

Group Iterative Spectrum Thresholding for Super-Resolution Sparse Spectral Selection

Yiyuan She, Jiangping Wang, Huanghuang Li, and Dapeng Wu, *Fellow, IEEE*

Abstract—Recently, sparsity-based algorithms are proposed for super-resolution spectrum estimation. However, to achieve adequately high resolution in real-world signal analysis, the dictionary atoms have to be close to each other in frequency, thereby resulting in a coherent design. The popular convex compressed sensing methods break down in presence of high coherence and large noise. We propose a new regularization approach to handle model collinearity and obtain parsimonious frequency selection simultaneously. It takes advantage of the pairing structure of sine and cosine atoms in the frequency dictionary. A probabilistic spectrum screening is also developed for fast computation in high dimensions. A data-resampling version of high-dimensional Bayesian Information Criterion is used to determine the regularization parameters. Experiments show the efficacy and efficiency of the proposed algorithms in challenging situations with small sample size, high frequency resolution, and low signal-to-noise ratio.

Keywords: spectral estimation, sparsity, super-resolution, non-convex optimization, iterative thresholding, model selection, spectral screening.

I. INTRODUCTION

The problem of spectral estimation studies how signal power is distributed over frequencies, and has rich applications in speech coding, radar & sonar signal processing and many other areas. Suppose a discrete-time real-valued signal is observed at finite time points contaminated with i.i.d. Gaussian noise. In common with all spectral models, we assume the signal can be represented as a linear combination of sinusoids, and aim to recover the spectrum of the signal at a desired resolution. However, the problem becomes very challenging when the required frequency resolution is high. In particular, the number of the frequency levels at the desired resolution can be (much) greater than the sample size, referred to as super-resolution spectral estimation. For such discrete-time signals of finite length, the classical methods based on Fourier analysis [1] or least-squares periodogram (LSP) [2], [3] suffer from power leakage and have very limited spectral resolution [1]. Some more recent algorithms, such as Burg [1], MUSIC [4] and RELAX [5] only alleviate the issue to some extent.

We assume that the signal is sparse in the frequency-domain, i.e., the number of its sinusoidal components is small relative

to the sample size, referred to as the *spectral sparsity*. It is a realistic assumption in many applications (e.g., astronomy [6] and radar signal processing [7]), and makes it possible to apply the revolutionary **compressed sensing** (CS) technique. In [6], Chen and Donoho proposed the basis pursuit (BP) to handle overcomplete dictionaries and unevenly sampled signals. A number of similar works followed, see, e.g., [8]–[15] and the references therein.

We point out two crucial facts that cannot be ignored in super-resolution spectrum reconstruction. (a) When the desired frequency resolution is very high, neighboring dictionary atoms become very similar and thus necessarily result in high coherence or collinearity. As is well known in the literature, the popular convex l_1 technique as used in the BP yields inconsistent frequency selection and suboptimal rates in estimation and prediction under such coherent setups [16]–[20]. (b) The grouping structure of the sinusoidal components is an essential feature in spectrum recovery: if frequency f is absent in the signal, the coefficients for $\cos(2\pi ft)$ and $\sin(2\pi ft)$ should *both* be zero.

In this paper we investigate super-resolution spectral recovery from a statistical perspective and propose a **group iterative spectrum thresholding (GIST)** framework to tackle the aforementioned challenges. GIST allows for (possibly nonconvex) shrinkage estimation and can exploit the pairing structure. Interestingly, we find that neither the l_1 nor the l_0 regularization is satisfactory for spectrum estimation, and advocate a hybrid $l_0 + l_2$ type shrinkage estimation. Theoretical analysis shows that the new regularization essentially removes the stringent coherence requirement and can accommodate much lower SNR and higher coherence. Furthermore, a GIST variant provides a screening technique for supervised dimension reduction to deal with applications in ultrahigh dimensions. The rest of this paper is organized as follows. We formulate the problem from a statistical point of view and briefly survey the literature in Section II. In Section III, we propose the GIST framework—in more details, a novel form of regularization, a generic algorithm for fitting group nonconvex penalized models, a data-resampling based model selection criterion, and a probabilistic spectral screening for fast computation. Experimental results are shown in Section IV. We summarize the conclusions in Section V. The technical details are left to the Appendices.

II. MODEL SETUP AND THE SUPER-RESOLUTION CHALLENGE

In this section, we introduce the problem of super-resolution spectrum estimation and review some existing methods from

Yiyuan She is with Dept. of Statistics, Florida State University, Tallahassee, FL 32306. Jiangping Wang is with Beckman Institute and Dept. of Electrical & Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801. Huanghuang Li and Dapeng Wu (correspondence author) are with Dept. of Electrical & Computer Engineering, University of Florida, Gainesville, FL 32611. Email of Dapeng Wu: wu@ece.ufl.edu. This work was supported in part by NSF grants CCF-1117012 and CCF-1116447. Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

a statistical point of view. Let $\mathbf{y} = [y(t_n)]_{1 \leq n \leq N}$ be a **real**-valued signal contaminated with i.i.d. Gaussian noise $N(0, \sigma^2)$. (We focus on real signals in this paper but our methodology carries over to complex-valued signals; see Section V.) The sampling time sequence $\{t_n\}_{1 \leq n \leq N}$ is *not* required to be uniform (cf. [6]). In order to achieve super-resolution spectral recovery, an *overcomplete* frequency dictionary must be applied. Concretely, we use a grid of evenly spaced frequencies $f_k = f_{\max} \cdot k/D$ for $k = 0, 1, \dots, D$ to construct the sine and cosine frequency predictors, i.e., $\cos(2\pi t f_k)$ and $\sin(2\pi t f_k)$. Let \mathcal{F} denote the set of nonzero frequencies $\{f_1, \dots, f_D\}$. The upper band limit f_{\max} can be $(2 \min_{1 \leq n \leq N} (t_n - t_{n-1}))^{-1}$ or estimated based on the spectral window [21]. The cardinality of the dictionary controls the frequency resolution given by f_{\max}/D . The true spectra of the signal are assumed to be discrete for convenience, because the quantization error can always be reduced by increasing the value of D . The signal can be represented by

$$y_n = y(t_n) = \sum_{k=0}^D A_k \cos(2\pi f_k t_n + \phi_k) + e_n, 1 \leq n \leq N, \quad (1)$$

where A_k, ϕ_k are unknown, and the noise $\{e_n\}_{n=1}^N$ are i.i.d. Gaussian with zero mean and unknown variance σ^2 . Traditionally, $D \leq N$. But in super resolution spectral analysis, D can take a much larger value than N . It still results in a well-defined problem because only a few A_k are nonzero under the *spectral sparsity* assumption.

From $A_k \cos(2\pi f_k t_n + \phi_k) = A_k \cos(\phi_k) \cos(2\pi f_k t_n) - A_k \sin(\phi_k) \sin(2\pi f_k t_n) = a_k \cos(2\pi f_k t_n) + b_k \sin(2\pi f_k t_n)$ with $a_k = A_k \cos \phi_k, b_k = -A_k \sin \phi_k$, we introduce two column vectors

$$\begin{aligned} \mathbf{X}^{\cos}(f) &\triangleq [\cos(2\pi t_n f)]_{1 \leq n \leq N}, \\ \mathbf{X}^{\sin}(f) &\triangleq [\sin(2\pi t_n f)]_{1 \leq n \leq N}, \end{aligned}$$

and define the predictor matrix

$$\mathbf{X} \triangleq [\mathbf{X}^{\cos}(f_1), \dots, \mathbf{X}^{\cos}(f_D), \mathbf{X}^{\sin}(f_1), \dots, \mathbf{X}^{\sin}(f_D)]. \quad (2)$$

(Some redundant or useless predictors can be removed in concrete problems, see (4).) Denote the coefficient vector by $\boldsymbol{\beta} \in \mathbb{R}^{2D}$ and the intercept (zero frequency component) by α . Now the model can be formulated as a linear regression

$$\mathbf{y} = \alpha + \mathbf{X}\boldsymbol{\beta} + \mathbf{e}, \quad (3)$$

where $\boldsymbol{\beta}$ is sparse and $\mathbf{e} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$. In super-resolution analysis, $D \gg N$, giving a small-sample-size-high-dimensional design. Linear analysis such as Fourier transform fails for such an underdetermined system.

As a demonstration, we consider a noisy ‘TwinSine’ signal at frequencies 0.25 Hz and 0.252 Hz with 100 observations. Obviously, the frequency resolution needs to be as fine as 0.002 HZ to perceive and distinguish the two sinusoidal components with different coefficients. We set $f_{\max} = 1/2$, and thus $2D$ must be at least 500 – much larger than the sample size. The concrete design matrix (without the intercept)

is given by

$$\mathbf{X} = \begin{bmatrix} \cos(\pi \frac{1}{D} t_1) & \dots & \cos(\pi \frac{D}{D} t_1) & \sin(\pi \frac{1}{D} t_1) & \dots & \sin(\pi \frac{D-1}{D} t_1) \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \cos(\pi \frac{1}{D} t_N) & \dots & \cos(\pi \frac{D}{D} t_N) & \sin(\pi \frac{1}{D} t_N) & \dots & \sin(\pi \frac{D-1}{D} t_N) \end{bmatrix}. \quad (4)$$

The last sine atom disappears because all t_n are integers. This yields a super-resolution spectral estimation problem.

There are many algorithms for identifying the spectrum of a discrete-time signal. But not all of them can super-resolve. From a modeling perspective, we classify them as nonsparse methods and sparse methods. Most classical methods (e.g., [2], [3], [21]) are nonsparse and assume no knowledge on the power spectra. For super-resolution spectrum estimation, they may seriously broaden the main lobes and introduce side lobes. *In this paper, we focus on sparse methods.* As aforementioned, one popular assumption for solving underdetermined systems is signal sparsity: the number of present frequency components is small relative to the number of samples. The problem is still NP hard because the frequency location of the truly relevant sinusoidal components is unknown and the number of candidate components can be very large. In fact, the frequency grid used for constructing the dictionary can be made arbitrarily fine by the customer.

Early attempts to enforce sparsity effects include greedy or exhaustive searches [22], [23] and genetic algorithms with a sparsity constraint [24]. Harikumar [25] computes the maximally sparse solutions under a constraint on the fitting error. A breakthrough is due to Chen & Donoho who proposed the basis pursuit (BP) for spectrum estimation [6]. A number of similar works followed [8]–[11]. BP is able to superresolve for unevenly sampled signals. In our notation, the noiseless version of BP solves the convex optimization problem $\min \|\boldsymbol{\beta}\|_1$ s.t. $\alpha + \mathbf{X}\boldsymbol{\beta} = \mathbf{y}$. The noisy versions can be defined similarly, in a penalty/constraint form. The l_1 -norm provides the tightest convex relaxation to the l_0 -norm and achieves a sparse spectral representation of the signal within feasible time and cost.

In recent years, the power and limitation of this convex relaxation have been systematically studied in a large body of compressed sensing literature. In short, to guarantee good statistical performance in either prediction, estimation, or model selection, the coherence of the system must be low, in terms of, e.g., *mutual coherence* conditions [16], *restricted isometry property* (RIP) [17] and *irrepresentable conditions* [19] among others. For example, the RIP of order s requires that for any index set $I \subset \mathcal{F}$ with $|I| = s$, there exists an RIP constant $\delta_s \geq 0$ such that $(1 - \delta_s)\|\mathbf{v}\|_2^2 \leq \|\mathbf{X}_I \mathbf{v}\|_2^2 \leq (1 + \delta_s)\|\mathbf{v}\|_2^2, \forall \mathbf{v} \in \mathbb{R}^s$; when δ_s is small, any s predictors in \mathbf{X} are approximately orthogonal. In theory, to guarantee l_1 's effectiveness in statistical accuracy, frequency selection consistency, and algorithmic stability, such RIP constants have to be small, e.g., $\delta_{3S} + 3\delta_{4S} < 2$ in a noisy setup, where $S = \|\boldsymbol{\beta}\|_0$ [17]. Similarly, the mutual coherence, defined as the maximum absolute value of the off-diagonal elements in the scaled Gram matrix $\mathbf{X}^T \mathbf{X}/N$, has to be as low as $O(1/S)$ [16]. Such theoretical results clearly indicate that the super-resolution challenge *cannot* be **fully** addressed by the l_1 -norm

based methods, because many similar sinusoidal components may arise in the dictionary and bring in high coherence.

To enhance the sparsity of the BP, Blumensath & Davies proposed the iterative hard thresholding (IHT) [12], [13]. See [14], [15] for some approximation methods. Intuitively, nonconvex penalties can better approximate the l_0 -norm and yield sparser estimates than the convex l_1 -penalty. On the other hand, we find that when the signal-to-noise ratio (SNR) is low and/or the coherence is high, the l_0 penalization may give an **over-sparse** spectral estimate and miss certain true frequency components. The high miss rates are due to the fact that the l_0 regularization is through (hard) thresholding only, offering no shrinkage at all for nonzero coefficients. Therefore, it tends to kill too many predictors to achieve the appropriate extent of shrinkage especially when the SNR is low. An inappropriate nonconvex penalty may seriously mask true signal components. This issue will be examined in the next section.

III. GIST FRAMEWORK

This section examines the super-resolution spectrum estimation in details. The complete group iterative spectrum thresholding (GIST) framework is introduced at the end.

A. A novel regularization form

In this subsection, we study a group penalized least-squares model and investigate the appropriate type of regularization.

The BP finds a solution to an underdetermined linear system with the minimum l_1 norm. When the signal is corrupted by noise as in (3), the following l_1 -penalized linear model is more commonly used:

$$\frac{1}{2}\|\mathbf{y} - \alpha - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_1, \quad (5)$$

where λ is a regularization parameter to provide a trade-off between the fitting error and solution sparsity. The intercept or zero frequency component α is not subject to any penalty. To include more sparsity-enforcing penalties, we consider a more general problem in this paper which minimizes

$$\frac{1}{2}\|\mathbf{y} - \alpha - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \sum_{k=1}^{2D} P(|\beta_k|; \lambda) =: F(\boldsymbol{\beta}; \lambda), \quad (6)$$

where $P(\cdot; \lambda)$ is a univariate penalty function parameterized by λ and is possibly **nonconvex**.

Some structural information can be further incorporated in spectrum estimation. From the derivation of (3), $A_k = 0$ implies $\beta_k = \beta_{D+k} = 0$, i.e., the sine and cosine predictors at f_k vanish *simultaneously*. The pairing structure shows it is more reasonable to impose the so-called group sparsity or block sparsity [26]–[28] on $\{(\beta_k, \beta_{D+k})\}_{1 \leq k \leq D}$ rather than the unstructured sparsity on $\{\beta_k\}_{1 \leq k \leq 2D}$. The group penalized model with the model design (2) minimizes

$$\frac{1}{2}\|\mathbf{y} - \alpha - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \sum_{k=1}^D P\left(\sqrt{\beta_k^2 + \beta_{D+k}^2}; \lambda\right) =: F(\boldsymbol{\beta}; \lambda). \quad (7)$$

(In the problem with the design matrix given by (4), the last sine predictor disappears and thus we always set β_{2D} to be 0.) The penalty function P is the same as before and is allowed to be nonconvex. For ease in computation, the first term in (6) and (7) will be replaced by $\frac{1}{2}\|\mathbf{y} - \alpha - \mathbf{X}\boldsymbol{\beta}\|_2^2/C$ for some C large enough; see the comment after Theorem 1.

A crucial problem is then to determine the appropriate form of P for regularization purposes. The popular l_1 -penalty $P_1(t; \lambda) = \lambda|t|$ may result in insufficient sparsity and relatively large prediction error, as shown in Section IV. There is still much room for improvement in super-resolution spectral estimation. Before we proceed, it is worth pointing out that there are two objectives involved in this task

Objective 1 (O1): *accurate* prediction of the signal at any new time point in the time domain;

Objective 2 (O2): *parsimonious* spectral representation of the signal in the Fourier domain.

O1+O2 complies with **Occam's razor** principle—the simplest way to explain the data is the best. A perfect approach must reflect both concerns to produce a stable sparse model with good generalizability.

From the perspective of **O2**, the l_0 -norm constructs an ideal penalty

$$P_0(t; \lambda) = \frac{\lambda^2}{2} 1_{t \neq 0}, \quad (8)$$

where the indicator function $1_{t \neq 0}$ is 1 when $t \neq 0$ and 0 otherwise. Yet it is discrete and strongly nonconvex. Interestingly, given any model matrix, the class of penalties $aP_H(t; \lambda/\sqrt{a})$ for any $a \geq 1$ mimics the behavior of (8), where P_H , referred to as the *hard-penalty*, is defined by

$$P_H(t; \lambda) = \begin{cases} -t^2/2 + \lambda|t|, & \text{if } |t| < \lambda \\ \lambda^2/2, & \text{if } |t| \geq \lambda. \end{cases} \quad (9)$$

Based on [29], we can show that all penalties, including the continuous penalty (9) ($a = 1$) and the discrete penalty (8) ($a = \infty$), result in the same global minima in optimization. Fig. 1 illustrates the penalty family in a neighborhood around 0.




Fig. 1: The nonconvex ‘hard’ penalty family ($a \geq 1$) in a neighborhood around 0. All penalties lead to the same Θ -estimators. The discrete l_0 -penalty P_0 corresponds to $a = \infty$. The one with the smallest curvature is given by P_H with $a = 1$.

A different type of regularization is desirable for objective **O1**. Even if all truly relevant sinusoidal components could be successfully located, these atoms are not necessarily far apart in the frequency domain, and thus collinearity may occur. In statistical signal processing, Tikhonov regularization is an effective means to deal with the singularity issue which seriously affects estimation and prediction accuracy. It is in the form of an l_2 -norm penalty

$$P_R(t; \eta) = \frac{1}{2} \eta t^2, \quad (10)$$

also known as the ridge penalty in statistics. The necessity and benefit of introducing such shrinkage in multidimensional estimation date back to the famous James-Stein estimator [30]. Even for the purpose of detection, **O1** plays an important role because most parameter tuning methods are designed to reduce prediction error.

Taking into account both concerns, we advocate the following hybrid *hard-ridge* (**HR**) penalty as a fusion of (9) and (10):

$$P_{HR}(t; \lambda, \eta) = \begin{cases} -\frac{1}{2} t^2 + \lambda |t|, & \text{if } |t| < \frac{\lambda}{1+\eta} \\ \frac{1}{2} \eta t^2 + \frac{1}{2} \frac{\lambda^2}{1+\eta}, & \text{if } |t| \geq \frac{\lambda}{1+\eta}. \end{cases} \quad (11)$$

The hard portion induces sparsity for small coefficients, while the ridge portion, representing Tikhonov regularization, helps address the coherence of the design and compensates for noise and collinearity. In the following subsections, we will show that such defined hard-ridge penalty also allows for ease in optimization and has better frequency selection performance.

Finally, we point out the difference between HR and the elastic net [31] which adds an additional ridge penalty in the lasso problem (5). However, this $l_1 + l_2$ penalty, i.e., $\lambda_1 \|\beta\|_1 + \lambda_2^2 \|\beta\|_2^2 / 2$, may over-shrink the model (referred to as the *double-shrinkage* effect [31]) and can not enforce higher level of sparsity than the l_1 -penalty. In contrast, using a q -function trick [32], it is shown that P_{HR} results in the same estimator as the ' $l_0 + l_2$ ' penalty

$$P(t; \lambda, \eta) = \frac{1}{2} \frac{\lambda^2}{1+\eta} 1_{t \neq 0} + \frac{1}{2} \eta t^2. \quad (12)$$

The ridge part does not affect the nondifferential behavior of the l_0 -norm at zero, and there is no double-shrinkage effect for nonzero coefficient estimates.

B. GIST fitting algorithm

We discuss how to fit the group penalized model (7) for a wide class of penalty functions. We assume both \mathbf{X} and \mathbf{y} have been centered so that the intercept term vanishes in the model. Our main tool to tackle the computational challenge is the class of Θ -estimators [33]. Let $\Theta(\cdot; \lambda)$ be an arbitrarily given threshold function (with λ as the parameter) which is odd, monotone, and a unbounded shrinkage rule (see [33] for the rigorous definition) with λ as the parameter. A group Θ -estimator is defined to be a solution to

$$\beta = \vec{\Theta}(\beta + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta); \lambda). \quad (13)$$

Here, for any $\xi \in \mathbb{R}^{2D}$, $\vec{\Theta}(\xi; \lambda)$ is a $2D$ -dimensional vector ξ' satisfying

$$[\xi'_k, \xi'_{k+D}] = [\xi_k, \xi_{k+D}] \Theta(\|[\xi_k, \xi_{k+D}]\|_2; \lambda) / \|[\xi_k, \xi_{k+D}]\|_2,$$

for $1 \leq k \leq D$. In the simpler case when no grouping is assumed, the Θ -estimator equation (13) reduces to

$$\beta = \Theta(\beta + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta); \lambda). \quad (14)$$

A Θ -estimator is necessarily a P -penalized estimator provided that

$$P(t; \lambda) - P(0; \lambda) = \int_0^{|t|} (\sup\{s : \Theta(s; \lambda) \leq u\} - u) du + q(t; \lambda) \quad (15)$$

holds for some nonnegative $q(\cdot; \lambda)$ satisfying $q(\Theta(s; \lambda); \lambda) = 0$ for any $s \in \mathbb{R}$ [29]. Based on this result, we can compute P -penalized estimators by solving (13) for an appropriate Θ .

Algorithm 1 GIST-fitting algorithm.

given \mathbf{X} (design matrix, normalized), \mathbf{y} (centered), λ (regularization parameter(s)), Θ (thresholding rule), ω (relaxation parameter), and Ω (maximum number of iterations).

1) $\mathbf{X} \leftarrow \mathbf{X}/\tau_0$, $\mathbf{y} \leftarrow \mathbf{y}/\tau_0$, with $\tau_0 \geq \|\mathbf{X}\|_2$ (spectral norm).

2) Let $j \leftarrow 0$ and $\beta^{(0)}$ be an initial estimate, say, $\mathbf{0}$.

while $\|\beta^{(j+1)} - \beta^{(j)}\|$ is not small enough or $j \leq \Omega$ **do**

3.1) $\xi^{(j+1)} \leftarrow (1 - \omega)\xi^{(j)} + \omega(\beta^{(j)} + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta^{(j)}))$
if $j > 0$, and $\xi^{(j+1)} \leftarrow \beta^{(j)} + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta^{(j)})$ if $j = 0$;

GROUP FORM:

3.2a) $l_k^{(j+1)} \leftarrow \sqrt{(\xi_k^{(j+1)})^2 + (\xi_{k+D}^{(j+1)})^2}$, $1 \leq k \leq D$.

3.2b) If $l_k^{(j+1)} \neq 0$, $\beta_k^{(j+1)} \leftarrow \xi_k^{(j+1)} \Theta(l_k^{(j+1)}; \lambda) / l_k^{(j+1)}$
and $\beta_{k+D}^{(j+1)} \leftarrow \xi_{k+D}^{(j+1)} \Theta(l_k^{(j+1)}; \lambda) / l_k^{(j+1)}$. Otherwise
 $\beta_k^{(j+1)} = \beta_{k+D}^{(j+1)} = 0$.

NON-GROUP FORM:

3.2') $\beta^{(j+1)} \leftarrow \Theta(\xi^{(j+1)}; \lambda)$;

end while

deliver $\hat{\beta} = \beta^{(j+1)}$.

We establish the convergence of Algorithm 1 in the following theorem. For simplicity, assume that there is no intercept term in the model (which is reasonable when \mathbf{X} and \mathbf{y} have both been centered), and $\tau_0 = 1 > \|\mathbf{X}\|_2$. Let $\Sigma = \mathbf{X}^T \mathbf{X}$. Construct an energy function for any $\gamma, \zeta, \beta, \xi \in \mathbb{R}^{2D}$ as follows

$$\begin{aligned} G(\gamma, \zeta, \beta, \xi) &= \frac{1}{2} \|\mathbf{X}\gamma - \mathbf{y}\|_2^2 + P(\gamma; \lambda) + \frac{\omega}{2} (\gamma - \beta)^T (\mathbf{I} - \Sigma) (\gamma - \beta) \\ &+ \frac{(1-\omega)^2}{2\omega} (\zeta - \xi)^T (\mathbf{I} - \Sigma)^{-1} (\zeta - \xi) + \frac{1-\omega}{2} [\gamma + (\mathbf{I} - \Sigma)^{-1} \mathbf{X}^T \mathbf{y} \\ &- (\mathbf{I} - \Sigma)^{-1} \xi]^T (\mathbf{I} - \Sigma) [\gamma + (\mathbf{I} - \Sigma)^{-1} \mathbf{X}^T \mathbf{y} - (\mathbf{I} - \Sigma)^{-1} \xi] \\ &+ \frac{1-\omega}{2} [\zeta - (\mathbf{I} - \Sigma)\beta - \mathbf{X}^T \mathbf{y}]^T (\mathbf{I} - \Sigma)^{-1} [\zeta - (\mathbf{I} - \Sigma)\beta - \mathbf{X}^T \mathbf{y}] \\ &- \frac{1-\omega}{2} [\xi - (\mathbf{I} - \Sigma)\beta - \mathbf{X}^T \mathbf{y}]^T (\mathbf{I} - \Sigma)^{-1} [\xi - (\mathbf{I} - \Sigma)\beta - \mathbf{X}^T \mathbf{y}], \end{aligned} \quad (16)$$

with the non-group and group versions of $P(\gamma; \lambda)$ being $\sum_k P(|\gamma_k|; \lambda)$ and $\sum_{k=1}^D P\left(\sqrt{\gamma_k^2 + \gamma_{D+k}^2}; \lambda\right)$, respectively. $G(\gamma, \zeta, \beta, \xi)$ is always greater than or equal to the objective

function $F(\gamma)$ as defined in (7) or (6). This energy function can be used to prove the convergence of the iterates to a Θ -estimator.

Theorem 1: For any $0 < \omega \leq 1$ and a thresholding rule Θ satisfying (15), under the continuity assumption \mathfrak{A} in Appendix A, Algorithm 1 in either group form or non-group form converges, and the iterates $(\beta^{(j)}, \xi^{(j)})$ satisfy

$$\begin{aligned} & G(\beta^{(j+1)}, \xi^{(j+1)}, \beta^{(j+1)}, \xi^{(j+1)}; \lambda) \\ & \leq G(\beta^{(j)}, \xi^{(j)}, \beta^{(j)}, \xi^{(j)}; \lambda) - \delta_1 - \delta_2, \end{aligned} \quad (17)$$

where $\delta_1 = \frac{1-\omega}{2\omega}(\xi^{(j+1)} - \xi^{(j)})^T(\mathbf{I} - \Sigma)^{-1}(\xi^{(j+1)} - \xi^{(j)})$ and $\delta_2 = \frac{1}{2\omega}[\omega(\mathbf{I} - \Sigma)(\beta^{(j)} - \beta^{(j+1)}) + (1-\omega)(\xi^{(j)} - \xi^{(j+1)})]^T(\mathbf{I} - \Sigma)^{-1}[\omega(\mathbf{I} - \Sigma)(\beta^{(j)} - \beta^{(j+1)}) + (1-\omega)(\xi^{(j)} - \xi^{(j+1)})]$. Furthermore, any limit point β° of $\{\beta^{(j)}\}$ is a group (or non-group) Θ -estimator that satisfies (13) (or (14)), and the sequence $G(\beta^{(j)}, \xi^{(j)}, \beta^{(j)}, \xi^{(j)}; \lambda)$ decreases to the limit $F(\beta^\circ; \lambda)$ with F defined in (7) (or (6)).

Applying the theorem to Algorithm 1, we know the nongroup form solves the optimization problem $\frac{1}{2}\|\mathbf{y} - \mathbf{X}\beta\|_2^2/\tau_0^2 + \sum_{k=1}^{2D} P(|\beta_k|; \lambda)$, and the group form solves $\frac{1}{2}\|\mathbf{y} - \mathbf{X}\beta\|_2^2/\tau_0^2 + \sum_{k=1}^D P(\sqrt{\beta_k^2 + \beta_{D+k}^2}; \lambda)$ for any arbitrarily given \mathbf{X} , \mathbf{y} . Algorithm 1 is justified for computing a penalized spectrum estimate associated with P , provided that a proper Θ can be found to satisfy (15).

The P - Θ strategy covers most commonly used penalties, either in group form or non-group form. We give some examples below. (i) When Θ is the soft-thresholding, the P -function according to (15) is the l_1 -norm penalty used in BP, and the non-group version of our algorithm reduces to the iterative soft thresholding [34]. The group l_1 penalty (called the group lasso [26]) is more suitable for frequency selection, and can be handled by Algorithm 1 as well. (ii) When Θ is the hard-thresholding, for $q(\cdot; \lambda) \equiv 0$ we get the hard-penalty (9), and for $q(t; \lambda) = \frac{(\lambda - |t|)^2}{2} \mathbf{1}_{0 < |t| < \lambda}$ we get the l_0 -penalty (8). The algorithm, in non-group form, corresponds to the iterative hard thresholding [12], [13]. (iii) Finally, if we define Θ to be the hard-ridge thresholding:

$$\Theta_{HR}(t; \lambda, \eta) = \begin{cases} 0, & \text{if } |t| < \lambda \\ \frac{t}{1+\eta}, & \text{if } |t| \geq \lambda. \end{cases} \quad (18)$$

then $P_{\Theta_{HR}}$ is the hard-ridge penalty (11). Setting $q(t; \lambda, \eta) = \frac{1+\eta}{2}(|t| - \lambda)^2 \mathbf{1}_{0 < |t| < \lambda}$, we successfully reach the $l_0 + l_2$ penalty (12). See [29] for more examples, such as SCAD, l_p ($0 < p < 1$), and elastic net.

Algorithm 1 includes a relaxation parameter ω , which is an effective means to accelerate the convergence. See the recent work by Maleki & Donoho [35]. (Our relaxation form is novel and is of Type I based on [36]). In practice, we set $\omega = 2$, and the number of iterations can be reduced by about 40% in comparison to nonrelaxation form.

C. Statistical analysis

Although the l_1 regularization is popular (see, e.g., BP [6]), in the following we show that the HR penalty has better selection power and can remove the stringent coherence

assumption and can accommodate lower SNRs. We focus on the group form based on the discussion in Section III.

Let \mathcal{F} be the entire frequency set covered by the dictionary. For the design matrix defined in (2), $\mathcal{F} = \{f_1, \dots, f_D\}$. Given any frequency $f \in \mathcal{F}$, we use \mathbf{X}_f to denote the submatrix of \mathbf{X} formed by the sine and cosine frequency atoms at f , and β_f the corresponding coefficient vector. If $I \subset \mathcal{F}$ is an index set, \mathbf{X}_I and β_I are defined similarly. In general, \mathbf{X}_f is of size $N \times 2$ and β_f 2×1 (but not always—cf. (4)). Given any coefficient vector β , we introduce

$$z(\beta) = \{f \in \mathcal{F} : \|\beta_f\|_2 = 0\}, nz(\beta) = \{f \in \mathcal{F} : \|\beta_f\|_2 \neq 0\} \quad (19)$$

to characterize the frequency selection outcome. In particular, we write $z^* = z(\beta^*)$, $nz^* = nz(\beta^*)$, associated with the true coefficient vector β^* , and let $p_{nz^*} = |nz^*|$ be the number of frequencies present in the true signal, and $p_{z^*} = |z^*|$ the number of irrelevant frequencies.

We introduce two useful quantities κ and μ . Recall $\Sigma = \mathbf{X}^T \mathbf{X}$ and $\tau_0^2 = \|\Sigma\|_2 = \mu_{\max}(\Sigma)$ (the largest eigenvalue of Σ). Given $I \subset \mathcal{F}$, let $\Sigma_{I,I'} = \mathbf{X}_I^T \mathbf{X}_{I'}$ and $\Sigma_I = \mathbf{X}_I^T \mathbf{X}_I$. In this subsection, we assume the design matrix has been column-normalized such that the 2-norm of every column is \sqrt{N} . Let $\Sigma^{(s)} = \Sigma/N$. Define

$$\mu := \mu_{\min}(\Sigma_{nz^*,nz^*}^{(s)}), \text{ and } \kappa := \max_{f \in z^*} \|\Sigma_{f,nz^*}^{(s)}\|_2 / \sqrt{p_{nz^*}},$$

where μ_{\min} denotes the smallest eigenvalue and $\|\cdot\|_2$ refers to the spectral norm. ($\Sigma_{f,nz^*}^{(s)}$ is of size $2 \times 2p_{nz^*}$ typically.) Intuitively, κ measures the ‘mean’ correlation between the relevant frequency atoms and the irrelevant atoms. When κ is high, the coherence of the dictionary is necessarily high. Denote by P_1 the probability that with soft-thresholding being applied, there exists at least one estimate $\hat{\beta}$ from **Algorithm 1** such that $nz(\hat{\beta}) = nz^*$. P_{02} is similarly defined for hard-ridge thresholding. Theorem 2 bounds these two probabilities.

Theorem 2: Assume $\mu > 0$.

(i) Let Θ be the soft-thresholding. Under the assumption that $\kappa < \mu/p_{nz}$ and λ is chosen such that $\min_{f \in nz^*} \|\beta_f^*\|_2 \geq \frac{\lambda \sqrt{p_{nz^*}}}{N\mu/\tau_0^2}$, we have

$$1 - P_1 \leq \frac{e}{4} \left(\frac{p_{z^*} M^2}{e^{M^2/4}} + \frac{p_{nz^*} L^2}{e^{L^2/4}} \right), \quad (20)$$

where $M := \frac{\lambda \tau_0^2}{\sigma \sqrt{N}} (1 - \frac{\kappa p_{nz^*}}{\mu})$ and $L := (\min_{f \in nz^*} \|\beta_f^*\|_2 - \frac{\lambda \tau_0^2 \sqrt{p_{nz^*}}}{N\mu}) \frac{\sqrt{N\mu}}{\sigma}$.

(ii) Let Θ be the hard-ridge thresholding. Assume λ, η are chosen such that $\kappa \leq \frac{1}{\eta} \frac{\lambda(\mu N + \eta \tau_0^2)}{\|\beta_{nz^*}^*\|_2 \sqrt{p_{nz^*}}}$, $\iota := \min_{f \in nz^*} \|[(\Sigma_{nz^*} + \eta \mathbf{I})^{-1} \Sigma_{nz^*} \beta_{nz^*}^*]_f\|_2 \geq \frac{\lambda}{1+\eta}$, and $\eta \leq \mu N / \tau_0^2$. Then

$$1 - P_{02} \leq \frac{e}{4} \left(\frac{p_{z^*} M'^2}{e^{M'^2/4}} + \frac{p_{nz^*} L'^2}{e^{L'^2/4}} \right), \quad (21)$$

where $M' := \frac{1}{\sigma \sqrt{N}} (\lambda \tau_0^2 - \frac{\eta \tau_0^2}{\mu N + \eta \tau_0^2} \kappa \sqrt{p_{nz^*}} \|\beta_{nz^*}^*\|_2)$ and $L' := (\iota - \frac{\lambda}{1+\eta}) \frac{\sqrt{\mu N + \eta \tau_0^2} / \sqrt{\mu N}}{\sigma}$.

Seen from (20) and (21), both inconsistent detection probabilities are small. It is worth mentioning that in practice, we

find the value of η is usually small, which, however, effectively handles singularity/collinearity in comparison to $\eta = 0$, as supported by the literature (e.g., [37]). In the following, we make a comparison of the assumptions and probability bounds. The setup of $p_{z^*} \gg N \gg p_{nz^*}$ is of particular interest, which means the number of truly present frequencies is small relative to the sample size but the number of irrelevant frequencies is overwhelmingly large. The κ -conditions characterize coherence accommodation, while the conditions on $\min_{f \in nz^*} \|\beta_f^*\|_2$ and ι describe how small the minimum signal strength can be. (i) For the l_1 penalty, $\kappa < \mu/p_{nz^*}$ is a version of the irrepresentable conditions and cannot be relaxed in general [19]. In contrast, for the $l_0 + l_2$, the bound for κ becomes large when η is small, and so the stringent coherence requirement can be essentially removed! (ii) When η is small in the hard-ridge thresholding, the noiseless ridge estimator $(\Sigma_{nz^*} + \eta \mathbf{I})^{-1} \Sigma_{nz^*} \beta_{nz^*}^*$ is close to $\beta_{nz^*}^*$, but the minimum signal strength can be much lower than that of the l_1 , due to the fact that $N\mu/\tau_0^2 = \mu_{\min}(\Sigma_{nz^*, nz^*}^{(s)})/\mu_{\max}(\Sigma^{(s)}) \leq 1 \leq 1 + \eta$ and in particular, the disappearance of $\sqrt{p_{nz^*}}$. (iii) Finally, for small values of η , $M' > M$, $L' > L$, and so $l_0 + l_2$ has a better chance to recover the whole spectra correctly.

Remark. Including the ridge penalty in regularization is helpful to enhance estimation and prediction accuracy, especially when the frequency resolution is quite high and the true signal is multi-dimensional. Even when the purpose is selection alone, it is meaningful because most tuning strategies of λ are prediction error (generalization error) based.

D. Model comparison criterion

This part studies the problem of how to choose proper regularization parameters for any given data (\mathbf{X}, \mathbf{y}) . In (7), the general parameter λ provides a statistical bias-variance trade-off in regularizing the model, and ought to be tuned in a data-driven manner. In common with most researchers (say [11], [38], [39]), we first specify a grid $\Lambda = \{\lambda_1, \dots, \lambda_l, \dots, \lambda_L\}$, then run Algorithm 1 for every λ in the grid to get a *solution path* $\hat{\beta}(\lambda_l)$, $1 \leq l \leq L$, and finally, use a *model comparison criterion* to find the optimal estimate $\hat{\beta}_{opt}$. The commonly used model comparison criteria are Akaike information criterion (AIC), Bayesian information criterion (BIC), and cross-validation (CV). But we found none of them is satisfactory in the high-dimensional super-resolution spectral estimation.

Ideally, in a data-rich situation, one would divide the whole dataset into a training subset denoted by $(\mathbf{X}^{trn}, \mathbf{y}^{trn})$ and a validation subset $(\mathbf{X}^{val}, \mathbf{y}^{val})$. For any $\lambda \in \Lambda$, train the model on $(\mathbf{X}^{trn}, \mathbf{y}^{trn})$ and evaluate the prediction accuracy on the validation subset by, say, $\|\mathbf{y}^{val} - \mathbf{X}^{val} \hat{\beta}(\lambda)\|_2^2$. However, this data-splitting approach is only reasonable when the validation subset is large enough to approximate the true prediction error. It cannot be used in our problem due to insufficiency of observations. A popular data-reusing method in small samples is the \mathcal{K} -fold CV. Divide the dataset into \mathcal{K} folds. Let $(\mathbf{X}^{(\kappa)}, \mathbf{y}^{(\kappa)})$ denote the κ th subset, and $(\mathbf{X}^{(-\kappa)}, \mathbf{y}^{(-\kappa)})$ denote the remaining data. To obtain the CV error at any $\lambda_l \in \Lambda$, one needs to fit \mathcal{K} penalized models. Concretely, setting $\mathbf{X} = \mathbf{X}^{(-\kappa)}$ and $\mathbf{y} = \mathbf{y}^{(-\kappa)}$ as the training data, solve

the penalized problem associated with λ_l , the estimate represented by $\hat{\beta}^{(-\kappa)}(\lambda_l)$. Then calculate the validation error on $(\mathbf{X}^{(\kappa)}, \mathbf{y}^{(\kappa)})$: $\text{cv-err}(\lambda_l, \kappa) = \|\mathbf{y}^{(\kappa)} - \mathbf{X}^{(\kappa)} \hat{\beta}^{(-\kappa)}(\lambda_l)\|_2^2$. The summarized CV error, $\text{cv-err}(\lambda_l) = \sum_{\kappa=1}^{\mathcal{K}} \text{cv-err}(\lambda_l, \kappa)/N$, serves as the comparison criterion. After the optimal λ_{opt} is determined, we refit the model on the global dataset to get $\hat{\beta}_{opt}$.

However, when a nonconvex penalty is applied, the above plain CV has an inherent drawback: the \mathcal{K} trained models at a common value of λ_l may not be comparable, and thus averaging their validation errors may make little sense. The reasons are twofold. (i) The regularization parameter λ appears in a Lagrangian form optimization problem (cf. (6) or (7)). In general, the optimal λ to guarantee good selection and estimation must be a function of both the true coefficient vector β^* and the data (\mathbf{X}, \mathbf{y}) . Notice that in the trainings of \mathcal{K} -fold CV, (\mathbf{X}, \mathbf{y}) changes. The same value of λ may have different regularization effects for different training datasets although β^* remains the same. Fig. 2 shows the numbers of nonzero coefficient estimates under the l_0 penalization in 5-fold CV—they are never consistent at any fixed value of λ ! (ii) The solution path $\hat{\beta}(\lambda)$ associated with a nonconvex penalty is generally discontinuous in λ . Fig. 3 plots the l_0 solution path for the default TwinSine signal. Even a small change in λ may result in a totally different estimate and zero-nonzero pattern. In consideration of both (i) and (ii), cross-validating λ is not a proper tuning strategy in our problem.




Fig. 2: The numbers of nonzero coefficients in 5-fold CV with respect to λ . The 5 CV trainings yield (sometimes quite) different models at the same value of λ .

To resolve the training inconsistency, we advocate a generic selective cross validation (SCV) for parameter tuning in sparsity-inducing penalties. First the sparsity algorithm is run on the *entire* dataset to get a solution path $\hat{\beta}(\lambda_l)$, $l = 1, \dots, L$. Every estimate $\hat{\beta}(\lambda_l)$ determines a candidate model with the predictor set given by $nz_l = nz(\hat{\beta}(\lambda_l)) = \{f_k \in \mathcal{F} : \hat{\beta}_k^2 + \hat{\beta}_{k+D}^2 \neq 0\}$. Next, we *cross-validate* nz_l (instead of




Fig. 3: The l_0 -penalized solution path $\hat{\beta}(\lambda)$ is discontinuous in λ . For clarity, only 5 frequency paths (chosen at random) are shown.

λ) to evaluate the goodness-of-fit of each candidate model. In this way, all \mathcal{X} trainings are restricted to the *same* subset of predictors. Concretely, for penalties without l_2 shrinkage, such as the l_0 -penalty, $\hat{\beta}^{(-\kappa)}(\lambda_l)$ is the unpenalized regression estimate fitted on $(\mathbf{y}^{(-\kappa)}, \mathbf{X}_{nz_l}^{(-\kappa)})$, while for penalties with l_2 shrinkage, such as the $l_0 + l_2$ -penalty, $\hat{\beta}^{(-\kappa)}(\lambda_l)$ is the ridge regression estimate fitted on $(\mathbf{y}^{(-\kappa)}, \mathbf{X}_{nz_l}^{(-\kappa)})$ (cf. Theorem 1), i.e., $\hat{\beta}^{(-\kappa)}(\lambda_l) = ((\mathbf{X}_{nz_l}^{(-\kappa)})^T \mathbf{X}_{nz_l}^{(-\kappa)} + \eta \mathbf{I})^{-1} (\mathbf{X}_{nz_l}^{(-\kappa)})^T \mathbf{y}^{(-\kappa)}$. Finally, the total SCV error is summarized by $\text{SCV}(\lambda_l) = \sum_{\kappa=1}^{\mathcal{X}} \|\mathbf{y}^{(\kappa)} - \mathbf{X}^{(\kappa)} \hat{\beta}^{(-\kappa)}(\lambda_l)\|_2^2$.

Motivated by the work of [38], we add a high-dimensional BIC correction term to define the model comparison criterion: $\text{SCV-BIC}(\lambda_l) = \text{SCV}(\lambda_l) + \text{DF}(\hat{\beta}(\lambda_l)) \log N$, where DF is the degrees of freedom function. When the true signal has a parsimonious representation in the frequency domain, i.e., the number of present frequencies is very small, such a correction is necessary—see [38] for a further theoretical justification. For the l_0 or l_1 penalty, DF is approximately the number of nonzero components in the estimate; for the $l_0 + l_2$ penalty, $\text{DF}(\hat{\beta}(\lambda_l))$ is given by $\text{Tr}((\mathbf{X}_{nz_l}^T \mathbf{X}_{nz_l} + \eta \mathbf{I})^{-1} \mathbf{X}_{nz_l}^T \mathbf{X}_{nz_l})$ [37]. The optimal estimate $\hat{\beta}_{\text{opt}}$ is chosen from the original solution path $\{\hat{\beta}(\lambda_l)\}_{l=1}^L$ by minimizing $\text{SCV-BIC}(\lambda_l)$.

We point out that in SCV, the sparsity algorithm is only required to run on the whole dataset to generate one solution path, while CV needs \mathcal{X} such solution paths. SCV is more efficient in computation.

E. Probabilistic spectra screening

Computational complexity is another major challenge in super-resolution studies. In Algorithm 1, each iteration step involves only matrix-vector multiplications and component-wise thresholding operations. Both have low complexity and can be vectorized. The total number of flops is no more than

$(4DN + 8D)\Omega$, which is linear in D . In our experiments, $\Omega = 200$ suffices and thus the complexity of Algorithm 1 is $O(DN)$. (Restricting attention to uniformly sampled data and frequency atoms in the dictionary construction, we can use the Fast Fourier transform (FFT) in computation to reduce the complexity to $O(D \log D)$, as pointed out by an anonymous reviewer, see [40] and Section IV.) On the other hand, with a superbly high resolution dictionary (where D is very large), dimension reduction is still desirable to further reduce the computational cost.

This is indeed possible under the spectral sparsity assumption, where the number of true components is supposed to be much smaller than N . One may reduce the dimension from $2D$ to ϑN (say $\vartheta = 0.5$) before running the formal algorithm. If the ϑN candidate predictors are wisely chosen, the truly relevant atoms will be included with high probability and the performance sacrifice in selection/estimation will be mild. Hereinafter, we call ϑ the *candidate ratio*. A well designed screening algorithm should not be very sensitive to ϑ as long as it is reasonably large. Significant decrease in computational time can be achieved after this supervised dimension reduction.

We propose an iterative probabilistic screening by adapting Algorithm 1 for dimension reduction. This has the benefit that the screening principle is consistent with the fitting criterion. We recommend using the hard-ridge thresholding and the associated Algorithm 2 is stated below.

The differences in comparison to Algorithm 1 lie in (3.2b) and (3.2'), where a dynamic threshold is constructed in performing the hard-ridge thresholding. We next show that this screening version still has convergence guarantee. Similar to Theorem 1, assume $\tau_0 = 1 > \|\mathbf{X}\|_2$. Let G be the same energy function constructed in (16) with P given by (11) or (12). For simplicity, suppose $m := \vartheta N \in \mathbb{N}$. Theorem 3 shows that Algorithm 2 solves an l_0 -constrained problem.

Theorem 3: For any $0 < \omega \leq 1$, the sequence of iterates $(\beta^{(j)}, \xi^{(j)})$ from Algorithm 2 has the same function value decreasing property (17) for the energy function G , and $\beta^{(j)}$ satisfies $nz(\beta^{(j)}) \leq m$. In addition, under $\eta > 0$ and the notice-occurring assumption \mathfrak{B} in Appendix C, the sequence of $\beta^{(j)}$ has a unique limit point β° which corresponds to the ridge estimate restricted to $\mathbf{X}_{nz(\beta^\circ)}$ with $|nz(\beta^\circ)| \leq m$.

We can use SCV to tune η or simply set η at a small value (say $1e-2$). In practice, the screening can proceed in a progressive fashion to avoid greedy selection: we use a varying sequence of $m^{(j)}$ that decreases to ϑN in Step 3.1), and add ‘squeezing’ operations after Step 3.2) or 3.2’): $\mathbf{d} \leftarrow \{f \in \mathcal{F} : \|\beta_f^{(j+1)}\|_2 \neq 0\}$, $\beta^{(j+1)} \leftarrow \beta^{(j+1)}[\mathbf{d}]$, $\mathbf{X} \leftarrow \mathbf{X}[\mathbf{d}]$ (group version), or $\mathbf{d} \leftarrow \{k : 1 \leq k \leq 2D, \beta_k^{(j+1)} \neq 0\}$, $\beta^{(j+1)} \leftarrow \beta^{(j+1)}[\mathbf{d}]$, $\mathbf{X} \leftarrow \mathbf{X}[\mathbf{d}]$ (non-group version). We have found that empirically, the sigmoidal decay cooling schedule $m^{(j)} = \lceil 2D / (1 + \exp(\alpha j)) \rceil$ with $\alpha = 0.01$ achieves good balance between selection and efficiency.

GIST-Screening works decently in super-resolution spectral analysis seen from the experiments: after dimension reduction the true signal components are included with high probability and the computational cost can be significantly reduced.

Algorithm 2 GIST-Screening algorithm.

given \mathbf{X} (design matrix, normalized), \mathbf{y} (centered), η (l_2 shrinkage parameter), ϑ (candidate ratio—ratio of new dimension to sample size), ω (relaxation parameter), and $\tilde{\Omega}$ (maximum number of iterations). (For simplicity, assume ϑN is an integer.)

1) $\mathbf{X} \leftarrow \mathbf{X}/\tau_0$, $\mathbf{y} \leftarrow \mathbf{y}/\tau_0$, with $\tau_0 \geq \|\mathbf{X}\|_2$ (spectral norm).

2) Let $j \leftarrow 0$ and $\beta^{(0)}$ be an initial estimate say $\mathbf{0}$.

while $\|\beta^{(j+1)} - \beta^{(j)}\|$ is not small enough or $j \leq \tilde{\Omega}$ **do**

3.1) $\xi^{(j+1)} \leftarrow (1 - \omega)\xi^{(j)} + \omega(\beta^{(j)} + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta^{(j)}))$

if $j > 0$ and $\xi^{(j+1)} \leftarrow \beta^{(j)} + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta^{(j)})$ if $j = 0$, and set $m^{(j)} = \vartheta N$;

GROUP FORM:

3.2a) $l_k^{(j+1)} \leftarrow \sqrt{(\xi_k^{(j+1)})^2 + (\xi_{k+D}^{(j+1)})^2}$, $1 \leq k \leq D$

3.2b) Let λ be the median of the $m^{(j)}$ th largest and $(m^{(j)} + 1)$ th largest elements in $\{l_k^{(j+1)}\}$. For each

$k : 1 \leq k \leq D$, if $l_k^{(j+1)} \neq 0$, set $[\beta_k^{(j+1)}, \beta_{k+D}^{(j+1)}] \leftarrow [\xi_k^{(j+1)}, \xi_{k+D}^{(j+1)}] \Theta_{HR}(l_k^{(j+1)}; \lambda, \eta) / l_k^{(j+1)}$; set $\beta_k^{(j+1)} = \beta_{k+D}^{(j+1)} = 0$ otherwise.

NON-GROUP FORM:

3.2') $\beta^{(j+1)} \leftarrow \Theta_{HR}(\xi^{(j+1)}; \lambda, \eta)$, where λ is the median of the $m^{(j)}$ th largest component and the $(m^{(j)} + 1)$ th largest component of $|\beta^{(j+1)}|$;

end while

deliver Remaining dimensions after screening: $\{f \in \mathcal{F} : \|\beta_f^{(j+1)}\|_2 \neq 0\}$ (group version) or $\{k : 1 \leq k \leq 2D, \beta_k^{(j+1)} \neq 0\}$ (non-group version).

An interesting observation is that with $\beta^{(0)} = \mathbf{0}$, the first iteration step of Algorithm 2 ranks the frequencies based on $\mathbf{X}^T \mathbf{y}$. In other words, the correlation between the signal \mathbf{y} and each dictionary atom is examined separately, to determine the candidate dimensions, see [41]. Of course, this type of single frequency analysis is merely marginal and does not amount to joint modeling, the resulting crude ranking not suitable for super resolution problems due to the existence of many correlated frequency predictors. Algorithm 2 iterates and avoids such greediness.

GIST screening is pretty flexible and useful, even if sparsity is not desired. It can be applied at any given value of ϑ (possibly greater than 1) and yields a meaningful result for super-resolution spectral problems.

F. GIST framework

We introduce the complete GIST framework to solve the spectral estimation problem. Fig. 4 shows the flowchart outline.

1) *Dictionary Construction and Normalization*: We construct an overcomplete dictionary through (2) with sufficiently high resolution. Then standardize the data, by (a) centering \mathbf{y} and (b) normalizing each predictor column in \mathbf{X} to have mean 0 and variance 1. After the standardization, all predictors are equally extended in the predictor space.

- 2) *GIST Spectrum Screening*: This step can greatly reduce the computational complexity. We perform the iterative probabilistic screening to remove a number of nuisance frequency components and keep ϑN candidate predictors with $\vartheta < 1$ (say $\vartheta = 0.5$) to achieve supervised dimension reduction. See Section III-E for details.
- 3) *Model Fitting*: For each given value of the regularization parameter in a predefined grid, run the iterative group-thresholding algorithm developed in Section III-B to obtain a local optimum to (7). All such solutions are collected to form a solution path parameterized by the regularization parameters.
- 4) *Model Selection*: An optimal solution $\hat{\beta}_{opt}$ is selected from the solution path based on a data-resampling version of high-dimensional BIC (Section III-D).
- 5) *Spectrum Recovery*: The signal can be reconstructed from the coefficient estimate. The amplitudes are estimated by $A(f_k) = \sqrt{\hat{\beta}_{opt,k}^2 + \hat{\beta}_{opt,D+k}^2}$, $1 \leq k \leq D$.

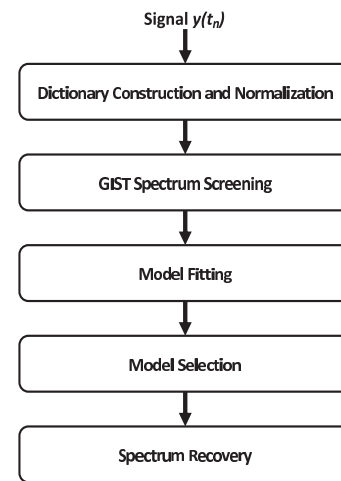


Fig. 4: The flowchart of the GIST framework for solving the spectral estimation problem.

IV. EXPERIMENTS

We conduct simulation experiments to show the performance of GIST fitting Algorithm 1 in sparse spectral estimation, and the power of GIST screening Algorithm 2 in fast computation (with little performance loss in frequency detection).

A. Simulation setup

Consider a discrete real-valued signal given by

$$y(t_n) = \sum_{f_k \in \mathcal{N}z^*} A_k \cos(2\pi f_k t_n + \phi_k) + e(t_n), \quad (22)$$

where $e(t_n)$ is white Gaussian noise with variance σ^2 . $N = 100$ training samples are observed at time $t_n = n$, $1 \leq n \leq N$. The spectrum frequency dictionary is constructed by setting the maximum frequency $f_{\max} = 0.5$ Hz, resolution level $\delta = 0.02$ Hz, and the number of frequency bins $D = f_{\max}/\delta = 250$ (and thus 500 atoms). Using the notation in Section

III-C (cf. (19)), we set $nz^* = \{0.248, 0.25, 0.252, 0.398, 0.4\}$, the associated amplitudes A_k and phases ϕ_k given by $[2, 4, 3, 3.5, 3]$ and $[\pi/4, \pi/6, \pi/3, \pi/5, \pi/2]$, respectively. We vary the noise level by $\sigma^2 = 1, 4, 8$ to study the algorithmic performance with respect to SNR.

Due to random fluctuation, reporting frequency identification for *one* particular simulation dataset is meaningless. Instead, we simulated each model 50 times to enhance stability, where at each run $e(t_n)$ are i.i.d. following $\mathcal{N}(0, \sigma^2)$.

Our simulations were performed in MATLAB R2010b and Win7 Professional 32-bit OS, on a desktop with an Intel(R) Core(TM)2 Quad 2.66 GHz processor and 4GB memory.

B. Experimental Results

1) *Comparison with some existing methods:* To compare with the advocated **group hard-ridge GIST** (or GIST for short), we implemented BP [6], IAA-APES (or IAA for short) [42], SPICE [43], LZA-F [14], CG-SLIM (or SLIM for short) [40]. To make a fair and realistic comparison, we used a common stopping criterion: the number of iterations reaches 200 or the change in β is less than $1e-4$. In GIST, we set $\vartheta = 0.25$ to give the cardinality bound in screening, and used SCV-BIC for parameter tuning. The algorithmic parameters in the other methods took default values suggested in the literature. (For example, the q parameter in SLIM is chosen to be 1, as recommended and used in the numerical examples of [40].) Figs. 5 and 6 show the frequency identification rates in 50 simulation runs for each of the methods under $\sigma^2 = 1$ and $\sigma^2 = 8$, respectively. That is, given each algorithm, we plotted the percentage of identifying f_k or $\hat{\beta}_{f_k} \neq 0$ in all runs, for every f_k in the dictionary. The blue solid lines show such identification rates, while the red dotted lines (with star marks at 100%) label the true frequencies. The plot for $\sigma^2 = 4$ is similar to Fig. 6; we do not show it here due to the page limit. We also included the running time (averaged over 50 runs) in Table I to reflect the computational cost.

TABLE I: Average runtime in seconds of different algorithms, with varying values of σ^2 at 1, 4, and 8.

	$\sigma^2 = 1$	$\sigma^2 = 4$	$\sigma^2 = 8$
BP	0.80	0.72	0.71
LZA-F	1.64	1.88	1.97
IAA	1.18	1.09	1.13
SLIM	3.77	3.71	3.70
SLIM with FFT	0.10	0.10	0.10
SPICE	4.31	4.23	4.22
GIST	1.40	1.40	1.39

IAA, SPICE, LZA-F, and SLIM are *not* capable of producing inherently sparse estimates. One must make a somewhat ad-hoc choice of the cutoff value τ to discern the present frequencies. We set $\tau = 1e - 2$ in performing such post-truncation. It behaved better than $\tau = 1e - 3$ or $\tau = 1e - 4$ in experimentation (which gave similar yet worse detection performance).

BP, though super fast, missed the frequency components at 0.25 and 0.4 all the time. An improvement is offered by the CG-SLIM which makes use of the group l_1 regularization. In [40], CG-SLIM is recommended to run for only 20 iteration steps (whereas the simulated signals there had very mild noise contamination, with $\sigma^2 = 0.001$). Here, we increased the maximum number of iterations to 200 for better identification, without sacrificing much efficiency. Otherwise CG-SLIM gave much poorer spectrum recovery in experiments.

SLIM is free of parameter tuning, because from a Bayesian perspective SLIM estimates the noise variance σ^2 in addition to the coefficient vector β . Unfortunately, we found all the variance estimates from SLIM were severely biased downward—for example, for $\sigma^2 = 8$, the mean $\hat{\sigma}^2$ in 50 runs is about $3e - 5$. This is perhaps the reason why SLIM failed in super-resolution recovery: with such a small σ^2 estimate, the threshold level tends to be very low, and thus SLIM always overselects. Seen from the figures, SLIM results in many spurious frequencies, some arising in more than 60 percent of the datasets. IAA is even worse and does not seem to have the ability to super-resolve.

It is observed that LZA-F may seriously mask the true components. In addition, with moderate/large noise contamination, we found that LZA-F may be unstable and produce huge errors. Because the design of LZA-F is to approximate the l_0 regularization, we substituted the hard-thresholding for Θ in GIST, which solves the exact l_0 -penalized problem. However, the high miss rates of the l_0 -type regularization are still commonly seen, and the resulting models are often over-sparse. To give an explanation of this under-selection, notice that the l_0 regularization either kills or keeps, thereby offering no shrinkage at all for nonzero coefficients. To attain the appropriate extent of shrinkage especially when the noise is not too small, it has to kill more predictors than necessary. As a conclusion, inappropriate nonconvex penalties may seriously mask true signal components.

In our experiments, SPICE performs well. GIST is much better and shows more concentrated signal power at the true frequencies. It produces very few spurious frequencies, and in terms of computation, it is much more efficient than SPICE (Table I). GIST adapts to SNR and is both stable and scalable.

Finally, a recent proposal of using the FFT for matrix-vector multiplication [40] was shown to be very effective: for SLIM, the average running time dropped from about 3.7 seconds to 0.1 seconds. The computational trick can be applied to all of the methods discussed here. However, it restricts to uniformly sampled data with Fourier dictionaries. We did not use the FFT implementation for the other methods. (GIST algorithms and analyses are general and do not have such restrictions, see Section III and Section V.)

2) *Probabilistic spectral screening:* We examine the performance of the GIST-screening Algorithm 2 in this experiment. The candidate ratio ϑ determines the dimensions (ϑN) of the reduced predictor space. Therefore, the lower the value of ϑ , the more efficient the computation, but also the higher the risk of mistakenly removing some true components. Our screening technique turns out to be pretty successful: even if we choose ϑN to be as small as 25 (which can be even lower),

it never misses any true frequency component. Fig. 7 shows the frequency location of 100, 50, and 25 remaining atoms, respectively, in GIST screening. The selected frequencies are non-uniform, and the density near the true spectra is much higher.

Next, we make a much more challenging problem by modifying the signal to have 10 present frequency components at 0.24, 0.242, \dots , 0.282 and large noise variance $\sigma^2 = 10$. Fig. 8 shows both the detection miss rates and the computation time, averaged over 50 runs. The miss rate is the mean of $|\{i : \hat{\beta}_i^* \neq 0, \beta_i^* = 0\}| / |\{i : \beta_i^* \neq 0\}|$ in all simulations, where $|\cdot|$ is the cardinality of a set. The plotted time is the total running time of both GIST screening and model fitting and selection. The empirical experience is that GIST-screening is safe when ϑN is roughly 3 times greater than the number of truly relevant atoms. It reduces the computation complexity significantly with little performance lost.

3) *Misspecified resolution level*: In super-resolution spectral selection, the frequency resolution level δ used in dictionary construction is customized by users. This requires the knowledge of a lower bound on frequency spacing. We are particularly interested in the performance of GIST when δ is misspecified in reference to the truth.

In this experiment, we set the signal frequencies at 0.2476, 0.2503, 0.2528, 0.3976, 0.4008, with amplitudes A_k and phases ϕ_k unchanged. Clearly, the ideal frequency resolution to resolve this signal should be no more than 0.0001 Hz.

We chose $\delta = 0.002$, 20 times as large as the required resolution. The results are nearly identical to Figs. (5f), (6f) (not shown due to space limitation). The crude resolution specification makes GIST unable to recover the true frequencies. On the other hand, the most frequently identified frequencies are 0.248, 0.25, 0.252, 0.398, 0.4, and a comparison shows that this is the best approximation in the given frequency grid. (For example, 0.398Hz is the closest frequency in the grid $\{0, 0.02, \dots, 0.396, 0.398, \dots, 0.5\}$ to 0.3976Hz.) This phenomenon is also seen in many other experiments: GIST gives the best possible identification to approximate the true frequencies, with the quantization error determined by the resolution level.

V. CONCLUSIONS

We have presented a sparsity-based GIST framework to tackle the super-resolution challenge in spectral estimation. It is able to handle nonconvex penalties and take the pairing structure of sine and cosine atoms into account in regularizing the model. The $l_0 + l_2$ type hard-ridge penalty was shown to be able to dramatically improve the popular convex l_1 penalty as well as the nonconvex l_0 penalty. Its variant, the iterative probabilistic spectrum screening, can be used for supervised dimension reduction and fast computation. In parameter tuning, the SCV criterion overcomes the training inconsistency issue of the plain CV and is much more computationally efficient. GIST can be applied to unevenly sampled signals (in which case the sampling time sequence $\{t_n\}_{1 \leq n \leq N}$ is not uniform) with guaranteed convergence (cf. Theorem 1 and Theorem 3).

It is worth mentioning that in our algorithm design and theoretical analyses, the only use of the Fourier frequency dictionary was to extract the atom grouping manner. Our methodology carries over to any type of dictionary as arising in signal processing, wavelets, and statistics. For example, although we focused on real-valued signals in the paper, for complex-valued signals, say, $\mathbf{y} = [y(t_n)] \in \mathbb{C}^{N \times 1}$ observed at t_n ($1 \leq n \leq N$) and the candidate frequency grid given by f_k ($1 \leq k \leq D$), a complex dictionary \mathbf{X} can be constructed as $[\exp(i2\pi f_k t_n)] \in \mathbb{C}^{N \times D}$ in place of (2), with $\boldsymbol{\beta} = [\beta_1, \dots, \beta_D]^T \in \mathbb{C}^D$. The group penalized model then minimizes $\frac{1}{2} \|\mathbf{y} - \alpha - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \sum_{k=1}^D P(\|\beta_k\|_2; \lambda)$ where $\|\beta_k\|_2 = \sqrt{\text{Re}(\beta_k)^2 + \text{Im}(\beta_k)^2}$, simply the complex norm of β_k . It is straightforward to extend all our algorithms and analyses to this problem. On the other hand, for real-valued signals, the formulation using the sine-cosine predictor matrix in (2) does not involve any imaginary/complex number processing in implementation.

Some future research topics include the extension of GIST to non-Gaussian and/or multivariate signals.

APPENDIX A PROOF OF THEOREM 1

We show the result for the group form only. The proof for the non-group form is similar and simpler. The following **continuity assumption** is made throughout the proof:

Assumption 2: Θ is continuous at any point in the closure of $\{\boldsymbol{\xi}^{(j)}\}$.

For continuous thresholding rules such as soft-thresholding, this regularity condition always holds. Practically used thresholding rules (such as hard-thresholding) have few discontinuity points and such discontinuities rarely occur in any real application. For $\omega = 1$, see [29] for the proof details. In the following, we assume $0 < \omega < 1$.

Note that G is quadratic and convex in $\boldsymbol{\beta}$, $\boldsymbol{\xi}$, and $\boldsymbol{\zeta}$, but possibly nonconvex and nonsmooth in $\boldsymbol{\gamma}$.

Lemma A.1: Given an arbitrary thresholding rule Θ , let P be any function satisfying

$$P(\theta; \lambda) - P(0; \lambda) = P_\Theta(\theta; \lambda) + q(\theta; \lambda)$$

where $P_\Theta(\theta; \lambda) \triangleq \int_0^{|\theta|} (\sup\{s : \Theta(s; \lambda) \leq u\} - u) du$, $q(\theta; \lambda)$ is nonnegative and $q(\Theta(t; \lambda)) = 0$ for all t . Then, the minimization problem

$$\min_{\boldsymbol{\beta}} \frac{1}{2} \|\mathbf{y} - \boldsymbol{\beta}\|_2^2 + P(\|\boldsymbol{\beta}\|_2; \lambda)$$

has a unique optimal solution given by $\hat{\boldsymbol{\beta}} = \vec{\Theta}(\mathbf{y}; \lambda)$ for every \mathbf{y} provided that $\Theta(\cdot; \lambda)$ is continuous at $\|\mathbf{y}\|_2$.

See [29] for its proof.

Given $\boldsymbol{\beta}$ and $\boldsymbol{\xi}$, the problem of minimizing G over $(\boldsymbol{\gamma}, \boldsymbol{\zeta})$ can be simplified to (detail omitted)

$$\min_{\boldsymbol{\gamma}} \frac{1}{2} \|\boldsymbol{\gamma} - \omega(\mathbf{I} - \boldsymbol{\Sigma})\boldsymbol{\beta} - \omega\mathbf{X}^T\mathbf{y} - (1 - \omega)\boldsymbol{\xi}\|_2^2 + P(\boldsymbol{\gamma}; \lambda),$$

and

$$\min_{\boldsymbol{\zeta}} \frac{1}{2} \frac{1 - \omega}{\omega} [\boldsymbol{\zeta} - \omega(\mathbf{I} - \boldsymbol{\Sigma})\boldsymbol{\beta} - \omega\mathbf{X}^T\mathbf{y} - (1 - \omega)\boldsymbol{\xi}]^T (\mathbf{I} - \boldsymbol{\Sigma})^{-1} [\boldsymbol{\zeta} - \omega(\mathbf{I} - \boldsymbol{\Sigma})\boldsymbol{\beta} - \omega\mathbf{X}^T\mathbf{y} - (1 - \omega)\boldsymbol{\xi}].$$

Based on Lemma A.1, the optimal solutions are

$$\begin{cases} \gamma_{opt} &= \bar{\Theta}(\omega(\mathbf{I} - \Sigma)\boldsymbol{\beta} + \omega\mathbf{X}^T\mathbf{y} + (1 - \omega)\boldsymbol{\xi}; \lambda) \\ \zeta_{opt} &= \omega(\mathbf{I} - \Sigma)\boldsymbol{\beta} + \omega\mathbf{X}^T\mathbf{y} + (1 - \omega)\boldsymbol{\xi}. \end{cases}$$

Therefore, we obtain

$$G(\boldsymbol{\beta}^{(j+1)}, \boldsymbol{\xi}^{(j+1)}, \boldsymbol{\beta}^{(j)}, \boldsymbol{\xi}^{(j)}; \lambda) \leq G(\boldsymbol{\beta}^{(j)}, \boldsymbol{\xi}^{(j)}, \boldsymbol{\beta}^{(j)}, \boldsymbol{\xi}^{(j)}; \lambda) - \frac{1 - \omega}{2\omega}(\boldsymbol{\xi}^{(j+1)} - \boldsymbol{\xi}^{(j)})^T(\mathbf{I} - \Sigma)^{-1}(\boldsymbol{\xi}^{(j+1)} - \boldsymbol{\xi}^{(j)}). \quad (23)$$

On the other hand, given γ and ζ , G can be expressed as a quadratic form in $\boldsymbol{\beta}$ that is positive definite. The same fact holds for $\boldsymbol{\xi}$. It can be computed that

$$\begin{cases} \nabla G_{\boldsymbol{\beta}} = \omega(\mathbf{I} - \Sigma)(\boldsymbol{\beta} - \gamma) + (1 - \omega)(\boldsymbol{\xi} - \zeta) \\ \nabla G_{\boldsymbol{\xi}} = \frac{1 - \omega}{\omega}(\mathbf{I} - \Sigma)^{-1}[\omega(\mathbf{I} - \Sigma)(\boldsymbol{\beta} - \gamma) + (1 - \omega)(\boldsymbol{\xi} - \zeta)], \end{cases}$$

from which it follows that G can be written as $\frac{1}{2}[\omega(\mathbf{I} - \Sigma)(\boldsymbol{\beta} - \gamma) + (1 - \omega)(\boldsymbol{\xi} - \zeta)]^T \omega^{-1}(\mathbf{I} - \Sigma)^{-1}[\omega(\mathbf{I} - \Sigma)(\boldsymbol{\beta} - \gamma) + (1 - \omega)(\boldsymbol{\xi} - \zeta)]$ in addition to the terms involving only γ and ζ . Hence $\boldsymbol{\beta}_{opt} = \boldsymbol{\beta}$ and $\boldsymbol{\xi}_{opt} = \boldsymbol{\xi}$ (though not unique) achieve the minimum. We obtain $G(\boldsymbol{\beta}^{(j+1)}, \boldsymbol{\xi}^{(j+1)}, \boldsymbol{\beta}^{(j+1)}, \boldsymbol{\xi}^{(j+1)}; \lambda) \leq G(\boldsymbol{\beta}^{(j+1)}, \boldsymbol{\xi}^{(j+1)}, \boldsymbol{\beta}^{(j)}, \boldsymbol{\xi}^{(j)}; \lambda) - \frac{1}{2\omega}[\omega(\mathbf{I} - \Sigma)(\boldsymbol{\beta}^{(j)} - \boldsymbol{\beta}^{(j+1)}) + (1 - \omega)(\boldsymbol{\xi}^{(j)} - \boldsymbol{\xi}^{(j+1)})]^T(\mathbf{I} - \Sigma)^{-1}[\omega(\mathbf{I} - \Sigma)(\boldsymbol{\beta}^{(j)} - \boldsymbol{\beta}^{(j+1)}) + (1 - \omega)(\boldsymbol{\xi}^{(j)} - \boldsymbol{\xi}^{(j+1)})]$. Combining this with (23) yields (17).

Assume a subsequence $\boldsymbol{\beta}^{(j_l)} \rightarrow \boldsymbol{\beta}^\circ$ as $l \rightarrow \infty$. Because

$$G(\boldsymbol{\beta}^{(j_l)}, \boldsymbol{\xi}^{(j_l)}, \boldsymbol{\beta}^{(j_l)}, \boldsymbol{\xi}^{(j_l)}) - G(\boldsymbol{\beta}^{(j_{l+1})}, \boldsymbol{\xi}^{(j_{l+1})}, \boldsymbol{\beta}^{(j_{l+1})}, \boldsymbol{\xi}^{(j_{l+1})}) \rightarrow 0,$$

we have $\boldsymbol{\xi}^{(j_l)} - \boldsymbol{\xi}^{(j_{l+1})} \rightarrow 0$ and thus $(\boldsymbol{\beta}^{(j_l)} - \boldsymbol{\beta}^{(j_{l+1})}) \rightarrow 0$. That is, $(1 - \omega)\boldsymbol{\xi}^{(j_l)} + \omega(\boldsymbol{\beta}^{(j_l)} + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(j_l)})) - \boldsymbol{\xi}^{(j_l)} \rightarrow 0$ and $\bar{\Theta}(\boldsymbol{\xi}^{(j_l)}; \lambda) - \boldsymbol{\beta}^{(j_l)} \rightarrow 0$. From $\mathbf{X}^T(\mathbf{y} - \mathbf{X}\boldsymbol{\beta}^{(j_l)}) - \boldsymbol{\xi}^{(j_l)} \rightarrow 0$ and the continuity assumption, $\boldsymbol{\beta}^\circ$ is a group Θ -estimate satisfying (13), and $\lim_{j \rightarrow \infty} G(\boldsymbol{\beta}^{(j)}, \boldsymbol{\xi}^{(j)}, \boldsymbol{\beta}^{(j)}, \boldsymbol{\xi}^{(j)}) = F(\boldsymbol{\beta}^\circ)$.

APPENDIX B PROOF OF THEOREM 2

Recall that \mathcal{F} denotes the frequency set covered by the dictionary \mathbf{X} and we assume all column norms of \mathbf{X} are \sqrt{N} (or the diagonal entries of $\Sigma = \mathbf{X}^T\mathbf{X}$ are equal to N).

Applying Theorem 1, we can characterize any group l_1 estimate $\hat{\boldsymbol{\beta}}$ from Algorithm 1 by

$$\hat{\boldsymbol{\beta}} = \bar{\Theta}(\hat{\boldsymbol{\beta}} + \mathbf{X}^T\mathbf{y}/\tau_0^2 - \Sigma\hat{\boldsymbol{\beta}}/\tau_0^2; \lambda) \quad (24)$$

with $\bar{\Theta}$ being the soft-thresholding function. Let $\mathbf{s} = \mathbf{s}(\boldsymbol{\beta})$ denote a function of $\boldsymbol{\beta}$ satisfying

$$\|\mathbf{s}_f\|_2 \leq 1, \forall f \in z(\boldsymbol{\beta}), \quad \mathbf{s}_f = \boldsymbol{\beta}_f / \|\boldsymbol{\beta}_f\|_2, \forall f \in nz(\boldsymbol{\beta}), \quad (25)$$

and $\mathbf{s}(\boldsymbol{\beta}_f) := [\mathbf{s}(\boldsymbol{\beta})]_f$. We have $\|\mathbf{s}(\boldsymbol{\beta}_f)\|_2 \leq 1, \forall f \in \mathcal{F}$. (In the group l_1 case, \mathbf{s} is a subgradient of $\sum_{f \in \mathcal{F}} \|\boldsymbol{\beta}_f\|_2$.) Then (24) reduces to $\hat{\boldsymbol{\beta}} + \lambda\mathbf{s}(\hat{\boldsymbol{\beta}}) = \hat{\boldsymbol{\beta}} + \mathbf{X}^T\mathbf{y}/\tau_0^2 - \Sigma\hat{\boldsymbol{\beta}}/\tau_0^2$ or

$$\Sigma\hat{\boldsymbol{\beta}} = \mathbf{X}^T\mathbf{y} - \lambda\tau_0^2\mathbf{s}(\hat{\boldsymbol{\beta}}), \quad (26)$$

for some \mathbf{s} satisfying (25).

Lemma B.1: Assume Σ_{nz^*} is nonsingular. Then (26) is equivalent to

$$\begin{cases} \mathbf{S}_{z^*}\hat{\boldsymbol{\beta}}_{z^*} = \mathbf{X}_{z^*}^T\mathbf{e} + \lambda\tau_0^2\Sigma_{z^*,nz^*}\Sigma_{nz^*}^{-1}\mathbf{s}(\hat{\boldsymbol{\beta}}_{nz^*}) - \lambda\tau_0^2\mathbf{s}(\hat{\boldsymbol{\beta}}_{z^*}) \\ \hat{\boldsymbol{\beta}}_{nz^*} = \boldsymbol{\beta}_{nz^*}^* + \Sigma_{nz^*}^{-1}(\mathbf{X}_{nz^*}^T\mathbf{e} - \lambda\tau_0^2\mathbf{s}(\hat{\boldsymbol{\beta}}_{nz^*})) - \Sigma_{nz^*}^{-1}\Sigma_{z^*,nz^*}^T\hat{\boldsymbol{\beta}}_{z^*} \end{cases} \quad (27)$$

where $\mathbf{S}_{z^*} := \Sigma_{z^*} - \Sigma_{z^*,nz^*}\Sigma_{nz^*}^{-1}\Sigma_{nz^*,z^*}$, and $\mathbf{X}_{z^*}^T := \mathbf{X}_{z^*}^T - \Sigma_{z^*,nz^*}\Sigma_{nz^*}^{-1}\mathbf{X}_{nz^*}^T$.

The proof details are given in [33].

Lemma B.2: Suppose $\begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \mathbf{V}\right)$, where \mathbf{V} is a correlation matrix. Then for any M , $P(z_1^2 + z_2^2 > M^2) \leq P(\xi > M^2/2)$ with $\xi \sim \chi^2(2)$.

From $\|\mathbf{V}\|_2 \leq \|\mathbf{V}\|_F \leq 2$, $2\mathbf{I} - \mathbf{V}$ is positive semi-definite. Let z'_1, z'_2 be independent standard Gaussian random variables. We get $P(z_1^2 + z_2^2 > M^2) \leq P((z'_1\sqrt{2})^2 + (z'_2\sqrt{2})^2 > M^2) = P(\xi > M^2/2)$ from Anderson's inequality [44].

Lemma B.3: Suppose $\xi \sim \chi^2(2)$. Then for any M , $P(\xi > 2M^2) \leq M^2e^{-(M^2-1)}$.

See, e.g., [45] for a proof of this χ^2 tail bound.

Let $\mathbf{X}_f^T = \mathbf{X}_f^T - \Sigma_{f,nz^*}\Sigma_{nz^*}^{-1}\mathbf{X}_{nz^*}^T, \forall f \in z^*$. From Lemma B.1, we have $P_1 \geq P(A \cap V)$, with $A := \{\|\mathbf{X}_f^T\mathbf{e} + \lambda\tau_0^2\Sigma_{f,nz^*}\Sigma_{nz^*}^{-1}\mathbf{s}(\hat{\boldsymbol{\beta}}_{nz^*})\|_2 \leq \lambda\tau_0^2, \forall \mathbf{s}$ satisfying (25), $\forall f \in z^*\}$, $V := \{\|\Sigma_{nz^*}^{-1}\mathbf{X}_{nz^*}^T\mathbf{e}\|_2 + \lambda\tau_0^2\|\Sigma_{nz^*}^{-1}\mathbf{s}(\hat{\boldsymbol{\beta}}_{nz^*})\|_2 < \|\boldsymbol{\beta}_f^*\|_2, \forall \mathbf{s}$ satisfying (25), $\forall f \in nz^*\}$. Therefore, $1 - P_1 \leq P(A^c \cup V^c) \leq P(A^c) + P(V^c)$.

From the definition of κ , $\|\Sigma_{f,nz^*}\Sigma_{nz^*}^{-1}\mathbf{s}(\hat{\boldsymbol{\beta}}_{nz^*})\|_2 \leq \kappa\sqrt{p_{nz^*}}\|\Sigma_{nz^*}^{-1}\|_2N\sqrt{p_{nz^*}} = \kappa p_{nz^*}/\mu, \forall f \in z^*$. It follows that

$$P(A^c) \leq p_{z^*}P(\|e'_f\|_2 \geq (1 - \kappa p_{nz^*}/\mu)\lambda\tau_0^2/(\sigma\sqrt{N})) \quad (28)$$

where $e'_f = \mathbf{X}_f^T\mathbf{e}/(\sigma\sqrt{N}) \sim N(0, \mathbf{S}_{z^*}/N)$ because $\mathbf{X}_{z^*}^T\mathbf{X}_{z^*} = \mathbf{S}_{z^*}$. Define $M := (1 - \kappa p_{nz^*}/\mu)\lambda\tau_0^2/(\sigma\sqrt{N})$. Based on Lemma B.2 and Lemma B.3 and the fact that the diagonal entries of $\mathbf{S}_{z^*} = \Sigma_{z^*} - \Sigma_{z^*,nz^*}\Sigma_{nz^*}^{-1}\Sigma_{nz^*,z^*}$ are all less than or equal to N , we obtain a bound for (28):

$$P(A^c) \leq \frac{e}{4}p_{z^*}M^2e^{-M^2/4}. \quad (29)$$

Next we bound $P(V^c)$. Suppose the spectral decomposition of Σ_{nz^*} is given by $\mathbf{U}\mathbf{D}\mathbf{U}^T$ with the i th row of \mathbf{U} given by \mathbf{u}_i^T , then we can represent $\Sigma_{nz^*}^{-1}$ as $[\mathbf{u}_i^T\mathbf{D}^{-1}\mathbf{u}_i]$, and thus $\text{diag}(\Sigma_{nz^*}^{-1}) \leq 1/(N\mu)$. Moreover, from $\|\Sigma_{nz^*}^{-1}\mathbf{s}\|_2 \leq \sqrt{p_{nz^*}}/(N\mu)$, $\|\Sigma_{nz^*}^{-1}\mathbf{s}_f\|_2 \leq \sqrt{p_{nz^*}}/(N\mu), \forall f \in nz^*$. Introduce $e'' = \sqrt{\mu N}\Sigma_{nz^*}^{-1}\mathbf{X}_{nz^*}^T\mathbf{e}/(\sigma) \sim N(0, \mu N\Sigma_{nz^*}^{-1})$. From the last two lemmas,

$$\begin{aligned} P(V^c) &\leq p_{nz^*}P(\|e''_f\|_2 \geq (\min_{f \in nz^*} \|\boldsymbol{\beta}_f^*\|_2 - \frac{\lambda\tau_0^2\sqrt{p_{nz^*}}}{\mu N})\frac{\sqrt{\mu N}}{\sigma}) \\ &\leq \frac{e}{4}p_{nz^*}L^2e^{-L^2/4}, \end{aligned}$$

where $L := (\min_{f \in nz^*} \|\boldsymbol{\beta}_f^*\|_2 - \lambda\tau_0^2\sqrt{p_{nz^*}}/(\mu N))\frac{\sqrt{\mu N}}{\sigma}$.

The proof of (21) for the hard-ridge thresholding follows similar lines. First, define $\mathbf{s}(\boldsymbol{\beta}; \lambda, \eta)$ as

$$\|\mathbf{s}_f\|_2 \leq 1, \forall f \in z(\boldsymbol{\beta}) \quad \text{and} \quad \mathbf{s}_f = \frac{\eta}{\lambda}\boldsymbol{\beta}_f, \forall f \in nz(\boldsymbol{\beta}). \quad (30)$$

Then similar to (26) we have

$$\Sigma \hat{\beta} = \mathbf{X}^T \mathbf{y} - \lambda \tau_0^2 \mathbf{s}(\hat{\beta}; \lambda, \eta), \quad (31)$$

Let $\mathbf{X}_{z^*}''^T := \mathbf{X}_{z^*}^T - \Sigma_{z^*,nz^*} \Sigma_{nz^*}^{-1} [\mathbf{I} - \eta(\Sigma_{nz^*} + \eta \mathbf{I})^{-1}] \mathbf{X}_{nz^*}^T$. To bound P_{02} , from Lemma B.1 and (30), we write (31) as

$$\begin{cases} \lambda \tau_0^2 \mathbf{s}(\hat{\beta}_{z^*}; \lambda, \eta) + \mathbf{S}_{z^*} \hat{\beta}_{z^*} = \\ \mathbf{X}_{z^*}''^T \mathbf{e} + \eta \tau_0^2 \Sigma_{z^*,nz^*} (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1} \beta_{nz^*}^*, \\ \hat{\beta}_{nz^*} = (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1} \Sigma_{nz^*} \beta_{nz^*}^* + \\ (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1} \mathbf{X}_{nz^*}^T \mathbf{e}, \end{cases}$$

where we used $\Sigma_{nz^*}^{-1} (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1} \Sigma_{nz^*} = (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1}$. It follows that

$$P_{02} \geq P(\exists \mathbf{s} \text{ satisfying (30) s.t. } \hat{\beta}_{z^*} = \mathbf{0} \text{ and } \|\hat{\beta}_f\|_2 \geq \lambda/(1+\eta), \forall f \in nz^*) \geq P(A \cap V)$$

with

$$\begin{aligned} A &:= \{ \|\mathbf{X}_f''^T \mathbf{e} + \eta \tau_0^2 \Sigma_{f,nz^*} (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1} \beta_{nz^*}^*\|_2 \leq \lambda \tau_0^2, \\ &\quad \forall f \in z^* \}, \\ V &:= \{ \|\Sigma_{nz^*}^{-1} (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1} \Sigma_{nz^*} \beta_{nz^*}^*\|_f \\ &\quad + \|(\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1} \mathbf{X}_{nz^*}^T \mathbf{e}\|_f \geq \frac{\lambda}{1+\eta}, \forall f \in nz^* \}. \end{aligned}$$

For $\mathbf{e}' = \mathbf{X}_{z^*}''^T \mathbf{e}/(\sigma\sqrt{N})$ and $\mathbf{e}'' = \frac{\mu N + \eta \tau_0^2}{\sqrt{\mu N}} (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1} \mathbf{X}_{nz^*}^T \mathbf{e}/\sigma$, their covariance matrices are computed as

$$(\Sigma_{z^*} - \Sigma_{z^*,nz^*} [\mathbf{I} - \eta^2 (\Sigma_{nz^*} + \eta \mathbf{I})^{-2}] \Sigma_{nz^*}^{-1} \Sigma_{z^*,nz^*}^T) / N$$

and

$$\frac{(\mu N + \eta \tau_0^2)^2}{\mu N} (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1} \Sigma_{nz^*} (\Sigma_{nz^*} + \eta \tau_0^2 \mathbf{I})^{-1},$$

respectively. Furthermore, it is not difficult to see that their diagonal entries are bounded by 1 (under $\eta \leq \mu N / \tau_0^2$). Therefore,

$$\begin{aligned} 1 - P_{02} &\leq P(\exists f \in z^* \text{ s.t. } \|\mathbf{e}'_f\|_2 \geq \frac{1}{\sigma\sqrt{N}} (\lambda \tau_0^2 - \frac{\eta \tau_0^2 \kappa \sqrt{p_{nz^*}} \|\beta_{nz^*}^*\|_2}{\mu N + \eta \tau_0^2})) \\ &\quad + P(\exists f \in nz^* \text{ s.t. } \|\mathbf{e}''_f\|_2 \geq (\iota - \frac{\lambda}{1+\eta}) \frac{\mu N + \eta \tau_0^2}{\sqrt{\mu N} \sigma}) \\ &\leq \frac{e}{4} p_{z^*} M'^2 e^{-M'^2/4} + \frac{e}{4} p_{z^*} L'^2 e^{-L'^2/4}, \end{aligned}$$

where $M' := \frac{1}{\sigma\sqrt{N}} (\lambda \tau_0^2 - \frac{\eta \tau_0^2}{\mu N + \eta \tau_0^2} \kappa \sqrt{p_{nz^*}} \|\beta_{nz^*}^*\|_2)$ and $L' := (\iota - \frac{\lambda}{1+\eta}) \frac{\mu N + \eta \tau_0^2}{\sqrt{\mu N} \sigma}$.

APPENDIX C PROOF OF THEOREM 3

We show the proof for the group form only. The proof in the non-group case is similar (and simpler). First, we introduce a group quantile thresholding rule $\tilde{\Theta}^\#(\cdot; m, \eta)$ as a variant of the hard-ridge thresholding. Given $1 \leq m \leq |\mathcal{F}|$ and $\eta \geq 0$, $\tilde{\Theta}^\#(\cdot; m, \eta) : \mathbf{a} \in \mathbb{R}^{2D} \rightarrow \mathbf{b} \in \mathbb{R}^{2D}$ is defined as follows: $\mathbf{b}_f = \mathbf{a}_f / (1 + \eta)$ if $\|\mathbf{a}_f\|_2$ is among the m largest norms in the set of $\{\|\mathbf{a}_f\|_2 : f \in \mathcal{F}\}$, and $\mathbf{b}_f = \mathbf{0}$ otherwise. In the case

of ties, a random tie breaking rule is used. With the notation, $\beta^{(j+1)} = \tilde{\Theta}^\#(\xi^{(j+1)}; m, \eta)$.

From the algorithm, $nz(\beta^{(j)}) \leq m$ is obvious. To prove the function value decreasing property, we introduce the following lemma.

Lemma C.1: $\hat{\beta} = \tilde{\Theta}^\#(\xi; m, \eta)$ is a globally optimal solution to

$$\min_{\beta} \frac{1}{2} \|\xi - \beta\|_2^2 + \frac{\eta}{2} \|\beta\|_2^2 =: f_0(\beta; \eta) \quad \text{s.t. } nz(\beta) \leq m. \quad (32)$$

Let $I \subset \mathcal{F}$ with $|I| = m$. Assuming $\beta_{I^c} = \mathbf{0}$, we get the optimal solution $\hat{\beta}$ with $\hat{\beta}_I = \xi_I / (1 + \eta)$. It follows that $f_0(\hat{\beta}; \eta) = \frac{1}{2} \|\xi\|_2^2 - \frac{1}{2(1+\eta)} \sum_{f \in I} \|\xi_f\|_2^2$. Hence the group quantile thresholding $\tilde{\Theta}^\#(\xi; m, \eta)$ yields a global minimizer.

Based on this lemma, (17) can be proved following the lines of Appendix A. Details are omitted.

Now suppose that $\eta > 0$ and the following **no tie occurring assumption** holds:

Assumption B: No ties occur in performing $\tilde{\Theta}^\#(\xi; m, \eta)$ for any ξ in the closure of $\{\xi^{(j)}\}$, i.e., either $\|\xi_{(m)}\|_2 > \|\xi_{(m+1)}\|_2$ or $\|\xi_{(m)}\|_2 = \|\xi_{(m+1)}\|_2 = 0$ occurs, where $\|\xi_{(m)}\|_2$ and $\|\xi_{(m+1)}\|_2$ are the m -th and $(m+1)$ -th largest norms in $\{\|\xi_f\|_2 : f \in \mathcal{F}\}$.

From $P(\beta^{(j+1)}; \eta) \leq G(\beta^{(j+1)}, \xi^{(j+1)}, \beta^{(j+1)}, \xi^{(j+1)}; \eta) \leq G(\beta^{(1)}, \xi^{(1)}, \beta^{(1)}, \xi^{(1)}; \eta)$ and $\eta > 0$, $\beta^{(j)}$ is uniformly bounded. Let β° be a limit point of $\beta^{(j)}$ satisfying $\beta^\circ = \lim_{l \rightarrow \infty} \beta^{(j_l)}$. Then from $G(\beta^{(j_l)}, \xi^{(j_l)}, \beta^{(j_l)}, \xi^{(j_l)}) - G(\beta^{(j_l+1)}, \xi^{(j_l+1)}, \beta^{(j_l+1)}, \xi^{(j_l+1)}) \rightarrow 0$, we have $\xi^{(j_l)} - \xi^{(j_l+1)} \rightarrow 0$ and thus $\beta^{(j_l)} - \beta^{(j_l+1)} \rightarrow 0$. That is, $(1 - \omega)\xi^{(j_l)} + \omega(\beta^{(j_l)} + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta^{(j_l)})) - \xi^{(j_l)} \rightarrow 0$ and $\tilde{\Theta}^\#(\xi^{(j_l)}; m, \eta) - \beta^{(j_l)} \rightarrow 0$. We obtain $\beta^\circ = \lim_{l \rightarrow \infty} \tilde{\Theta}^\#(\beta^{(j_l)} + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta^{(j_l)}); m, \eta)$. Because the limit of $\beta^{(j_l)} + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta^{(j_l)})$ exists, it is easy to show that β° is a fixed point of

$$\beta = \tilde{\Theta}^\#(\beta + \mathbf{X}^T(\mathbf{y} - \mathbf{X}\beta); m, \eta).$$

Let $nz^\circ = nz(\beta^\circ)$. By the definition of $\tilde{\Theta}^\#$, $\beta_{nz^\circ}^\circ = \beta_{nz^\circ}^\circ / (1 + \eta) + \mathbf{X}_{nz^\circ}^T(\mathbf{y} - \mathbf{X}\beta^\circ) / (1 + \eta)$, and thus $\eta \beta_{nz^\circ}^\circ + \mathbf{X}_{nz^\circ}^T(\mathbf{y} - \mathbf{X}\beta_{nz^\circ}^\circ) = \mathbf{0}$. But this is the KKT equation of the convex optimization problem

$$\min_{\gamma} \frac{1}{2} \|\mathbf{y} - \mathbf{X}_I \gamma\|_2^2 + \frac{\eta}{2} \|\gamma\|_2^2 \quad (33)$$


with $I = nz^\circ$ given. Therefore, β° is the ridge regression estimate restricted to \mathbf{X}_{nz° . Note that $\eta > 0$ guarantees its uniqueness given I .

Next, based on Ostrowski's convergence theorem, the boundedness of $\beta^{(j)}$ and $\lim \|\beta^{(j)} - \beta^{(j+1)}\| = 0$ imply that the set of limit points of $\beta^{(j)}$ (denoted by L) must be connected. On the other hand, the set of all restricted ridge regression estimates (denoted by R) is finite. Therefore, $\lim \beta^{(j)} = \beta^\circ$. The convergence of $\xi^{(j)}$ is guaranteed as well.

REFERENCES

- [1] P. Stoica and R. Moses, *Spectral Analysis of Signals*. Pearson/Prentice Hall, 2005.

- [2] P. Stoica, J. Li, and H. He, "Spectral analysis of nonuniformly sampled data: a new approach versus the periodogram," *IEEE Transactions on Signal Processing*, vol. 57, no. 3, pp. 843–858, 2009.
- [3] N. Lomb, "Least-squares frequency analysis of unequally spaced data," *Astrophysics and space science*, vol. 39, no. 2, pp. 447–462, 1976.
- [4] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, 1986.
- [5] J. Li and P. Stoica, "Efficient mixed-spectrum estimation with applications to target feature extraction," *IEEE Transactions on Signal Processing*, vol. 44, no. 2, pp. 281–295, 2002.
- [6] S. Chen and D. Donoho, "Application of basis pursuit in spectrum estimation," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 3, 1998, pp. 1865–1868.
- [7] J. Li, P. Stoica, and E. Corporation, *MIMO radar signal processing*. Wiley Online Library, 2009.
- [8] S. Bourguignon, H. Carfantan, and J. Idier, "A sparsity-based method for the estimation of spectral lines from irregularly sampled data," *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 4, p. 575, 2007.
- [9] E. Candes, M. Wakin, and S. Boyd, "Enhancing sparsity by reweighted l_1 minimization," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5, pp. 877–905, 2008.
- [10] J. Fuchs and B. Delyon, "Minimal L_1 -norm reconstruction function for oversampled signals: applications to time-delay estimation," *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1666–1673, 2002.
- [11] S. Bourguignon, H. Carfantan, and T. Böhm, "Sparspec: a new method for fitting multiple sinusoids with irregularly sampled data," *Astron. Astrophys.*, vol. 462, no. 1, pp. 379–387, 2007.
- [12] T. Blumensath and M. Davies, "Iterative hard thresholding for compressed sensing," *Applied and Computational Harmonic Analysis*, vol. 27, no. 3, pp. 265–274, 2009.
- [13] —, "Normalized iterative hard thresholding: Guaranteed stability and performance," *IEEE Journal on Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 298–309, 2009.
- [14] M. Hyder and K. Mahata, "An l_0 norm based method for frequency estimation from irregularly sampled data," in *Proceedings of IEEE ICASSP*, 2010, pp. 4022–4025.
- [15] —, "An improved smoothed approximation algorithm for sparse representation," *IEEE Transactions on Signal Processing*, vol. 58, no. 4, pp. 2194–2205, april 2010.
- [16] D. Donoho, M. Elad, and V. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Transactions on Information Theory*, vol. 52, no. 1, pp. 6–18, 2006.
- [17] E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [18] C. Zhang and J. Huang, "The sparsity and bias of the Lasso selection in high-dimensional linear regression," *Annals of Statistics*, vol. 36, no. 4, pp. 1567–1594, 2008.
- [19] P. Zhao and B. Yu, "On model selection consistency of lasso," *Journal of Machine Learning Research*, vol. 7, no. 2, pp. 2541–2563, 2006.
- [20] E. J. Candès and Y. Plan, "Near-ideal model selection by l_1 minimization," *Ann. Statist.*, vol. 37, no. 5A, pp. 2145–2177, 2009.
- [21] J. Scargle, "Studies in astronomical time series analysis. II-Statistical aspects of spectral analysis of unevenly spaced data," *The Astrophysical Journal*, vol. 263, pp. 835–853, 1982.
- [22] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [23] B. Natarajan, "Sparse approximate solutions to linear systems," *SIAM Journal on Computing*, vol. 24, no. 2, pp. 227–234, 1995.
- [24] J. Holland, "Genetic algorithms," *Scientific American*, vol. 267, no. 1, pp. 66–72, 1992.
- [25] G. Harikumar and Y. Bresler, "A new algorithm for computing sparse solutions to linear inverse problems," in *Proceedings of IEEE ICASSP*, 1996, pp. 1331–1334.
- [26] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2005.
- [27] Y. Eldar, P. Kuppinger, and H. Bölcskei, "Block-sparse signals: Uncertainty relations and efficient recovery," *IEEE Transactions on Signal Processing*, vol. 58, no. 6, pp. 3042–3054, Jun. 2010.
- [28] T.-H. Li, *Time Series with Mixed Spectra*. Chapman and Hall/CRC, 2013.
- [29] Y. She, "An iterative algorithm for fitting nonconvex penalized generalized linear models with grouped predictors," *Computational Statistics & Data Analysis*, vol. 56, pp. 2976–2990, 2012.
- [30] W. James and C. Stein, "Estimation with quadratic loss," in *Proceedings of the 4th Berkeley Symposium on Mathematical Statistics and Probability, Vol. 1*. University of California Press, 1961, pp. 361–379.
- [31] H. Zou and T. Hastie, "Regularization and variable selection via the elastic net," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 2, pp. 301–320, 2005.
- [32] Y. She and A. B. Owen, "Outlier detection using nonconvex penalized regression," *Journal of the American Statistical Association*, vol. 106, no. 494, pp. 626–639, 2011.
- [33] Y. She, "Thresholding-based iterative selection procedures for model selection and shrinkage," *Electronic Journal of Statistics*, vol. 3, pp. 384–415, 2009.
- [34] I. Daubechies, M. DeFrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Communications on Pure and Applied Mathematics*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [35] A. Maleki and D. L. Donoho, "Optimally tuned iterative reconstruction algorithms for compressed sensing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 330–341, 2010.
- [36] Y. She, "Sparse regression with exact clustering," *Electronic Journal of Statistics*, vol. 4, pp. 1055–1096, 2010.
- [37] M. Y. Park and T. Hastie, "L1-regularization path algorithm for generalized linear models," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 69, no. 4, pp. 659–677, 2007.
- [38] J. Chen and Z. Chen, "Extended Bayesian information criterion for model selection with large model space," *Biometrika*, vol. 95, pp. 759–771, 2008.
- [39] X. Tan, W. Roberts, J. Li, and P. Stoica, "Sparse learning via iterative minimization with application to mimo radar imaging," *IEEE Transactions on Signal Processing*, vol. 59, no. 3, pp. 1088–1101, 2011.
- [40] D. Vu, L. Xu, M. Xue, and J. Li, "Nonparametric missing sample spectral analysis and its applications to interrupted sar," *IEEE Journal of Selected Topics in Signal Processing*, vol. 6, no. 1, pp. 1–14, 2012.
- [41] J. Fan and J. Lv, "Sure independence screening for ultrahigh dimensional feature space," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 70, no. 5, pp. 849–911, 2008.
- [42] T. Yardibi, J. Li, P. Stoica, M. Xue, and A. B. Baggeroer, "Source localization and sensing: A nonparametric iterative adaptive approach based on weighted least squares," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 46, no. 1, pp. 425–443, 2010.
- [43] P. Stoica, P. Babu, and J. Li, "New method of sparse parameter estimation in separable models and its use for spectral analysis of irregularly sampled data," *IEEE Transactions on Signal Processing*, vol. 59, no. 1, pp. 35–47, 2011.
- [44] T. W. Anderson, "The integral of a symmetric unimodal function over a symmetric convex set and some probability inequalities," *Proc. Amer. Math. Soc.*, vol. 6, pp. 170–176, 1955.
- [45] L. Cavalier, G. K. Golubev, D. Picard, and A. B. Tsybakov, "Oracle inequalities for inverse problems," *Ann. Statist.*, vol. 30, no. 3, pp. 843–874, 2002, dedicated to the memory of Lucien Le Cam.



Yiyuan She received B.S. in Mathematics and M.S. in Computer Science from Peking University in 2000 and 2003, respectively, and received Ph.D. in Statistics from Stanford University in 2008. He is currently an Assistant Professor in the Department of Statistics at Florida State University. His research interests include high-dimensional statistics, machine learning, robust statistics, statistics computing, and bioinformatics.



Jiangping Wang (S'10) received the BE degree in electronic engineering from Tsinghua University, Beijing in 2008, and MS degree in electrical and computer engineering from University of Florida, Gainesville in 2010. He is currently a PhD candidate in the Department of Electrical and Computer Engineering at the University of Illinois Urbana-Champaign (UIUC), and is also affiliated with Beckman Institute at UIUC. His research interests include signal and image processing, computer vision and machine learning.



Huanghuang Li is a Ph.D. student in Electrical and Computer Engineering Department at University of Florida, Gainesville, FL. Prior to that, he received his M.S. in signal processing and B.S. in mathematics from Peking University, Beijing, China.



Dapeng Wu (S'98-M'04-SM06-F'13) received Ph.D. in Electrical and Computer Engineering from Carnegie Mellon University, Pittsburgh, PA, in 2003. Since 2003, he has been on the faculty of Electrical and Computer Engineering Department at University of Florida, Gainesville, FL, where he is a professor. His research interests are in the areas of networking, communications, signal processing, computer vision, and machine learning. He received University of Florida Research Foundation Professorship Award in 2009, AFOSR Young Investigator Program (YIP)

Award in 2009, ONR Young Investigator Program (YIP) Award in 2008, NSF CAREER award in 2007, the IEEE Circuits and Systems for Video Technology (CSVT) Transactions Best Paper Award for Year 2001, and the Best Paper Awards in IEEE GLOBECOM 2011 and International Conference on Quality of Service in Heterogeneous Wired/Wireless Networks (QShine) 2006.

Currently, he serves as an Associate Editor for IEEE Transactions on Circuits and Systems for Video Technology. He is the founder of IEEE Transactions on Network Science and Engineering. He was the founding Editor-in-Chief of Journal of Advances in Multimedia between 2006 and 2008, and an Associate Editor for IEEE Transactions on Wireless Communications and IEEE Transactions on Vehicular Technology between 2004 and 2007. He has served as Technical Program Committee (TPC) Chair for IEEE INFOCOM 2012. He has served as Chair for the Award Committee, and Chair of Mobile and wireless multimedia Interest Group (MobIG), Technical Committee on Multimedia Communications, IEEE Communications Society. He was a member of Multimedia Signal Processing Technical Committee, IEEE Signal Processing Society from Jan. 1, 2009 to Dec. 31, 2012.

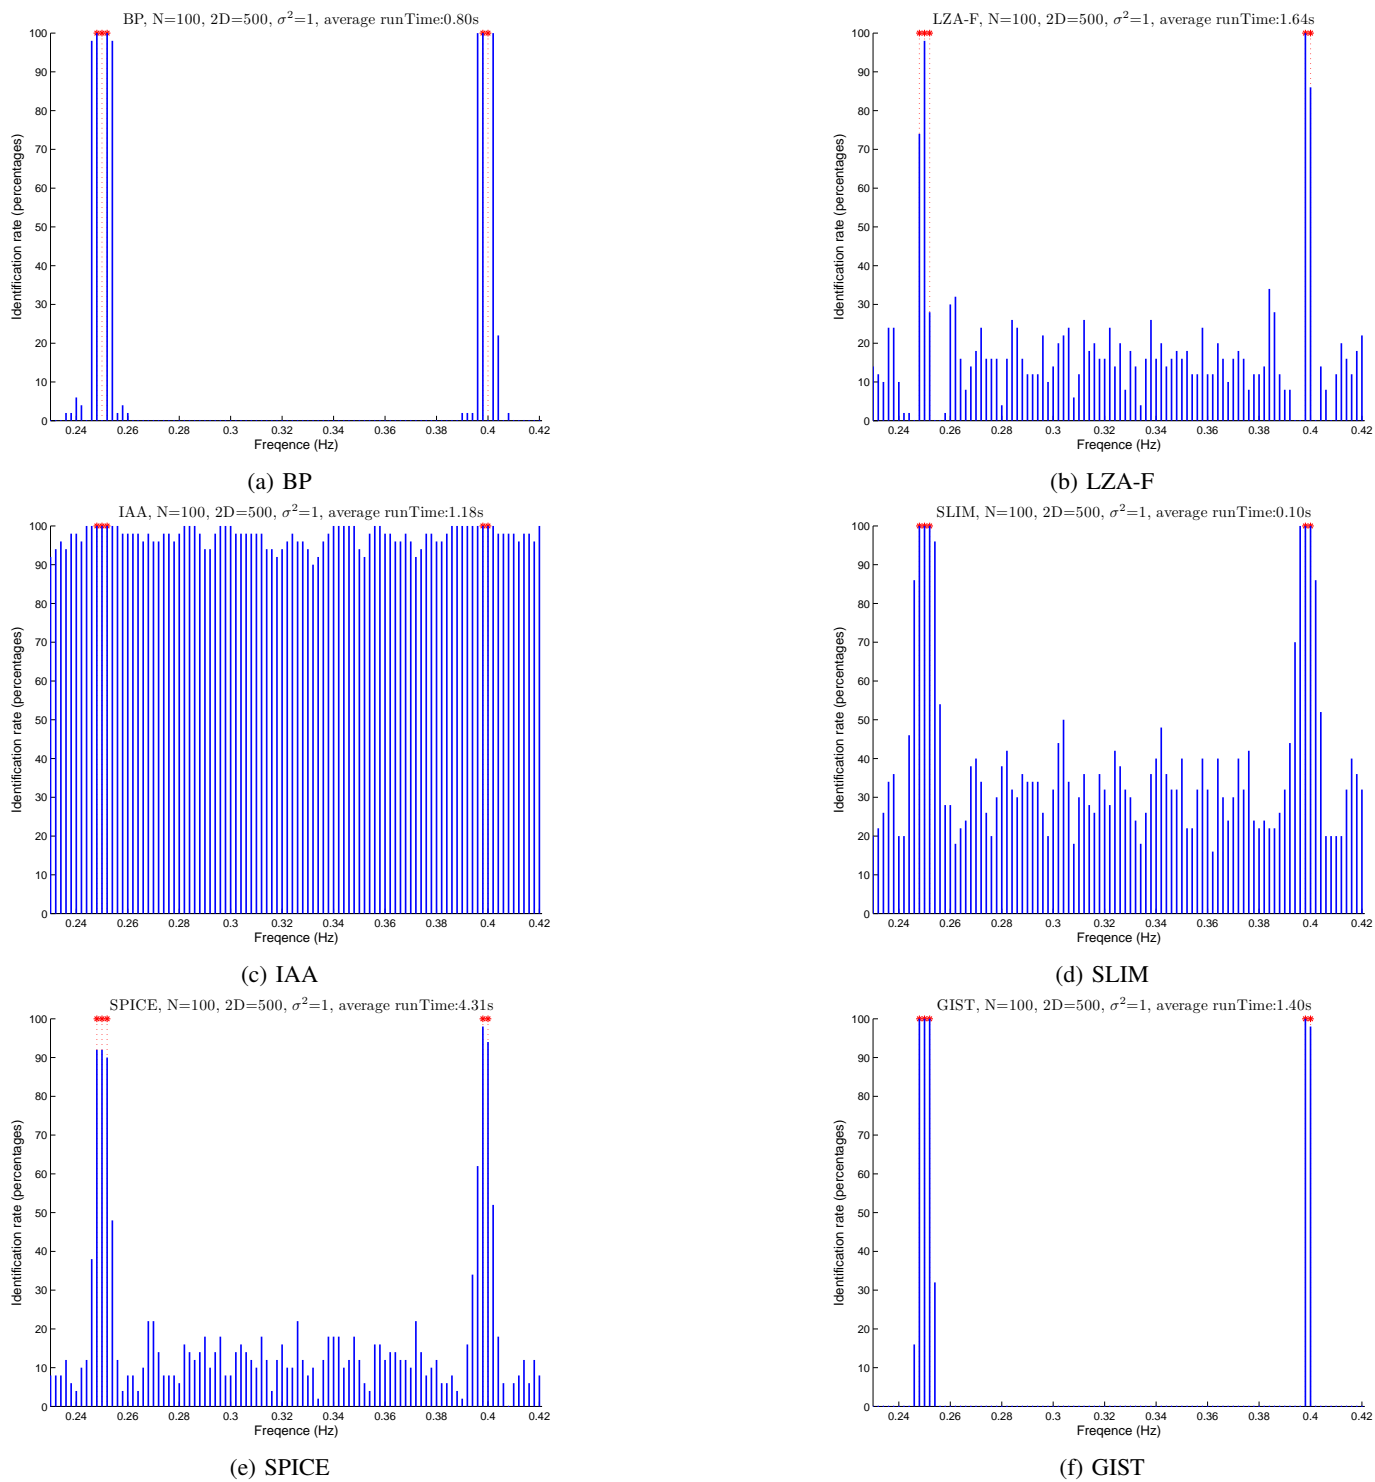


Fig. 5: Frequency identification rates with $\sigma^2 = 1$ in 50 simulation runs, using BP, LZA-F, IAA, SLIM, SPICE, and GIST.

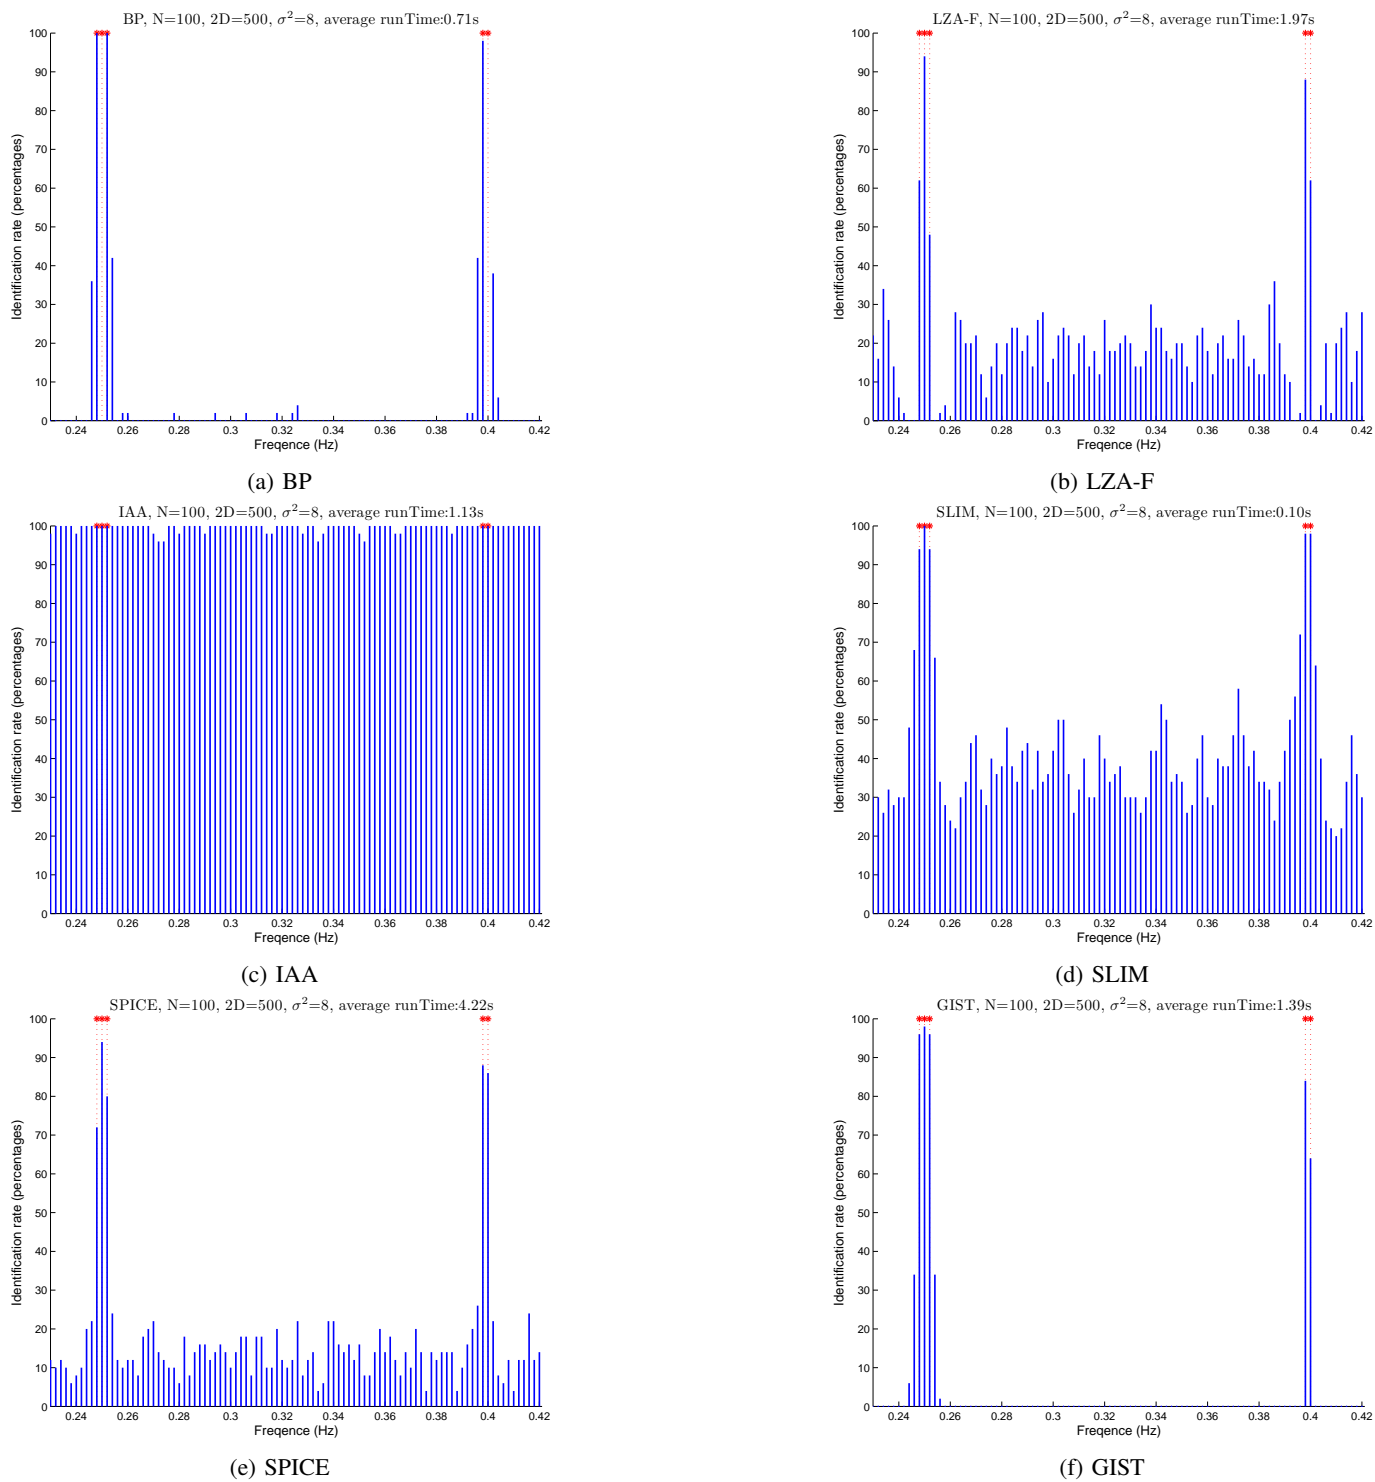


Fig. 6: Frequency identification rates with $\sigma^2 = 8$ in 50 simulation runs, using BP, LZA-F, IAA, SLIM, SPICE, and GIST.

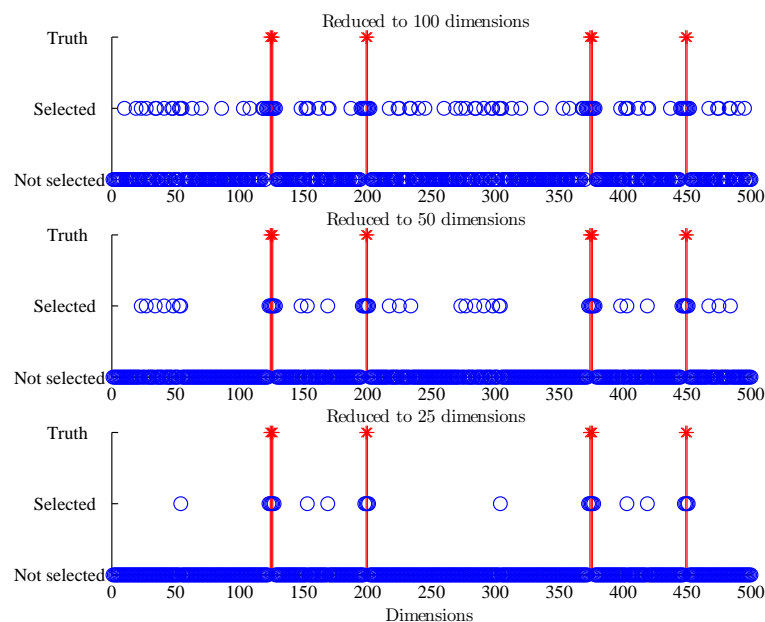


Fig. 7: Locations of the remaining frequency atoms after GIST screening with $\vartheta N = 100, 50, 25$. The true frequencies are indicated by red lines and stars.




Fig. 8: Performance of GIST screening on a hard problem with 10 present frequency components and large noise variance $\sigma^2 = 10$. The left panel shows the miss rates, while the right panel shows the total computational time (including the GIST fitting time thereafter); both x -axes represent ϑN , the dimensions to be kept after screening.