u_{j'}, \mathbf{X}_i \rangle$ are independent. We now use (8) over all u_j and get that \mathcal{D}_γ holds with probability of at most,

$$\left(\frac{6\sqrt{dn}\gamma}{\sigma} \right)^s.$$

Finally, choosing $\gamma = \Omega\left(\frac{\sigma\delta^{\frac{1}{s}}}{\sqrt{dn}}\right)$ completes the proof. \square

7. Coupling the solutions

In this section, we show that when the total number of linear segments in a small interval is excessively large, the optimal solution $w[\lambda]$ can be coupled with the optimal solution $v[\lambda]$ of the constrained Lasso problem of (2).

Sign changes. For a given fixed $\lambda_0 > 0$ and $\nu > 0$, let us denote by $\zeta(i)$ the number of times that the generalized sign of w_i changes as λ increases from λ_0 to $\lambda_0 + \nu$. Thus, the total number of linear segments in the interval $[\lambda_0, \lambda_0 + \nu]$ is at least $\sum_{i=1}^d \zeta(i)$. We prove the following lemmas related to the sign changes.

Lemma 13 (Number of sign changes). *For any integer $N > 0$, any $\lambda_0, \nu > 0$, if $\sum_{i=1}^d \zeta(i) \geq N$, then there exists at least $\log_3(N)$ many indices $j \in [d]$ such that $\zeta(j) \geq 1$.*

Proof. According to Lemma 2, each linear segment is associated with a *unique* sign pattern in $\{-1, 0, 1\}^d$. Since there are N segments, the pigeon hole principle implies that there must exist at least $\log_3(N)$ many coordinates of $w[\lambda]$ that change their sign in this interval. \square

From the Lipschitzness of $w[\lambda]$ and $u[\lambda]$, we also obtain the following lemma.

Lemma 14 (Sign change \Rightarrow small weight). *For $i \in [d]$, if $\zeta(i) \geq 1$, then following properties hold: $|w_i[\lambda_0]| \leq \mathcal{L}_w \nu$ and $\|u_i[\lambda_0] - \lambda_0\| \leq (\mathcal{L}_u + 1)\nu$.*

Proof. Since $\zeta(i) \geq 1$, we know that there exists $\lambda \in [\lambda_0, \lambda_0 + \nu]$ such that $w_i[\lambda] = 0$ and $|u_i[\lambda]| = \lambda$. Using Lipschitzness of $w[\lambda]$ we get $|w_i[\lambda_0] - w_i[\lambda]| = |w_i[\lambda_0]| \leq \mathcal{L}_w \nu$. For $u_i[\lambda_0]$ we have,

$$\begin{aligned} \|u_i[\lambda_0] - \lambda_0\| &\leq \|u_i[\lambda_0] - |u_i[\lambda]|\| + \||u_i[\lambda]| - \lambda\| \\ &\quad + |\lambda - \lambda_0| \leq (\mathcal{L}_u + 1)\nu \end{aligned}$$

which concludes the proof. \square

We use $\tilde{\mathcal{L}}_u$ in the sequel as a shorthand for $\mathcal{L}_u + 1$. Based on the two lemmas above we readily get the following corollary.

Corollary 15. *For any integer $N > 0$ and $\nu, \lambda_0 \geq 0$, if $\sum_i \zeta(i) \geq N$, then there exists a subset $\mathcal{S} \subseteq [d]$ of cardinality at least $\log_3(N)$ such that $\forall i \in \mathcal{S}$,*

$$|w_i[\lambda_0]| \leq \mathcal{L}_w \nu \text{ and } \|u_i[\lambda_0] - \lambda_0\| \leq \tilde{\mathcal{L}}_u \nu.$$

Rare events. Let \mathcal{S} be defined as in Corollary 15, we next show that if the size of \mathcal{S} is too large, then certain rare couplings would take place. Thus, the size of \mathcal{S} is likely to be small with high probability. Throughout the rest of the paper we *overload* notation and denote by $\mathbf{X}_{\mathcal{S}} \in \mathbb{R}^{n \times d}$ the matrix where each column \mathbf{X}_i , for $i \notin \mathcal{S}$, is replaced with the zero vector. The matrix $\mathbf{X}_{\bar{\mathcal{S}}}$ is defined analogously. Note that by definition, $\mathbf{X}_{\mathcal{S}} + \mathbf{X}_{\bar{\mathcal{S}}} = \mathbf{X}$.

Lemma 16 (Size of \mathcal{S}). *For any fixed $\lambda_0, \nu > 0$, and a set $\mathcal{S} \subseteq [d]$, let $\mathbf{X}_{\bar{\mathcal{S}}} \in \mathbb{R}^{n \times d}$ be defined as above. Let $v_{\bar{\mathcal{S}}}[\lambda_0]$ be the minimizer,*

$$v_{\bar{\mathcal{S}}}[\lambda_0] = \arg \min_{v \in \mathbb{R}^d} \frac{1}{2} \|\mathbf{X}_{\bar{\mathcal{S}}} v - y\|_2^2 + \lambda_0 \|v\|_1$$

Assume that the properties of Corollary 15 hold for a set \mathcal{S} . Then, for every $j \in \mathcal{S}$ the following inequality holds,

$$\begin{aligned} &\left| \mathbf{X}_j^\top (\mathbf{X}_{\bar{\mathcal{S}}} v_{\bar{\mathcal{S}}}[\lambda_0] - y) \right| - \lambda_0 \\ &\leq 2\sqrt{|\mathcal{S}|} \|\mathbf{X}\|_2^2 \left(\frac{\mathcal{L}_w \|\mathbf{X}\|_2^2 \nu}{\alpha^2} + \tilde{\mathcal{L}}_u \nu \right). \end{aligned}$$

We refer to this event as $\mathcal{E}_{\mathcal{S}}^{(\tau)}$ with parameter $\tau = 2\sqrt{|\mathcal{S}|} \|\mathbf{X}\|_2^2 \left(\frac{\mathcal{L}_w \|\mathbf{X}\|_2^2 \nu}{\alpha^2} + \tilde{\mathcal{L}}_u \nu \right)$.

Proof. We know that the i 'th coordinate of $v_{\bar{\mathcal{S}}}[\lambda_0]$ is zero for all $i \in \mathcal{S}$. Therefore, we need to focus solely on the set of vectors v which are in

$$\mathcal{K} \stackrel{\text{def}}{=} \{v \in \mathbb{R}^d \mid \forall i \in \mathcal{S}, v_i = 0\}.$$

Since by definition $\mathbf{X}_{\mathcal{S}} = \mathbf{X} - \mathbf{X}_{\bar{\mathcal{S}}}$ we can rewrite the original objective as,

$$\frac{1}{2} \|\mathbf{X}_{\bar{\mathcal{S}}} w + \mathbf{X}_{\mathcal{S}} w - y\|_2^2 + \lambda_0 \|w\|_1.$$

Let $w_{\bar{\mathcal{S}}}[\lambda_0] \in \mathbb{R}^d$ be a vector whose i 'th coordinate is the i 'th coordinate of $w[\lambda_0]$ for $i \notin \mathcal{S}$ and is zero otherwise and let $w_{\mathcal{S}}[\lambda_0] = w - w_{\bar{\mathcal{S}}}[\lambda_0]$. From the optimality of $w[\lambda_0]$, we know that $w_{\bar{\mathcal{S}}}[\lambda_0]$ is the minimizer of

$$h(w) = \frac{1}{2} \|\mathbf{X}_{\bar{\mathcal{S}}} w + \mathbf{X}_{\mathcal{S}} w_{\mathcal{S}}[\lambda_0] - y\|_2^2 + \lambda_0 \|w\|_1 \quad \text{s.t. } w \in \mathcal{K}.$$

Let $g(w) \stackrel{\text{def}}{=} \frac{1}{2} \|\mathbf{X}_{\bar{\mathcal{S}}} w - y\|_2^2 + \lambda_0 \|w\|_1$. Expanding terms we get that for every $w \in \mathcal{K}$,

$$\begin{aligned} h(w) - g(w) &= \frac{1}{2} \|\mathbf{X}_{\mathcal{S}} w_{\mathcal{S}}[\lambda_0]\|_2^2 - \langle \mathbf{X}_{\mathcal{S}} w_{\mathcal{S}}[\lambda_0], y \rangle \\ &\quad + \langle \mathbf{X}_{\mathcal{S}} w_{\mathcal{S}}[\lambda_0], \mathbf{X}_{\bar{\mathcal{S}}} w \rangle, \end{aligned}$$

and the gradient of $h(w) - g(w)$ satisfies $\|\nabla(h(w) - g(w))\|_2 \leq \left\| \mathbf{X}_{\bar{\mathcal{S}}}^\top \mathbf{X}_{\mathcal{S}} w_{\mathcal{S}}[\lambda_0] \right\|_2 \leq \|\mathbf{X}\|_2^2 \|w_{\mathcal{S}}[\lambda_0]\|_2$. From our assumption that Corollary 15 holds for \mathcal{S} , we know that $\|w_{\mathcal{S}}[\lambda_0]\|_\infty \leq \mathcal{L}_w \nu$ which in turn implies that $\|w_{\mathcal{S}}[\lambda_0]\|_2 \leq \sqrt{|\mathcal{S}|} \mathcal{L}_w \nu$.

We can now apply Lemma 11 w.r.t $g(w)$ and $h(w) - g(w)$ to conclude that

$$\|v_{\bar{\mathcal{S}}}[\lambda_0] - w_{\bar{\mathcal{S}}}[\lambda_0]\|_2 \leq \|\mathbf{X}\|_2^2 \frac{\sqrt{|\mathcal{S}|} \mathcal{L}_w \nu}{\alpha^2}.$$

Therefore, for every $j \in \mathcal{S}$ the following holds

$$\begin{aligned} &\left| \mathbf{X}_j^\top (\mathbf{X}_{\bar{\mathcal{S}}} v_{\bar{\mathcal{S}}}[\lambda_0] - y) \right| - \lambda_0 \\ &\leq \left| \mathbf{X}_j^\top (\mathbf{X}_{\bar{\mathcal{S}}} w_{\bar{\mathcal{S}}}[\lambda_0] - y) \right| - \lambda_0 + \|\mathbf{X}\|_2^4 \frac{\sqrt{|\mathcal{S}|} \mathcal{L}_w \nu}{\alpha^2} \\ &\leq \left| \mathbf{X}_j^\top (\mathbf{X}_{\bar{\mathcal{S}}} w[\lambda_0] - y) \right| - \lambda_0 \\ &\quad + \|\mathbf{X}\|_2^2 \left(\frac{\sqrt{|\mathcal{S}|} \|\mathbf{X}\|_2^2 \mathcal{L}_w \nu}{\alpha^2} + \sqrt{|\mathcal{S}|} \|w_{\mathcal{S}}[\lambda_0]\|_\infty \right) \\ &\leq 2\sqrt{|\mathcal{S}|} \|\mathbf{X}\|_2^2 \left(\frac{\mathcal{L}_w \|\mathbf{X}\|_2^2 \nu}{\alpha^2} + \tilde{\mathcal{L}}_u \nu \right), \end{aligned}$$

which concludes the proof. \square

Using again that $\|\mathbf{X}\|_2 \stackrel{c}{\leq} 1$ we obtain

$$\tau \stackrel{c}{\leq} \nu \sqrt{|\mathcal{S}|} \left(\frac{\mathcal{L}_w}{\alpha^2} + \tilde{\mathcal{L}}_u \right). \quad (9)$$

This means that of gradient of objective scales as the product of the root of the size of \mathcal{S} and the length of the interval ν .

Bounding the probability of bad events Next we show that for every fixed set \mathcal{S} of sufficiently large cardinality and sufficiently small ν , $\mathcal{E}_{\mathcal{S}}^{(\tau)}$ holds with very small probability if \mathbf{X} satisfies the smoothness assumption.

Lemma 17 (Smoothing). *For any fixed λ_0 , $\nu > 0$, $\tau \geq 0$, and $\mathcal{S} \subseteq [d]$, let us decompose \mathbf{G} into $\mathbf{G} = \mathbf{G}_{\bar{\mathcal{S}}} + \mathbf{G}_{\mathcal{S}}$ as in Lemma 16. Then, the following inequality holds,*

$$\Pr_{\mathbf{G}_{\mathcal{S}}} \left[\mathcal{E}_{\mathcal{S}}^{(\tau)} | \mathbf{G}_{\bar{\mathcal{S}}} \right] \leq \left(\frac{e\tau\sqrt{n}}{\sigma \|\mathbf{X}_{\bar{\mathcal{S}}} v_{\bar{\mathcal{S}}}[\lambda_0] - y\|_2} \right)^{|\mathcal{S}|}.$$

Proof. Consider the vector $v_{\bar{\mathcal{S}}}[\lambda_0]$ defined as in Lemma 16. We know that $v_{\bar{\mathcal{S}}}[\lambda_0]$ depends only on $\mathbf{G}_{\bar{\mathcal{S}}}$ but not on $\mathbf{G}_{\mathcal{S}}$. Therefore, for any fixed \mathbf{G}_0 , by conditioning on $\mathbf{G}_{\bar{\mathcal{S}}} = \mathbf{G}_0$, we get that for all $j \in \mathcal{S}$

$$\begin{aligned} & \Pr_{\mathbf{G}_{\mathcal{S}}} \left[\left| \left\| \mathbf{X}_j^\top (\mathbf{X}_{\bar{\mathcal{S}}} v_{\bar{\mathcal{S}}}[\lambda_0] - y) \right\| - \lambda_0 \right| \leq \tau \right] \\ &= \Pr_{\mathbf{G}_{\mathcal{S}}} \left[\left| \left\| \left((\mathbf{X}^h)_j + \mathbf{G}_j \right)^\top (\mathbf{X}_{\bar{\mathcal{S}}} v_{\bar{\mathcal{S}}}[\lambda_0] - y) \right\| - \lambda_0 \right| \leq \tau \right] \\ &= \Pr_{\mathbf{G}_j} \left[\left| \left\| \left((\mathbf{X}^h)_j + \mathbf{G}_j \right)^\top (\mathbf{X}_{\bar{\mathcal{S}}} v_{\bar{\mathcal{S}}}[\lambda_0] - y) \right\| - \lambda_0 \right| \leq \tau \right] \\ &\leq \frac{e\tau\sqrt{n}}{\sigma \|\mathbf{X}_{\bar{\mathcal{S}}} v_{\bar{\mathcal{S}}}[\lambda_0] - y\|_2}. \end{aligned}$$

Since the inequality holds for all $j \in \mathcal{S}$ and any \mathbf{G}_0 the proof is completed. \square

8. Proof of Theorem 6

Recall that we assume that the target vector is of unit norm $\|y\|_2 = 1$. We slightly overload notation and denote by $\alpha(\mathbf{X})$ the smallest right singular value of \mathbf{X} . From the optimality of $w[\lambda]$, we know that

$$\frac{1}{2} \|\mathbf{X}w[\lambda] - y\|_2^2 + \lambda \|w[\lambda]\|_1 \leq \frac{1}{2} \|y\|_2^2 \leq \frac{1}{2},$$

which implies that $\|\mathbf{X}w[\lambda] - y\|_2 \leq 1$. From necessary conditions for optimality we also get,

$$\mathbf{X}^\top (\mathbf{X}w[\lambda] - y) = u[\lambda].$$

Thus, we have

$$\|u[\lambda]\|_2 = \|\mathbf{X}^\top (\mathbf{X}w[\lambda] - y)\|_2 \leq \|\mathbf{X}\|_2 = O(1).$$

Therefore, there exists a constant λ_{\max} such that implies that $w[\lambda] = 0$ for $\lambda \geq \lambda_{\max}$.

We next employ the randomness of \mathbf{G} . Consider a fixed α_0 and examine the event that there exists one set \mathcal{S} of size at

least s such that $\mathcal{E}_{\mathcal{S}}^{(\tau)}$ is true, then it holds that,

$$\begin{aligned} & \Pr \left[\left(\exists \mathcal{S} : |\mathcal{S}| = s, \mathcal{E}_{\mathcal{S}}^{(\tau)} \right) \cap \mathcal{D}_\gamma \cap \alpha(\mathbf{X}) \geq \alpha_0 \right] \\ &\leq \sum_{\mathcal{S}_0 \subseteq [d], |\mathcal{S}_0| = s} \Pr \left[\left(\mathcal{E}_{\mathcal{S}_0}^{(\tau)}, \mathcal{S} = \mathcal{S}_0 \right) \cap \mathcal{D}_\gamma \cap \alpha(\mathbf{X}) \geq \alpha_0 \right] \\ &\leq \sum_{\mathcal{S}_0 \subseteq [d], |\mathcal{S}_0| = s} \Pr \left[\mathcal{E}_{\mathcal{S}_0}^{(\tau)} \cap \mathcal{D}_\gamma \cap \alpha(\mathbf{X}) \geq \alpha_0 \right] \\ &\leq \binom{d}{s} \left(\frac{e\tau\sqrt{n}}{\sigma\gamma} \right)^s \leq \left(\frac{e\tau d\sqrt{n}}{\sigma\gamma} \right)^s, \end{aligned}$$

where we used the definition of \mathcal{D}_γ and the fact that $\alpha \stackrel{\text{def}}{=} \alpha(\mathbf{X}) \geq \alpha_0$ to obtain the last inequality. We now set

$$\tau = O \left((\delta\nu)^{1/s} \frac{\sigma\gamma}{d\sqrt{n}} \right),$$

which in turn implies that (see (9)), $\nu^{1-1/s} \stackrel{c}{>} \frac{\delta^{1/s}\sigma\gamma}{d\sqrt{ns}(\mathcal{L}_w/\alpha_0^2 + \mathcal{L}_u)} \Rightarrow \nu \stackrel{c}{>} \left(\frac{\delta^{1/s}\sigma\gamma}{d\sqrt{ns}(\mathcal{L}_w/\alpha_0^2 + \mathcal{L}_u)} \right)^{\frac{s}{s-1}}$. and obtain

$$\Pr \left[\left(\exists \mathcal{S} : |\mathcal{S}| = s, \mathcal{E}_{\mathcal{S}}^{(\tau)} \text{ holds} \right) \cap \mathcal{D}_\gamma \cap \alpha(\mathbf{X}) \geq \alpha_0 \right] \leq \delta\nu.$$

Since we have at most $1/\nu$ many intervals, taking union bound over all linear segments we get that

$$\Pr \left[\left(N(\mathcal{P}) \geq \frac{3^s}{\nu} \right) \cap \mathcal{D}_\gamma \cap (\alpha(\mathbf{X}) \geq \alpha_0) \right] \leq \delta$$

Finally, for properly chosen γ and α_0 we also obtain $\Pr[-\mathcal{D}_\gamma \cup (\alpha(\mathbf{X}) < \alpha_0)] \leq 2\delta$ which completes the proof. \square

To recap, there exists a *universal* constant c such that for *all* $s \in [d]$ the complexity of the Lasso path is bounded above by,

$$c 3^s \left(\frac{\sqrt{snd} \left(\frac{\mathcal{L}_w}{\alpha^2} + \mathcal{L}_u \right)}{\delta^2 \sigma \gamma_s} \right)^{\frac{s}{s-1}}.$$

We now use the bounds on \mathcal{L}_w , \mathcal{L}_u , and α , yielding,

$$\begin{aligned} & \frac{\mathcal{L}_w}{\alpha^2} + \mathcal{L}_u \stackrel{c}{<} \frac{d^{4.5}}{\delta^4 \sigma^4}, \quad \gamma_s \stackrel{c}{<} \frac{\sigma \delta^{\frac{1}{s}}}{\sqrt{dn}} \\ & \Rightarrow |\mathcal{P}| \stackrel{c}{<} 3^s \left(\frac{\sqrt{sn} d^6}{\delta^{6+\frac{1}{s}} \sigma^6} \right)^{\frac{s}{s-1}}. \end{aligned}$$

By choosing $s = O \left(\log \left(\frac{nd}{\delta\sigma} \right) \right)$ we get that

$$|\mathcal{P}| \stackrel{c}{<} n^{1.1} \left(\frac{d}{\delta\sigma} \right)^6.$$




Figure 1. Path complexity as a function of dimension for different levels of smoothing. The theoretical non-smoothed ($\sigma = 0$) complexity is exponential in the size of the problem.




Figure 2. Path complexity in a regression task that predicts the value of a pixel from its neighboring pixels using the MNIST dataset.

9. Experiments

We performed two sets of experiments. Our path following implementation uses Python with **Float128** for high accuracy computations. In the first set of experiments, we start with the exponential complexity construction for $\mathbf{X}^h \in \mathbb{R}^{d \times d}$ from (Mairal & Yu, 2012), which has $(3^d + 1)/2$ many line segments. We artificially added to each entry of \mathbf{X}^h i.i.d. Gaussian noise of mean zero and variance σ^2 . In this setting, the largest value of the entries of \mathbf{X}^h is 1 and y is an all-one vector. We show the effect of dimension d and smoothing σ on $N(\mathcal{P})$. We report the average over 100 random choices for smoothing per \mathbf{X}^h . As can be seen from the figure below, even for a tiny amount of entry-wise noisy of 10^{-10} , the number of linear segments dramatically shrinks. We also include a full table of results, where $1/\text{SNR}$ denotes $-\log_{10}(\sigma)$.

For the next experiment we use the MNIST data set. We randomly selected $n = 1000$ images from the data set. We constructed the data matrix $\mathbf{X} \in \mathbb{R}^{n \times d^2}$ such that the i 'th row of \mathbf{X} is a randomly chosen patch from the i 'th image of size $d \times d$. We cast the center pixel of patch i as the target y_i and discard the pixel from \mathbf{X} . Thus, the regression task amounts to predicting the center pixel y_i using its surrounding pixels $\mathbf{X}^{(i)}$. We plot the relation between the size of the patch and the path complexity $N(\mathcal{P})$. Each point in the graph is the average over 100 random samples of patches and images. As can be seen, when the amount of noise in the data is fixed and governed by the data acquisition process (the minimum pixel value is 0 and maximum is 255 for MNIST), the number of linear segments increases barely faster than linearly in the dimension.

10. Conclusions

We proved that the smoothed complexity of the Lasso's regularization path is pragmatically polynomial in the input

| $1/\text{SNR}$ | $d = 4$ | $d = 5$ | $d = 6$ | $d = 7$ | $d = 8$ | $d = 9$ | $d = 10$ |
|----------------|---------|---------|---------|---------|---------|---------|----------|
| 0 | 6 | 8 | 10 | 12 | 13 | 15 | 18 |
| 1 | 8 | 9 | 10 | 13 | 16 | 14 | 17 |
| 2 | 10 | 13 | 13 | 15 | 17 | 21 | 24 |
| 3 | 16 | 18 | 18 | 20 | 25 | 24 | 26 |
| 4 | 21 | 26 | 27 | 27 | 30 | 33 | 33 |
| 5 | 31 | 41 | 41 | 44 | 46 | 45 | 47 |
| 6 | 36 | 50 | 55 | 58 | 59 | 66 | 67 |
| 7 | 41 | 71 | 81 | 91 | 89 | 91 | 98 |
| 8 | 41 | 91 | 118 | 136 | 133 | 134 | 138 |
| 9 | 41 | 110 | 148 | 181 | 184 | 189 | 192 |
| 10 | 41 | 122 | 205 | 256 | 259 | 286 | 268 |
| 11 | 41 | 122 | 276 | 364 | 407 | 410 | 411 |
| 12 | 41 | 122 | 354 | 467 | 538 | 560 | 566 |
| 13 | 41 | 122 | 365 | 642 | 704 | 795 | 831 |
| 14 | 41 | 122 | 365 | 872 | 1088 | 1141 | 1165 |
| 15 | 41 | 122 | 365 | 978 | 1404 | 1601 | 1694 |
| 16 | 41 | 122 | 365 | 1094 | 1814 | 2162 | 2416 |
| 17 | 41 | 122 | 365 | 1094 | 2478 | 3046 | 3343 |
| 18 | 41 | 122 | 365 | 1094 | 3030 | 3894 | 4345 |
| 19 | 41 | 122 | 365 | 1094 | 3281 | 5323 | 6048 |
| 20 | 41 | 122 | 365 | 1094 | 3281 | 7137 | 8592 |
| ∞ | 41 | 122 | 365 | 1094 | 3281 | 9842 | 29525 |

Table 1. Path complexity in a worst-case synthetic setting. The theoretical non-smoothed complexity is exponential in the size of the problem.

size. Our analysis contrasts worst case settings for which the Lasso's complexity is known to be exponential. To illustrate the key idea, we provided analysis when smoothing each entry in the data matrix by adding small amount of Gaussian noise. Although not presented here, our analysis carries over to settings in which the smoothing is performed using other distributions which can be sub-Gaussian, sub-exponential, or even non i.i.d. so long as the rows of \mathbf{G} are statistically independent.

The nature of smoothed analysis usually imposes a large polynomial factor (Spielman & Teng, 2001), as we do not make any additional assumptions on the hidden matrix \mathbf{X}^h . However, we believe that the polynomial degree in our results can be further reduced. For example, in our proof, we used the fact that the condition number of \mathbf{X} is $\sim \frac{1}{d}$, which is close to being ill-conditioned. For well-conditioned matrices the polynomial bound can be improved to $O(d^{2.1}n^{1.1})$. This reduction is also valid in settings when $n \geq 2d$ (see e.g. (Rudelson & Vershynin, 2010) for different behaviors of the condition number of Gaussian random matrices for $n = d$ and $n \geq 2d$). Furthermore, when $n \geq 2d$, we can also improve γ_s to $\Omega(1)$, hence our polynomial dependency can be further reduced to $O(d^{1.6}n^{0.6})$. A final improvement may stem from \mathcal{L}_w and \mathcal{L}_u . We actually proved that $\left\| \frac{dw[\lambda]}{d\lambda} \right\|_2 = O\left(\frac{\sqrt{d}}{a^2}\right)$, while we only need to use the infinity norm $\left\| \frac{dw[\lambda]}{d\lambda} \right\|_\infty$. We leave these improvements and further generalizations to future research.

References

- Arthur, David, Manthey, Bodo, and Röglin, Heiko. K-means has polynomial smoothed complexity. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pp. 405–414. IEEE, 2009.
- Blum, Avrim and Dunagan, John. Smoothed analysis of the perceptron algorithm for linear programming. In *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 905–914, 2002.
- Dubiner, Moshe and Singer, Yoram. Entire relaxation path for maximum entropy problems. In *Proc. of the 2011 Conference on Empirical Methods in Natural Language Processing, EMNLP 2011*, pp. 941–948, 2011.
- Efron, Bradley, Hastie, Trevor, Johnstone, Iain, Tibshirani, Robert, et al. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.
- Giesen, Joachim, Jaggi, Martin, and Laue, Sören. Approximating parameterized convex optimization problems. In *European Symposium on Algorithms*, pp. 524–535. Springer, 2010.
- Mairal, Julien and Yu, Bin. Complexity analysis of the lasso regularization path. In *Proceedings of the 29th International Conference on International Conference on Machine Learning*, pp. 1835–1842. Omnipress, 2012.
- Osborne, Michael R, Presnell, Brett, and Turlach, Berwin A. A new approach to variable selection in least squares problems. *IMA Journal of Numerical Analysis*, 20(3): 389–403, 2000.
- Rosset, Saharon and Zhu, Ji. Piecewise linear regularized solution paths. *The Annals of Statistics*, pp. 1012–1030, 2007.
- Rudelson, Mark and Vershynin, Roman. Non-asymptotic theory of random matrices: extreme singular values. In *Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (In 4 Volumes) Vol. I: Plenary Lectures and Ceremonies Vols. II–IV: Invited Lectures*, pp. 1576–1602. World Scientific, 2010.
- Sankar, Arvind, Spielman, Daniel A, and Teng, Shang-Hua. Smoothed analysis of the condition numbers and growth factors of matrices. *SIAM Journal on Matrix Analysis and Applications*, 28(2):446–476, 2006.
- Spielman, Daniel and Teng, Shang-Hua. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pp. 296–305. ACM, 2001.
- Spielman, Daniel A and Teng, Shang-Hua. Smoothed analysis: an attempt to explain the behavior of algorithms in practice. *Communications of the ACM*, 52(10):76–84, 2009.
- Tibshirani, Robert. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- Tibshirani, Ryan J. and Taylor, Jonathan. Degrees of freedom in lasso problems. *The Annals of Statistics*, 40(2): 1198–1232, 2012.