
Efficient Designs of SLOPE Penalty Sequences in Finite Dimension

Yiliang Zhang

Zhiqi Bu

University of Pennsylvania

Abstract

In linear regression, SLOPE is a new convex analysis method that generalizes the Lasso via the sorted ℓ_1 penalty: larger fitted coefficients are penalized more heavily. This magnitude-dependent regularization requires an input of penalty sequence $\boldsymbol{\lambda}$, instead of a scalar penalty as in the Lasso case, thus making the design extremely expensive in computation. In this paper, we propose two efficient algorithms to design the possibly high-dimensional SLOPE penalty, in order to minimize the mean squared error. For Gaussian data matrices, we propose a first order Projected Gradient Descent (PGD) under the Approximate Message Passing regime. For general data matrices, we present a zero-th order Coordinate Descent (CD) to design a sub-class of SLOPE, referred to as the k -level SLOPE. Our CD allows a useful trade-off between the accuracy and the computation speed. We demonstrate the performance of SLOPE with our designs via extensive experiments on synthetic data and real-world datasets.

1 Introduction

In sparse linear regression, we aim to find an accurate sparse estimator $\hat{\boldsymbol{\beta}}$ of the unknown truth $\boldsymbol{\beta}$ from

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{w}.$$

Here the response $\mathbf{y} \in \mathbb{R}^n$, the data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$, the true parameter $\boldsymbol{\beta} \in \mathbb{R}^p$ and the noise $\mathbf{w} \in \mathbb{R}^n$. Specifically, in high dimension where $p > n$, ordinary linear regression fails to find a unique solution and L_1 -related regularization is usually introduced to achieve

sparse estimators, including the Lasso [21], elastic net [27], (sparse) group Lasso [24], adaptive Lasso [26] and the recent SLOPE [5]:

$$\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda}) = \arg \min_{\mathbf{b}} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{b}\|_2^2 + \sum_{i=1}^p \lambda_i |b|_{(i)} \quad (1.1)$$

Here $\sum_{i=1}^p \lambda_i |b|_{(i)}$ is the *sorted* ℓ_1 norm of \mathbf{b} governed by the penalty vector $\boldsymbol{\lambda} \in \mathbb{R}^p$ with $\lambda_1 \geq \dots \geq \lambda_p \geq 0$, and $|b|_{(i)}$ is the ordered statistics of absolute values $|b_i|$ such that $|b|_{(1)} \geq \dots \geq |b|_{(p)}$. We pause here to remark that the Lasso is a sub-case of SLOPE when $\lambda_1 = \dots = \lambda_p$, since there is no need to sort and hence sorted ℓ_1 norm is simply the ℓ_1 norm. Generally speaking, the sorting step in the norm allows SLOPE to work in a way similar to the taxation, assigning larger thresholds to larger fitted coefficients.

Many desirable properties have been proven for SLOPE. For example, SLOPE is a convex optimization that can be solved by existing gradient methods, such as the subgradient descent and the proximal gradient descent; SLOPE achieves minimax estimation properties without requiring knowledge of the sparsity degree of $\boldsymbol{\beta}$ [17]; SLOPE controls the false discovery rate in the case of independent predictors. However, understanding the SLOPE problem is difficult. Questions such as what posterior distribution does SLOPE solution follow, can we characterize statistics (e.g. the false discovery rate and true positive rate) from SLOPE exactly, whether SLOPE has better estimation error than the Lasso, are not answered until recently [4, 6, 12]. Still, the substantial difficulty imposed by the sorted penalty impedes the general application of SLOPE for two reasons. From the practical point of view, tuning a \mathbb{R}^p penalty can be extremely costly for large p (e.g. in high dimensional regression or over-parameterized neural networks) and naive methods that work for the Lasso, such as the grid search, renders not pragmatic. From a theoretical perspective, the sorted norm is complicated in that the effect of thresholding of SLOPE is *non-separable* and *data-dependent*, unlike the Lasso, thus making the analysis much involved.

In this paper, we further exploit the advantage of

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

the data-depending penalty in SLOPE and investigate, from the estimation error perspective, how to design the SLOPE penalty sequence to achieve better performance.

We give a computationally efficient framework to design the SLOPE penalty sequence $\boldsymbol{\lambda} \in \mathbb{R}^p$ which corresponds to an estimator $\hat{\boldsymbol{\beta}}(\boldsymbol{\lambda})$ that minimizes the estimation error. To be more specific, we derive the gradient of penalty for SLOPE under the Approximate Message Passing (AMP) regime [1, 2, 9, 10] and propose the k -level SLOPE for the general data matrices. In words, k -level SLOPE is a sub-class of SLOPE, where the p elements in $\{\lambda_i\}$ have only k unique values. Under this definition, the general SLOPE is p -level SLOPE and the Lasso is indeed 1-level SLOPE. Additionally, k -level SLOPE is a sub-class of $(k+1)$ -level SLOPE, and larger k leads to better performance but requires longer computation time. As a result, by choosing an appropriate k , we can establish a trade-off between speed and accuracy. We illustrate in various experiments that such a trade-off is of practical use as even a small k may improve the performance non-trivially.

1.1 Notations

We start by introducing the proximal operator of SLOPE,

$$\text{prox}_{J_\theta}(\mathbf{y}) := \underset{\mathbf{b}}{\text{argmin}} \frac{1}{2} \|\mathbf{y} - \mathbf{b}\|^2 + J_\theta(\mathbf{b}), \quad (1.2)$$

where $J_\theta(\mathbf{b}) := \sum_{i=1}^p \theta_i |b_{(i)}|$ and the proximal operator indeed solves (1.1) with an identity data matrix. This operator is the building block that is iteratively applied to derive the SLOPE estimator in the proximal gradient descent (ISTA [8]) and in FISTA [3]. We note that there is no closed form of $\text{prox}_{J_\theta}(\mathbf{x})$ but it can be efficiently computed as in [5, Algorithm 3]. Next we denote the mean squared error (MSE) between two vectors in \mathbb{R}^m as $\text{MSE}(\mathbf{u}, \mathbf{v}) := \|\mathbf{u} - \mathbf{v}\|^2$. Two performance measures that are investigated in this work are the prediction error, $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$, and the estimation error, $\text{MSE}(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}})$.

2 SLOPE penalty design under AMP regime

2.1 Computing the gradients with respect to the penalty

We introduce a special regime of the AMP for SLOPE [6], within which the SLOPE estimator can be asymptotically exactly characterized. A similar regime is the case when Convex Gaussian Min-max Theorem

(CGMT) [7, 18, 19, 20] applies, which shares similar assumptions as those of AMP. We then derive the gradient of $\text{MSE}(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}})$ with respect to the penalty $\boldsymbol{\lambda}$ and optimize our penalty design iteratively. Generally speaking, AMP is a class of gradient-based optimization algorithms that mainly work on independent Gaussian random data matrices, offering both a sequence of estimators that converges to the true minimizer and a distributional characterization of the latter (see [6, Theorem 3] and [12, Theorem 1]). Here we present the five assumptions of the SLOPE AMP [6]:

- The data matrix \mathbf{X} has independent and identically-distributed (i.i.d.) gaussian entries that have mean 0 and variance $1/n$.
- The signal $\boldsymbol{\beta}$ has elements that are i.i.d. and follow Π , with $\mathbb{E}(\Pi^2 \max\{0, \log \Pi\}) < \infty$.
- The noise \mathbf{w} is elementwise i.i.d. and follows W , with $\sigma_w^2 := \mathbb{E}(W^2) < \infty$.
- The vector $\boldsymbol{\lambda}(p) = (\lambda_1, \dots, \lambda_p)$ is elementwise i.i.d. and follows Λ , with $\mathbb{E}(\Lambda^2) < \infty$.
- The ratio n/p reaches a constant $\delta \in (0, \infty)$ in the large system limit, as n and $p \rightarrow \infty$.

Under these assumptions, Theorem 3 in [6] provides an asymptotic characterization of $\hat{\boldsymbol{\beta}}$, which can be informally interpreted as

$$\hat{\boldsymbol{\beta}} \approx \text{prox}_{J_{\boldsymbol{\alpha}\tau}}(\boldsymbol{\beta} + \tau\mathbf{Z}) \quad (2.1)$$

in which $(\boldsymbol{\alpha}, \tau)$ are the unique solutions of two key equations, namely the calibration and the state evolution in the AMP (or CGMT) regime (see [6, 12]):

$$\boldsymbol{\lambda} = \boldsymbol{\alpha}\tau \left(1 - \frac{1}{n} \mathbb{E} \|\text{prox}_{J_{\boldsymbol{\alpha}\tau}}(\boldsymbol{\beta} + \tau\mathbf{Z})\|_0^* \right) \quad (2.2)$$

$$\tau^2 = \sigma_w^2 + \frac{1}{\delta p} \mathbb{E} \|\text{prox}_{J_{\boldsymbol{\alpha}\tau}}(\boldsymbol{\beta} + \tau\mathbf{Z}) - \boldsymbol{\beta}\|^2 \quad (2.3)$$

Here we assume the noise \mathbf{w} has variance σ_w^2 , $\|\cdot\|_0^*$ is a modified ℓ_0 norm that counts the unique non-zero absolute values in a vector and $\mathbf{Z} \in \mathbb{R}^p$ is a vector in which each element is i.i.d. standard normal. We denote $\delta := \lim_p n/p$ as the aspect ratio or sampling ratio and $\epsilon := \lim_p |\{j : \beta_j \neq 0\}|/p$.

Using (2.1), we observe that to minimize $\text{MSE}(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}})$ is equivalent to finding desirable $(\boldsymbol{\alpha}, \tau)$. We now introduce some properties that are useful in deriving the desirable $\boldsymbol{\lambda}$, which uniquely defines $(\boldsymbol{\alpha}, \tau)$. By [6, Proposition 2.3], the calibration (2.2) describes a bijective, monotone and parallel mapping Λ^1 between $\boldsymbol{\alpha}$ and $\boldsymbol{\lambda}$

¹For a given $\boldsymbol{\alpha}$, we can use (2.3) to obtain a unique $\tau(\boldsymbol{\alpha})$ and leverage (2.2) to obtain a corresponding penalty vector $\boldsymbol{\lambda}(\boldsymbol{\alpha})$.

[6, Proposition 2.3], which allows us to work with α easily instead of λ . By [6, Theorem 1], the state evolution (2.3) can be solved via a fixed point recursion, which converges to the unique solution $\tau(\alpha)$ monotonically under any initial condition.

Under AMP region, our strategy is to design $\lambda \in \mathbb{R}^p$ in SLOPE that, by quoting [6, Corollary 3.2], minimizes:

$$\text{plim} \|\hat{\beta} - \beta\|^2/p = \delta(\tau^2 - \sigma_w^2)$$

where plim is the probability limit. Hence minimizing $\text{MSE}(\beta, \hat{\beta})$ is in fact equivalent to minimizing τ , which depends on α and leads to differentiating (2.3) against each of α_i for $i \in [p]$. In what follows, we view the scalar τ as a function of the penalty α given the prior. Next, we use the gradient information to descend (with the projection elaborated in Algorithm 1) till convergence. Once the minimizer α is obtained, we leverage the calibration (2.2) to map to the corresponding $\lambda(\alpha)$.

In what follows, we shorthand $\text{prox}_{J_b}(\mathbf{a})$ by using $\boldsymbol{\eta}(\mathbf{a}; \mathbf{b})$. In particular $\text{prox}_{J_{\alpha\tau}}(\beta + \tau\mathbf{Z})$ is denoted by $\boldsymbol{\eta}$ and we let η_j represent its j -th element. We define a set $I_j := \{k : |\eta_k| = |\eta_j|\}$, which will be used in characterization of gradients. We also define an inverse mapping for ranking of indices: $\sigma : \{1, \dots, p\} \rightarrow \{1, \dots, p\}$ such that $\sigma(i) = j$ representing $|\boldsymbol{\eta}|_{(i)} = |\eta_j|$. Consider a toy example $\boldsymbol{\eta} = (-2, -4, 3, 1)$, then the ranking of magnitudes is (3, 1, 2, 4) whose inverse gives: $\sigma(1) = 2$, $\sigma(2) = 3$, $\sigma(3) = 1$ and $\sigma(4) = 4$. This mapping is useful in assigning the penalties to coefficients in $\hat{\beta}$ due to the sorting procedure.

We state the following theorem to give a concrete form of gradients $\partial\tau/\partial\alpha_i$, which is used in the projected gradient descent (PGD) in Algorithm 2.

Theorem 1. *The gradients satisfy*

$$\frac{\partial\tau}{\partial\alpha_i} = \mathbb{E} \frac{1}{|I_{\sigma(i)}|D(\alpha, \tau)} \sum_{j \in I_{\sigma(i)}} (\eta_j - \beta_j) \text{sign}(\eta_j)\tau \quad (2.4)$$

where $D(\alpha, \tau)$ is a negative constant that is independent of index i .

Here the expectation is taken with respect to \mathbf{Z} in $\boldsymbol{\eta}$, which in turn also affects $I_{\sigma(i)}$. The detailed form of $D(\alpha, \tau)$ in the denominator and the proof of Theorem 1 can be found in Appendix B, where we also claim that $D(\alpha, \tau)$ is always negative. In practice, we can either set the step size s_t as constant or simply set $D = 1$ to save computation time. We remark that, using a constant step size and $\mathbb{E} \left(\sum_{j \in I_{\sigma(i)}} (\eta_j - \beta_j) \text{sign}(\eta_j)\tau \right) / |I_{\sigma(i)}|$ as the gradient is equivalent to using a time-dependent $s_t = s \cdot D(\alpha^t, \tau_t)$ and the actual gradient $\frac{\partial\tau}{\partial\alpha_i}$.

2.2 Projection onto non-negative decreasing vectors

We notice that α (and λ) must be non-negative and decreasing, hence the vanilla gradient descent is unsuitable for this constrained optimization problem of α . We design a projected gradient descent (PGD) in the following. To do so, we first give Algorithm 1 to compute the projection and establish the correctness of the algorithm in Theorem 2.

Let \mathcal{S} denote the set of non-negative and decreasing vectors in \mathbb{R}^p (i.e. $\lambda \in \mathcal{S} \Rightarrow \lambda_i \geq \lambda_{i+1} \geq 0, \forall i$). Define the projection on to \mathcal{S} as

$$\Pi_{\mathcal{S}}(\gamma) = \underset{\gamma' \in \mathcal{S}}{\text{argmin}} \frac{1}{2} \|\gamma - \gamma'\|_2^2. \quad (2.5)$$

Algorithm 1 ProjectOnS ($\Pi_{\mathcal{S}}$)

Input: Arbitrary sequence $\gamma = (\gamma_1, \dots, \gamma_p)$

for $i = 1, \dots, p$ **do**

\triangleright Identify the shortest sub-sequence $\{\gamma_j, \dots, \gamma_i\}$ whose average is smaller than its left neighbor (with $\gamma_0 = \infty$):

$$\frac{1}{i-j+1} \sum_{k=j}^i \gamma_k \leq \gamma_{j-1}$$

\triangleright Assign the average value to such sub-sequence for $(\gamma_j, \dots, \gamma_i)$:

$$\gamma_j, \dots, \gamma_i \leftarrow \frac{1}{i-j+1} \sum_{k=j}^i \gamma_k$$

Output: $\max\{\gamma, 0\}$ \triangleright Element-wise truncation

We show that Algorithm 1 indeed finds the minimizer of (2.5) and provide the proof in Appendix C.1.

Theorem 2. *Given an arbitrary $\gamma \in \mathbb{R}^p$ as input, Algorithm 1 outputs the projection of γ on \mathcal{S} , that is, $\Pi_{\mathcal{S}}(\gamma)$.*

On high level, the proof consists of two parts. In the first part we provide a detailed characterization of $\Pi_{\mathcal{S}}(\gamma)$ by partitioning the index sequence $\{1, \dots, p\}$ into a number of carefully selected sub-sequences. We prove that within each sub-sequence, $\Pi_{\mathcal{S}}(\gamma)$ takes the same value at each index, and such value is exactly the average of the sub-sequence γ 's values at these indices. In the second part, we prove that Algorithm 1 indeed finds such sub-sequences and thus operates in a way that matches the goal of the projection $\Pi_{\mathcal{S}}(\gamma)$. The final truncation of the averaged sequence at 0 is a trivial method to guarantee the non-negativity.




Figure 1: \mathbf{X} i.i.d. $\mathcal{N}(0, 1/n)$, $n = 300$, $p = 1000$, Π is Bernoulli(ϵ), $\delta = n/p = 0.3$, $\epsilon = 0.5$, $\sigma_w = 0$. Top: A sample run of PGD that finds the minimizing α , *not* the minimizing $\hat{\beta}$ for SLOPE in [6]. Additionally, we plot two gradient descent methods with 0.9 momentum: the projected heavy ball (PHB) and the projected Nesterov accelerated gradient. All methods use the fine-tuned Lasso penalty as their starting points. Top-middle: Red dashed line is Lasso MSE path (each point corresponds to one Lasso penalty, as λ varies from 0 to large values); other lines are different SLOPE MSE for a single SLOPE penalty. BH here stands for “Benjamini and Hochberg”. Bottom-middle: Best α (right y-axis) and best λ (left y-axis) sequences found by PGD. Bottom-right: Histogram of best λ (bottom) and α (upper) sequences found by PGD.

2.3 Projected Gradient Descents

Now that we have the gradients in Theorem 1 and the projection in Algorithm 1, the projected gradient descent is straight-forward. In each iteration we first conduct gradient descent using Theorem 1 and transform the sequence from α regime to λ regime. The transformation Λ is defined in Section 2.1 using the calibration and the state evolution in AMP. Then we project the gradient onto the constrained space \mathcal{S} and transform it back to α regime. The procedure is summarized in Algorithm 2.

Algorithm 2 Projected Gradient Descent (PGD)

Input: initial α^0 , step size $\{s_t\}$

for $t = 1, \dots, T$ **do**

\triangleright Gradient descent on α and transform to λ regime

$$\gamma^{t+1} = \Lambda(\alpha^t - s_t \nabla_{\alpha}(\tau(\alpha^t)))$$

\triangleright Project onto \mathcal{S}

$$\lambda^{t+1} = \text{ProjectOnS}(\gamma^{t+1})$$

\triangleright Transform back to α regime

$$\alpha^{t+1} = \Lambda^{-1}(\lambda^{t+1})$$

Output: α^{T+1}

We highlight that Algorithm 2 is only one form of PGD. In fact, with a concrete form of the gradients, we can use any off-the-shelf first-order optimizer to find α iteratively. Some examples include projected versions of stochastic gradient descent, Heavy Ball method [16], Nesterov accelerated gradient descent [15], Adagrad [11], AdaDelta [25] and Adam [13]. We include some of these optimizers in Figure 1.

To understand the convergence behavior of PGD, we need to study the convexity of the domain \mathcal{S} and the objective function τ . Clearly \mathcal{S} is convex by simply applying the definition. Unfortunately, $\tau(\alpha)$ for SLOPE is in general non-convex: even in the Lasso AMP regime, it is shown that $\tau(\alpha)$ is only a quasi-convex function of α [14, Theorem 3.3]. We note that some non-convex problems may enjoy desirable properties such as having unique global minimum or local minima do not exist. As for SLOPE, the analysis on quasi-convexity of $\tau(\alpha)$ has not been well established but in practice, we do not observe any local minimum.

Remarkably, the gradient information that we use distinguishes our work from [12]. We pause a bit and compare our approach with theirs, as they work under very similar assumptions as our AMP regime (in fact, both AMP and CGMT regimes agree asymptotically). Instead of optimizing directly on τ , they propose to optimize the proximal operator η in the functional space [12, Proposition 3]: for a fixed candidate

τ , they use the finite approximation with 2048 grids to solve a functional optimization, whose minimum is $\mathcal{L}(\tau)$. Next, they check the feasibility of the candidate τ by whether $\mathcal{L}(\tau) \leq \delta(\tau^2 - \sigma_w^2)$. Lastly, a binary search is conducted to find the optimal τ (smallest feasible τ) and the optimal design can be derived from the corresponding η . In summary, this approach took a detour by using a zeroth-order optimization algorithm, as the authors did not search over $\boldsymbol{\lambda}$ (or $\boldsymbol{\alpha}$) directly. Our first-order algorithm overcomes the seemingly unwieldy computation burden, especially in the high dimension when p is very large.

2.4 Transforming from α to λ

Once we find the desirable $\boldsymbol{\alpha}$ with Algorithm 2, the calibration (2.2) allows us to convert $\boldsymbol{\alpha}$ to $\boldsymbol{\lambda}$ in the original SLOPE problem. We demonstrate in Figure 1 and Figure 2 that SLOPE can outperform the best-tuned Lasso significantly. In Figure 1, SLOPE reduces $\text{MSE}(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}})$ from 0.473 by Lasso to 0.350 by SLOPE, a 26% improvement in the estimation error. In fact, we observe from Figure 2 that SLOPE is even comparable to Minimum Mean Squared Error (MMSE; proposed by [1]) estimator, which produces the lowest MSE possible. We emphasize that our result does not contradict [23] which states that, under some conditions, the Lasso is the optimal SLOPE. We note that the condition in [23, Theorem 2] does not hold for large systems: the premise of Lasso being optimal is that the Lasso achieves exact recovery, which requires $n \sim p \log p$ (see [22]). Therefore, in our setting where $n/p \rightarrow \delta$, the Lasso is incapable of achieving the exact recovery nor outperforming general SLOPE.

3 k -level SLOPE

In this section we propose a method, described in Algorithm 3, that works on the general linear model. I.e. our method works on arbitrary data $\boldsymbol{X}, \boldsymbol{y}, \boldsymbol{\beta}, \boldsymbol{w}$ and does not require $n/p \rightarrow \delta$ when $n, p \rightarrow \infty$.

In contrast to the AMP regime, we directly search on $\boldsymbol{\lambda}$ without implicitly using $\boldsymbol{\alpha}$ and we do not try to use the gradient information. To avoid searching in the high dimension of $\boldsymbol{\lambda}$ space, we propose to restrict that the penalty $\boldsymbol{\lambda}$ only contains k different non-negative values, which is denoted by $(\lambda_1, \dots, \lambda_k; S_1, \dots, S_{k-1})$. Here λ_i denotes the penalty magnitude and S_i represents the splitting index in $[p]$, where the penalty magnitudes change, i.e. $S_i - S_{i-1}$ entries in $\boldsymbol{\lambda}$ take the value λ_i . We note that λ_i is decreasing in i while S_i is increasing, guaranteeing that $\boldsymbol{\lambda}$ satisfies the assumption of SLOPE. As an example in \mathbb{R}^5 , $\boldsymbol{\lambda} = (7, 5, 1; 2, 3) = (7, 7, 5, 1, 1)$. We name this restricted SLOPE problem as the k -level SLOPE and

design the $(2k - 1)$ degree of freedom penalty $\boldsymbol{\lambda}$, so as to only search in the reduced dimension $k \ll p$.

Notice that the original SLOPE is the p -level SLOPE and the Lasso is the 1-level SLOPE. We note that k -level SLOPE is always a sub-case of $(k + 1)$ -level SLOPE. Therefore intuitively, by allowing k to take values other than 1 and p , we can trade off the difficulty of designing the penalty and the accuracy gain by employing more penalty levels. We demonstrate that empirically, the trade-off is surprisingly encouraging: even 2 or 3 levels of penalty is sufficient to exploit the benefit of SLOPE.

3.1 Practical penalty design for k -level SLOPE

We emphasize that in the general regime beyond AMP and CGMT, we cannot access the gradient information nor the functional optimization in [12] for two reasons: the true $\boldsymbol{\beta}$ distribution is not known in real data and the data matrix \boldsymbol{X} is general (not i.i.d. with a specific variance). To design the k -level SLOPE penalty in the real-world datasets, we propose the Coordinate Descent (CD, Algorithm 3)² and compare to the PGD in Algorithm 2 under the AMP and CGMT regimes in Figure 2.

Algorithm 3 Coordinate Descent (CD)

Input: initial $\boldsymbol{\lambda}, \text{MSE}_{old} = \infty$, level k
while $\text{MSE} < \text{MSE}_{old}$ **do**
 Set $\text{MSE}_{old} = \text{MSE}$
 for $i \in \{1, \dots, k\}$ **do**
 ▷ Search on magnitudes λ_i
 $\lambda_i = \underset{\lambda_i \in (\lambda_{i+1}, \lambda_{i-1})}{\text{argmin}} \text{MSE}(\lambda_1, \dots; S_1, \dots)$
 for $i \in \{1, \dots, k - 1\}$ **do**
 ▷ Search on splits S_i
 $S_i = \underset{S_i \in (S_{i-1}, S_{i+1})}{\text{argmin}} \text{MSE}(\lambda_1, \dots; S_1, \dots)$
Output: $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_k; S_1, \dots, S_{k-1})$

We highlight some details of Algorithm 3 that make it efficient and practical. First of all, Algorithm 3 directly works on $\boldsymbol{\lambda}$ instead of $\boldsymbol{\alpha}$ (the calibration is generally unavailable). Second, the projection is not needed as in Algorithm 2 since $\boldsymbol{\lambda}$ is decreasing and non-negative by our definition of the search domain. Third, Algorithm 3 is flexible in the following sense: (1) we can choose any order of coordinates to successively minimize the error, e.g. by $\lambda_1, S_1, \lambda_2, S_2, \dots$; (2) we can use any zeroth-order search method such

²We slightly abuse the notation of MSE to mean either the estimation error (only available in synthetic data) or the prediction error.

as the grid search or the binary search for the magnitudes and splits.

4 Experiments

In this section, we justify the effectiveness of k -level SLOPE on various synthetic and real datasets, on linear and logistic regression tasks. For the sake of implementation consistency, we adopt R package SLOPE to run both Lasso and SLOPE in experiments. Empirically, we remark that PGD and k -level CD are both significantly fast in all experimental settings, taking only a few minutes to converge even for $p = 1000$.

4.1 Synthetic datasets

Independent case

In this experiment we investigate the performance of k -level SLOPE in the AMP regime: data matrix \mathbf{X} is i.i.d. $\mathcal{N}(0, 1/n)$, $n = 300$. The signal distribution Π is Gaussian-Bernoulli with probability 0.5 being standard normal and 0 otherwise. We work on a high-dimensional setting where $\delta = n/p = 0.3$ and hence $\mathbf{X} \in \mathbb{R}^{300 \times 1000}$. The noise σ_w is 0. We observe from Figure 2 that employing more levels of penalty is beneficial and fast, suggesting that even a small k may be sufficient to reduce the errors significantly.

Dependent case

Different from the AMP regime in which each entry in the design matrix is i.i.d. gaussian, we study the performance of 2-level SLOPE in a synthetic dataset with features strongly correlated with each other. We include three other methods: Lasso and SLOPE with two other designs: Benjamini Hochberg design and the MR design proposed in [4]. The data \mathbf{X} is generated from an ARMA(1,1) model:

$$X_t = \varepsilon_t + 0.8X_{t-1} + 0.8\varepsilon_{t-1} \quad (4.1)$$

where X_t denote the t -th feature and ε_t follows i.i.d. $\mathcal{N}(0, 1)$. We set Π , the asymptotic distribution of β , to be i.i.d. Gaussian-Binomial: $\beta_i \sim \mathbf{BZ}$ with $\mathbf{B} \sim B(5, 0.3)$ and \mathbf{Z} being standard normal. In terms of the dimension, we study two cases (1) $n = 20$, $p = 50$; (2) $n = 200$, $p = 500$. 10-fold cross-validation $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$ are calculated for both the Lasso and SLOPE. We highlight that, different than the previous section, we investigate the prediction error instead of the estimation error here. Curves for $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$ with different iterations in both cases are shown in Figure 3. In the first case, using grid search, the optimal prediction $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$ given by Lasso is 0.128 while the optimal prediction $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$ given by




Figure 2: Top: the result of a single run for MMSE AMP, Lasso AMP, 2-level SLOPE AMP (by CD) and the p -level SLOPE AMP (by PGD). Bottom: averaged result over 10 independent runs of different k -level SLOPE by CD.

2-level SLOPE (using Algorithm 3) is 0.083. Prediction errors of SLOPE with other two penalty sequences are also under 0.1, but worse than that of 2-level SLOPE. We observe a 35% improvement on prediction error when using 2-level SLOPE for this case of small sample size and dimension, compared with Lasso. In the second case, the optimal prediction $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$ given by Lasso and other two SLOPEs are no smaller than 0.2, while that of 2-level SLOPE is 0.186, giving a 7.5% reduction in the prediction error.

4.2 Real datasets for linear and logistic regression

To further demonstrate the utility of k -level SLOPE in practice, we apply the model to real datasets, where $\text{MSE}(\beta, \hat{\beta})$ is intractable, and focus on the prediction $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$. In this experiment, again we compare the performance of 2-level SLOPE in a linear regression setting with three other methods we studied in Section 4.1. The dataset we adopt is atherosclerosis cardiovascular disease (ASCVD), which records medical information of 236 patients and their corresponding ASCVD risk score (outcome variable). We select 1000 features out of 4216 features, which have the largest correlation with the outcome variable. We conduct 20-fold cross-validation and calculate the cross-validation

prediction $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$. Using grid search, the optimal prediction $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$ given by Lasso is 0.528. Interestingly, prediction $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$ given by other two SLOPEs are worse than that of Lasso while that given by 2-level SLOPE (using Algorithm 3) is 0.489. This result clearly demonstrates the outperformance of k -level SLOPE compared to Lasso and SLOPE using other penalty sequences. A curve for $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$ with different iterations is shown in Figure 3.

We further extend the idea of k -level SLOPE in logistic regression and justify the results on Alzheimer’s Disease Neuroimaging Initiative (ADNI) gene dataset. The dataset contains over 19000 genomic features of 649 patients, along with a binary disease status (normal or ill). We select the first 300 patients in the original dataset and 500 features out of the total features, which has the largest correlation with the outcome variable. We conduct 10-fold cross-validation and calculate the cross-validation prediction accuracy. Using grid search, the optimal prediction accuracy given by Lasso is 0.62. The optimal prediction accuracy given by 2-level SLOPE (using Algorithm 3) is 0.66.

5 Discussion

In this work, we propose a framework to flexibly and efficiently design the SLOPE penalty sequence. Under the AMP setting, our first-order PGD approach is capable of finding the effective penalty sequence with reasonable computation budget. The key is to use the gradient with respect to the penalty instead of using zeroth-order search as previous works have proposed. In the practical world beyond the AMP setting, via various experiments, we illustrate that the proposed k -level SLOPE with penalty sequence determined by Algorithm 3 can provide decent results. Although Algorithm 3 loses the access to the first-order information when compared to Algorithm 2, the universal ability to search good penalty is desirable for practical use, as we can view the algorithm as a dimension-reduction trick. In many cases even 2-level SLOPE, the simplest k -level SLOPE (other than Lasso), can outperform the Lasso in accuracy as well as the (p -level) SLOPE in computation speed. Additionally, our framework indeed generalizes to other high-dimensional penalty designs. Some direct extensions include group SLOPE and weighted Lasso.

Much room is left for future study. From a theoretical perspective, the quasi-convexity of $\tau(\alpha)$ in AMP setting is still not well studied. The asymptotic $\text{MSE}(\boldsymbol{\beta}, \hat{\boldsymbol{\beta}})$ (i.e. Equation (7) in [14]) is shown to be quasi-convex in Lasso case. However, no such theoretical property has been shown for SLOPE. If the quasi-convexity indeed holds true for SLOPE AMP, then we can guaran-




Figure 3: $\text{MSE}(\mathbf{y}, \hat{\mathbf{y}})$ in linear regression cases. SLOPE-MR: SLOPE using penalty sequence suggested in [4]; SLOPE-BH: SLOPE using Benjamini–Hochberg penalty sequence in R function ‘SLOPE’. Top: Synthetic data with \mathbf{X} i.i.d. drawn from ARMA(1,1) model (4.1), $n = 20$, $p = 50$. Middle: \mathbf{X} i.i.d. drawn from (4.1) with $n = 200$, $p = 500$. Bottom: the results of ASCVD dataset.

tee that the minimizing λ by PGD is indeed the global minimizer and thus claim our design is optimal.

It would also be interesting to develop PGD (based on AMP regime) for k -level SLOPE, i.e. using gradient descent to find the optimal magnitudes and splits. One could then derive a theoretical trade-off curve between the minimum τ and each k , similarly to Figure 2 bottom subplot. This would suggest a proper choice of k for our k -level SLOPE.

From a practical perspective, we anticipate that k -level SLOPE can also be explored in various applications that already employ the Lasso, such as the matrix completion, the compressed sensing and the neural network regularization.

References

- [1] Mohsen Bayati and Andrea Montanari. “The dynamics of message passing on dense graphs, with applications to compressed sensing”. In: *IEEE Transactions on Information Theory* 57.2 (2011), pp. 764–785.
- [2] Mohsen Bayati and Andrea Montanari. “The LASSO risk for Gaussian matrices”. In: *IEEE Transactions on Information Theory* 58.4 (2011), pp. 1997–2017.
- [3] Amir Beck and Marc Teboulle. “A fast iterative shrinkage-thresholding algorithm for linear inverse problems”. In: *SIAM journal on imaging sciences* 2.1 (2009), pp. 183–202.
- [4] Pierre C Bellec, Guillaume Lecué, Alexandre B Tsybakov, et al. “Slope meets lasso: improved oracle bounds and optimality”. In: *The Annals of Statistics* 46.6B (2018), pp. 3603–3642.
- [5] Małgorzata Bogdan et al. “SLOPE—adaptive variable selection via convex optimization”. In: *The annals of applied statistics* 9.3 (2015), p. 1103.
- [6] Zhiqi Bu et al. “Algorithmic Analysis and Statistical Estimation of SLOPE via Approximate Message Passing”. In: *IEEE Transactions on Information Theory* 67.1 (2020), pp. 506–537.
- [7] Michael Celentano, Andrea Montanari, and Yuting Wei. “The Lasso with general Gaussian designs with applications to hypothesis testing”. In: *arXiv preprint arXiv:2007.13716* (2020).
- [8] Ingrid Daubechies, Michel Defrise, and Christine De Mol. “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint”. In: *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences* 57.11 (2004), pp. 1413–1457.
- [9] David L Donoho, Arian Maleki, and Andrea Montanari. “Message passing algorithms for compressed sensing: I. motivation and construction”. In: *2010 IEEE information theory workshop on information theory (ITW 2010, Cairo)*. IEEE. 2010, pp. 1–5.
- [10] David L Donoho, Arian Maleki, and Andrea Montanari. “Message-passing algorithms for compressed sensing”. In: *Proceedings of the National Academy of Sciences* 106.45 (2009), pp. 18914–18919.
- [11] John Duchi, Elad Hazan, and Yoram Singer. “Adaptive subgradient methods for online learning and stochastic optimization.” In: *Journal of machine learning research* 12.7 (2011).
- [12] Hong Hu and Yue M Lu. “Asymptotics and optimal designs of SLOPE for sparse linear regression”. In: *2019 IEEE International Symposium on Information Theory (ISIT)*. IEEE. 2019, pp. 375–379.
- [13] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [14] Ali Mousavi, Arian Maleki, Richard G Baraniuk, et al. “Consistent parameter estimation for LASSO and approximate message passing”. In: *The Annals of Statistics* 46.1 (2018), pp. 119–148.
- [15] Yurii Nesterov. “A method for unconstrained convex minimization problem with the rate of convergence $O(1/k^2)$ ”. In: *Doklady an ussr*. Vol. 269. 1983, pp. 543–547.
- [16] Boris T Polyak. “Some methods of speeding up the convergence of iteration methods”. In: *USSR Computational Mathematics and Mathematical Physics* 4.5 (1964), pp. 1–17.
- [17] Weijie Su, Emmanuel Candes, et al. “SLOPE is adaptive to unknown sparsity and asymptotically minimax”. In: *The Annals of Statistics* 44.3 (2016), pp. 1038–1068.
- [18] Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. “Precise Error Analysis of Regularized M -Estimators in High Dimensions”. In: *IEEE Transactions on Information Theory* 64.8 (2018), pp. 5592–5628.
- [19] Christos Thrampoulidis, Samet Oymak, and Babak Hassibi. “Regularized linear regression: A precise analysis of the estimation error”. In: *Proceedings of Machine Learning Research* 40 (2015), pp. 1683–1709.
- [20] Christos Thrampoulidis, Samet Oymak, and Babak Hassibi. “The Gaussian min-max theorem in the presence of convexity”. In: *arXiv preprint arXiv:1408.4837* (2014).

- [21] Robert Tibshirani. “Regression shrinkage and selection via the lasso”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 58.1 (1996), pp. 267–288.
- [22] Martin J Wainwright. “Sharp thresholds for High-Dimensional and noisy sparsity recovery using ℓ_1 -Constrained Quadratic Programming (Lasso)”. In: *IEEE transactions on information theory* 55.5 (2009), pp. 2183–2202.
- [23] Shuaiwen Wang, Haolei Weng, and Arian Maleki. “Does SLOPE outperform bridge regression?” In: *arXiv preprint arXiv:1909.09345* (2019).
- [24] Ming Yuan and Yi Lin. “Model selection and estimation in regression with grouped variables”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68.1 (2006), pp. 49–67.
- [25] Matthew D Zeiler. “Adadelata: an adaptive learning rate method”. In: *arXiv preprint arXiv:1212.5701* (2012).
- [26] Hui Zou. “The adaptive lasso and its oracle properties”. In: *Journal of the American statistical association* 101.476 (2006), pp. 1418–1429.
- [27] Hui Zou and Trevor Hastie. “Regularization and variable selection via the elastic net”. In: *Journal of the royal statistical society: series B (statistical methodology)* 67.2 (2005), pp. 301–320.