Pattern Recognition Letters
journal homepage: www.elsevier.com

Revisiting Convexity-Preserving Signal Recovery with the Linearly Involved GMC Penalty

Xiaoqian Liu[a,**], Eric C. Chi[b]

[a]Department of Statistics, North Carolina State University, Raleigh, NC 27695, USA
[b]Department of Statistics,Rice University, Houston, TX 77005, USA

Article history:
Received 15 July 2020
Received in final form
Accepted
Available online

Communicated by

Keywords: Convexity-preserving nonconvex strategy, Generalized minimax concave penalty, Linearly involved convexity-preserving model, Feasibility problem, Saddle-point problem

ABSTRACT

The generalized minimax concave (GMC) penalty is a newly proposed regularizer that can maintain the convexity of the objective function. This paper deals with signal recovery with the linearly involved GMC penalty. First, we propose a new method to set the matrix parameter in the penalty via solving a feasibility problem. The new method possesses appealing advantages over the existing method. Second, we recast the linearly involved GMC model as a saddle-point problem and use the primal-dual hybrid gradient (PDHG) algorithm to compute the solution. Another important work in this paper is that we provide guidance on the tuning parameter selection by proving desirable properties of the solution path. Finally, we apply the linearly involved GMC penalty to 1-D signal recovery and matrix regression. Numerical results show that the linearly involved GMC penalty can obtain better recovery performance and preserve the signal structure more successfully in comparison with the total variation (TV) regularizer.

1. Introduction

The task of recovering a sparse representation is ubiquitous in signal and image recovery and is often formulated as the following optimization problem:

$$\text{minimize } F(\mathbf{x}) = \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \mu\psi(\mathbf{x}), \tag{1}$$

where $\mathbf{y} \in \mathbb{R}^m$ is the vector of observations, $\mathbf{A} \in \mathbb{R}^{m\times n}$ is a matrix, and $\psi : \mathbb{R}^n \mapsto \mathbb{R}$ is a sparsity promoting penalty function. The nonnegative tuning parameter μ trades off the emphasis on the least squares data-fidelity term and the sparsity level in $\mathbf{x}$. Many applications in signal processing fall under the framework of solving (1), for instance, signal denoising (Chen et al., 1998) and signal smoothing (Selesnick, 2015). In statistics, (1) corresponds to the sparse linear regression problem, where $\mathbf{y}$ is the response vector, $\mathbf{A}$ is the design matrix, and ψ is the regularizer to promote sparsity in the vector of coefficients $\mathbf{x}$. A commonly used regularizer is the l_1-norm $\psi(\mathbf{x}) = \|\mathbf{x}\|_1$, which is a convex function. Convexity ensures that all local minimizers are global. Nonetheless, the l_1-norm regularizer also introduces estimation bias, namely the solution to (1) tends to underestimate large magnitude components of $\mathbf{x}$ (Chen et al., 1998; Tibshirani, 1996). To deal with this issue, several nonconvex sparse regularizers have been proposed to mitigate the estimation bias and consequently more accurately recover the signal. However, with a nonconvex regularizer ψ, the convexity of F is no longer guaranteed. Therefore, (1) could have multiple suboptimal local minimizers which are undesirable.

To overcome the drawbacks of nonconvex optimization but enjoy the less biased estimation of nonconvex regularizers, convexity-preserving nonconvex strategies were introduced in Blake and Zisserman (1987); Nikolova (1998); Nikolova et al. (2010), and further investigated in Bayram (2015); Selesnick (2017b); Zou et al. (2018). The basic idea of convexity-preserving nonconvex strategies is to use a nonconvex regularizer but keep the optimization problem convex. Selesnick (2017a) introduced the generalized minimax concave (GMC) penalty, which can maintain the convexity of F under a modest regularity condition. The GMC penalty is defined as:

**Corresponding author.
e-mail: xliu62@ncsu.edu (Xiaoqian Liu)

$$\psi_{\mathbf{B}}(\mathbf{x}) = \|\mathbf{x}\|_1 - \min_{\mathbf{v}\in\mathbb{R}^n}\left\{\|\mathbf{v}\|_1+\frac{1}{2}\|\mathbf{B}(\mathbf{x}-\mathbf{v})\|_2^2\right\}, \tag{2}$$

where $\mathbf{B} \in \mathbb{R}^{m\times n}$ is a parameter for $\psi_{\mathbf{B}}$. The objective function F in (1) is convex if $\mathbf{A}^\mathsf{T}\mathbf{A} - \mu\mathbf{B}^\mathsf{T}\mathbf{B}$ is positive semidefinite.

Abe et al. (2019) proposed a linearly involved convexity-preserving model by composing a linear operator with the GMC penalty. The linearly involved GMC model is:

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}}\, F_{\mathbf{L}}(\mathbf{x}) = \frac{1}{2}\|\mathbf{y}-\mathbf{A}\mathbf{x}\|_2^2 + \mu\psi_{\mathbf{B}}\circ\mathbf{L}(\mathbf{x}), \tag{3}$$

where $\mathbf{L} \in \mathbb{R}^{l\times n}$ is a linear operator and $\mathbf{B} \in \mathbb{R}^{m\times l}$ is a matrix parameter. Here we call $\psi_{\mathbf{B}}\circ\mathbf{L}$ the linearly involved GMC penalty which includes the original GMC penalty (2) as a special case by taking $\mathbf{L}$ as the identity operator. When the signal is supposed to be piece-wise constant, a typical choice for $\mathbf{L}$ is the first-order difference operator defined as

$$\mathbf{D}_n^{(1)} = \begin{bmatrix} -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \end{bmatrix} \in \mathbb{R}^{(n-1)\times n}. \tag{4}$$

Abe et al. (2019) states that problem (3) maintains its convexity when $\mathbf{B}$ satisfies the following convexity-preserving condition:

$$\mathbf{A}^\mathsf{T}\mathbf{A} - \mu\mathbf{L}^\mathsf{T}\mathbf{B}^\mathsf{T}\mathbf{B}\mathbf{L} \succeq \mathbf{O}_n. \tag{5}$$

There are basically three issues in the context of the GMC penalization method. First, guidance on setting the matrix parameter $\mathbf{B}$ for the GMC family of penalties is still limited. Selesnick (2017a) chose to set $\mathbf{B} = \sqrt{\theta/\mu}\mathbf{A}$ with $\theta \in (0,1)$ for the original GMC penalty. Abe et al. (2019) provided a method based on Schur complement to find a feasible matrix $\mathbf{B}$ for the linearly involved GMC penalty, which only applies to $\mathbf{L}$ with full row rank. Two recent papers on matrix-parametric nonconvex penalization methods also discussed how to set the matrix parameter in their models. Lanza et al. (2019) developed a general parametric nonconvex nonseparable regularizer for image processing and they produced the matrix $\mathbf{B}$ by rescaling the eigenvalues of $\mathbf{A}^\mathsf{T}\mathbf{A}$. Selesnick et al. (2020) introduced the generalized Moreau envelope for the nonconvex total variation regularization and they set the matrix $\mathbf{B}$ as a filter with $\mathbf{L} = \mathbf{D}_n^{(1)}$, but prior information and additional parameters need to be specified in that case. Second, computation of the GMC family of optimization problems warrants further exploration. Selesnick (2017a) applied the forward-backward (FB) algorithm to compute the solution of the original GMC model. Lanza et al. (2019) and Selesnick et al. (2020) also used the FB algorithm and its variant to solve the corresponding optimization problems. Unfortunately, the FB algorithm is no longer applicable to the linearly involved GMC problem due to the existence of the linear operator. Third, discussion on how to select the crucial tuning parameter μ in the current literature is rare. All the references cited above manually set μ, which is not efficient or robust in practice. These existing issues motivate the work in this paper.

This paper focuses on the linearly involved convexity-preserving model (3) and makes the following contributions:

1. Firstly, and most importantly, we propose a new method to set a matrix $\mathbf{B}$ satisfying condition (5). We focus on producing the matrix $\mu\mathbf{B}^\mathsf{T}\mathbf{B}$, instead of $\mathbf{B}$. We formulate the task of producing $\mu\mathbf{B}^\mathsf{T}\mathbf{B}$ as a feasibility problem and present two algorithms to solve the problem. The advantage of our method is that it applies to $\mathbf{L}$ of any structure and can find different feasible $\mathbf{B}$ matrices by setting different initial values for the algorithms.

2. We provide a fast algorithm, primal-dual hybrid gradient (PDHG), to solve problem (3).

3. We prove theoretical properties of the solution path of (3). These properties in turn expedite the computation in the numerical experiments and provide guidance on selecting the tuning parameter μ.

Note that since the original GMC penalty (2) is a special case of the linearly involved GMC penalty (3), all the contributions in this work also apply to the original GMC method.

The rest of this paper is organized as follows. In section 2, we introduce our method to compute matrix $\mathbf{B}$, or $\mu\mathbf{B}^\mathsf{T}\mathbf{B}$ to be exact, to ensure the convexity of problem (3). We briefly review the CQ and ADMM algorithms and then derive specific instances of each of them to compute a feasible $\mathbf{B}$. In section 3, we pose problem (3) as a saddle-point problem and apply the PDHG algorithm to compute the solution. In section 4, we characterize the solution path theoretically. In section 5, we present two numerical examples to show the advantage of our new method of computing $\mathbf{B}$ as well as to highlight performance gains of the linearly involved GMC model in signal recovery. In section 6, we summarize our work and discuss future directions. Proofs are included in the appendix.

2. Algorithms for computing matrix B

Throughout, vectors are denoted by boldface lowercase letters (e.g. $\mathbf{x} \in \mathbb{R}^n$) and matrices by boldface capital letters (e.g. $\mathbf{A} \in \mathbb{R}^{m\times n}$). $\mathbf{O}_{m\times n} \in \mathbb{R}^{m\times n}$ and $\mathbf{O}_n \in \mathbb{R}^{n\times n}$ stands for the zero matrices and $\mathbf{I}_n \in \mathbb{R}^{n\times n}$ the identity matrix. For a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, $\mathbf{A}^\dagger \in \mathbb{R}^{n\times m}$ represents the Moore–Penrose inverse of $\mathbf{A}$, and $\text{vec}(\mathbf{A})$ denotes the column-major vectorization of $\mathbf{A}$, namely the vector obtained by stacking the columns of $\mathbf{A}$ on top of each other. We denote the i-th row of $\mathbf{A}$ by $\mathbf{A}_{i:}$ and the j-th column of $\mathbf{A}$ by $\mathbf{A}_{:j}$. The inner product between two matrices is defined as $\langle\mathbf{A},\mathbf{B}\rangle = \text{vec}(\mathbf{A})^\mathsf{T}\text{vec}(\mathbf{B})$. For a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n\times n}$, $\lambda_j(\mathbf{A})$ indicates the j-th largest eigenvalue of $\mathbf{A}$ for $1 \le j \le n$. $\mathbf{A} \succeq \mathbf{O}_n$ means $\mathbf{A}$ is positive semidefinite and $\mathbf{A} \succ \mathbf{O}_n$ indicates the positive definiteness of $\mathbf{A}$.

For model (3), the convexity-preserving condition (5) imposes a constraint on the matrix parameter $\mathbf{B}$, or equivalently on $\mu\mathbf{B}^\mathsf{T}\mathbf{B}$. We first focus our attention on finding a matrix $\mathbf{Z} = \mu\mathbf{B}^\mathsf{T}\mathbf{B} \in \mathbb{R}^{l\times l}$ satisfying condition (5). It is straightforward to verify that identifying a matrix $\mathbf{B}$ satisfying condition (5) is equivalent to identifying a pair of symmetric matrices $\mathbf{Z}$ and $\mathbf{D}$ satisfying the following three conditions

$$\mathbf{Z} \succeq \mathbf{O}_l, \quad \mathbf{L}^\mathsf{T}\mathbf{Z}\mathbf{L} = \mathbf{D}, \quad \text{and} \quad \mathbf{A}^\mathsf{T}\mathbf{A} - \mathbf{D} \succeq \mathbf{O}_n. \tag{6}$$

The singular value decomposition of a matrix $\mathbf{Z}$ satisfying (6) can be used to construct a matrix $\mathbf{B}$. We will see, however, that we do not need to explicitly construct $\mathbf{B}$. We next discuss two algorithms to obtain a feasible $\mathbf{Z}$ matrix.

2.1. CQ algorithm

Let C and Q be nonempty closed convex sets in $\mathbb{R}^N$ and $\mathbb{R}^M$, respectively, and $\mathbf{H} \in \mathbb{R}^{M\times N}$. The split feasibility problem seeks to find a point $\mathbf{z}$ such that $\mathbf{z} \in C$ and $\mathbf{Hz} \in Q$, if such a $\mathbf{z}$ exists. We can formulate the split feasibility problem as the following constrained optimization problem

$$\text{minimize } f(\mathbf{z}) = \frac{1}{2}\|\mathbf{Hz} - \mathcal{P}_Q(\mathbf{Hz})\|_2^2 \quad \text{subject to } \mathbf{z} \in C. \tag{7}$$

Noting that the objective function f is differentiable with gradient $\nabla f(\mathbf{z}) = \mathbf{H}^\mathsf{T}\mathbf{Hz} - \mathbf{H}^\mathsf{T}\mathcal{P}_Q(\mathbf{Hz})$, Byrne (2002) proposed the CQ algorithm to solve the split feasibility problem by solving (7) with projected gradient descent.

We now formulate the task of computing a matrix $\mathbf{Z}$ satisfying (6) as a split feasibility problem. Let C be the set of matrices $\mathbf{Z} \in \mathbb{R}^{l\times l}$ that are positive semidefinite and Q be the set of symmetric matrices $\mathbf{D} \in \mathbb{R}^{n\times n}$ such that $\lambda_j(\mathbf{A}^\mathsf{T}\mathbf{A} - \mathbf{D}) \geq \sigma_\theta$ for $1 \leq j \leq n$, where $\sigma_\theta = (1-\theta)\lambda_n(\mathbf{A}^\mathsf{T}\mathbf{A}) \geq 0$ and $\theta \in (0, 1)$. If $\mathbf{Z}$ is a solution to the following split feasibility problem

$$\text{Find } \mathbf{Z} \in C \text{ with } \mathbf{L}^\mathsf{T}\mathbf{ZL} \in Q, \tag{8}$$

then $\mathbf{Z}$ satisfies (6). Note that we use $\sigma_\theta = (1-\theta)\lambda_n(\mathbf{A}^\mathsf{T}\mathbf{A})$, instead of 0, to lower bound the eigenvalues of $\mathbf{A}^\mathsf{T}\mathbf{A} - \mathbf{L}^\mathsf{T}\mathbf{ZL}$ to control the convexity of model (3) by the so-called convexity-preserving parameter θ.

It is well known that C is a closed and convex set. The proof of the closedness and convexity of Q is given in the appendix. Since both C and Q are closed convex sets, we can apply the CQ algorithm to compute a target $\mathbf{Z}$ matrix. We rewrite problem (8) in terms of the vectorization of $\mathbf{Z}$, denoted by $\mathbf{z}$, as follows:

$$\text{Find } \mathbf{z} \in C \text{ with } \mathbf{Hz} \in Q, \tag{9}$$

where $\mathbf{H} = \mathbf{L}^\mathsf{T} \otimes \mathbf{L}^\mathsf{T}$ and $\otimes$ denotes the Kronecker product between two matrices. This follows from the fact that $\text{vec}(\mathbf{L}^\mathsf{T}\mathbf{ZL}) = \mathbf{L}^\mathsf{T} \otimes \mathbf{L}^\mathsf{T}\,\text{vec}(\mathbf{Z})$. For notational simplicity, we overload our definition of the sets C and Q in (8) and (9), namely $\mathbf{z} \in C$ is equivalent to $\mathbf{Z} \in C$. We solve problem (9) with the CQ algorithm, specifically an accelerated version of the CQ algorithm (Qu and Xiu, 2005).

The CQ algorithm requires projections onto the sets C and Q. The projection of $\mathbf{Z}$ onto C is given by $\mathcal{P}_C(\mathbf{Z}) = \mathbf{U}\mathbf{T}(\mathbf{\Lambda}, 0)\mathbf{U}^\mathsf{T}$, where $\mathbf{U\Lambda U}^\mathsf{T}$ is the eigendecomposition of $\mathbf{Z}$ and $\mathbf{T}(\mathbf{\Lambda}, \sigma)$ is the diagonal matrix whose ith diagonal entry is the projection of λ_i onto the interval $[\sigma, \infty)$, namely $\max\{\lambda_i, \sigma\}$. The projection onto Q is given in proposition 1.

Proposition 1. *Let* $\mathbf{D}$ *be a symmetric matrix and the matrix* $\mathbf{A}^\mathsf{T}\mathbf{A} - \mathbf{D}$ *have eigendecomposition* $\mathbf{U\Lambda U}^\mathsf{T}$. *Then the projection of* $\mathbf{D}$ *onto* Q *is given by*

$$\mathcal{P}_Q(\mathbf{D}) = \mathbf{A}^\mathsf{T}\mathbf{A} - \mathbf{U}\mathbf{T}(\mathbf{\Lambda}, \sigma_\theta)\mathbf{U}^\mathsf{T}.$$

Algorithm 1 summarizes the CQ algorithm for computing $\mathbf{Z}$ in (8). Here, $\gamma \in (0, \frac{2}{\lambda_1(\mathbf{H})})$ with $\mathbf{H} = \mathbf{L}^\mathsf{T} \otimes \mathbf{L}^\mathsf{T}$.

Algorithm 1 CQ algorithm for computing matrix $\mathbf{Z}$

```
Initialize Z_0 ∈ R^{l×l} and set z_0 = vec(Z_0)
Initialize β, γ ∈ (0, 1) and α_0 ∈ (0, ∞)
α_min ← min(α_0, γ)
for k = 1, 2, ... do
    α_k ← α_0
    z^+ ← P_C(z_k − α_k ∇f(z_k))
    while f(z^+) ≥ γ f(z_k) − ∇f(z_k)^T (z_k − z^+) do
        α_k ← max(β α_k, α_min)
        z^+ ← P_C(z_k − α_k ∇f(z_k))
    end while
    z_{k+1} ← z^+
end for
```

2.2. ADMM algorithm

The ADMM algorithm (or specifically, parallel projections) proposed by Boyd et al. (2011) is a classical method to solve the convexity feasibility problem, which is to find a point in the intersection of a collection of nonempty closed convex sets.

Let us define three sets in the space $\Theta = \mathbb{R}^{l\times l} \times \mathbb{R}^{n\times n}$.

$$\begin{aligned} G &= \left\{\{\mathbf{Z}, \mathbf{D}\} \in \Theta : \mathbf{D} = \mathbf{L}^\mathsf{T}\mathbf{ZL}\right\}, \\ C' &= \{\{\mathbf{Z}, \mathbf{D}\} \in \Theta : \mathbf{Z} \in C\}, \\ Q' &= \{\{\mathbf{Z}, \mathbf{D}\} \in \Theta : \mathbf{D} \in Q\}. \end{aligned}$$

The sets C and Q are the ones defined in (8). It is straightforward to verify that G, C' and Q' are nonempty, closed and convex. Then finding a matrix $\mathbf{Z}$ satisfying (6) is equivalent to finding a point in the intersection of C', G and Q'. The ADMM algorithm requires projections onto C', G and Q'. The projection onto G is given in Proposition 2.

Proposition 2. *Let* $\mathbf{H} = \mathbf{L}^\mathsf{T} \otimes \mathbf{L}^\mathsf{T}$. *Then the projection of* $\{\mathbf{Z}, \mathbf{D}\} \in \Theta$ *onto* G *is given by*

$$\begin{aligned} \mathcal{P}_G(\mathbf{Z}, \mathbf{D}) &= \left\{\hat{\mathbf{Z}}, \mathbf{L}^\mathsf{T}\hat{\mathbf{Z}}\mathbf{L}\right\} \\ \hat{\mathbf{z}} &= \left(\mathbf{H}^\mathsf{T}\mathbf{H} + \mathbf{I}_{l^2}\right)^{-1}\left(\mathbf{H}^\mathsf{T}\mathbf{d} + \mathbf{z}\right), \end{aligned}$$

where $\mathbf{z} = \text{vec}(\mathbf{Z}), \mathbf{d} = \text{vec}(\mathbf{D})$, *and* $\hat{\mathbf{z}} = \text{vec}(\hat{\mathbf{Z}})$.

We close this section by noting that both CQ and ADMM have no assumption on the structure of $\mathbf{L}$, meaning that the two algorithms are very general and applicable to any $\mathbf{L}$. This overcomes the drawback of the existing method to compute $\mathbf{B}$ in Abe et al. (2019) that it only applies to $\mathbf{L}$ with full row rank.

3. Algorithm for the linearly involved GMC model

The linearly involved GMC model (3) can be written as the following saddle-point problem

$$\min_{\mathbf{x}\in\mathbb{R}^n} \max_{\mathbf{v}\in\mathbb{R}^l} \; f(\mathbf{x}) + \mathbf{v}^\mathsf{T}\mathbf{Px} - g(\mathbf{v}), \tag{10}$$

where $f(\mathbf{x}) = \frac{1}{2}\|\mathbf{y} - \mathbf{Ax}\|_2^2 - \frac{\mu}{2}\|\mathbf{BLx}\|_2^2 + \mu\|\mathbf{Lx}\|_1$ is convex under the convexity-preserving condition (5), $g(\mathbf{v}) = \frac{\mu}{2}\|\mathbf{Bv}\|_2^2 + \mu\|\mathbf{v}\|_1$ is convex in general, and $\mathbf{P} = \mu\mathbf{B}^\mathsf{T}\mathbf{BL} = \mathbf{ZL}$ is a linear operator.

We propose using the PDHG (Goldstein et al., 2013, 2015) algorithm to solve (10) to take advantage of the special structure of the method we have introduced for computing $\mathbf{B}$. Note that using the matrix $\mathbf{Z}$, rather than $\mathbf{B}$, avoids redundant computations. The GMC penalty (2) depends on $\mathbf{B}^\mathsf{T}\mathbf{B}$, not $\mathbf{B}$ itself (Selesnick, 2017a). Thus problem (3) only depends on $\mu\mathbf{B}^\mathsf{T}\mathbf{B}$, which is exactly the $\mathbf{Z}$ matrix computed by the CQ or the ADMM algorithm. We get further computational gains by using an adaptive PDHG algorithm to speed up the convergence (Goldstein et al., 2013, 2015).

4. Properties of the solution path

A key practical question is how to choose the tuning parameter μ or, in other words, perform model selection. The extended Bayesian Information Criterion (eBIC) proposed in Chen and Chen (2008, 2012) is commonly used for tuning parameter selection in fitting structured sparse models as in (3). In this paper, we calculate eBIC on a grid of values of μ and select the μ corresponding to the smallest value of eBIC. As we need to solve a sequence of optimization problems, we next describe two properties of the solution to the optimization problem (3) as a function of the tuning parameter μ that can help expedite our model selection step. We write a minimizer to (3) as $\mathbf{x}^\star(\mu)$ to denote this dependency.

The first property establishes under modest assumptions that $\mathbf{x}^\star(\mu)$ not only exists and is unique but is also a continuous function of μ. In other words, $\mathbf{x}^\star(\mu)$ traces out a continuous solution path as μ varies. Note that Proposition 8 in Lanza et al. (2019) claims the existence and uniqueness of the solution to a convex nonconvex model, which includes model (3) in this paper. We emphasize, however, that our main contribution is actually the continuity of the solution path, which is not discussed in Lanza et al. (2019) but quite important for model selection.

Theorem 1. *Suppose* $\mathbf{A}^\mathsf{T}\mathbf{A} - \mu\mathbf{L}^\mathsf{T}\mathbf{B}^\mathsf{T}\mathbf{B}\mathbf{L} \succ \mathbf{O}_n$, *then the solution* $\mathbf{x}^\star(\mu)$ *to* (3) *exists, is unique, and is continuous in* μ.

The existence and uniqueness of $\mathbf{x}^\star(\mu)$ indicate that problem (3) is well-posed. The continuity of $\mathbf{x}^\star(\mu)$ suggests a homotopy strategy for reducing computation time in computing $\mathbf{x}^\star(\mu_i)$ for $i = 1, \ldots, N$ and $\mu_1 \leq \mu_2 \leq \mu_3 \leq \cdots \leq \mu_N$. Suppose we have already computed $\mathbf{x}^\star(\mu_1), \ldots, \mathbf{x}^\star(\mu_i)$ and wish to compute $\mathbf{x}^\star(\mu_{i+1})$. Since $\mathbf{x}^\star(\mu)$ is continuous in μ, we can initialize, or warm start, our PHDG algorithm at $\mathbf{x}^\star(\mu_i)$ which will likely be closer to $\mathbf{x}^\star(\mu_{i+1})$ than a random initialization.

The second property identifies a maximum value of μ of interest for performing model selection. One may intuitively expect that for sufficiently large μ that a solution to (3) coincides with a solution to the constrained optimization problem

$$\text{minimize } \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 \quad \text{subject to} \quad \mathbf{L}\mathbf{x} = \mathbf{0}. \tag{11}$$

This is indeed the case.

Theorem 2. *Let* $\tilde{\mathbf{x}}$ *be a solution to* (11). *Then* $F_\mathbf{L}(\mathbf{x})$ *in* (3) *is minimized by* $\tilde{\mathbf{x}}$ *for all* μ *greater than*

$$\mu_0 = \|\mathbf{A}^\mathsf{T}(\mathbf{A}\tilde{\mathbf{x}} - \mathbf{y})\|_2 / \sigma_{\min}(\mathbf{L}),$$

where $\sigma_{\min}(\mathbf{L})$ *is the smallest singular value of* $\mathbf{L}$.

For example, if $\mathbf{L}$ is the first-order difference operator, Theorem 2 tells us that $\mathbf{x}^\star(\mu) = \frac{\mathbf{y}^\mathsf{T}\mathbf{A}\mathbf{1}}{\mathbf{1}^\mathsf{T}\mathbf{A}^\mathsf{T}\mathbf{A}\mathbf{1}}\mathbf{1}$ for $\mu > \mu_0$.

The practical value of Theorem 2 is that when choosing μ, it is sufficient to only compute $\mathbf{x}^\star(\mu)$ for μ between 0 and μ_0 to sample the full dynamic range between least smoothed and most smoothed values of $\mathbf{x}^\star(\mu)$. Computing μ_0 is straightforward since computing $\tilde{\mathbf{x}}$ requires solving a linear system. This follows from the fact that (11) is an equality constrained convex quadratic program.

We close this section by noting that the two properties of $\mathbf{x}^\star(\mu)$ are inherent to the optimization problem (3) and consequently independent of how $\mathbf{x}^\star(\mu)$ is computed.

5. Numerical experiments

We present two simulated examples to investigate the performance of our new method to compute the matrix $\mathbf{B}$ and the effectiveness of the linearly involved GMC penalty. The first example is a 1-D signal recovery problem where $\mathbf{L}$ is the first-order difference operator (4). In this case, $\mathbf{L}$ has full row rank, therefore both the CQ and ADMM algorithms as well as the existing method in Abe et al. (2019) can be applied to compute $\mathbf{B}$. We will show that CQ and ADMM enjoy some advantages over the existing method. The second example is a matrix regression problem where $\mathbf{L}$ has more rows than columns. As a result, the existing method in Abe et al. (2019) is no longer applicable, but both CQ and ADMM can be applied without any limitation. In both examples, we present the recovery performance of the linearly involved GMC method in different aspects.

5.1. 1-D signal recovery

This example applies the linearly involved GMC penalty on 1-D signal recovery by considering $\mathbf{L}$ as the first-order difference operator. The true signal $\boldsymbol{x} \in \mathbb{R}^{128}$ is piece-wise constant and generated by the `MakeSignal` function in the R package *rwavelet*. The observed data $\boldsymbol{y} \in \mathbb{R}^{100}$ is generated by $\boldsymbol{y} = \mathbf{A}\boldsymbol{x} + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \in \mathbb{R}^{100}$ is additive white Gaussion noise with unit variance, and entries of $\mathbf{A}$ are i.i.d. standard normal.

We first focus on the computation of the matrix $\mathbf{B}$. Since $\mathbf{L}$ has full row rank, we use three methods to produce a feasible matrix $\mathbf{B}$ for the linearly involved GMC penalty: the existing method in Abe et al. (2019), the CQ and ADMM algorithms in this paper. For a fixed $\theta = 0.9$, we set the initial values $\mathbf{Z}_0 = \theta(\mathbf{L}^\dagger)^\mathsf{T}\mathbf{A}^\mathsf{T}\mathbf{A}\mathbf{L}^\dagger$ and $\mathbf{D}_0 = \mathbf{L}^\mathsf{T}\mathbf{Z}_0\mathbf{L}$ for CQ and ADMM. It turns out that both CQ and ADMM obtain the same matrix $\mathbf{B}$ as that from the method in Abe et al. (2019).

After obtaining a feasible matrix $\mathbf{B}$, we then move to recover the piece-wise constant signal by the linearly involved GMC penalty. We measure the accuracy of the recovered signal $\hat{\mathbf{x}}$ by the mean squared error (MSE) defined as the average of $\text{SE} = \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2$ over 20 independent replicates. We compare the performance of the linearly involved GMC with the TV regularization.

We first discuss the dependency of the recovery performance on the tuning parameter μ. Figure 1 displays the log(MSE) as a function of $\log(\mu)$. It can be seen that (i) the linearly involved GMC can achieve better accuracy (a lower MSE) than the TV

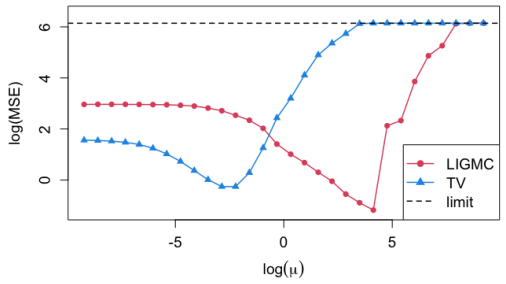

Figure 1. log(**MSE**) **as a function of** log(μ)**. The limit of** log(**MSE**) **is obtained by the** $\bar{\mathbf{x}}$ **in Theorem 2.**

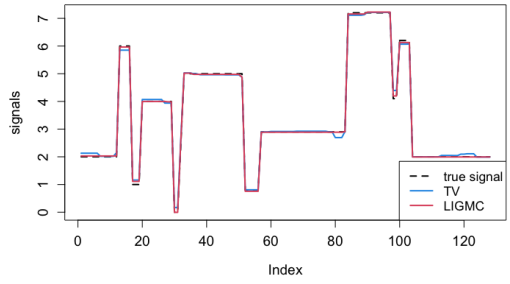

Figure 2. The true piece-wise constant signal (in black) and the recovered signals from the TV regularization (in blue) and the linearly involved GMC method (in red).

method; (ii) to obtain the best recovery, the linearly involved GMC prefers a larger μ than the TV method; (iii) log(MSE) indeed goes to the constant limit as predicted in Theorem 2.

Next, to further illustrate the effectiveness of the linearly involved GMC penalty, Figure 2 compares the recovered signals (selected by eBIC) from the linearly involved GMC method and the TV regularization. From Figure 2, we observe that (i) the linearly involved GMC can recover the "jumps" in the signal better than the TV method, which means that the linearly involved GMC estimator has less estimation bias on larger magnitude components of the signal difference; (ii) the linearly involved GMC recovers smoother signals over the constant segments than the TV method, which indicates that the linearly involved GMC tends to recover signals with sparser discrete first derivatives.

Finally, we emphasize the advantage of our method to compute the matrix $\mathbf{B}$. We have shown that given specific initial values, both CQ and ADMM can produce the same $\mathbf{B}$ computed from the existing method in Abe et al. (2019). Since the feasible matrix $\mathbf{B}$ is not necessarily unique, we can get different feasible $\mathbf{B}$ matrices by varying the initializations for CQ and ADMM. For instance, in this experiment, if we set $\mathbf{Z}_0 = \theta\mathbf{B}_0^\top\mathbf{B}_0$ and $\mathbf{D}_0 = \mathbf{L}^\top\mathbf{Z}_0\mathbf{L}$ where entries of $\mathbf{B}_0$ are i.i.d. standard normally distributed, we can get two different $\mathbf{B}$ matrices from the CQ and ADMM algorithms, and both of them are different from the one that is produced by the existing method. When we plug the three different $\mathbf{B}$ matrices into the linearly involved GMC model, we find that the matrix $\mathbf{B}$ computed by the ADMM algorithm gives a slightly lower MSE than the other two. Due to space limitations, we have not included plots of the MSE or recovered signals here. Nonetheless, it is still an open question how to choose the matrix $\mathbf{B}$ subject to the convexity preserving condition to maximize accuracy in the recovered signal. Our algorithms, which have the flexibility of producing different feasible $\mathbf{B}$ matrices by setting different initial values, can provide a way to explore alternative choices of $\mathbf{B}$ that can give insight into how to best design $\mathbf{B}$.

5.2. *Matrix regression*

In this experiment, we consider the linearly involved GMC method for matrix regression (Wang et al., 2017). Matrix regression is a generalization of the classical linear regression, where covariates are matrices. A typical matrix regression model is

$$y = \langle \mathbf{X}, \mathbf{S} \rangle + \epsilon, \tag{12}$$

where y is the response, $\mathbf{X} \in \mathbb{R}^{p\times q}$ is the matrix covariate, $\mathbf{S} \in \mathbb{R}^{p\times q}$ is the coefficient matrix, and ϵ is the additive white Gaussian noise. Zhou and Li (2014) investigated a class of regularized matrix regression methods based on spectral regularization to estimate signal $\mathbf{S}$ of a low rank structure. When the signal $\mathbf{S}$ is piecewiese constant, we can apply the TV regularizer on both rows and columns of $\mathbf{S}$, which leads to

$$\underset{\mathbf{S}\in\mathbb{R}^{p\times q}}{\text{minimize}}\ \frac{1}{2}\sum_{i=1}^{m}(\mathbf{y}_i - \langle \mathbf{X}_i, \mathbf{S}\rangle)^2 + \mu\left(\sum_{j=1}^{q}\|\mathbf{D}_p^{(1)}\mathbf{S}_{:j}\|_1 + \sum_{k=1}^{p}\|\mathbf{D}_q^{(1)}\mathbf{S}_{i:}\|_1\right), \tag{13}$$

where $\mathbf{D}_p^{(1)}$ and $\mathbf{D}_q^{(1)}$ are the first-order difference operators. We can rewrite (13) as a regularized least squares problem

$$\underset{\mathbf{s}\in\mathbb{R}^{pq}}{\text{minimize}}\ \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{s}\|_2^2 + \mu\|\mathbf{L}\mathbf{s}\|_1, \tag{14}$$

where $\mathbf{y} \in \mathbb{R}^m$ is the response vector, $\mathbf{s} = \text{vec}\,(\mathbf{S})$ is the vectorization of $\mathbf{S}$, and $\mathbf{A} \in \mathbb{R}^{m\times pq}$ whose i-th row is $\text{vec}\,(\mathbf{X}_i)$. Here the linear operator $\mathbf{L}$ is given by

$$\mathbf{L} = \begin{pmatrix} \mathbf{I}_q \otimes \mathbf{D}_p^{(1)} \\ \mathbf{D}_q^{(1)} \otimes \mathbf{I}_p \end{pmatrix}$$

Obviously, we can apply the linearly involved GMC penalty on this regularized matrix regression problem (14) to recover the signal $\mathbf{S}$. An important point here is that the number of rows of $\mathbf{L}$, $p(q-1) + q(p-1)$, is larger than its number of columns, pq. Consequently, $\mathbf{L}$ cannot have full row rank thus the existing method in Abe et al. (2019) to compute $\mathbf{B}$ for the linearly involved GMC cannot be used in this case.

We use two different image signals, as shown in Figure 3. For each case, the true signal $\mathbf{S} \in \mathbb{R}^{20\times 20}$ has binary entries 1 (black) and -1 (white). The vector of observations $\mathbf{y} \in \mathbb{R}^{80}$ is generated by $y_i = \text{vec}(\mathbf{X}_i)^\top \text{vec}(\mathbf{S}) + \epsilon_i$, where $\mathbf{X}_i$ are covariate matrices whose elements are standard normal and $\boldsymbol{\epsilon} \in \mathbb{R}^{80}$ is the additive white Gaussian noise with standard deviation σ.

Regarding the computation of the matrix $\mathbf{B}$ for the linearly involved GMC penalty, both CQ and ADMM are applicable here since they both have no restriction on the structure of $\mathbf{L}$. We again set $\theta = 0.9$ and $\mathbf{Z}_0 = \theta(\mathbf{L}^\dagger)^\top\mathbf{A}^\top\mathbf{A}\mathbf{L}^\dagger$, $\mathbf{D}_0 = \mathbf{L}^\top\mathbf{Z}_0\mathbf{L}$ to initialize CQ and ADMM and finally obtain the same $\mathbf{B}$ from the two algorithms.

In terms of the recovery performance, we use the squared error defined as $\text{SE} = \|\hat{\mathbf{S}} - \mathbf{S}\|_{\text{F}}^2$ as a metric of the accuracy of the recovered image $\hat{\mathbf{S}}$. We conduct the same analysis of how the

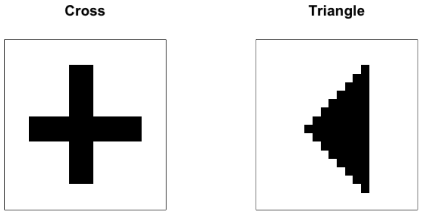

Figure 3. The true image signals.

MSE varies with the tuning parameter μ for a fixed noise level $\sigma = 1$. The conclusions from both the cross and triangle signals are similar with what we obtained in Section 5.1. That is, the linearly involved GMC method reaches a lower MSE at a larger value of μ in comparison with the TV method. Our main goal in this simulation example is to investigate how the linearly involved GMC penalty and the TV regularizer perform at different noise levels. To that end, we set a sequence of values for the standard deviation of noise $\sigma \in \{0.1, 0.5, 1, 2, 3\}$. At each noise level, we set a grid of values of μ to obtain the solution paths from the two methods and apply the eBIC for model selection. Figure 4 displays the recovery performance of the two different methods at different noise levels. The results show that (i) the linearly involved GMC method consistently outperforms the TV method at different noise levels for both cross and triangle signals; (ii) both methods perform worse as the noise level increases, but the performance of the TV method is affected by the noise level more severely.

6. Discussion

For the linearly involved GMC model (3), we proposed a new method to produce the matrix parameter $\mathbf{B}$ satisfying condition (5). That is, we first produce a feasible matrix $\mathbf{Z} = \mu\mathbf{B}^{\mathsf{T}}\mathbf{B}$ and then compute $\mathbf{B}$ by the singular value decomposition of $\mathbf{Z}$. We formulated computing $\mathbf{Z}$ as solving a feasibility problem and offered two algorithms, CQ and ADMM, to solve the problem. Then, we cast problem (3) as a saddle-point problem and adopted the PDHG algorithm to compute the solution. We also provided guidance on the tuning parameter selection for model (3) via theoretical properties of its solution path.

We presented two numerical experiments on 1-D signal recovery and matrix regression where the linear operator $\mathbf{L}$ has different structures. The simulation results showed that both CQ and ADMM work well to produce a feasible matrix $\mathbf{B}$. Unlike the existing method in Abe et al. (2019), they do not have any requirement on the structure of the linear operator $\mathbf{L}$ and have the advantage of producing different $\mathbf{B}$ by setting different initial values. With respect to recovery performance, the linearly involved GMC penalty achieved better accuracy than the TV regularizer in both examples. Moreover, the linearly involved GMC method maintained this advantage in recovery accuracy across different noise levels.

There are several directions that we leave for further study. The matrix parameter in the GMC penalization method plays an important role in its recovery performance. One limitation of our work is on the discussion of choice of $\mathbf{B}$ or how to initialize the CQ and ADMM algorithms. Determining what type

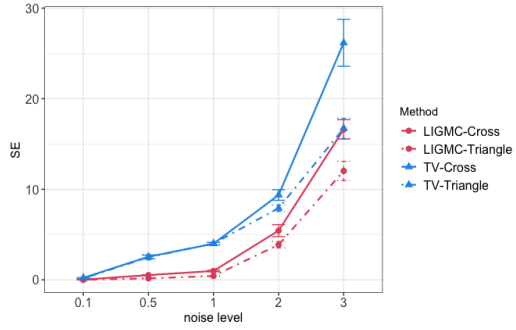

Figure 4. Average performance plus/minus one standard error based on 20 simulation replicates for each method

of $\mathbf{B}$ gives better recovery warrants further investigation. Additionally, the statistical properties, e.g. consistency and support recovery, of the linearly involved GMC estimator are left for future work. Finally, we note that the GMC penalization approach can be extended to other sparse learning problems in both signal processing and statistical learning.

Acknowledgments

This work was partially supported by the National Science Foundation [DMS-1752692].

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A

A.1. Proof of the closedness and convexity of Q set

The matrix set $Q = \{\mathbf{D} \in \mathbb{R}^{n\times n} | \mathbf{D}$ is symmetric with $\lambda_j(\mathbf{A}^{\mathsf{T}}\mathbf{A} - \mathbf{D}) \geq \sigma_\theta$ for $1 \leq j \leq n\}$. Here λ_j indicates the j-th largest eigenvalue and $\sigma_\theta = (1-\theta)\lambda_n(\mathbf{A}^{\mathsf{T}}\mathbf{A}) \geq 0$ with $\theta \in (0, 1)$ is a specified number.

First, to show Q is closed, suppose a sequence of matrices $(\mathbf{D}_k)_{k\in\mathbf{N}} \in Q$, and suppose $\mathbf{D}_k \to \mathbf{D}$ as $k \to \infty$, that is $\|\mathbf{D}_k - \mathbf{D}\|_{\mathrm{F}} \to 0$ as $k \to \infty$. We only need to show $\mathbf{D} \in Q$.

Let $\mathbf{M}_k = \mathbf{A}^{\mathsf{T}}\mathbf{A} - \mathbf{D}_k$, then we have $\lambda_j(\mathbf{M}_k) \geq \sigma_\theta$ for all $1 \leq j \leq n$. Denote $\mathbf{M} = \mathbf{A}^{\mathsf{T}}\mathbf{A} - \mathbf{D}$, then $\mathbf{M}_k \to \mathbf{M}$ since $\mathbf{D}_k \to \mathbf{D}$ as $k \to \infty$. Since $\mathbf{M} \in \mathbb{R}^{n\times n}$ is symmetric, it has eigen decompostion and there exists a unit Eulidean length vector $\mathbf{u}$ such that $\mathbf{u}^{\mathsf{T}}\mathbf{M}\mathbf{u} = \lambda_n(\mathbf{M})$. Let $g : \mathbb{R}^{n\times n} \to \mathbb{R}$, $g(\mathbf{X}) = \mathbf{u}^{\mathsf{T}}\mathbf{X}\mathbf{u}$, then $g(\mathbf{M}) = \lambda_n(\mathbf{M})$, and $\sigma_\theta \leq \lambda_n(\mathbf{M}_k) \leq \mathbf{u}^{\mathsf{T}}\mathbf{M}_k\mathbf{u} = g(\mathbf{M}_k)$. Note than g is an affine thus a continuous function, so $\lim_{k\to\infty} g(\mathbf{M}_k) = g(\mathbf{M}) = \lambda_n(\mathbf{M})$. Therefore, we get $\lambda_n(\mathbf{M}) \geq \sigma_\theta$, thus $\mathbf{D} \in Q$.

The convexity of Q can be proved by the definition of convex set. We leave it as an exercise for the reader.

A.2. Proof of Proposition 1

The projection of $\mathbf{D}$ onto the Q, $\hat{\mathbf{D}}$, is the solution of the optimization problem

$$\begin{aligned}\underset{\mathbf{X}\in Q}{\text{minimize}}\ \|\mathbf{D}-\mathbf{X}\|_{\mathrm{F}}^2 &= \underset{\mathbf{X}\in Q}{\text{minimize}}\ \|(\mathbf{A}^\mathsf{T}\mathbf{A}-\mathbf{D})-(\mathbf{A}^\mathsf{T}\mathbf{A}-\mathbf{X})\|_{\mathrm{F}}^2 \\ &= \underset{\mathbf{A}^\mathsf{T}\mathbf{A}-\mathbf{Y}\in Q}{\text{minimize}}\ \|\mathbf{M}-\mathbf{Y}\|_{\mathrm{F}}^2, \end{aligned} \tag{A.1}$$

where $\mathbf{M}=\mathbf{A}^\mathsf{T}\mathbf{A}-\mathbf{D}$. Since $\mathbf{Y}$ is symmetric, we can express $\mathbf{Y}$ with its eigendecomposition $\mathbf{Y}=\mathbf{V}\mathbf{\Sigma}\mathbf{V}^\mathsf{T}$. Moreover, $\lambda_j(\mathbf{Y})\geq\sigma_\theta$ since $\mathbf{A}^\mathsf{T}\mathbf{A}-\mathbf{Y}\in Q$. Then $\|\mathbf{M}-\mathbf{Y}\|_{\mathrm{F}}^2=\|\mathbf{V}^\mathsf{T}\mathbf{M}\mathbf{V}-\mathbf{\Sigma}\|_{\mathrm{F}}^2$. Let $\mathbf{B}=\mathbf{V}^\mathsf{T}\mathbf{M}\mathbf{V}$ and $\mathbf{B}_{ij}$ be the (i,j) element, then $\|\mathbf{M}-\mathbf{Y}\|_{\mathrm{F}}^2=\sum_{i\neq j}\mathbf{B}_{ij}^2+\sum_{i=1}^n(\mathbf{B}_{ii}-\mathbf{\Sigma}_{ii})^2$. Therefore, $\mathbf{Y}$ is the solution of (A.1) if and only if

$$\begin{cases}\mathbf{B}_{ij}=0, & 1\leq i\neq j\leq n\\ \mathbf{\Sigma}_{ii}=\max(\mathbf{B}_{ii},\sigma_\theta), & 1\leq i\leq n\end{cases}$$

Suppose $\mathbf{M}$ has the eigen decomposition $\mathbf{M}=\mathbf{U}\mathbf{\Lambda}\mathbf{U}^\mathsf{T}$, plugging it into $\mathbf{B}$, we get $\mathbf{V}^\mathsf{T}\mathbf{U}\mathbf{\Lambda}\mathbf{U}^\mathsf{T}\mathbf{V}$ is diagonal. Hence $\mathbf{U}=\mathbf{V}$ and $\mathbf{B}=\mathbf{\Lambda}$. So the solution $\mathbf{Y}=\mathbf{V}\mathbf{\Sigma}\mathbf{V}^\mathsf{T}=\mathbf{U}\mathbf{T}(\mathbf{\Lambda},\sigma_\theta)\mathbf{U}^\mathsf{T}$, where $\mathbf{T}(\mathbf{\Lambda},\sigma_\theta)$ is the diagonal matrix whose ith diagonal entry is $\max(\mathbf{\Lambda}_{ii},\sigma_\theta)$. Therefore, the projection of $\mathbf{D}$ onto Q is

$$\mathcal{P}_Q(\mathbf{D})=\mathbf{A}^\mathsf{T}\mathbf{A}-\mathbf{Y}=\mathbf{A}^\mathsf{T}\mathbf{A}-\mathbf{U}\mathbf{T}(\mathbf{\Lambda},\sigma_\theta)\mathbf{U}^\mathsf{T}.$$

A.3. Proof of Proposition 2

The projection of $(\mathbf{Z},\mathbf{D})$ onto the G, $\mathcal{P}_G(\mathbf{Z},\mathbf{D})$, is the solution of the optimization problem

$$\underset{(\mathbf{X},\mathbf{Y})\in G}{\text{minimize}}\ \|\mathbf{Z}-\mathbf{X}\|_{\mathrm{F}}^2+\|\mathbf{D}-\mathbf{Y}\|_{\mathrm{F}}^2=\underset{\mathbf{X}\in\mathbb{R}^{l\times l},\mathbf{Y}=\mathbf{L}^\mathsf{T}\mathbf{X}\mathbf{L}}{\text{minimize}}\ \|\mathbf{Z}-\mathbf{X}\|_{\mathrm{F}}^2+\|\mathbf{D}-\mathbf{L}^\mathsf{T}\mathbf{X}\mathbf{L}\|_{\mathrm{F}}^2. \tag{A.2}$$

Denote $\mathbf{z}=\mathrm{vec}(\mathbf{Z})$, $\mathbf{d}=\mathrm{vec}(\mathbf{D})$, $\mathbf{x}=\mathrm{vec}(\mathbf{X})$, and $\mathbf{H}=\mathbf{L}^\mathsf{T}\otimes\mathbf{L}^\mathsf{T}$, then (A.2) is equivalent to

$$\underset{\mathbf{x}\in\mathbb{R}^{l^2},\ \mathbf{y}=\mathbf{H}\mathbf{x}}{\text{minimize}}\ g(\mathbf{x})=\|\mathbf{z}-\mathbf{x}\|_2^2+\|\mathbf{d}-\mathbf{H}\mathbf{x}\|_2^2. \tag{A.3}$$

Since g is differentiable, $(\mathbf{x},\mathbf{y})$ solves (A.3) if and only if $\mathbf{x}$ is the solution to the following linear system of equations

$$(\mathbf{H}^\mathsf{T}\mathbf{H}+\mathbf{I}_{l^2})\mathbf{x}=\mathbf{H}^\mathsf{T}\mathbf{d}+\mathbf{z},$$

and $\mathbf{y}=\mathbf{H}\mathbf{x}$. Transforming the vector solutions into matrices, we get

$$\mathcal{P}_G(\mathbf{Z},\mathbf{D})=(\hat{\mathbf{Z}},\mathbf{L}^\mathsf{T}\hat{\mathbf{Z}}\mathbf{L}),$$

with $\mathrm{vec}(\hat{\mathbf{Z}})=\left(\mathbf{H}^\mathsf{T}\mathbf{H}+\mathbf{I}_{l^2}\right)^{-1}\left(\mathbf{H}^\mathsf{T}\mathbf{d}+\mathbf{z}\right)$.

A.4. Proof of Theorem 1

Recall that a coercive and strictly convex function has a unique minimizer on its domain. We rewrite $F_{\mathbf{L}}$ in (3) as

$$F_{\mathbf{L}}(\mathbf{x})=\frac{1}{2}\|\mathbf{y}-\mathbf{A}\mathbf{x}\|_2^2-\frac{\mu}{2}\|\mathbf{B}\mathbf{L}\mathbf{x}\|_2^2+\mu\|\mathbf{L}\mathbf{x}\|_1+\mu\max_{\mathbf{v}}g(\mathbf{x},\mathbf{v}), \tag{A.4}$$

where $g(\mathbf{x},\mathbf{v})=\langle\mathbf{L}^\mathsf{T}\mathbf{B}^\mathsf{T}\mathbf{B}\mathbf{v},\mathbf{x}\rangle-\frac{1}{2}\|\mathbf{B}\mathbf{v}\|_2^2-\|\mathbf{v}\|_1$ is affine in $\mathbf{x}$. The last term in (A.4) is convex as it is the pointwise maximum of a set of convex functions. Hence, $F_{\mathbf{L}}$ is strongly convex thus strictly convex when $\mathbf{A}^\mathsf{T}\mathbf{A}-\mu\mathbf{L}^\mathsf{T}\mathbf{B}^\mathsf{T}\mathbf{B}\mathbf{L}\succ\mathbf{O}_n$. It is straightforward to verify that $\lim_{\|x\|\to\infty}F_{\mathbf{L}}(\mathbf{x})=+\infty$, so $F_{\mathbf{L}}$ is coercive. Hence the solution to (3) exists and is unique.

Next, we prove the continuity of the solution path. Here, we write $F_{\mathbf{L}}$ as

$$F_{\mathbf{L}}(\mathbf{x})=\frac{1}{2}\|\mathbf{y}-\mathbf{A}\mathbf{x}\|_2^2+\mu\left(\|\mathbf{L}\mathbf{x}\|_1-\min_{\mathbf{v}}\|\mathbf{v}\|_1+\frac{1}{2}\|\mathbf{B}(\mathbf{L}\mathbf{x}-\mathbf{v})\|_2^2\right).$$

For $\forall(\mathbf{x},\mu)\in(\mathbb{R}^n,\mathbb{R}^+)$, $F_{\mathbf{L}}$ is a jointly continuous function of $(\mathbf{x},\mu)$ by the convexity of $F_{\mathbf{L}}$, and

$$\begin{aligned}&\|\mathbf{L}\mathbf{x}\|_1-\min_{\mathbf{v}}\|\mathbf{v}\|_1+\frac{1}{2}\|\mathbf{B}(\mathbf{L}\mathbf{x}-\mathbf{v})\|_2^2\\ &\geq\|\mathbf{L}\mathbf{x}\|_1-\|\mathbf{L}\mathbf{x}\|_1+\frac{1}{2}\|\mathbf{B}(\mathbf{L}\mathbf{x}-\mathbf{L}\mathbf{x})\|_2^2=0.\end{aligned}$$

Therefore, $F_{\mathbf{L}}$ is non-decreasing in μ. Hence, for any subinterval $[a,b]\subseteq(0,+\infty)$ and for $\forall\mu_0\in[a,b]$,

$$F_{\mathbf{L}}(\mathbf{x}(\mu_0),a)\leq F_{\mathbf{L}}(\mathbf{x}(\mu_0),\mu_0)\leq F_{\mathbf{L}}(\mathbf{0},\mu_0)\leq F_{\mathbf{L}}(\mathbf{0},b).$$

Since $F_{\mathbf{L}}(\mathbf{x},a)$ is coercive in $\mathbf{x}$, $S=\{\mathbf{x}:F_{\mathbf{L}}(\mathbf{x},a)\leq F_{\mathbf{L}}(\mathbf{0},b)\}$ is compact. Hence $\mathbf{x}(\mu_0)\in S$ is bounded for $\forall\mu_0\in[a,b]$. Suppose $\mathbf{x}(\mu)$ is not continuous at some point μ_0. Choose a,b such that $\mu_0\in[a,b]$, then there exists $\epsilon_0>0$ and a sequence $\{\mu_n\}_{n\in\mathbb{N}}\in[a,b]$ such that $\mu_n\to\mu_0$ but $\|\mathbf{x}(\mu_n)-\mathbf{x}(\mu_0)\|_2\geq\epsilon_0$ for $\forall n\in\mathbb{N}$. We have seen that for each n, $x(\mu_n)\in S$, a compact set, hence $\mathbf{x}(\mu_n)$ is a bounded sequence. So there exists a subsequence $\mathbf{x}(\mu_{n_k})$, such that $\mathbf{x}(\mu_{n_k})\to\tilde{\mathbf{x}}\in S$, and

$$F_{\mathbf{L}}(x(\mu_{n_k}),\mu_{n_k})\leq F_{\mathbf{L}}(x(\mu_0),\mu_{n_k}). \tag{A.5}$$

By the joint continuity of $F_{\mathbf{L}}$, taking limit on both sides of (A.5), we get $F_{\mathbf{L}}(\tilde{\mathbf{x}},\mu_0)\leq F_{\mathbf{L}}(\mathbf{x}(\mu_0),\mu_0)$. By the uniqueness of the solution path, $\tilde{\mathbf{x}}=\mathbf{x}(\mu_0)$, which contradicts with $\|\mathbf{x}(\mu_n)-\mathbf{x}(\mu_0)\|_2\geq\epsilon_0$ for $\forall n\in\mathbb{N}$. Therefore, the solution path $\mathbf{x}(\mu)$ is continuously dependent on μ.

A.5. Proof of Theorem 2

For the proof of Theorem 2, we need an auxiliary lemma, denoted as Lemma 1 in the following.

Lemma 1. *Suppose f is a convex and differentiable function, if the solution set $C=\arg\min_{\mathbf{x}:\mathbf{L}\mathbf{x}=\mathbf{0}}f(\mathbf{x})\neq\emptyset$, then $\forall\mathbf{x}^*\in C$, $\langle\nabla f(\mathbf{x}^*),\boldsymbol{\theta}\rangle=0$ holds for $\forall\ \boldsymbol{\theta}$ such that $\mathbf{L}\boldsymbol{\theta}=\mathbf{0}$. In particular, the conclusion holds for $f(\mathbf{x})=\frac{1}{2}\|\mathbf{y}-\mathbf{A}\mathbf{x}\|_2^2$.*

Proof of Lemma 1. Take $\mathbf{x}^*\in C$, suppose there exists one $\boldsymbol{\theta}$ such that $\mathbf{L}\boldsymbol{\theta}=\mathbf{0}$, but $\langle\nabla f(\mathbf{x}^*),\boldsymbol{\theta}\rangle\neq0$. Without lose of generalization, suppose $\langle\nabla f(\mathbf{x}^*),\boldsymbol{\theta}\rangle<0$, then the directional derivative of f at $\mathbf{x}^*$ is

$$d_{\boldsymbol{\theta}}f(\mathbf{x}^*)=\lim_{h\to0}\frac{f(\mathbf{x}^*+h\boldsymbol{\theta})-f(\mathbf{x}^*)}{h}=\langle\nabla f(\mathbf{x}^*),\boldsymbol{\theta}\rangle<0,$$

which is in contradiction with $\mathbf{x}^*\in C$. Therefore, $\langle\nabla f(\mathbf{x}^*),\boldsymbol{\theta}\rangle=0$ for $\forall\boldsymbol{\theta}$ such that $\mathbf{L}\boldsymbol{\theta}=\mathbf{0}$.

In particular, $f(\mathbf{x}) = \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$ is convex and differentiable. And the solution set

$$\begin{aligned} C &= \arg\min_{\mathbf{x}:\mathbf{L}\mathbf{x}=\mathbf{0}} \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 = \arg\min_{\alpha\in\mathbb{R}} \frac{1}{2}\|\mathbf{y} - \alpha\mathbf{A}\mathbf{1}\|_2^2 \\ &= \arg\min_{\alpha\in\mathbb{R}} \frac{1}{2}\|\mathbf{A}\mathbf{1}\|_2^2\alpha^2 - \langle \mathbf{y}, \mathbf{A}\mathbf{1}\rangle\alpha + \frac{1}{2}\|\mathbf{y}\|_2^2 \end{aligned}$$

is always nonempty. Therefore, for $\forall\ \mathbf{x}^* \in \arg\min_{\mathbf{x}:\mathbf{L}\mathbf{x}=\mathbf{0}} \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$, and for $\forall\ \boldsymbol{\theta}$ such that $\mathbf{L}\boldsymbol{\theta} = \mathbf{0}$, we have

$$\langle \nabla f(\mathbf{x}^*), \boldsymbol{\theta}\rangle = \langle \mathbf{A}^\mathsf{T}(\mathbf{A}\mathbf{x}^* - \mathbf{y}), \boldsymbol{\theta}\rangle = 0.$$

Proof of Theorem 2. A point $\tilde{\mathbf{x}}$ furnishes a global minimum of the convex function $F_\mathbf{L}(\mathbf{x})$ if and only if all forward directional derivatives $d_{\boldsymbol{\theta}}F_\mathbf{L}(\mathbf{x})$ at $\tilde{\mathbf{x}}$ are nonnegative. Let $f(\mathbf{x}) = \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$ and $m(\mathbf{x}) = \min_\mathbf{v}\|\mathbf{v}\|_1 + \frac{1}{2}\|\mathbf{B}(\mathbf{L}\mathbf{x} - \mathbf{v})\|_2^2$, then

$$\begin{aligned} d_{\boldsymbol{\theta}}F_\mathbf{L}(\tilde{\mathbf{x}}) &= \lim_{h\to 0} \frac{F_\mathbf{L}(\tilde{\mathbf{x}} + h\boldsymbol{\theta}) - F_\mathbf{L}(\tilde{\mathbf{x}})}{h} \\ &= \langle \nabla f(\tilde{\mathbf{x}}), \boldsymbol{\theta}\rangle + \mu \lim_{h\to 0} \frac{\|\mathbf{L}(\tilde{\mathbf{x}} + h\boldsymbol{\theta})\|_1 - \|\mathbf{L}\tilde{\mathbf{x}}\|_1}{h} - \mu d_{\boldsymbol{\theta}}m(\tilde{\mathbf{x}}). \end{aligned} \tag{A.6}$$

Suppose $\mathbf{L}\tilde{\mathbf{x}} = \mathbf{0}$, we first show that

$$d_{\boldsymbol{\theta}}m(\tilde{\mathbf{x}}) = \lim_{h\to 0} \frac{m(\tilde{\mathbf{x}} + h\boldsymbol{\theta}) - m(\tilde{\mathbf{x}})}{h} = 0.$$

Note that $m(\tilde{\mathbf{x}} + h\boldsymbol{\theta}) = \|\tilde{\mathbf{v}}\|_1 + \frac{1}{2}\|\mathbf{B}(\mathbf{L}(h\boldsymbol{\theta}) - \tilde{\mathbf{v}})\|_2^2$ where $\tilde{\mathbf{v}}$ satisfies $\mathbf{B}^\mathsf{T}\mathbf{B}(h\mathbf{L}\boldsymbol{\theta} - \tilde{\mathbf{v}}) \in \partial\|\tilde{\mathbf{v}}\|_1$ by Fermat's rule. If $\tilde{\mathbf{v}} = \mathbf{0}$, then $\|\mathbf{B}^\mathsf{T}\mathbf{B}(h\mathbf{L}\boldsymbol{\theta} - \tilde{\mathbf{v}})\|_\infty = \|h\mathbf{B}^\mathsf{T}\mathbf{B}\mathbf{L}\boldsymbol{\theta}\|_\infty \leq 1$ for all sufficiently small h. Therefore, $m(\tilde{\mathbf{x}} + h\boldsymbol{\theta}) = \frac{h}{2}\|\mathbf{B}\mathbf{L}\boldsymbol{\theta}\|_2^2$ for all sufficiently small h. It is straightforward to verify that $m(\tilde{\mathbf{x}}) = 0$. Hence, $d_{\boldsymbol{\theta}}m(\tilde{\mathbf{x}}) = 0$ for all $\boldsymbol{\theta}$ by L'Hopital's rule.

Therefore, (A.6) becomes

$$d_{\boldsymbol{\theta}}F_\mu(\tilde{\mathbf{x}}) = \langle \nabla f(\tilde{\mathbf{x}}), \boldsymbol{\theta}\rangle + \mu\|\mathbf{L}\boldsymbol{\theta}\|_1.$$

We now show that there exists some μ_0 such that for $\forall \mu > \mu_0$, $F_\mathbf{L}$ is minimized by $\tilde{\mathbf{x}}$, namely $d_{\boldsymbol{\theta}}F_\mathbf{L}(\tilde{\mathbf{x}}) \geq 0$ for all $\boldsymbol{\theta}$. Since $\mathbf{L}$ is not of full column rank, we can express $\mathbf{L}$ via the singular value decomposition as

$$\mathbf{L} = \begin{bmatrix}\mathbf{U}_1 & \mathbf{U}_2\end{bmatrix}\begin{bmatrix}\mathbf{D} & \mathbf{O} \\ \mathbf{O} & \mathbf{O}\end{bmatrix}\begin{bmatrix}\mathbf{V}_1 & \mathbf{V}_2\end{bmatrix}^\mathsf{T}.$$

Take $\forall \boldsymbol{\theta} \in \mathbb{R}^n$ which can be written as a linear combination of the columns of $\mathbf{V}_1$ and $\mathbf{V}_2$, namely $\boldsymbol{\theta} = \mathbf{V}_1\boldsymbol{\alpha}_1 + \mathbf{V}_2\boldsymbol{\alpha}_2$. Then $\mathbf{L}\boldsymbol{\theta} = \mathbf{U}_1\mathbf{D}\boldsymbol{\alpha}_1$ and $\mathbf{L}\mathbf{V}_2\boldsymbol{\alpha}_2 = \mathbf{0}$. Therefore,

$$\begin{aligned} d_\theta F_\mathbf{L}(\tilde{\mathbf{x}}) &= \langle \nabla f(\tilde{\mathbf{x}}), \mathbf{V}_1\boldsymbol{\alpha}_1\rangle + \langle \nabla f(\tilde{\mathbf{x}}), \mathbf{V}_2\boldsymbol{\alpha}_2\rangle + \mu\|\mathbf{U}_1\mathbf{D}\boldsymbol{\alpha}_1\|_1 \\ &= \langle \nabla f(\tilde{\mathbf{x}}), \mathbf{V}_1\boldsymbol{\alpha}_1\rangle + \mu\|\mathbf{U}_1\mathbf{D}\boldsymbol{\alpha}_1\|_1 \quad \text{(by Lemma 1)} \\ &\geq \mu\|\mathbf{U}_1\mathbf{D}\boldsymbol{\alpha}_1\|_2 - \|\nabla f(\tilde{\mathbf{x}})\|_2\|\mathbf{V}_1\boldsymbol{\alpha}_1\|_2 \\ &= \mu\|\mathbf{D}\boldsymbol{\alpha}_1\|_2 - \|\nabla f(\tilde{\mathbf{x}})\|_2\|\boldsymbol{\alpha}_1\|_2 \\ &\geq \mu\sigma_{\min}\|\boldsymbol{\alpha}_1\|_2 - \|\nabla f(\tilde{\mathbf{x}})\|_2\|\boldsymbol{\alpha}_1\|_2 \\ &= (\mu\sigma_{\min} - \|\nabla f(\tilde{\mathbf{x}})\|_2)\|\mathbf{V}_1^\mathsf{T}\boldsymbol{\theta}\|_2 \quad (\text{since } \boldsymbol{\alpha}_1 = \mathbf{V}_1^\mathsf{T}\boldsymbol{\theta}) \end{aligned}$$

where $\sigma_{\min}$ is the smallest singular value of $\mathbf{L}$. Note that the first inequality is obtained by the Cauchy-Schwarz inequality and the fact that $\|\mathbf{z}\|_1 \geq \|\mathbf{z}\|_2$ for all $\mathbf{z}$.

In conclusion, we have $d_{\boldsymbol{\theta}}F_\mathbf{L}(\tilde{\mathbf{x}}) \geq 0$ for all $\boldsymbol{\theta}$, when $\mu > \mu_0 = \frac{\|\nabla f(\tilde{\mathbf{x}})\|_2}{\sigma_{\min}} = \frac{\|\mathbf{A}^\mathsf{T}(\mathbf{A}\tilde{\mathbf{x}} - \mathbf{y})\|_2}{\sigma_{\min}}$.

References

Abe, J., Yamagishi, M., Yamada, I., 2019. Convexity-edge-preserving signal recovery with linearly involved generalized minimax concave penalty function. ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) , 4918–4922.

Bayram, I., 2015. On the convergence of the iterative shrinkage/thresholding algorithm with a weakly convex penalty. IEEE Trans. Signal Process. 64, 1597–1608.

Blake, A., Zisserman, A., 1987. Visual Reconstruction. MIT press.

Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. Inverse Problems 3, 1–122.

Byrne, C., 2002. Iterative oblique projection onto convex sets and the split feasibility problem. Inverse Problems 18, 441–453.

Chen, J., Chen, Z., 2008. Extended bayesian information criteria for model selection with large model spaces. Biometrika 95, 759–771.

Chen, J., Chen, Z., 2012. Extended bic for small-n-large-p sparse glm. Statistica Sinica 22, 759–771.

Chen, S.S., Donoho, D.L., Saunders, M.A., 1998. Atomic decomposition by basis pursuit. SIAM Journal on Scientific Computing 20, 33–61.

Goldstein, T., Li, M., Yuan, X., Esser, E., Baraniuk, R., 2013. Adaptive primal-dual hybrid gradient methods for saddle-point problems. arXiv preprint arXiv:1305.0546 .

Goldstein, T., Li, M., Yuan, X., Esser, E., Baraniuk, R., 2015. Adaptive primal-dual splitting methods for statistical learning and image processing. Advances in Neural Information Processing Systems , 2089–2097.

Lanza, A., Morigi, S., Selesnick, I., Sgallari, F., 2019. Sparsity-inducing nonconvex nonseparable regularization for convex image processing. SIAM Journal on Imaging Sciences 12, 1099–1134.

Nikolova, M., 1998. Estimation of binary images by minimizing convex criteria. Proc. IEEE Int. Conf. Image Processing (ICIP) 2, 108–112.

Nikolova, M., Ng, M.K., Tam, C.P., 2010. Fast nonconvex nonsmooth minimization methods for image restoration and reconstruction. IEEE Trans. Signal Process 19, 3073–3088.

Qu, B., Xiu, N., 2005. A note on the cq algorithm for the split feasibility problem. Inverse Problems 21, 1655–1665.

Selesnick, I., 2017a. Sparse regularization via convex analysis. IEEE Transactions on Signal Processing 65, 4481–4494.

Selesnick, I., 2017b. Total variation denoising via the moreau envelope. EEE Signal Process. Lett. 24, 216–220.

Selesnick, I., Lanza, A., Morigi, S., Sgallari, F., 2020. Non-convex total variation regularization for convex denoising of signals. Journal of Mathematical Imaging and Vision , 1–17.

Selesnick, I.W., 2015. Sparsity-assisted signal smoothing, in: Excursions in Harmonic Analysis, Volume 4. Springer, pp. 149–176.

Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. J. R. Statist. Soc. B 24, 267–288.

Wang, X., Zhu, H., Initiative, A.D.N., 2017. Generalized scalar-on-image regression models via total variation. Journal of the American Statistical Association 112, 1156–1168.

Zhou, H., Li, L., 2014. Regularized matrix regression. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 76, 463–483.

Zou, J., Shen, M., Zhang, Y., Li, H., Liu, G., Ding, S., 2018. Total variation denoising with non-convex regularizers. IEEE Access. 7, 4422–4431.