

Why Gabor Frames? Two Fundamental Measures of Coherence and Their Role in Model Selection

Waheed U. Bajwa, Robert Calderbank, and Sina Jafarpour

(Invited Paper)

Abstract: The problem of model selection arises in a number of contexts, such as subset selection in linear regression, estimation of structures in graphical models, and signal denoising. This paper studies non-asymptotic model selection for the general case of arbitrary (random or deterministic) design matrices and arbitrary nonzero entries of the signal. In this regard, it generalizes the notion of *incoherence* in the existing literature on model selection and introduces two fundamental measures of coherence—termed as the worst-case coherence and the average coherence—among the columns of a design matrix. It utilizes these two measures of coherence to provide an in-depth analysis of a simple, model-order agnostic one-step thresholding (OST) algorithm for model selection and proves that OST is feasible for exact as well as partial model selection as long as the design matrix obeys an easily verifiable property, which is termed as the *coherence property*. One of the key insights offered by the ensuing analysis in this regard is that OST can successfully carry out model selection even when methods based on convex optimization such as the lasso fail due to the rank deficiency of the submatrices of the design matrix. In addition, the paper establishes that if the design matrix has reasonably small worst-case and average coherence then OST performs near-optimally when either (i) the energy of any nonzero entry of the signal is close to the average signal energy per nonzero entry or (ii) the signal-to-noise ratio in the measurement system is not too high. Finally, two other key contributions of the paper are that (i) it provides bounds on the average coherence of Gaussian matrices and Gabor frames, and (ii) it extends the results on model selection using OST to low-complexity, model-order agnostic recovery of sparse signals with arbitrary nonzero entries. In particular, this part of the analysis in the paper implies that an Alltop Gabor frame together with OST can successfully carry out model selection and recovery of sparse signals irrespective of the phases of the nonzero entries even if the number of nonzero entries scales almost linearly with the number of rows of the Alltop Gabor frame.

Index Terms: Basis pursuit, coherence property, compressed sensing, Gabor frames, hard thresholding, incoherence condition, irrepresentable condition, lasso, marginal regression, matching pur-

suit, model selection, sparse signals, sparsity pattern recovery, statistical orthogonality condition, variable selection.

I. INTRODUCTION

A. Background

In many information processing and statistics problems involving high-dimensional data, the “curse of dimensionality” can often be broken by exploiting the fact that real-world data tend to live in low-dimensional manifolds. This phenomenon is exemplified by the important special case in which a data vector $\beta \in \mathbb{C}^p$ satisfies $\|\beta\|_0 \doteq \sum_{i=1}^p 1_{\{|\beta_i| > 0\}} \leq k \ll p$ and is observed according to the linear measurement model $y = X\beta + \eta$. Here, X is an $n \times p$ (real- or complex-valued) matrix called the *measurement* or *design matrix*, while $\eta \in \mathbb{C}^n$ represents noise in the measurement system. In this problem, the assumption that the data vector β is “ k -sparse” allows one to operate in the so-called “compressed” setting, $k < n \ll p$, thereby enabling tasks that might be deemed prohibitive otherwise because of either technological or computational constraints.

Fundamentally, given a measurement vector $y = X\beta + \eta$ in the compressed setting, there are three complementary—but nonetheless distinct—questions that one needs to answer:

[Estimation] Under what conditions can we obtain a reliable estimate of a k -sparse β from y ?

[Regression] Under what conditions can we reliably approximate $X\beta$ corresponding to a k -sparse β from y ?

[Model selection] Under what conditions can we reliably recover the locations of the nonzero entries of a k -sparse β (in other words, the model $\mathcal{S} \doteq \{i \in \{1, \dots, p\} : |\beta_i| > 0\}$) from y ?

A number of researchers have attempted to address the estimation and the regression question over the past several years. In many application areas, however, the model-selection question is equally—if not more—important than the other two questions. In particular, the problem of model selection (sometimes also known as *variable selection* or *sparsity pattern recovery*) arises indirectly in a number of contexts, such as subset selection in linear regression [1], estimation of structures in graphical models [2], and signal denoising [3]. In addition, solving the model-selection problem in some (but not all) cases also enables one to solve the estimation and/or the regression problem.

B. Main Contributions

Model selection: One of the primary objectives of this paper is to study the problem of polynomial time, model-order agnostic model selection in a compressed setting for the general case

Manuscript received June 03, 2010; approved for publication by Helmut Bölcskei, Guest Editor, July 02, 2010.

This paper was presented in part at the IEEE International Symposium on Information Theory, Austin, TX, June 2010.

This work was supported in part by NSF under grant DMS 0914892, by ONR under grant N00173-06-1-G006, and by AFOSR under grants FA9550-05-1-0443 and FA9550-09-1-0422.

W. U. Bajwa is with the Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544, email: wbajwa@math.princeton.edu.

R. Calderbank is with the Department of Electrical Engineering and the Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544, email: calderbk@math.princeton.edu.

S. Jafarpour is with the Department of Computer Science, Princeton University, Princeton, NJ 08544, email: sina@cs.princeton.edu.

Algorithm 1 The OST algorithm for model selection**Input:** An $n \times p$ matrix X , a vector $y \in \mathbb{C}^n$, and a threshold $\lambda > 0$ **Output:** An estimate $\hat{S} \subset \{1, \dots, p\}$ of the true model S

$$f \leftarrow X^H y$$

{Form signal proxy}

$$\hat{S} \leftarrow \{i \in \{1, \dots, p\} : |f_i| > \lambda\}$$

{Select model via OST}

of arbitrary (random or deterministic) design matrices and arbitrary nonzero entries of the signal. In order to accomplish this task, we introduce in the paper two fundamental measures of coherence among the (normalized) columns $\{x_i \in \mathbb{C}^n\}$ of the $n \times p$ design matrix X , namely,¹

- *Worst-case coherence*: $\mu(X) \doteq \max_{i,j:i \neq j} |\langle x_i, x_j \rangle|$, and

- *Average coherence*: $\nu(X) \doteq \frac{1}{p-1} \max_i \left| \sum_{j:j \neq i} \langle x_i, x_j \rangle \right|$.

Roughly speaking, worst-case coherence—which seems to have been introduced in the related literature in [4]—is a similarity measure between the columns of a design matrix: The smaller the worst-case coherence, the less similar the columns. On the other hand, average coherence—which was first introduced in a conference version of this paper [5]—is a measure of the spread of the columns of a design matrix within the n -dimensional unit ball: The smaller the average coherence, the more spread out the column vectors.

Our main contribution in the area of model selection is that we make use of these two measures of coherence to propose and analyze a model-order agnostic threshold for the *one-step thresholding* (OST) algorithm (see Algorithm 1) for model selection. Specifically, we characterize in Section II both the exact and the partial model-selection performance of OST in a *non-asymptotic* setting in terms of μ and ν . In particular, we establish in Section II that if $\mu(X) \asymp n^{-1/2}$ and $\nu(X) \lesssim n^{-1}$ then OST—despite being computationally primitive—can perform near-optimally for the case when either (i) the energy of any nonzero entry of β is not too far away from the average signal energy per nonzero entry $\|\beta\|_2^2/k$ or (ii) the signal-to-noise ratio (SNR) in the measurement system is not too high.² Equally importantly, in contrast to some of the existing literature on model selection, this analysis in the paper holds for arbitrary values of the nonzero entries of β and it does not require the $n \times k$ submatrices of the design matrix X to have full column rank.

Sparse-signal recovery: The second main objective of this paper is to study the problem of low-complexity, model-order agnostic recovery of k -sparse signals with arbitrary nonzero entries in the noiseless case. In this regard, our main contribution in the area of sparse-signal recovery is that we make use of a recent result by Tropp [6] in Section III to extend our results on model selection to recovery of k -sparse signals using OST (see Algorithm 2). In particular, we establish in Section IV that Ga-

bor frames—which are collections of time- and frequency-shifts of a nonzero seed vector (sequence) in \mathbb{C}^n —can potentially be used together with OST to exactly recover *most* k -sparse signals with arbitrary nonzero entries as long as $k \lesssim \mu^{-2}/\log n$ and the energy of any nonzero entry of β is not too far away from $\|\beta\|_2^2/k$. This result then applies immediately to Gabor frames generated from the Alltop sequence [7]. Specifically, since Gabor frames generated from the Alltop sequence have worst-case coherence $\mu = \frac{1}{\sqrt{n}}$ for any prime $n \geq 5$ [8], this result implies that an Alltop Gabor frame together with OST successfully recovers most k -sparse signals *irrespective* of the values of the nonzero entries of β as long as $k \lesssim n/\log n$ and the energy of any nonzero entry of β is not too far away from $\|\beta\|_2^2/k$.

C. Relationship to Previous Work

The problems of model selection and sparse-signal recovery in general and the use of OST (also known as *simple thresholding* [9] and *marginal regression* [10]) to solve these problems in particular have a rich history in the literature. In the context of model selection in the compressed setting, Mallows's C_p selection procedure [11] and the *Akaike information criterion* (AIC) [12]—both of which essentially attempt to solve a complexity-regularized version of the least-squares criterion—are considered to be seminal works, and are known to perform well empirically as well as theoretically; see, e.g., [13] and the references therein. These two procedures have been modified by numerous researchers over the years in order to improve their performance—the most notable variants being the *Bayesian information criterion* (BIC) [14] and the *risk inflation criterion* (RIC) [15]. Solving model-selection procedures such as C_p , AIC, BIC, and RIC, however, is known to be an NP-hard problem [16] even if the true model order k is made available to these procedures.

In order to overcome the computational intractability of these model-selection procedures, several methods based on convex optimization have been proposed by various researchers in recent years. Among these proposed methods, the lasso [17] has arguably become the standard tool for model selection, which can be partly attributed to the theoretical guarantees provided for the lasso in [2], [18]–[20]. In particular, the results reported in [2] and [18] establish that the lasso asymptotically identifies the correct model under certain conditions on the design matrix X and the sparse vector β . Later, Wainwright in [19] strengthens the results of [2] and [18], and makes explicit the dependence of exact model selection using the lasso on the smallest (in magnitude) nonzero entry of β . However, apart from the fact that the results reported in [2], [18], and [19] are for exact model selection and are only asymptotic in nature, the main limitation of these works is that explicit verification of the conditions (such as the *irrepresentable condition* of [18] and the *incoherence con-*

¹Here, and throughout the rest of this paper, we assume without loss of generality that X has unit ℓ_2 -norm columns. This is because deviations to this assumption can always be accounted for by appropriately scaling the entries of the data vector β instead.

²Recall “Big- O ” notation: $f(n) = O(g(n))$ (alternatively, $f(n) \lesssim g(n)$) if $\exists c_o > 0, n_o : \forall n \geq n_o, f(n) \leq c_o g(n)$, $f(n) = \Omega(g(n))$ (alternatively, $f(n) \gtrsim g(n)$) if $g(n) = O(f(n))$, and $f(n) = \Theta(g(n))$ (alternatively, $f(n) \asymp g(n)$) if $g(n) \lesssim f(n) \lesssim g(n)$.

Algorithm 2 The OST algorithm for sparse-signal recovery**Input:** An $n \times p$ matrix X , a vector $y \in \mathbb{C}^n$, and a threshold $\lambda > 0$ **Output:** An estimate $\hat{\beta} \in \mathbb{C}^p$ of the true sparse signal β

$\hat{\beta} \leftarrow \mathbf{0}$	{Initialize}
$f \leftarrow X^H y$	{Form signal proxy}
$\mathcal{I} \leftarrow \{i \in \{1, \dots, p\} : f_i > \lambda\}$	{Select indices via OST}
$\hat{\beta}_{\mathcal{I}} \leftarrow (X_{\mathcal{I}})^{\dagger} y$	{Recover signal via least-squares}

dition of [19]) that a generic design matrix X needs to satisfy is computationally intractable for $k \gtrsim \mu^{-1}$. The most general (and non-asymptotic) model-selection results using the lasso for arbitrary design matrices have been reported in [20]. Specifically, Candès and Plan have established in [20] that the lasso correctly identifies most models with probability $1 - O(p^{-1})$ under certain conditions on the smallest nonzero entry of β provided: (i) The spectral norm (the largest singular value) and the worst-case coherence of X are not too large, and (ii) the values of the nonzero entries of β are independent and statistically symmetric around zero. Despite these recent theoretical triumphs of the lasso, it is still desirable to study alternative solutions to the problem of polynomial time, model-order agnostic model selection in a compressed setting. This is because:³

- 1) Lasso requires the minimum singular values of the submatrices of X corresponding to the true models to be bounded away from zero [2], [18]–[20]. While this is a plausible condition for the case when one is interested in estimating β , it is arguable whether this condition is necessary for the case of model selection.
- 2) The current literature on model selection using the lasso lacks guarantees beyond $k \gtrsim \mu^{-1}$ for the case of generic design matrices and arbitrary nonzero entries. In particular, given an arbitrary design matrix X , [2], [18]–[20] do not provide any guarantees beyond $k \gtrsim \sqrt{n}$ for even the simple case of $\beta \in \mathbb{R}_+^p$.
- 3) The computational complexity of the lasso for generic design matrices tends to be $O(p^3 + np^2)$ [10]. This makes the lasso computationally demanding for large-scale model-selection problems.

Recently, a few researchers have raised somewhat similar concerns about the lasso and revisited the much older (and oft-forgotten) method of thresholding for model selection [10], [22]–[24], which has computational complexity of $O(np)$ only and which is known to be nearly optimal for $p \times p$ orthonormal design matrices [25]. Algorithmically, this makes our approach to model selection similar to that of [10], [22]–[24]. Nevertheless, the OST algorithm presented in this paper differs from [10], [22]–[24] in five key aspects:

- 1) *Model-order agnostic model selection:* Unlike [10], [22]–[24], the OST algorithm presented in this paper is completely agnostic to both the true model order k and any estimate of k .
- 2) *Generic design matrices and arbitrary nonzero entries:*

The results reported in this paper hold for arbitrary (random or deterministic) design matrices and do not assume any statistical prior on the values of the nonzero entries of β even when k scales linearly with n . In contrast, [23] only studies the problem of Gaussian design matrices whereas the most influential results reported in [10], [22], and [24] assume that the values of the nonzero entries of β are independent and statistically symmetric around zero.

- 3) *Verifiable sufficient conditions:* In contrast to [10], [22]–[24], we relate the model-selection performance of OST to two global parameters of X , namely, μ and ν , which are trivially computable in polynomial time: $\mu(X) = \|X^H X - I\|_{\max}$ and $\nu(X) = \frac{1}{p-1} \|(X^H X - I)\mathbf{1}\|_{\infty}$.
- 4) *Non-asymptotic theory:* Similar to [10], [23], and [24], the analysis in this paper can be used to establish that OST achieves (asymptotically) consistent model selection under certain conditions. However, the results reported in this paper are completely non-asymptotic in nature (with explicit constants) and thereby shed light on the rate at which OST achieves consistent model selection.
- 5) *Partial model selection:* In addition to the exact model-selection performance of OST, we also characterize in the paper its partial model-selection performance. In this regard, we establish that the *universal threshold* proposed in Section II for OST guarantees $\hat{\mathcal{S}} \subset \mathcal{S}$ with high probability and we quantify the cardinality of the estimate $\hat{\mathcal{S}}$. On the other hand, both [22] and [23] study only exact model selection, whereas [10] and [24] study approximate (though not partial) model selection only for Gaussian design matrices [10] and assuming Gaussian (resp. statistical) priors on the nonzero entries of β [24] (resp. [10]).

We conclude this discussion of model selection by making three important remarks. First, to the best of our knowledge, Donoho in [9, Theorem 7.2] reported some of the earliest known results for thresholding in the compressed setting. Nevertheless, the conclusion drawn in [9] was that thresholding is feasible for model selection as long as $k \lesssim \mu^{-1}$, the so-called “square root bottleneck.” Second, the structure of OST and the model-order agnostic threshold of this paper enable us to carry out *localized model selection*. Specifically, if one is provided at the time of recovery with a set \mathcal{T} such that $\mathcal{T} \supset \mathcal{S}$ then the threshold proposed in this paper enables one to carry out model selection using the submatrix $X_{\mathcal{T}}$ instead of X , thereby reducing the complexity of OST from $O(np)$ to $O(n|\mathcal{T}|)$. Third, the results reported in this paper hold for any $n \leq p$ and, in particular, the universal threshold proposed here for model selection reduces to the universal threshold proposed by Donoho and Johnstone [25] for $p \times p$ orthonormal design matrices. In this sense, some of the results

³During the course of revising this paper we also became aware of [21], which proposes a thresholded variant of *basis pursuit* [3] for sparsity pattern recovery using Gaussian design matrices. However, the results reported in [21] are limited because of similar issues and because of the requirement that the magnitude of the smallest nonzero entry of β be known to the algorithm.




Fig. 1. A Venn diagram used to illustrate the major difference between the BP-based recovery guarantees and the OST-based recovery guarantees for k -sparse unimodal signals in \mathbb{R}_+^p measured using Alltop Gabor frames. The OST algorithm is guaranteed to recover $\beta \in \Sigma_1 - \Gamma_B$. But BP, unlike OST, is only guaranteed to recover $\beta \in \Sigma_2$ in this case.

reported in [25] can also be thought of as special instances of the results reported in this paper.

Finally, in the context of sparse-signal recovery in the compressed setting, there exists now a large body of literature that studies this problem under the rubric of *compressed sensing* [26]. However, convex optimization procedures such as basis pursuit (BP) [3], Dantzig selector [27], and lasso—although known for their ability to recover sparse signals under a variety of conditions—are ill-suited for large-scale problems because of their computational complexity. On the other hand, low-complexity iterative algorithms such as matching pursuit [28], subspace pursuit [29], compressive sampling matching pursuit (CoSaMP) [30], and iterative hard thresholding [31], and combinatorial algorithms based on group testing such as heavy hitters on steroids (HHS) pursuit [32] and Fourier samplers [33], [34] have been shown to perform well either only for some special classes of design matrices [32]–[34] or for design matrices that satisfy the *restricted isometry property* (RIP) [35]. Nevertheless, explicitly verifying that X satisfies the RIP of order $k \lesssim \mu^{-1}$ is computationally intractable; in particular, since we have from the Welch bound [36] that $\mu^{-1} \lesssim \sqrt{n}$ for $p \gg 1$, the guarantees provided in [29]–[31] for the case of generic design matrices at best hold only for k -sparse signals with $k \lesssim \sqrt{n}$.

In contrast, and motivated by the need to have verifiable sufficient conditions for low-complexity algorithms and arbitrary values of the nonzero entries of β even when $k \lesssim \sqrt{n}$, we extend in Section III our results on model selection using OST and characterize the performance of Algorithm 2 in terms of three global parameters of the design matrix X : $\mu(X)$, $\nu(X)$, and $\|X\|_2$. In particular, a highlight of this part of the paper is that we partially strengthen the results of Pfander *et al.* [37] and Herman and Strohmer [38] by establishing that Gabor frames generated from the Alltop sequence can be used along with OST to recover most k -sparse signals belonging to certain classes even when $k \lesssim \sqrt{n}$. It is worth pointing out here that both [37] and [38] also establish that Alltop Gabor frames can recover most k -sparse signals—albeit using BP—even when $k \lesssim \sqrt{n}$. Nevertheless, the basic difference between [37] and [38], and the work presented here is that [37] and [38] require the phases of the nonzero entries of β to be statistically independent and uniformly distributed on the unit torus whereas we do not assume any statistical prior on the values of the nonzero entries of β . Note in particular that, just like the lasso result in [20], the results reported in [37] and [38] for Alltop Gabor frames conse-

quently do not provide any guarantees beyond $k \lesssim \sqrt{n}$ for even the simple case of $\beta \in \mathbb{R}_+^p$. This difference between the BP-based recovery guarantees presented in [37] and [38] (which are essentially based on [39]) and the OST-based recovery guarantees provided in this paper is also illustrated using a Venn diagram in Fig. 1 for *unimodal signals* (defined as: $|\beta_i| \approx c$ for some arbitrary $c > 0$ and for all $i \in S$).

D. Notation

The following notation is used throughout the rest of this paper. We use lowercase letters to denote scalars and vectors, while we use uppercase letters to denote matrices. We also use $\mathbf{0}$, $\mathbf{1}$, and I to denote the all-zeros vector, the all-ones vector, and the identity matrix, respectively. In addition, we use $\|v\|_p$ to denote the usual ℓ_p -norm of a vector v , while we use A^\dagger , $\|A\|_2$, and $\|A\|_{\max}$ to denote the Moore–Penrose pseudoinverse, the spectral norm, and the maximum magnitude of any entry of a matrix A , respectively. Further, we use $(\cdot)^T$ and $(\cdot)^H$ to denote the operations of transposition and conjugate transposition, respectively, while we use $\langle \cdot, \cdot \rangle$ to denote inner product that is conjugate linear in the first argument. Finally, given a set \mathcal{I} , we use $v_{\mathcal{I}}$ to denote the part of a vector v corresponding to the indices in \mathcal{I} and $A_{\mathcal{I}}$ to denote the submatrix obtained by collecting the $|\mathcal{I}|$ columns of a matrix A corresponding to the indices in \mathcal{I} .

E. Organization

The rest of this paper is organized as follows. In Section II, we propose a model-order agnostic threshold for the OST algorithm and characterize both the exact and the partial model-selection performance of OST. In Section III, we extend our results on model selection and characterize the sparse-signal recovery performance of OST. In Section IV, we specialize the model-selection and the sparse-signal recovery results of the previous sections to Gabor frames. Finally, we provide proofs of the main results of this paper in Section V and conclude with a discussion of the limitations and extensions of our results in Section VI.

II. MODEL SELECTION USING ONE-STEP THRESHOLDING

A. Assumptions

Before proceeding with presenting our results on model selection using OST, we need to be mathematically precise about

our problem formulation. To this end, we begin by reconsidering the measurement model $y = X\beta + \eta$ and assume that X is an $n \times p$ real- or complex-valued design matrix having unit ℓ_2 -norm columns, $\beta \in \mathbb{C}^p$ is a k -sparse signal ($\|\beta\|_0 \leq k$), and $k < n \leq p$. Here, we allow X to be either a random or a deterministic design matrix, while we take η to be a complex additive white Gaussian noise vector. It is worth mentioning here though that Gaussianity of η is just a simplified assumption for the sake of this exposition; in particular, the results presented in this section are readily generalizable to other noise distributions as well as perturbations having bounded ℓ_2 -norms. Finally, the main assumption that we make here is that the true model $\mathcal{S} \doteq \{i \in \{1, \dots, p\} : |\beta_i| > 0\}$ is a uniformly random k -subset of $\{1, \dots, p\}$. In other words, we have a uniform prior on the support of the data vector β .

B. Main Results

Intuitively speaking, successful model selection requires the columns of the design matrix to be *incoherent*. In the case of the lasso, this notion of incoherence has been quantified in [18] and [19] in terms of the “irrepresentable condition” and the “incoherence condition,” respectively (see also [20]). In contrast to earlier work on model selection, however, we formulate this idea of incoherence in terms of the *coherence property*.

Definition 1 (The coherence property). An $n \times p$ design matrix X having unit ℓ_2 -norm columns is said to obey the coherence property if the following two conditions hold

$$\mu(X) \leq \frac{0.1}{\sqrt{2 \log p}} \quad \text{and} \quad (\text{CP-1})$$

$$\nu(X) \leq \frac{\mu}{\sqrt{n}}. \quad (\text{CP-2})$$

In words, (CP-1) roughly states that the columns of X are not too similar, while (CP-2) roughly states that the columns of X are somewhat distributed within the n -dimensional unit ball. Note that the coherence property is superior to other measures of incoherence such as the irrepresentable condition in two key aspects. First, it does not require the singular values of the submatrices of X to be bounded away from zero. Second, it can be easily verified in polynomial time since it simply requires checking that $\|X^H X - I\|_{\max} \leq (200 \log p)^{-1/2}$ and $\|(X^H X - I)\mathbf{1}\|_{\infty} \leq (p-1)n^{-1/2}\|X^H X - I\|_{\max}$.

Below, we describe the implications of the coherence property for both the exact and the partial model-selection performance of OST. Before proceeding further, however, it is instructive to first define some fundamental quantities pertaining to the problem of model selection as follows

$$\begin{aligned} \beta_{\min} &\doteq \min_{i \in \mathcal{S}} |\beta_i|, & \text{MAR} &\doteq \frac{\beta_{\min}^2}{\|\beta\|_2^2/k}, \\ \text{SNR}_{\min} &\doteq \frac{\beta_{\min}^2}{\mathbb{E}[\|\eta\|_2^2]/k}, & \text{SNR} &\doteq \frac{\|\beta\|_2^2}{\mathbb{E}[\|\eta\|_2^2]}. \end{aligned}$$

In words, β_{\min} is the magnitude of the smallest nonzero entry of β , while *minimum-to-average ratio* (MAR) [23] is the ratio of the energy in the smallest nonzero entry of β and the average signal energy per nonzero entry of β . Likewise, SNR_{\min} is the ratio of the energy in the smallest nonzero entry of β and the average

noise energy per nonzero entry, while SNR simply denotes the usual signal-to-noise ratio in the system. It is also worth pointing out here the relationship between SNR_{\min} and SNR and MAR; specifically, it is easy to see that $\text{SNR}_{\min} = \text{SNR} \text{MAR}$. We are now ready to state the first main result of this paper that concerns the performance of OST in terms of exact model selection.

Theorem 1 (Exact model selection using OST). Suppose that the design matrix X satisfies the coherence property and let η be distributed as $\mathcal{CN}(0, \sigma^2 I)$. Next, choose the threshold $\lambda = \max \left\{ \frac{1}{t} 10\mu\sqrt{n\text{SNR}}, \frac{1}{1-t}\sqrt{2} \right\} \sqrt{2\sigma^2 \log p}$ for any $t \in (0, 1)$. Then, if we write $\mu(X)$ as $\mu = c_1 n^{-1/\gamma}$ for some $c_1 > 0$ (which may depend on p) and $\gamma \in \{0\} \cup [2, \infty)$, the OST algorithm (Algorithm 1) satisfies $\Pr(\hat{\mathcal{S}} \neq \mathcal{S}) \leq 6p^{-1}$ provided $p \geq 128$ and the number of measurements satisfies

$$\begin{aligned} n &> \max \left\{ 2k \log p, \frac{c_2 k \log p}{\text{SNR}_{\min}}, \left(\frac{c_3 k \log p}{\text{MAR}} \right)^{\gamma/2} \right\} \\ &\equiv \max \left\{ 2k \log p, \frac{c_2 k \log p}{\text{SNR} \text{MAR}}, \left(\frac{c_3 k \log p}{\text{MAR}} \right)^{\gamma/2} \right\}. \end{aligned} \quad (1)$$

Here, the quantities $c_2, c_3 > 0$ are defined as $c_2 \doteq 16(1-t)^{-2}$ and $c_3 \doteq 800c_1^2 t^{-2}$, while the probability of failure is with respect to the true model \mathcal{S} and the noise vector η .

The proof of this theorem is provided in Section V. Note that the parameter ‘ t ’ in Theorem 1 can always be fixed a priori (say $t = 1/2$) without affecting the scaling relation in (1). In practice, however, t should be chosen so as to reduce the total number of measurements needed to ensure successful model selection; the optimal choice of t in this regard is $t_{\text{opt}} = \arg \min_t \left(\max \left\{ \frac{c_2 k \log p}{\text{SNR} \text{MAR}}, \left(\frac{c_3 k \log p}{\text{MAR}} \right)^{\gamma/2} \right\} \right)$. Notice also that Theorem 1 is best suited for applications where one is interested in quantifying the minimum number of measurements needed to guarantee exact model selection for a given class of signals. Alternatively, it might be the case in some other applications that the problem dimensions are fixed and one is instead interested in specifying the class of signals that leads to successful model selection. The following variant of Theorem 1 is best suited in such situations.

Theorem 2. Suppose that the design matrix X satisfies the coherence property and let the noise vector η be distributed as $\mathcal{CN}(0, \sigma^2 I)$. Next, let $p \geq 128$ and choose the threshold $\lambda = \max \left\{ \frac{1}{t} 10\mu\sqrt{n\text{SNR}}, \frac{1}{1-t}\sqrt{2} \right\} \sqrt{2\sigma^2 \log p}$ for any $t \in (0, 1)$. Then, the OST algorithm (Algorithm 1) satisfies $\Pr(\hat{\mathcal{S}} \neq \mathcal{S}) \leq 6p^{-1}$ as long as we have that $k \leq n/(2 \log p)$ and

$$\text{MAR} > \max \left\{ \frac{c_2 k \log p}{n \text{SNR}}, \frac{c'_3 k \log p}{\mu^{-2}} \right\}. \quad (2)$$

Here, $c_2 > 0$ is as defined in Theorem 1, $c'_3 > 0$ is defined as $c'_3 \doteq 800t^{-2}$, and the probability of failure is with respect to the true model \mathcal{S} and the noise vector η .

Note that the proof of Theorem 2 follows directly from the proof of Theorem 1. There are a few important remarks that need to be made at this point concerning the threshold proposed in Theorem 1 and Theorem 2 for the OST algorithm. First, it is

Algorithm 3 The SOST algorithm for model selection**Input:** An $n \times p$ matrix X , a vector $y \in \mathbb{C}^n$, and model order k **Output:** An estimate $\hat{\mathcal{S}} \subset \{1, \dots, p\}$ of the true model \mathcal{S}

$f \leftarrow X^H y$	{Form signal proxy}
$(\mathcal{I}, f_s) \leftarrow \text{SORT}(\{(1, \dots, p), f\})$	{Sort the signal proxy}
$\hat{\mathcal{S}} \leftarrow \mathcal{I}[1 : k]$	{Select model via OST}

easy to see that the proposed threshold is completely agnostic to the model order k and only requires knowledge of the SNR and the noise variance. Second, some of the bounds in the proof of Theorem 1 and extensive simulations suggest that the absolute constant 10 in the proposed threshold is somewhat conservative and can be reduced through the use of more sophisticated analytical tools (also see Section VI). Finally, while estimating the true model order k tends to be harder than estimating the SNR and the noise variance σ^2 in majority of the situations, it might be the case that estimating k is easier in some applications. It is better in such situations to work with a slight variant of the OST algorithm (see Algorithm 3) that relies on knowledge of the model order k instead and returns an estimate $\hat{\mathcal{S}}$ corresponding to the k largest (in magnitude) entries of $X^H y$. We characterize the performance of this algorithm—which we term as *sorted one-step thresholding* (SOST) algorithm—in terms of the following theorem.

Theorem 3 (Exact model selection using SOST). Suppose that the design matrix X satisfies the coherence property and let the noise vector η be distributed as $\mathcal{CN}(\mathbf{0}, \sigma^2 I)$. Next, write $\mu(X)$ as $\mu = c_1 n^{-1/\gamma}$ for some $c_1 > 0$ (which may depend on p) and $\gamma \in \{0\} \cup [2, \infty)$. Then, the SOST algorithm (Algorithm 3) satisfies $\Pr(\hat{\mathcal{S}} \neq \mathcal{S}) \leq 6p^{-1}$ as long as $p \geq 128$ and the number of measurements satisfies

$$n > \min_{t \in (0,1)} \max \left\{ 2k \log p, \frac{c_2 k \log p}{\text{SNR}_{\min}}, \left(\frac{c_3 k \log p}{\text{MAR}} \right)^{\gamma/2} \right\} \\ \equiv \min_{t \in (0,1)} \max \left\{ 2k \log p, \frac{c_2 k \log p}{\text{SNR} \text{ MAR}}, \left(\frac{c_3 k \log p}{\text{MAR}} \right)^{\gamma/2} \right\}. \quad (3)$$

Here, the quantities $c_2, c_3 > 0$ are as defined in Theorem 1, while the probability of failure is with respect to the true model \mathcal{S} and the noise vector η .

The proof of this theorem is just a slight variant of the proof of Theorem 1 and is therefore omitted here. A few remarks are in order now concerning OST and SOST. First, the computational complexity of SOST is comparable with that of OST since efficient sorting algorithms (such as heap sort) tend to have computational complexity of $O(p \log p)$ only. Second, (1) and (3) suggest that knowledge of the true model order k allows SOST to perform better than OST in situations where the threshold parameter t is fixed a priori (cf. Theorem 1). In this sense, SOST should be preferred over OST for exact model selection *provided* one has accurate knowledge of the true model order k . On the other hand, OST should be the algorithm of choice for model-selection problems where it is difficult to obtain a reliable estimate of the true model order. We conclude this discussion by rephrasing Theorem 3 for SOST along the lines of Theo-

rem 2 for OST.

Theorem 4. Suppose that the design matrix X satisfies the coherence property. Next, let $p \geq 128$ and let the noise vector η be distributed as $\mathcal{CN}(\mathbf{0}, \sigma^2 I)$. Then, the SOST algorithm (Algorithm 3) satisfies $\Pr(\hat{\mathcal{S}} \neq \mathcal{S}) \leq 6p^{-1}$ as long as we have $k \leq n/(2 \log p)$ and

$$\text{MAR} > \min_{t \in (0,1)} \left\{ \frac{c_2 k \log p}{n \text{SNR}}, \frac{c'_3 k \log p}{\mu^{-2}} \right\}. \quad (4)$$

Here, the quantities $c_2, c'_3 > 0$ are as defined in Theorem 2, while the probability of failure is with respect to the true model \mathcal{S} and the noise vector η .

The final result that we present in this section concerns the partial model-selection performance of OST. Specifically, note that our focus in this section has so far been on specifying conditions for either the number of measurements or the MAR of the signal that ensure exact model selection. In many real-world applications, however, the parameters of the problem are fixed and it is not always possible to ensure that either the number of measurements or the MAR of the signal satisfy the aforementioned conditions. A natural question to ask then is whether the OST algorithm completely fails in such circumstances or whether any guarantees can still be provided for its performance. We address this aspect of the OST algorithm in the following and show that, even if the MAR of β is very small, OST has the ability to identify the locations of the nonzero entries of β whose energies are greater than both the noise power and the average signal energy per nonzero entry. In order to make this notion mathematically precise, we first define the m th largest-to-average ratio (LAR_{*m*}) of β as the ratio of the energy in the m th largest (in magnitude) nonzero entry of β and the average signal energy per nonzero entry of β ; that is,

$$\text{LAR}_m \doteq \frac{|\beta_{(m)}|^2}{\|\beta\|_2^2/k}$$

where $\beta_{(m)}$ denotes the m th largest nonzero entry of β (note that $\text{MAR} \equiv \text{LAR}_k$). We are now ready to specify the partial model-selection performance of the OST algorithm.

Theorem 5 (Partial model selection using OST). Suppose that the design matrix X satisfies the coherence property. Next, let $p \geq 128$ and η be distributed as $\mathcal{CN}(\mathbf{0}, \sigma^2 I)$. Finally, fix a parameter $t \in (0,1)$ and choose the threshold $\lambda = \max \left\{ \frac{1}{t} 10\mu \sqrt{n \text{SNR}}, \frac{1}{1-t} \sqrt{2} \right\} \sqrt{2\sigma^2 \log p}$. Then, under the assumption that $k \leq n/(2 \log p)$, the OST algorithm (Algorithm 1) guarantees with probability exceeding $1 - 6p^{-1}$ that $\hat{\mathcal{S}} \subset \mathcal{S}$ and $|\mathcal{S} - \hat{\mathcal{S}}| \leq (k - M)$, where M is the largest

integer for which the following inequality holds

$$\text{LAR}_M > \max \left\{ \frac{c_2 k \log p}{n \text{SNR}}, \frac{c'_3 k \log p}{\mu^{-2}} \right\}. \quad (5)$$

Here, the quantities $c_2, c'_3 > 0$ are as defined in Theorem 2, while the probability of failure is with respect to the true model S and the noise vector η .

The proof of this theorem, which relies to a great extent on the proof of Theorem 1, is provided in Section V. We conclude this section by pointing out that no counterpart of Theorem 5 exists for the SOST algorithm since we can never have $\hat{S} \subset S$ in that case because of the nature of the algorithm.

C. Discussion

The results reported earlier can be best put into perspective by considering some specific model-selection problems that are commonly studied in the literature and juxtaposing our results with the ones reported in previous works. The rest of this section is devoted to such comparison purposes.

C.1 Gaussian Design Matrices

Matrices with independent and identically distributed (i.i.d.) $\mathcal{N}(0, 1/n)$ entries (i.e., Gaussian matrices) are perhaps the most widely assumed design matrices in the model-selection literature. In order to specialize our results to Gaussian design matrices, we first need to specify the worst-case coherence μ and the average coherence ν of i.i.d. Gaussian matrices. The first lemma that we have in this regard follows immediately from Proposition 5 in Appendix A through a simple union bound argument.

Lemma 1 (Worst-case coherence of Gaussian matrices). Let X be an $n \times p$ design matrix with i.i.d. $\mathcal{N}(0, 1/n)$ entries. Then, we have that $\mu(X) \leq \sqrt{(15 \log p)/n}$ with probability exceeding $1 - 2p^{-1}$ as long as $n \geq 60 \log p$.

Remark 1. A cautious reader might argue here that Lemma 1 only provides an upperbound on the worst-case coherence of Gaussian matrices. Nevertheless, the results (and the definition of the coherence property) presented earlier remain valid if one replaces $\mu(X)$ with an upperbound $\bar{\mu}(X)$ on $\mu(X)$.

Lemma 2 (Average coherence of Gaussian matrices). Let X be an $n \times p$ design matrix with i.i.d. $\mathcal{N}(0, 1/n)$ entries. Then, we have that $\nu(X) \leq \sqrt{15 \log p}/n$ with probability exceeding $1 - 2p^{-2}$ as long as $p > n \geq 60 \log p$.

Proof: The proof of this lemma is also a consequence of Proposition 5 in Appendix A. Specifically, fix an index $i \in \{1, \dots, p\}$ and define $\tilde{x}_i \doteq \sum_{j \neq i} x_{ji} / \sqrt{p-1}$. Then, it is easy to see that \tilde{x}_i is distributed as $\mathcal{N}(0, 1/n)$ and it is independent of x_i . Therefore, Proposition 5 in Appendix A implies through a simple union bound argument that $\max_i |\langle x_i, \tilde{x}_i \rangle| \leq \sqrt{(15 \log p)/n}$ with probability exceeding $1 - 2p^{-2}$ as long as $n \geq 60 \log p$. The proof now follows from the fact that $p > n$ and $\nu(X) = \max_i |\langle x_i, \tilde{x}_i \rangle| / \sqrt{p-1}$. \square

Lemma 1 and Lemma 2 establish that Gaussian design matrices satisfy the coherence property with high probability as long as $n \gtrsim (\log p)^2$. Theorem 1 (resp. Theorem 3) therefore implies that OST (resp. SOST) correctly identifies the exact model with probability exceeding $1 - O(p^{-1})$ as long as

$n \gtrsim \max \left\{ 1, \frac{1}{\text{SNR} \text{MAR}}, \frac{\log p}{\text{MAR}} \right\} k \log p$. In particular, this suggests that if either $\text{MAR}(\beta) = \Theta(1)$ or $\text{SNR} = O(1)$ then OST leads to successful model selection with high probability provided $n \gtrsim \max \left\{ 1, \frac{1}{\text{SNR} \text{MAR}} \right\} k \log p$.⁴ On the other hand, one of the best known results for model selection using the maximum likelihood algorithm requires that $n \gtrsim \max \left\{ \frac{k \log(p-k)}{\text{SNR} \text{MAR}}, k \log(p/k) \right\}$ [40] (also see [23], [41]). This establishes that OST (and its variants) performs near-optimally for Gaussian design matrices provided (i) the SNR in the measurement system is not too high or (ii) the energy of any nonzero entry of β is not too far away from the average energy $\|\beta\|_2^2/k$ and k scales sublinearly with p .

Remark 2. It is worth pointing out here that similar results can also be obtained for sub-Gaussian design matrices (i.e., matrices with entries given by i.i.d. bounded random variables, etc.) using standard concentration inequalities. Note also that the preceding discussion regarding Gaussian matrices strengthens the results of Fletcher *et al.* [23] concerning asymptotic (Gaussian) model selection using thresholding (cf. [23, Theorem 2]).

C.2 Lasso versus OST

Historically, OST (and its variants) is preferred over the lasso because of its low computational complexity. The results reported in this paper, however, bring forth another important aspect of OST (also see [10]): OST can lead to successful model selection even when the lasso fails. Specifically, note that the lasso solution is not even guaranteed to be unique if the minimum singular value of the submatrix of X corresponding to the true model is not bounded away from zero (see, e.g., [18], [19]). On the other hand, OST does not require the aforementioned condition for model selection. Note that this is in part due to the fact that model selection using the lasso is in fact a byproduct of signal reconstruction, whereas the OST results do not guarantee signal reconstruction without imposing additional constraints on X . In other words, we have established in the paper that model selection is inherently an easier problem than signal reconstruction.

Finally, it is worth comparing the model-selection performance of OST with that of the lasso for the cases when the lasso does succeed. In this regard, the most general result for model selection using the lasso states that if X is close to being a tight frame in the sense that $\|X\|_2 \approx \sqrt{p/n}$ then the lasso identifies the correct model with probability exceeding $1 - O(p^{-1})$ as long as (i) the nonzero entries of β are independent and statistically symmetric around zero, (ii) $k \lesssim n/\log p$, and (iii) $\text{MAR} \gtrsim \frac{k \log p}{n \text{SNR}}$ [20, Theorem 1.3]. On the other hand, assume now that the design matrix X has $\mu(X) \asymp n^{-1/2}$ and $\nu(X) \lesssim n^{-1}$; there indeed exist design matrices that satisfy these conditions (e.g., Gaussian matrices, as proved earlier, and Alltop Gabor frames, as proved in Section IV). We then have from Theorem 2 (resp. Theorem 4) that OST (resp. SOST) identifies the correct model with probability exceeding $1 - O(p^{-1})$ as long as $k \lesssim n/\log p$ and $\text{MAR} \gtrsim \max \left\{ \frac{1}{\text{SNR}}, 1 \right\} \frac{k \log p}{n}$. This suggests that, even for the cases in which the lasso succeeds,

⁴Here, we use the shorthand notation $f(n) \gtrsim g(n)$ (resp. $f(n) \lesssim g(n)$) to indicate that $f(n) \gtrsim g(n)$ (resp. $f(n) \lesssim g(n)$) modulo a logarithmic factor.

OST can be guaranteed to perform as well as the lasso in situations where either the energy of any nonzero entry of β is not too far away from the average energy ($\text{MAR} = \Theta(1)$) or the SNR is not too high ($\text{SNR} = O(1)$). Equally importantly, and in contrast to the lasso results reported in [20], OST is guaranteed to attain this performance *irrespective* of the values of the nonzero entries of the data vector β .

C.3 Near-Optimality of OST

We have concluded up to this point that—under certain conditions on MAR and SNR—the OST algorithm can perform as well as the lasso and it performs near-optimally for Gaussian design matrices. We conclude this discussion by arguing that the OST algorithm in fact performs near-optimally for *any* design matrix that satisfies $\mu(X) \asymp n^{-1/2}$ and $\nu(X) \lesssim n^{-1}$ as long as $\text{MAR} = \Theta(1)$ or $\text{SNR} = O(1)$.⁵ In order to accomplish this goal, we first recall the thresholding results obtained by Donoho and Johnstone [25]—which form the basis of ideas such as the wavelet denoising—for the case of $p \times p$ orthonormal design matrices. Specifically, it was established in [25] that if X is an orthonormal basis then hard thresholding the entries of $X^H y$ at $\lambda \asymp \sqrt{\sigma^2 \log p}$ results in oracle-like performance in the sense that one recovers (with high probability) the locations of all the nonzero entries of β that are above the noise floor.

Now the first thing to note regarding the results presented earlier is the intuitively pleasing nature of the threshold proposed for the OST algorithm. Specifically, assume that X is an orthonormal design and notice that, since $\mu(X) = 0$ in this case, the threshold $\lambda \asymp \max \left\{ \mu \sqrt{n \text{SNR}}, 1 \right\} \sqrt{\sigma^2 \log p}$ proposed earlier reduces to the threshold proposed in [25] and Theorem 5 guarantees that thresholding recovers (with high probability) the locations of all the nonzero entries of β that are above the noise floor: $\text{LAR}_m \gtrsim \frac{k \log p}{n \text{SNR}} \Rightarrow m \in \hat{S}$. Now consider instead design matrices that are not necessarily orthonormal but which satisfy $\mu(X) \asymp n^{-1/2}$ and $\nu(X) \lesssim n^{-1}$. Then, we have from Theorem 5 that OST identifies (with high probability) the locations of the nonzero entries of β whose energies are greater than both the noise power and the average signal energy per nonzero entry: $\text{LAR}_m \gtrsim \max \left\{ \frac{1}{\text{SNR}}, 1 \right\} \frac{k \log p}{n} \Rightarrow m \in \hat{S}$. In particular, under the assumption that either $\text{MAR} = \Theta(1)$ (and since $\text{MAR} \leq \text{LAR}_m$) or $\text{SNR} = O(1)$, this suggests that the OST in such situations performs in a near-optimal (oracle-like) fashion in the sense that it recovers (with high probability) the locations of all the nonzero entries of β that are above the noise floor *without* requiring the design matrix X to be an orthonormal basis.

III. RECOVERY OF SPARSE SIGNALS USING ONE-STEP THRESHOLDING

In this section, we extend our results on model selection using OST to model-order agnostic recovery of k -sparse signals.

⁵Note that it trivially follows from the Welch bound [36] that there exists no design matrix with $p > n \gg 1$ that satisfies $\mu(X) \asymp n^{-1/\gamma}$ with $\gamma < 2$. On the other hand, there does exist a large body of literature devoted to constructing matrices with $\mu(X) \asymp n^{-1/2}$ [42].

In doing so, we also strengthen the results of Schnass and Vandergheynst [22] for signal recovery using thresholding in at least three key aspects. First, we specify polynomial-time verifiable sufficient conditions under which recovery of k -sparse signals using OST succeeds. Second, the threshold that we specify for the OST algorithm (Algorithm 2) does not require knowledge of the model order k . Third, we do not impose a statistical prior on the nonzero entries of the data vector β . Note that, just like [22], we limit ourselves in this exposition to recovery of k -sparse signals in a noiseless setting; extensions of these results to noisy settings would be reported in a sequel to this paper. In other words, the measurement model that we study in this section is $y = X\beta$ and the goal is to recover the k -sparse β using OST under the assumption that the true model $\mathcal{S} \doteq \{i \in \{1, \dots, p\} : |\beta_i| > 0\}$ is a uniformly random k -subset of $\{1, \dots, p\}$.

A. Main Result

Intuitively speaking (and as noted in the discussion in Section II), the problem of sparse-signal recovery is inherently more difficult than the problem of model selection. We capture part of this intuitive notion in the following in terms of the *strong coherence property*.

Definition 2 (The strong coherence property). An $n \times p$ design matrix X having unit ℓ_2 -norm columns is said to obey the strong coherence property if the following hold

$$\mu(X) \leq \frac{1}{60e \log p} \quad \text{and} \quad (\text{SCP-1})$$

$$\nu(X) \leq \frac{\mu}{\sqrt{n}}. \quad (\text{SCP-2})$$

In order to better illustrate the difference between the coherence property and the strong coherence property, note that we have from Lemma 1 and Lemma 2 that Gaussian design matrices satisfy the coherence property with high probability as long as $n \gtrsim (\log p)^2$. On the other hand, Lemma 1 and Lemma 2 suggest that Gaussian design matrices satisfy the strong coherence property with high probability as long as $n \gtrsim (\log p)^4$. In other words, there are scaling regimes in which Gaussian matrices satisfy the coherence property but are not guaranteed to satisfy the strong coherence property. We are now ready to state the main result of this section that makes use of the notation developed earlier in Section II.

Theorem 6 (Sparse-signal recovery using OST). Suppose that the design matrix X satisfies the strong coherence property and choose the threshold $\lambda = 10\mu\|y\|_2 \sqrt{\frac{2 \log p}{1 - e^{-1/2}}}$ for any $p \geq 128$. Then, the OST algorithm (Algorithm 2) satisfies $\Pr(\hat{\beta} \neq \beta) \leq 6p^{-1}$ as long as

$$k \leq \min \left\{ \frac{p}{c_4^2 \|X\|_2^2 \log p}, \frac{\mu^{-2} \text{MAR}}{c_5^2 \log p} \right\}. \quad (6)$$

Here, the probability of failure is only with respect to the true model \mathcal{S} (locations of the nonzero entries of β), while $c_4, c_5 > 0$ are numerical constants given by $c_4 \doteq 37e$ and $c_5 \doteq 43$.

The proof of this theorem is provided in Section V. The significance of Theorem 6 can be best put into perspective by considering the case of the design matrix X being an approximately

tight frame in the sense that $\|X\|_2 \approx \sqrt{p/n}$; indeed, we have that Gaussian design matrices satisfy this condition with high probability [43] and that Gabor frames generated from any (unit-norm) nonzero vector satisfy $\|X\|_2 \equiv \sqrt{p/n} (= \sqrt{n})$ [44]. It then follows from Theorem 6 that if X satisfies the strong coherence property then OST exactly recovers any k -sparse vector β with high probability as long as $k \lesssim \mu^{-2} \text{MAR}$; in particular, if we assume that $\text{MAR} = \Theta(1)$ then this condition reduces to $k \lesssim \mu^{-2}$. On the other hand, low-complexity sparse-recovery algorithms such as subspace pursuit [29], CoSaMP [30], and iterative hard thresholding [31] all rely on the restricted isometry property (RIP) [35]. Therefore, the guarantees provided in [29]–[31] for the case of generic design matrices are limited to k -sparse signals that satisfy $k \lesssim \mu^{-1}$, which is much weaker than the $k \lesssim \mu^{-2}$ scaling claimed here.⁶ We conclude this section by pointing out that if one does have knowledge of the true model order then it can be shown through a slight variation of the proof of Theorem 6 that SOST (the sorted variant of the OST) can also recover sparse signals with high probability—the only difference in that case being that the constant c_5 in Theorem 6 gets replaced with a smaller constant $c'_5 \doteq \sqrt{800}$.

IV. WHY GABOR FRAMES?

Our focus in Section II and Section III has been on establishing that OST leads to successful model selection and sparse-signal recovery under certain conditions on three global parameters of the design matrix X : $\mu(X)$, $\nu(X)$, and $\|X\|_2$. As noted earlier, one particular class of design matrices that satisfies these conditions is the class of random sub-Gaussian matrices. In contrast, our focus in this section is on establishing that Gabor frames—which are collections of time- and frequency-shifts of a nonzero seed vector in \mathbb{C}^n —also tend to satisfy the aforementioned conditions on the matrix geometry. Note that Gabor frames constitute an important class of design matrices because of the facts that (i) Gabor frames are completely specified by a total of n numbers that describe the seed vector, (ii) multiplications with Gabor frames (and their adjoints) can be efficiently carried out using algorithms such as the *fast Fourier transform*, (iii) Gabor frames arise naturally in many important application areas such as communications, radar, and signal/image processing, and (iv) there exist deterministic constructions of Gabor frames that (as shown next) are nearly-optimal in terms of the requisite conditions on $\mu(X)$, $\nu(X)$, and $\|X\|_2$.

A. Geometry of Gabor Frames and Its Implications

A (finite) frame for \mathbb{C}^n is defined as any collection of $p \geq n$ vectors that span the n -dimensional Hilbert space \mathbb{C}^n [49]. Gabor frames for \mathbb{C}^n constitute an important class of frames, having applications in areas such as communications [50] and radar [38], that are constructed from time- and frequency-shifts of a nonzero seed vector in \mathbb{C}^n . Specifically, let $g \in \mathbb{C}^n$ be a unit-norm seed vector and define T to be an $n \times n$ time-shift

matrix that is generated from g as follows

$$T(g) \doteq \begin{bmatrix} g_1 & g_n & & g_2 \\ g_2 & g_1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & g_n \\ g_n & g_{n-1} & & g_1 \end{bmatrix} \quad (7)$$

where we write $T = T(g)$ to emphasize that T is a matrix-valued function on \mathbb{C}^n . Next, denote the collection of n samples of a discrete sinusoid with frequency $2\pi \frac{m}{n}$, $m \in \{0, \dots, n-1\}$ as $\omega_m \doteq [e^{j2\pi \frac{m}{n} 0} \dots e^{j2\pi \frac{m}{n} (n-1)}]^T$. Finally, define the corresponding $n \times n$ diagonal modulation matrices as $W_m = \text{diag}(\omega_m)$. Then, the Gabor frame generated from g is an $n \times n^2$ block matrix of the form

$$X = [W_0 T \quad W_1 T \quad \dots \quad W_{n-1} T]. \quad (8)$$

In words, columns of the Gabor frame X are given by downward circular shifts and modulations (frequency shifts) of the seed vector g . We are now ready to state the first main result concerning the geometry of Gabor frames, which follows directly from [44].

Proposition 1 (Spectral norm of Gabor frames [44]). Gabor frames generated from nonzero (unit-norm) seed vectors are tight frames; in other words, we have that $\|X\|_2 = \sqrt{n}$.

Recall from Theorem 6 and the subsequent discussion that design matrices with small spectral norms are particularly well-suited for recovery of k -sparse signals. In this regard, Proposition 1 implies that Gabor frames are the best that one can hope for in terms of the spectral norm. The next result that we prove concerns the average coherence of Gabor frames.

Theorem 7 (Average coherence of Gabor frames). Let X be a Gabor frame generated from a unit-norm seed vector $g \in \mathbb{C}^n$. Then, using the notation $g_{\max} \doteq \max_i |g_i|$ and $g_{\min} \doteq \min_i |g_i|$, the average coherence of X can be bounded from the above as follows

$$\nu(X) \leq \frac{n g_{\max} (\sqrt{n} - g_{\min}) + 1 - n g_{\min}^2}{n^2 - 1}. \quad (9)$$

Proof: In order to facilitate the proof of this theorem, we first map the indices of the columns of X from $\{1, \dots, n^2\}$ to $\mathcal{C} \doteq \{0, \dots, n-1\} \times \{0, \dots, n-1\}$ as follows

$$\kappa : i \mapsto \left((i \bmod n) - 1, \left\lfloor \frac{i-1}{n} \right\rfloor \right). \quad (10)$$

In words, $\kappa(i) = (\ell, m)$ signifies that the i th column of X corresponds to the $(\ell+1)$ th column of $W_m T$. Next, fix an index i (resp. $\kappa(i) = (\ell, m)$) and make use of the above reindexing to write

$$\begin{aligned} \sum_{\substack{j=1 \\ j \neq i}}^{n^2} \langle x_{\kappa(i)}, x_{\kappa(j)} \rangle &= \sum_{\substack{(\ell', m') \in \mathcal{C} \\ (\ell', m') \neq (\ell, m)}} \langle x_{\ell, m}, x_{\ell', m'} \rangle \\ &= \sum_{\substack{\ell'=0 \\ \ell' \neq \ell}}^{n-1} \sum_{m'=0}^{n-1} \langle x_{\ell, m}, x_{\ell', m'} \rangle + \sum_{\substack{m'=0 \\ m' \neq m}}^{n-1} \langle x_{\ell, m}, x_{\ell, m'} \rangle. \end{aligned} \quad (11)$$

⁶Note that the $k \lesssim \mu^{-1}$ claim is an easy consequence of the Geršgorin circle theorem [45]; see, for example, [39], [46]–[48].

Finally, note that we can explicitly write the columns of X using (8) for any $(\ell, m) \in \mathcal{C}$ as follows

$$\mathbf{x}_{\ell, m} \doteq [g_{(1-\ell)_n} e^{j2\pi \frac{m}{n} 0} \quad \dots \quad g_{(n-\ell)_n} e^{j2\pi \frac{m}{n} (n-1)}]^T \quad (12)$$

where we use the notation $g_{(q)_n}$ as a shorthand for $g_{q \bmod n}$.

The rest of the proof now follows from simple algebraic manipulations. Specifically, it is easy to see from (12) that the first term in (11) can be simplified as

$$\begin{aligned} & \sum_{\substack{\ell'=0 \\ \ell' \neq \ell}}^{n-1} \sum_{m'=0}^{n-1} \langle \mathbf{x}_{\ell, m}, \mathbf{x}_{\ell', m'} \rangle \\ &= \sum_{q=1}^n \sum_{\substack{\ell'=0 \\ \ell' \neq \ell}}^{n-1} g_{(q-\ell)_n}^* g_{(q-\ell')_n} \sum_{m'=0}^{n-1} e^{j2\pi \frac{q-1}{n} (m'-m)} \\ &= \sum_{q=2}^n \sum_{\substack{\ell'=0 \\ \ell' \neq \ell}}^{n-1} g_{(q-\ell)_n}^* g_{(q-\ell')_n} \sum_{m'=0}^{n-1} e^{j2\pi \frac{q-1}{n} (m'-m)} + \\ & \quad + n \sum_{\substack{\ell'=0 \\ \ell' \neq \ell}}^{n-1} g_{(1-\ell)_n}^* g_{(1-\ell')_n} \stackrel{(a)}{=} n g_{(1-\ell)_n}^* \sum_{\substack{\ell'=0 \\ \ell' \neq \ell}}^{n-1} g_{(1-\ell')_n} \quad (13) \end{aligned}$$

where (a) in the above expression is a consequence of the fact that $\sum_{m'=0}^{n-1} e^{j2\pi \frac{q-1}{n} (m'-m)} = 0$ for any fixed $q \in \{2, \dots, n\}$. Likewise, we can simplify the second term in (11) as follows

$$\begin{aligned} & \sum_{\substack{m'=0 \\ m' \neq m}}^{n-1} \langle \mathbf{x}_{\ell, m}, \mathbf{x}_{\ell, m'} \rangle \\ &= \sum_{q=1}^n g_{(q-\ell)_n}^* g_{(q-\ell)_n} \sum_{\substack{m'=0 \\ m' \neq m}}^{n-1} e^{j2\pi \frac{q-1}{n} (m'-m)} \\ &= \sum_{q=2}^n |g_{(q-\ell)_n}|^2 \sum_{\substack{m'=0 \\ m' \neq m}}^{n-1} e^{j2\pi \frac{q-1}{n} (m'-m)} + |g_{(1-\ell)_n}|^2 \sum_{\substack{m'=0 \\ m' \neq m}}^{n-1} 1 \\ &\stackrel{(b)}{=} - \sum_{q=2}^n |g_{(q-\ell)_n}|^2 + (n-1) |g_{(1-\ell)_n}|^2 \\ &= -1 + n |g_{(1-\ell)_n}|^2 \quad (14) \end{aligned}$$

where (b) follows from the fact that $\sum_{m' \neq m} e^{j2\pi \frac{q-1}{n} (m'-m)} = -1$ for any fixed $q \in \{2, \dots, n\}$.

To conclude the theorem, note from (11), (13), and (14) that we can write

$$\begin{aligned} & \max_{i \in \{1, \dots, n^2\}} \left| \sum_{\substack{j=1 \\ j \neq i}}^{n^2} \langle \mathbf{x}_i, \mathbf{x}_j \rangle \right| \\ &= \max_{\ell} \left| n g_{(1-\ell)_n}^* \sum_{\substack{\ell'=0 \\ \ell' \neq \ell}}^{n-1} g_{(1-\ell')_n} - 1 + n |g_{(1-\ell)_n}|^2 \right| \\ &\stackrel{(c)}{\leq} \max_{r \in \{1, \dots, n\}} \left| n g_r^* \sum_{\substack{s=1 \\ s \neq r}}^n g_s \right| + \max_{r \in \{1, \dots, n\}} | -1 + n |g_r|^2 | \end{aligned}$$

$$\begin{aligned} & \leq n \max_{r \in \{1, \dots, n\}} |g_r| \sum_{\substack{s=1 \\ s \neq r}}^n |g_s| + 1 - n g_{\min}^2 \\ &\stackrel{(d)}{\leq} n g_{\max} (\sqrt{n} - g_{\min}) + 1 - n g_{\min}^2. \quad (15) \end{aligned}$$

Here, (c) mainly follows from the triangle inequality and a simple reindexing argument, while (d) mainly follows from the Cauchy-Schwarz inequality since $\sum_{\substack{s=1 \\ s \neq r}}^n |g_s| = \|g\|_1 - |g_r| \leq \sqrt{n} - g_{\min}$. The proof of the theorem now follows by dividing the above expression by $n^2 - 1$. \square

In words, Theorem 7 states that the average coherence of Gabor frames cannot be too large. In particular, it implies that Gabor frames generated from unimodal (unit-norm) seed vectors (i.e., seed vectors characterized by $g_{\min} \asymp g_{\max} \asymp n^{-1/2}$) satisfy $\nu(X) \lesssim n^{-1}$. On the other hand, recall that the Welch bound [36] dictates that $\mu(X) \geq (n+1)^{-1/2}$ for Gabor frames. It is therefore easy to conclude from these two facts that Gabor frames generated from unimodal seed vectors are automatically guaranteed to satisfy the coherence property (resp. strong coherence property) as long as $\mu(X) \lesssim (\log p)^{-1/2}$ (resp. $\mu(X) \lesssim (\log p)^{-1}$). In the context of model selection and sparse-signal recovery, Theorem 7 therefore suggests that Gabor frames generated from unimodal seed vectors are the best that one can hope for in terms of the average coherence.

Finally, recall from the discussions in Section II and Section III that—among the class of matrices that satisfy the (strong) coherence property—design matrices with small worst-case coherence are particularly well-suited for model selection and sparse-signal recovery. In the context of Gabor frames, the goal then is to design unimodal seed vectors that yield Gabor frames with the smallest-possible worst-case coherence. This, however, is an active area of mathematical research and a number of researchers have looked at this problem in recent years; see, e.g., [8]. As such, we can simply leverage some of the existing research in this area in order to provide explicit constructions of Gabor frames that satisfy the (strong) coherence property with nearly-optimal worst-case coherence.

Specifically, let $n \geq 5$ be a prime number and construct a unimodal seed vector $g \in \mathbb{C}^n$ as follows

$$g = \left[\frac{1}{\sqrt{n}} e^{j2\pi \frac{0^3}{n}} \quad \frac{1}{\sqrt{n}} e^{j2\pi \frac{1^3}{n}} \quad \dots \quad \frac{1}{\sqrt{n}} e^{j2\pi \frac{(n-1)^3}{n}} \right]^T. \quad (16)$$

The sequence $\left\{ \frac{1}{\sqrt{n}} e^{j2\pi \frac{q^3}{n}} \right\}_{q=0}^{n-1}$ is termed as the *Alltop sequence* [7] in the literature. This sequence has the property that its autocorrelation decays very fast and, therefore, it is particularly well-suited for generating Gabor frames with small worst-case coherence. In particular, it was established recently in [8] that Gabor frames generated from the Alltop seed vector g given in (16) satisfy

$$\mu(X) \doteq \max_{i, j: i \neq j} |\langle \mathbf{x}_i, \mathbf{x}_j \rangle| \leq \frac{1}{\sqrt{n}}. \quad (17)$$

In addition, since we have that $g_{\min} = g_{\max} = n^{-1/2}$ for the Alltop seed vector, it is easy to check using Theorem 7 that the average coherence of Alltop Gabor frames satisfies

$\nu(X) \leq (n+1)^{-1} \leq \mu(X)/\sqrt{n}$. An immediate consequence of this discussion is that all the results reported in Section II and Section III in the context of model selection and sparse-signal recovery using OST apply directly to the case of Alltop Gabor frames. In particular, it follows from Theorem 6 that Alltop Gabor frames together with OST are guaranteed to recover most k -sparse signals—regardless of the statistical dependence across the nonzero entries of β —as long as $k \lesssim n$ and $\text{MAR} = \Theta(1)$. In contrast, the only other results available in the sparse-signal recovery literature for Alltop Gabor frames are based on the higher-complexity basis pursuit [3] and require the nonzero entries of β to be independent and statistically symmetric around zero for the case when $\sqrt{n} \lesssim k \lesssim n$ [37], [38].

V. PROOFS OF MAIN RESULTS

In this section, we provide detailed proofs of the main results reported in Section II and Section III. Before proceeding further, however, it is advantageous to develop some notation that will facilitate our forthcoming analysis. In this regard, recall that the true model \mathcal{S} is taken to be a uniformly random k -subset of $[p] \doteq \{1, \dots, p\}$. We can therefore write the data vector β under this assumption as concatenation of a random permutation matrix and a deterministic k -sparse vector. Specifically, let $\bar{z} \in \mathbb{C}^p$ be a *deterministic* k -sparse vector that we write (without loss of generality) as

$$\bar{z} \doteq \left(\underbrace{z_1, \dots, z_k}_{\doteq z \in \mathbb{C}^k}, \underbrace{0, \dots, 0}_{(p-k) \text{ times}} \right)^T \quad (18)$$

and let P_π be a $p \times p$ random permutation matrix; in other words,

$$P_\pi \doteq [e_{\pi_1} \quad e_{\pi_2} \quad \dots \quad e_{\pi_p}]^T \quad (19)$$

where e_j denotes the j th column of the canonical basis I and $\bar{\Pi} \doteq (\pi_1, \dots, \pi_p)$ is a random permutation of $[p]$. Then, the assumption that the model \mathcal{S} is a random subset of $[p]$ is equivalent to stating that the data vector β can be written as $\beta = P_\pi \bar{z}$. In other words, the measurement vector y can be expressed as

$$y = X\beta + \eta = X P_\pi \bar{z} + \eta = X_\Pi z + \eta \quad (20)$$

where $\Pi \doteq (\pi_1, \dots, \pi_k)$ denotes the first k elements of the random permutation $\bar{\Pi}$, X_Π denotes the $n \times k$ submatrix obtained by collecting the columns of X corresponding to the indices in Π , and the vector $z \in \mathbb{C}^k$ represents the k nonzero entries of β .

A. Proof of Theorem 1

The general road map for the proof of Theorem 1 is as follows. Below, we first introduce the notion of (k, ϵ, δ) -statistical orthogonality condition (StOC). We next establish the relationship between the StOC parameters and the worst-case and average coherence of X in Lemma 3 and Lemma 4. We then provide a proof of Theorem 1 by first showing that if X satisfies the StOC then OST recovers \mathcal{S} with high probability and then relating the results of Lemma 3 and Lemma 4 to the coherence property.

Definition 3 ((k, ϵ, δ) -Statistical orthogonality condition).

Let $\bar{\Pi} = (\pi_1, \dots, \pi_p)$ be a random permutation of $[p]$, and define $\Pi \doteq (\pi_1, \dots, \pi_k)$ and $\Pi^c \doteq (\pi_{k+1}, \dots, \pi_p)$ for any $k \in [p]$.

Then, the $n \times p$ (normalized) design matrix X is said to satisfy the (k, ϵ, δ) -statistical orthogonality condition if there exist $\epsilon, \delta \in [0, 1)$ such that the inequalities

$$\|(X_\Pi^H X_\Pi - I)z\|_\infty \leq \epsilon \|z\|_2 \quad (\text{StOC-1})$$

$$\|X_{\Pi^c}^H X_\Pi z\|_\infty \leq \delta \|z\|_2 \quad (\text{StOC-2})$$

hold for every *fixed* $z \in \mathbb{C}^k$ with probability exceeding $1 - \delta$ (with respect to the random permutation $\bar{\Pi}$).

Remark 3. Note that the StOC derives its name from the fact that if X is a $p \times p$ orthonormal matrix then it trivially satisfies the StOC for every $k \in [p]$ with $\epsilon = \delta = 0$. In addition, although we will not use this fact explicitly in the paper, it can be checked that if X satisfies (k, ϵ, δ) -StOC then it *approximately* preserves the ℓ_2 -norms of k -sparse signals with probability exceeding $1 - \delta$ as long as $k < \epsilon^{-2}$. Having defined StOC, our goal in the next two lemmas is to relate the StOC parameters k, ϵ , and δ to the worst-case and average coherence of the design matrix X .

Lemma 3. Let $\Pi = (\pi_1, \dots, \pi_k)$ denote the first k elements of a random permutation of $[p]$ and choose a parameter $a \geq 1$. Then, for any $\epsilon \in [0, 1)$, $k \leq \min \{\epsilon^2 \nu^{-2}, (1+a)^{-1}p\}$, and fixed $z \in \mathbb{C}^k$, we have

$$\Pr \left(\{X \text{ does not satisfy (StOC-1)}\} \right) \leq 4k \exp \left(- \frac{(\epsilon - \sqrt{k} \nu)^2}{16(2+a^{-1})^2 \mu^2} \right). \quad (21)$$

Proof: The proof of this lemma relies heavily on the so-called *method of bounded differences* (MOBD) [51]. Specifically, we begin by noting that $\|(X_\Pi^H X_\Pi - I)z\|_\infty = \max_i |\sum_{j \neq i} z_j \langle x_{\pi_i}, x_{\pi_j} \rangle|$. Therefore, for a fixed index i , and conditioned on the event $\mathcal{A}_{i'} \doteq \{\pi_i = i'\}$, we have the following equality from basic probability theory

$$\begin{aligned} \Pr \left(\left| \sum_{\substack{j=1 \\ j \neq i}}^k z_j \langle x_{\pi_i}, x_{\pi_j} \rangle \right| > \epsilon \|z\|_2 \middle| \mathcal{A}_{i'} \right) \\ = \Pr \left(\left| \sum_{\substack{j=1 \\ j \neq i}}^k z_j \langle x_{i'}, x_{\pi_j} \rangle \right| > \epsilon \|z\|_2 \middle| \mathcal{A}_{i'} \right). \end{aligned} \quad (22)$$

Next, in order to apply the MOBD to obtain an upper bound for (22), we first define a random $(k-1)$ -tuple $\Pi^{-i} \doteq (\pi_1, \dots, \pi_{i-1}, \pi_{i+1}, \dots, \pi_k)$ and then construct a Doob martingale $(M_0, M_1, \dots, M_{k-1})$ as follows

$$\begin{aligned} M_0 &= \mathbb{E} \left[\sum_{\substack{j=1 \\ j \neq i}}^k z_j \langle x_{i'}, x_{\pi_j} \rangle \middle| \mathcal{A}_{i'} \right], \quad \text{and} \\ M_\ell &= \mathbb{E} \left[\sum_{\substack{j=1 \\ j \neq i}}^k z_j \langle x_{i'}, x_{\pi_j} \rangle \middle| \pi_{1 \rightarrow \ell}^{-i}, \mathcal{A}_{i'} \right], \quad \ell = 1, \dots, k-1 \end{aligned}$$

where $\pi_{1 \rightarrow \ell}^{-i}$ denotes the first ℓ elements of Π^{-i} . The first thing to note here is that we have from the linearity of (conditional)

$$\begin{aligned}
|M_\ell(r) - M_\ell(s)| &= \left| \sum_{j \neq i} z_j \left(\mathbb{E}[\langle x_{i'}, x_{\pi_j} \rangle | \pi_{1 \rightarrow \ell-1}^{-i}, \pi_\ell^{-i} = r, \mathcal{A}_{i'}] - \mathbb{E}[\langle x_{i'}, x_{\pi_j} \rangle | \pi_{1 \rightarrow \ell-1}^{-i}, \pi_\ell^{-i} = s, \mathcal{A}_{i'}] \right) \right| \\
&\leq \sum_{j \neq i} |z_j| \underbrace{\left| \mathbb{E}[\langle x_{i'}, x_{\pi_j} \rangle | \pi_{1 \rightarrow \ell-1}^{-i}, \pi_\ell^{-i} = r, \mathcal{A}_{i'}] - \mathbb{E}[\langle x_{i'}, x_{\pi_j} \rangle | \pi_{1 \rightarrow \ell-1}^{-i}, \pi_\ell^{-i} = s, \mathcal{A}_{i'}] \right|}_{\doteq d_{\ell,j}} \\
&= \sum_{\substack{j \leq \ell+1 \\ j \neq i}} |z_j| |d_{\ell,j}| + \sum_{\substack{j > \ell+1 \\ j \neq i}} |z_j| |d_{\ell,j}|. \tag{25}
\end{aligned}$$

expectation

$$\begin{aligned}
|M_0| &= \left| \sum_{j \neq i} z_j \mathbb{E}[\langle x_{i'}, x_{\pi_j} \rangle | \mathcal{A}_{i'}] \right| \\
&\leq \sum_{j \neq i} |z_j| \left| \mathbb{E}[\langle x_{i'}, x_{\pi_j} \rangle | \mathcal{A}_{i'}] \right| \\
&\stackrel{(a)}{\leq} \sum_{j \neq i} |z_j| \left| \sum_{\substack{q=1 \\ q \neq i'}}^p \frac{1}{p-1} \langle x_{i'}, x_q \rangle \right| \\
&\stackrel{(b)}{\leq} \nu \|z\|_1 \leq \sqrt{k} \nu \|z\|_2 \tag{23}
\end{aligned}$$

where (a) follows from the fact that, conditioned on $\mathcal{A}_{i'}$, π_j has a uniform distribution over $[p] - \{i'\}$, while (b) is mainly a consequence of the definition of average coherence. In addition, if we use π_ℓ^{-i} to denote the ℓ th element of Π^{-i} and define

$$M_\ell(r) \doteq \mathbb{E} \left[\sum_{\substack{j=1 \\ j \neq i}}^k z_j \langle x_{i'}, x_{\pi_j} \rangle \middle| \pi_{1 \rightarrow \ell-1}^{-i}, \pi_\ell^{-i} = r, \mathcal{A}_{i'} \right] \tag{24}$$

for $\ell = 1, \dots, k-1$ then, since $(M_0, M_1, \dots, M_{k-1})$ is a Doob martingale, it can be easily verified that $|M_\ell - M_{\ell-1}|$ is upper bounded by $\sup_{r,s} [M_\ell(r) - M_\ell(s)]$ (see, e.g., [52]).

Now in order to upper bound $\sup_{r,s} [M_\ell(r) - M_\ell(s)]$, notice that we can bound $|M_\ell(r) - M_\ell(s)|$ as in (25). In addition, we have that for every $j > \ell+1, j \neq i$, the random variable π_j has a uniform distribution over $[p] - \{\pi_{1 \rightarrow \ell-1}^{-i}, r, i'\}$ when conditioned on $\{\pi_{1 \rightarrow \ell-1}^{-i}, \pi_\ell^{-i} = r, i'\}$, whereas π_j has a uniform distribution over $[p] - \{\pi_{1 \rightarrow \ell-1}^{-i}, s, i'\}$ when conditioned on $\{\pi_{1 \rightarrow \ell-1}^{-i}, \pi_\ell^{-i} = s, i'\}$. Therefore, we get $\forall j > \ell+1, j \neq i$,

$$|d_{\ell,j}| = \frac{1}{p-\ell-1} \left| \langle x_{i'}, x_r \rangle - \langle x_{i'}, x_s \rangle \right| \leq \frac{2\mu}{p-\ell-1}. \tag{26}$$

Similarly, it can be shown that $\sum_{\substack{j \leq \ell+1 \\ j \neq i}} |z_j| |d_{\ell,j}| \leq |z_{\ell+1}| 2\mu$ when $i \leq \ell$, $\sum_{\substack{j \leq \ell+1 \\ j \neq i}} |z_j| |d_{\ell,j}| \leq |z_\ell| 2\mu$ when $i = \ell+1$, and $\sum_{\substack{j \leq \ell+1 \\ j \neq i}} |z_j| |d_{\ell,j}| \leq (|z_\ell| + \frac{|z_{\ell+1}|}{p-\ell-1}) 2\mu$ when $i > \ell+1$. Consequently, regardless of the initial choice of i , we obtain

$$\sup_{r,s} [M_\ell(r) - M_\ell(s)]$$

$$\leq 2\mu \underbrace{\left(|z_\ell| + |z_{\ell+1}| + \frac{1}{p-\ell-1} \sum_{j > \ell+1} |z_j| \right)}_{\doteq d_\ell}. \tag{27}$$

We have now established that $(M_0, M_1, \dots, M_{k-1})$ is a (real- or complex-valued) bounded-difference martingale sequence with $|M_\ell - M_{\ell-1}| \leq 2\mu d_\ell$ for $\ell = 1, \dots, k-1$. Therefore, under the assumption that $k \leq \epsilon^2 \nu^{-2}$ and since it has been established in (23) that $|M_0| \leq \sqrt{k} \nu \|z\|_2$, it is easy to see that

$$\begin{aligned}
&\Pr \left(\left| \sum_{\substack{j=1 \\ j \neq i}}^k z_j \langle x_{i'}, x_{\pi_j} \rangle \right| > \epsilon \|z\|_2 \middle| \mathcal{A}_{i'} \right) \\
&\leq \Pr \left(|M_{k-1} - M_0| > \epsilon \|z\|_2 - \sqrt{k} \nu \|z\|_2 \middle| \mathcal{A}_{i'} \right) \\
&\stackrel{(c)}{\leq} 4 \exp \left(- \frac{(\epsilon - \sqrt{k} \nu)^2 \|z\|_2^2}{16\mu^2 \sum_{\ell=1}^{k-1} d_\ell^2} \right) \tag{28}
\end{aligned}$$

where (c) follows from the complex Azuma inequality for bounded-difference martingale sequences (see Lemma 5 in Appendix A). Further, it can be established through routine calculations from (27) that $\sum_{\ell=1}^{k-1} d_\ell^2 \leq (2 + a^{-1})^2 \|z\|_2^2$ since $k \leq p/(1+a)$. Combining all these facts together, we finally obtain that

$$\begin{aligned}
&\Pr \left(\|(X_\Pi^H X_\Pi - I)z\|_\infty > \epsilon \|z\|_2 \right) \\
&\stackrel{(d)}{\leq} k \Pr \left(\left| \sum_{\substack{j=1 \\ j \neq i}}^k z_j \langle x_{\pi_i}, x_{\pi_j} \rangle \right| > \epsilon \|z\|_2 \right) \\
&= k \sum_{i'=1}^p \Pr \left(\left| \sum_{\substack{j=1 \\ j \neq i}}^k z_j \langle x_{i'}, x_{\pi_j} \rangle \right| > \epsilon \|z\|_2 \middle| \mathcal{A}_{i'} \right) \Pr(\mathcal{A}_{i'}) \\
&\stackrel{(e)}{\leq} 4k \exp \left(- \frac{(\epsilon - \sqrt{k} \nu)^2}{16(2 + a^{-1})^2 \mu^2} \right) \tag{29}
\end{aligned}$$

where (d) follows from the union bound and the fact that the π_i 's are identically (though not independently) distributed, while (e) follows from (28) and the fact that π_i has a uniform distribution over $[p]$. \square

Lemma 4. Let $\Pi = (\pi_1, \dots, \pi_k)$ and $\Pi^c = (\pi_{k+1}, \dots, \pi_p)$ denote the first k and the last $(p-k)$ elements of a random

$$\begin{aligned}
|M_\ell(r) - M_\ell(s)| &= \left| \sum_j z_j \left(\mathbb{E}[\langle x_{i'}, x_{\pi_j} \rangle | \pi_{1 \rightarrow \ell-1}, \pi_\ell = r, \mathcal{A}_{i'}] - \mathbb{E}[\langle x_{i'}, x_{\pi_j} \rangle | \pi_{1 \rightarrow \ell-1}, \pi_\ell = s, \mathcal{A}_{i'}] \right) \right| \\
&\leq |z_\ell| |\langle x_{i'}, x_r \rangle - \langle x_{i'}, x_s \rangle| + \frac{|\langle x_{i'}, x_r \rangle - \langle x_{i'}, x_s \rangle|}{p - \ell - 1} \sum_{j > \ell} |z_j| \leq 2\mu \underbrace{\left(|z_\ell| + \frac{\sum_{j > \ell} |z_j|}{p - \ell - 1} \right)}_{\doteq d_\ell}. \quad (32)
\end{aligned}$$

permutation of $\llbracket p \rrbracket$, respectively, and choose a parameter $a \geq 1$. Then, for any $\epsilon \in [0, 1]$, $k \leq \min\{\epsilon^2 \nu^{-2}, (1+a)^{-1}p\}$, and fixed $z \in \mathbb{C}^k$, we have

$$\begin{aligned}
&\Pr\left(\{X \text{ does not satisfy (StOC-2)}\}\right) \\
&\leq 4(p-k) \exp\left(-\frac{(\epsilon - \sqrt{k}\nu)^2}{8(1+a^{-1})^2\mu^2}\right). \quad (30)
\end{aligned}$$

Proof: The proof of this lemma is very similar to that of Lemma 3 and also relies on the MOBD. To begin with, we note that $\|X_{\Pi^c}^H X_{\Pi} z\|_\infty = \max_{i \in \llbracket p-k \rrbracket} \left| \sum_j z_j \langle x_{\pi_i^c}, x_{\pi_j} \rangle \right|$, where $\llbracket p-k \rrbracket \doteq \{1, \dots, p-k\}$ and π_i^c denotes the i th element of Π^c . Then, for a fixed index $i \in \llbracket p-k \rrbracket$, and conditioned on the event $\mathcal{A}_{i'} \doteq \{\pi_i^c = i'\}$, we again have the following simple equality

$$\begin{aligned}
&\Pr\left(\left|\sum_{j=1}^k z_j \langle x_{\pi_i^c}, x_{\pi_j} \rangle\right| > \epsilon \|z\|_2 \middle| \mathcal{A}_{i'}\right) \\
&= \Pr\left(\left|\sum_{j=1}^k z_j \langle x_{i'}, x_{\pi_j} \rangle\right| > \epsilon \|z\|_2 \middle| \mathcal{A}_{i'}\right). \quad (31)
\end{aligned}$$

Next, as in the case of Lemma 3, we construct a Doob martingale sequence (M_0, M_1, \dots, M_k) as follows

$$\begin{aligned}
M_0 &= \mathbb{E}\left[\sum_{j=1}^k z_j \langle x_{i'}, x_{\pi_j} \rangle \middle| \mathcal{A}_{i'}\right] \quad \text{and} \\
M_\ell &= \mathbb{E}\left[\sum_{j=1}^k z_j \langle x_{i'}, x_{\pi_j} \rangle \middle| \pi_{1 \rightarrow \ell}, \mathcal{A}_{i'}\right], \quad \ell = 1, \dots, k
\end{aligned}$$

where $\pi_{1 \rightarrow \ell}$ now denotes the first ℓ elements of Π . Then, since π_j has a uniform distribution over $\llbracket p \rrbracket - \{i'\}$ when conditioned on $\mathcal{A}_{i'}$, we once again have the bound $|M_0| \leq \sqrt{k}\nu \|z\|_2$. Therefore, the only remaining thing that we need to show in order to be able to apply the complex Azuma inequality to the constructed martingale (M_0, M_1, \dots, M_k) is that $|M_\ell - M_{\ell-1}|$ is suitably bounded.

In this regard, we make use of the notation $M_\ell(r) \doteq \mathbb{E}\left[\sum_{j=1}^k z_j \langle x_{i'}, x_{\pi_j} \rangle \middle| \pi_{1 \rightarrow \ell-1}, \pi_\ell = r, \mathcal{A}_{i'}\right]$ and note that $|M_\ell(r) - M_\ell(s)|$ can be bounded as in (32), which implies that $\sup_{r,s} [M_\ell(r) - M_\ell(s)] \leq 2\mu d_\ell$, $\ell = 1, \dots, k$. Consequently, we have now established that (M_0, M_1, \dots, M_k) is a bounded-difference martingale with $|M_\ell - M_{\ell-1}| \leq 2\mu d_\ell$. Therefore,

since $k \leq \epsilon^2 \nu^{-2}$ and $|M_0| \leq \sqrt{k}\nu \|z\|_2$, we once again have from the complex Azuma inequality that

$$\begin{aligned}
&\Pr\left(\left|\sum_{j=1}^k z_j \langle x_{i'}, x_{\pi_j} \rangle\right| > \epsilon \|z\|_2 \middle| \mathcal{A}_{i'}\right) \\
&\leq \Pr\left(|M_k - M_0| > \epsilon \|z\|_2 - \sqrt{k}\nu \|z\|_2 \middle| \mathcal{A}_{i'}\right) \\
&\stackrel{(a)}{\leq} 4 \exp\left(-\frac{(\epsilon - \sqrt{k}\nu)^2}{8(1+a^{-1})^2\mu^2}\right) \quad (33)
\end{aligned}$$

where (a) follows by noting that $\sum_{\ell=1}^k d_\ell^2 \leq (1+a^{-1})^2 \|z\|_2^2$ since $k \leq p/(1+a)$. Combining all these facts together, we finally obtain the claimed result as follows

$$\begin{aligned}
&\Pr\left(\|X_{\Pi^c}^H X_{\Pi} z\|_\infty > \epsilon \|z\|_2\right) \\
&\stackrel{(b)}{\leq} (p-k) \Pr\left(\left|\sum_{j=1}^k z_j \langle x_{\pi_i^c}, x_{\pi_j} \rangle\right| > \epsilon \|z\|_2\right) \\
&\leq (p-k) \sum_{i'=1}^p \Pr\left(\left|\sum_{j=1}^k z_j \langle x_{i'}, x_{\pi_j} \rangle\right| > \epsilon \|z\|_2 \middle| \mathcal{A}_{i'}\right) \Pr(\mathcal{A}_{i'}) \\
&\stackrel{(c)}{\leq} 4(p-k) \exp\left(-\frac{(\epsilon - \sqrt{k}\nu)^2}{8(1+a^{-1})^2\mu^2}\right) \quad (34)
\end{aligned}$$

where (b) follows from the union bound and the fact that the π_i^c 's are identically (though not independently) distributed, while (c) follows from (33) and the fact that π_i^c has a uniform distribution over $\llbracket p \rrbracket$. \square

Note that Lemma 3 and Lemma 4 collectively prove through a simple union bound argument that an $n \times p$ design matrix X satisfies (k, ϵ, δ) -StOC for any $\epsilon \in [0, 1]$ with $\delta \leq 4p \exp\left(-\frac{(\epsilon - \sqrt{k}\nu)^2}{16(2+a^{-1})^2\mu^2}\right)$ for any $a \geq 1$ as long as we have that $k \leq \min\{\epsilon^2 \nu^{-2}, (1+a)^{-1}p\}$. We are now ready to provide a proof of Theorem 1.

Proof (Theorem 1): We begin by making use of the notation developed at the start of this section and writing the signal proxy $f = X^H y$ as $f = X^H X_{\Pi} z + X^H \eta$. Now, let $\Pi^c = (\pi_{k+1}, \dots, \pi_p)$ denote the last $(p-k)$ elements of Π and note that we need to show that $\|f_{\Pi^c}\|_\infty \leq \lambda$ and $\min_{i \in \{1, \dots, k\}} |f_{\pi_i}| > \lambda$ in order to establish that $\hat{S} = S$.

In this regard, we first assume that X satisfies (k, ϵ, δ) -StOC and define $\lambda_\epsilon \doteq \max\left\{\frac{1}{t}\epsilon \|z\|_2, \frac{1}{1-t}2\sqrt{\sigma^2 \log p}\right\}$ for any $t \in (0, 1)$. Next, it can be verified through Lemma 6 in Appendix A

that $\tilde{\eta} \doteq X^H \eta$ satisfies $\|\tilde{\eta}\|_\infty \leq 2\sqrt{\sigma^2 \log p}$ with probability exceeding $1 - 2(\sqrt{2\pi \log pp})^{-1}$. Now define the probability event

$$\mathcal{G} \doteq \left\{ \left\{ X \text{ satisfies (StOC-1) and (StOC-2)} \right\} \cap \bigcap \left\{ \|\tilde{\eta}\|_\infty \leq 2\sqrt{\sigma^2 \log p} \right\} \right\} \quad (35)$$

and notice that we have $\Pr(\mathcal{G}) > 1 - \delta - 2(\sqrt{2\pi \log pp})^{-1}$. Further, conditioned on the event \mathcal{G} , we have

$$\begin{aligned} \|f_{\Pi^c}\|_\infty &\stackrel{(a)}{\leq} \|X_{\Pi^c}^H X_{\Pi} z\|_\infty + \|X_{\Pi^c}^H \eta\|_\infty \\ &\stackrel{(b)}{\leq} \epsilon \|z\|_2 + 2\sqrt{\sigma^2 \log p} \stackrel{(c)}{\leq} \lambda_\epsilon \end{aligned} \quad (36)$$

where (a) follows from the triangle inequality, (b) is mainly a consequence of the conditioning on the event \mathcal{G} , and (c) follows from the definition of λ_ϵ . Next, we define $r = (X_{\Pi}^H X_{\Pi} - I)z$ and notice that, conditioned on the event \mathcal{G} , we have for any $i \in [k] \doteq \{1, \dots, k\}$ the following inequality

$$\begin{aligned} |f_{\pi_i}| &= |z_i + r_i + \tilde{\eta}_{\pi_i}| \geq |z_i| - \|r\|_\infty - \|\tilde{\eta}\|_\infty \\ &\stackrel{(d)}{\geq} \beta_{\min} - \epsilon \|z\|_2 - 2\sqrt{\sigma^2 \log p} \stackrel{(e)}{\geq} \beta_{\min} - \lambda_\epsilon. \end{aligned} \quad (37)$$

Here, (d) follows from the conditioning on \mathcal{G} , while (e) is a simple consequence of the choice of λ_ϵ . It can therefore be concluded from (36) and (37) that if X satisfies (k, ϵ, δ) -StOC and the OST algorithm uses the threshold λ_ϵ then we have $\Pr(\hat{\mathcal{S}} \neq \mathcal{S}) \leq \Pr(\mathcal{G}^c)$ as long as $\beta_{\min} > 2\lambda_\epsilon$.

Finally, to complete the proof of this theorem, we let $k \leq n/(2 \log p)$ and fix $\epsilon = 10\mu\sqrt{2 \log p}$. Then, the claim is that X satisfies (k, ϵ, δ) -StOC with $\delta \leq 4p^{-1}$. In order to establish this claim, we only need to ensure that the chosen parameters satisfy the assumptions of Lemma 3 and Lemma 4. In this regard, note that (i) $\epsilon < 1$ because of (CP-1), and (ii) $\sqrt{k} \nu \leq \frac{\epsilon}{9}$ because of the assumption that $k \leq n/(2 \log p)$ and (CP-2). Therefore, since the assumption $p \geq 128$ together with $k \leq n/(2 \log p)$ implies that $16(2 + a^{-1})^2 < 72$, we obtain $\exp\left(-\frac{(\epsilon - \sqrt{k} \nu)^2}{16(2 + a^{-1})^2 \mu^2}\right) \leq p^{-2}$. We can now combine this fact with the previously established facts to see that the threshold $\lambda = \max\left\{\frac{1}{t} 10\mu\sqrt{n \text{SNR}}, \frac{1}{1-t} \sqrt{2}\right\} \sqrt{2\sigma^2 \log p}$ guarantees that $\Pr(\hat{\mathcal{S}} \neq \mathcal{S}) \leq 6p^{-1}$ as long as $n \geq 2k \log p$ and $\beta_{\min} > 2\lambda$. Finally, note that

$$\beta_{\min} > \frac{1}{1-t} 4\sqrt{\sigma^2 \log p} \iff n > \frac{c_2 k \log p}{\text{SNR}_{\min}}$$

and

$$\beta_{\min} > \frac{1}{t} 20\mu\sqrt{2n\sigma^2 \log p \text{SNR}} \iff n > \left(\frac{c_3 k \log p}{\text{MAR}}\right)^{\gamma/2}.$$

This completes the proof of the theorem. \square

B. Proof of Theorem 5

Proof (Theorem 5): We begin by making use of the notation developed earlier in this section and conditioning on the event

\mathcal{G} defined in (35) with $\epsilon = 10\mu\sqrt{2 \log p}$. Then, it is easy to see from the proof of Theorem 1 that the estimate $\hat{\mathcal{S}}$ is a subset of \mathcal{S} because of the fact that $\|f_{\Pi^c}\|_\infty \leq \lambda$.

Next, assume without loss of generality that $z_i \equiv \beta_{(i)}$ and note from (37) that $|f_{\pi_i}| \geq |\beta_{(i)}| - \lambda$ for any $i \in \{1, \dots, k\}$. Then, since $\pi_i \in \hat{\mathcal{S}}$ if and only if $|f_{\pi_i}| > \lambda$, we have that $\beta_{(i)} > 2\lambda \Rightarrow \pi_i \in \hat{\mathcal{S}}$. Now define M to be the largest integer for which $\beta_{(M)} > 2\lambda$ holds and note that $\beta_{(M)} > 2\lambda \Rightarrow \beta_{(i)} > 2\lambda \Rightarrow \pi_i \in \hat{\mathcal{S}}$ for every $i \in \{1, \dots, M\}$, which in turn implies $|\mathcal{S} - \hat{\mathcal{S}}| \leq (k - M)$. Finally, note that

$$\beta_{(M)} > \frac{1}{1-t} 4\sqrt{\sigma^2 \log p} \iff \text{LAR}_M > \frac{c_2 k \log p}{n \text{SNR}}$$

and

$$\beta_{(M)} > \frac{1}{t} 20\mu\sqrt{2n\sigma^2 \log p \text{SNR}} \iff \text{LAR}_M > \frac{c'_3 k \log p}{\mu^{-2}}.$$

This completes the proof of the theorem since the event \mathcal{G} holds with probability exceeding $1 - 6p^{-1}$. \square

C. Proof of Theorem 6

The first key result that we will need to prove Theorem 6 is regarding the expected spectral norm of a random principal-submatrix of $(X^H X - I)$. The following result is mainly due to Tropp [6] and it was first presented in the following form by Candès and Plan in [20].

Proposition 2 ([6], [20]). Let $\bar{\Pi} = (\pi_1, \dots, \pi_p)$ be a random permutation of $[p]$ and define $\Pi \doteq (\pi_1, \dots, \pi_k)$ for any $k \in [p]$. Then, for $q = 2 \log p$, we have

$$\begin{aligned} &\left(\mathbb{E} \left[\|X_{\Pi}^H X_{\Pi} - I\|_2^q \right]\right)^{1/q} \\ &\leq 2^{1/q} \left(30\mu \log p + 13 \sqrt{\frac{2k \|X\|_2^2 \log p}{p}} \right) \end{aligned} \quad (38)$$

provided that $k \leq p/4 \|X\|_2^2$. Here, the expectation is with respect to the random permutation $\bar{\Pi}$.

Using this result, it is easy to obtain a probabilistic bound (with respect to $\bar{\Pi}$) on the minimum and maximum singular values of a random submatrix of X since, by the Markov inequality, we have that $\Pr\left(\|X_{\Pi}^H X_{\Pi} - I\|_2 \geq \varsigma\right) \leq \varsigma^{-q} \mathbb{E} \left[\|X_{\Pi}^H X_{\Pi} - I\|_2^q \right]$. The following result is simply a generalization of the corresponding result presented in [20].

Proposition 3 (Extreme singular values of a random submatrix). Let $\Pi = (\pi_1, \dots, \pi_k)$ denote the first k elements of a random permutation of $[p]$ and suppose that $\mu(X) \leq (c'_1 \log p)^{-1}$ and $k \leq p/(c_4^2 \|X\|_2^2 \log p)$ for numerical constants $c'_1 \doteq 60e$ and $c_4 \doteq 37e$. Then, we have that

$$\Pr\left(\|X_{\Pi}^H X_{\Pi} - I\|_2 \geq e^{-1/2}\right) \leq 2p^{-1}. \quad (39)$$

Note that Proposition 3 guarantees that, under certain conditions on $\mu(X)$ and k , every singular value of *most* $n \times k$ submatrices of X lies within $(\sqrt{1 - e^{-1/2}}, \sqrt{1 + e^{-1/2}})$. We are now ready

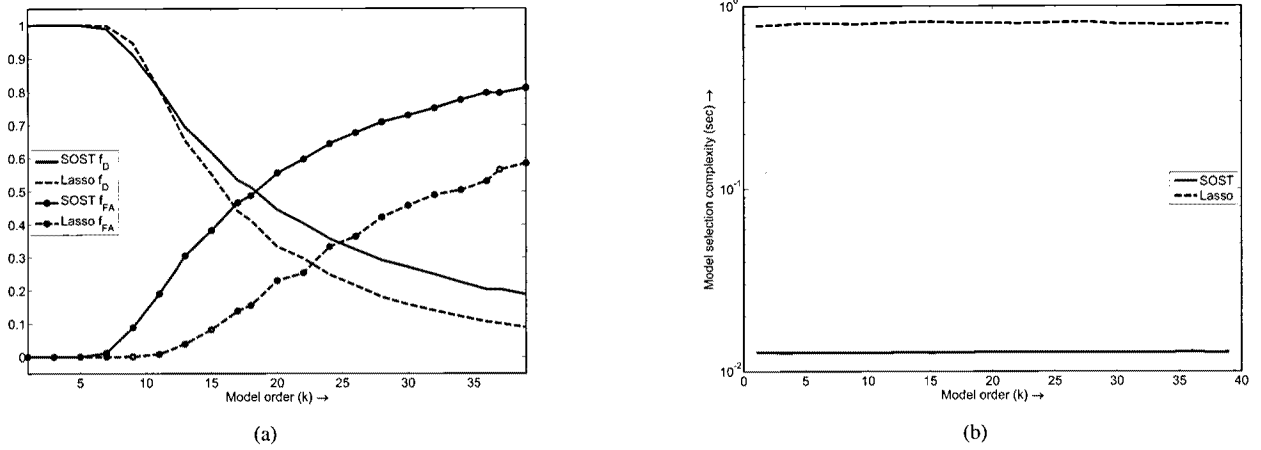


Fig. 2. Numerical comparisons between the performance of the SOST algorithm (Algorithm 3) and the lasso [17] using an Alltop Gabor frame. The $n \times p$ design matrix X has dimensions $n = 127$ and $p = n^2$, the MAR of the signals is 1, the SNR in the measurement system is 10 dB, and the noise variance is $\sigma^2 = 10^{-2}$. The (forward and adjoint) matrix–vector multiplications are carried out using the fast Fourier transform, while the lasso is solved using the SpaRSA package [53] with the regularization parameter set to $\tau = 2\sqrt{2\sigma^2 \log p}$ [20]. (a) Plots of the fraction of detections, defined as $f_D = |\mathcal{S} \cap \hat{\mathcal{S}}|/k$, and the fraction of false alarms, defined as $f_{FA} = (|\hat{\mathcal{S}}| - |\mathcal{S} \cap \hat{\mathcal{S}}|)/|\hat{\mathcal{S}}|$, versus the model order (averaged over 200 independent trials) for both SOST and the lasso. (b) Plots of the amount of time (averaged over 200 independent trials) that it takes SOST and the lasso to solve one model-selection problem versus the model order.

to provide a proof of Theorem 6 that relies on this key result as well as on Lemma 3 and Lemma 4.

Proof (Theorem 6): The proof of this theorem follows along somewhat similar lines as the proof of Theorem 1. Specifically, by making use of the notation developed at the start of this section, we write $f = X^H y \equiv X^H X_{\Pi} z$ and first argue that the set of indices $\mathcal{I} = \{i \in [p] : |f_i| > \lambda\}$ is the same as the true model \mathcal{S} with high probability. Then, we make use of the union bound and argue using Proposition 3 that $\hat{\beta} = \beta$ with high probability.

In this regard, recall that it was established in the proof of Theorem 1 using Lemma 3 and Lemma 4 that if X obeys the coherence property then it satisfies (k, ϵ, δ) -StOC with $\epsilon = 10\mu\sqrt{2\log p}$ and $\delta \leq 4p^{-1}$ as long as $k \leq n/(2\log p)$. This fact therefore implies that, under the assumptions of the theorem,⁷ the following inequalities hold with probability exceeding $1 - 4p^{-1}$:

$$\|f_{\Pi^c}\|_{\infty} = \|X_{\Pi^c}^H X_{\Pi} z\|_{\infty} \leq \epsilon \|z\|_2 \quad (40)$$

and

$$\begin{aligned} \min_{i \in \{1, \dots, k\}} |f_{\pi_i}| &\geq \beta_{\min} - \|(X_{\Pi}^H X_{\Pi} - I)z\|_{\infty} \\ &\geq \beta_{\min} - \epsilon \|z\|_2. \end{aligned} \quad (41)$$

Also note that, conditioned on $\mathcal{E} \doteq \{\|X_{\Pi}^H X_{\Pi} - I\|_2 < e^{-1/2}\}$, we can write

$$\sqrt{1 - e^{-1/2}} \|z\|_2 < \underbrace{\|X_{\Pi} z\|_2}_{\equiv y} < \sqrt{1 + e^{-1/2}} \|z\|_2. \quad (42)$$

Therefore, if we condition on the event \mathcal{E} then it trivially follows from the assumptions of the theorem and (40) and (41)

⁷Note that the assumptions of the theorem trivially guarantee the condition $k \leq n/(2\log p)$ since $\|X\|_2^2 \geq p/n$ from elementary linear algebra.

that $\mathcal{I} = \mathcal{S}$ with probability exceeding $1 - 4p^{-1}$ since (i) $\mathcal{I} \subset \mathcal{S}$ because $\|f_{\Pi^c}\|_{\infty} < \frac{\epsilon \|y\|_2}{\sqrt{1 - e^{-1/2}}} \equiv \lambda$ (cf. (40), (42)), and (ii) $\mathcal{I} \supset \mathcal{S}$ because $k \leq \mu^{-2} \text{MAR} / (c_4^2 \log p)$ implies that $\beta_{\min} - \epsilon \|z\|_2 > \lambda \Rightarrow \min_{i \in \{1, \dots, k\}} |f_{\pi_i}| > \lambda$ (cf. (41), (42)). Consequently, we conclude that $(X_{\mathcal{I}})^{\dagger} = (X_{\Pi}^H X_{\Pi})^{-1} X_{\Pi}^H$ with high probability when conditioned on the probability event \mathcal{E} , which in turn implies that $\hat{\beta}_{\mathcal{I}} = (X_{\mathcal{I}})^{\dagger} X_{\Pi} z \equiv \beta_{\mathcal{S}}$ with probability exceeding $1 - 4p^{-1}$ when conditioned on \mathcal{E} . The claim of the theorem now follows trivially from the union bound and the fact that $\Pr(\mathcal{E}^c) \leq 2p^{-1}$ from Proposition 3 since X satisfies the strong coherence property and $k \leq p/(c_3^2 \|X\|_2^2 \log p)$. \square

VI. CONCLUSIONS

In the modern statistics and signal processing literature, the lasso has arguably become the standard tool for model selection because of its computational tractability [17] and some recent theoretical guarantees [2], [18]–[20]. Nevertheless, it is desirable to study alternative solutions to the lasso since (i) it is still computationally expensive for massively large-scale inference problems (think of p in the millions), (ii) it lacks theoretical guarantees beyond $k \lesssim \mu^{-1}$ for the case of generic design matrices and arbitrary nonzero entries, and (iii) it requires the submatrices of the design matrix to have full rank, which seems reasonable for signal reconstruction but appears too restrictive for model selection.

In this paper, we have revisited two variants of the oft-forgotten but extremely fast OST algorithm for model selection. One of the key insights offered by the paper in this regard is that polynomial-time model selection can be carried out even when signal reconstruction (and thereby the lasso) fails. In addition, we have established in the paper that if the $n \times p$ design matrix X satisfies $\mu(X) \asymp n^{-1/2}$ and $\nu(X) \lesssim n^{-1}$ then OST can perform near-optimally for the case when either (i) the MAR of the signal is not too small or (ii) the SNR in the measurement




Fig. 3. Partial model-selection performance of the OST algorithm (averaged over 200 independent trials) corresponding to an Alltop Gabor frame in \mathbb{C}^{997} . The MAR of the signals in this experiment is 1, the SNR in the measurement system is 3 dB, and the noise variance is $\sigma^2 = 10^{-2}$.

system is not too high. It is worth pointing out here that some researchers in the past have observed that the SOST algorithm at times performs similar to or better than the lasso (see Fig. 2 for an illustration of this in the case of an Alltop Gabor frame in \mathbb{C}^{127}). One of our main contributions in this regard is that we have taken the mystery out of this observation and explicitly specified in the paper the four key parameters of the model-selection problem, namely, $\mu(X)$, $\nu(X)$, MAR, and SNR, that determine the non-asymptotic performance of the SOST algorithm for generic (random or deterministic) design matrices and data vectors having generic (random or deterministic) nonzero entries; also, see [10] for a comparison of our results with corresponding results recently reported in the literature.

The second main contribution of this paper—which completely sets it apart from existing work on thresholding for model selection—is that we have proposed and analyzed a model-order agnostic threshold for the OST algorithm. The significance of this aspect of the paper can be best understood by realizing that in real-world applications it is often easier to estimate the SNR and the noise variance in the system than to estimate the true model order. In particular, we have established in the paper that the threshold $\lambda = \max \left\{ \frac{1}{t} 10\mu\sqrt{n\text{SNR}}, \frac{1}{1-t}\sqrt{2} \right\} \sqrt{2\sigma^2 \log p}$ for $t \in (0, 1)$ enables the OST algorithm to carry out near-optimal partial model selection. It is worth pointing out here that this threshold is rather conservative in nature for small-scale problems (see (5)) and we believe that there is still a lot of room for improvement as far as reducing (or eliminating) some of the constants in the threshold is concerned. In particular, it is easy to see from the proof of Theorem 1 that the constant 10 in the threshold is mainly there due to a number of loose upperbounds; in fact, this constant was 24 in a conference version of this paper [5] and we believe that it can be reduced even further. Some of the numerical experiments that we have carried out in this regard also seem to lend credence to our belief. Specifically, Fig. 3 reports the results of one such experiment concerning partial model-selection performance of the OST algorithm in terms of the metrics of *fraction of detections*, $f_D \stackrel{\text{def}}{=} \frac{|\mathcal{S} \cap \hat{\mathcal{S}}|}{k}$, and

fraction of false alarms, $f_{FA} \stackrel{\text{def}}{=} \frac{|\hat{\mathcal{S}}| - |\mathcal{S} \cap \hat{\mathcal{S}}|}{|\hat{\mathcal{S}}|}$, averaged over 200 independent trials. In this experiment, the $n \times p$ design matrix X corresponds to an Alltop Gabor frame in \mathbb{C}^{997} , the noise variance is $\sigma^2 = 10^{-2}$, the MAR and the SNR are chosen to be 1 and 3 dB, respectively, and the initial threshold is set at $\lambda_s \stackrel{\text{def}}{=} \max \left\{ \frac{1}{t} c' \mu \sqrt{n\text{SNR}}, \frac{1}{1-t} \sqrt{2} \right\} \sqrt{2\sigma^2 \log p}$ with $t = (\sqrt{2} - 1)/\sqrt{2}$ and $c' = 2t$. It can be easily seen from Fig. 3 that OST successfully carries out partial model selection ($f_{FA} \equiv 0$) even when the threshold is set at $0.6\lambda_s$, which proves the somewhat conservative nature of the proposed threshold in terms of the constants.

Finally, the third main contribution of this paper is that we have extended our results on model selection using OST to low-complexity recovery of sparse signals. In particular, within the area of low-complexity algorithms for sparse-signal recovery (such as, matching pursuit [28], subspace pursuit [29], CoSaMP [30], and iterative hard thresholding [31]), we have for the first time specified polynomial-time verifiable sufficient conditions under which recovery of sparse signals having generic (random or deterministic) nonzero entries succeeds using generic (random or deterministic) design matrices. In addition, we have also provided a bound in the paper on the average coherence of generic Gabor frames and used this result to establish that an Alltop Gabor frame in \mathbb{C}^n can be used together with the OST algorithm to successfully carry out model selection and recovery of sparse signals irrespective of the phases of the nonzero entries even if the number of nonzero entries scales almost linearly with n .

APPENDIX

A. Concentration Inequalities

In this appendix, we collect the various concentration inequalities that are used throughout the paper.

Proposition 4 (The Azuma inequality [54]). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let (M_0, M_1, \dots, M_n) be a bounded difference, (real-valued) martingale sequence on $(\Omega, \mathcal{F}, \mathbb{P})$. That

is, $\mathbb{E}[M_i] = M_{i-1}$ and $|M_i - M_{i-1}| \leq b_i$ for every $i = 1, \dots, n$. Then, for every $\epsilon \geq 0$, we have

$$\Pr(|M_n - M_0| \geq \epsilon) \leq 2 \exp\left(-\frac{\epsilon^2}{2 \sum_{i=1}^n b_i^2}\right). \quad (43)$$

Proposition 5 (Inner product of independent Gaussian random vectors [48]). Let $x, y \in \mathbb{R}^n$ be two random vectors that are independently drawn from $\mathcal{N}(0, \sigma^2 I)$ distribution. Then, for every $\epsilon \geq 0$, we have

$$\Pr(|\langle x, y \rangle| \geq \epsilon) \leq 2 \exp\left(-\frac{\epsilon^2}{4\sigma^2(n\sigma^2 + \epsilon/2)}\right). \quad (44)$$

Since we are mainly concerned with complex-valued random variables in this paper, it is helpful to state a complex version of the Azuma inequality. The following lemma is an easy consequence of Proposition 4.

Lemma 5 (The complex Azuma inequality). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let (M_0, M_1, \dots, M_n) be a bounded difference, complex-valued martingale sequence on $(\Omega, \mathcal{F}, \mathbb{P})$. That is, $\mathbb{E}[M_i] = M_{i-1} \in \mathbb{C}$ and further $|M_i - M_{i-1}| \leq b_i$ for every $i = 1, \dots, n$. Then, for every $\epsilon \geq 0$, we have

$$\Pr(|M_n - M_0| \geq \epsilon) \leq 4 \exp\left(-\frac{\epsilon^2}{4 \sum_{i=1}^n b_i^2}\right). \quad (45)$$

Proof: To establish this lemma, first define $S_i \doteq \text{Re}(M_i)$ and $T_i \doteq \text{Im}(M_i)$. Further, notice that since $\mathbb{E}[M_i] = M_{i-1}$ and $|M_i - M_{i-1}| \leq b_i$, we equivalently have that: (i) $\mathbb{E}[S_i] = S_{i-1}$ and $|S_i - S_{i-1}| \leq b_i$, and (ii) $\mathbb{E}[T_i] = T_{i-1}$ and $|T_i - T_{i-1}| \leq b_i$. Therefore, we have that (S_0, S_1, \dots, S_n) and (T_0, T_1, \dots, T_n) are bounded difference, real-valued martingale sequences on $(\Omega, \mathcal{F}, \mathbb{P})$ and hence

$$\begin{aligned} \Pr(|M_n - M_0| \geq \epsilon) &\stackrel{(a)}{\leq} \Pr\left(|S_n - S_0| \geq \frac{\epsilon}{\sqrt{2}}\right) + \Pr\left(|T_n - T_0| \geq \frac{\epsilon}{\sqrt{2}}\right) \\ &\stackrel{(b)}{\leq} 4 \exp\left(-\frac{\epsilon^2}{4 \sum_{i=1}^n b_i^2}\right) \end{aligned} \quad (46)$$

where (a) follows from a simple union bounding argument and (b) follows from the Azuma inequality. \square

Lemma 6 (ℓ_∞ -Norm of the projection of a complex Gaussian vector). Let X be a real- or complex-valued $n \times p$ matrix having unit ℓ_2 -norm columns and let η be a $p \times 1$ vector having entries independently distributed as $\mathcal{CN}(0, \sigma^2)$. Then, for any $\epsilon > 0$, we have

$$\Pr(\|X^H \eta\|_\infty \geq \sigma \epsilon) < \frac{4p}{\sqrt{2\pi}} \frac{\exp(-\epsilon^2/2)}{\epsilon}. \quad (47)$$

Proof: Assume without loss of generality that $\sigma = 1$, since the general case follows from a simple rescaling argument. Let $x_1, \dots, x_p \in \mathbb{C}^n$ be the p columns of X and define

$$z_i \doteq x_i^H \eta, \quad i = 1, \dots, p. \quad (48)$$

Note that the z_i 's are identically (but not independently) distributed as $z_i \sim \mathcal{CN}(0, 1)$, which follows from the fact that $\eta_i \stackrel{i.i.d.}{\sim} \mathcal{CN}(0, 1)$ and the columns of X have unit ℓ_2 -norms. The rest of the proof is pretty elementary and follows from the facts that

$$\begin{aligned} \Pr(\|X^H \eta\|_\infty \geq \epsilon) &\stackrel{(a)}{\leq} p \Pr(|\text{Re}(z_1)|^2 + |\text{Im}(z_1)|^2 \geq \epsilon^2) \\ &\stackrel{(b)}{\leq} 2p \Pr(|\text{Re}(z_1)| \geq \frac{\epsilon}{\sqrt{2}}) \\ &\stackrel{(c)}{<} \frac{4p}{\sqrt{2\pi}} \frac{\exp(-\epsilon^2/2)}{\epsilon}. \end{aligned} \quad (49)$$

Here, (a) follows by taking a union bound over the event $\bigcup_i \{|z_i| \geq \epsilon\}$, (b) follows from taking a union bound over the event $\{|\text{Re}(z_1)| \geq \epsilon/\sqrt{2}\} \cup \{|\text{Im}(z_1)| \geq \epsilon/\sqrt{2}\}$ and noting that the real and imaginary parts of z_i 's are identically distributed as $\mathcal{N}(0, \frac{1}{2})$, and (c) mainly follows by upper bounding the complementary cumulative distribution function as $Q(\epsilon) < \frac{1}{\sqrt{2\pi}\epsilon} \exp(-\frac{1}{2}\epsilon^2)$ [55]. \square


REFERENCES

- [1] A. Miller, *Subset Selection in Regression*. Chapman and Hall, 1990.
- [2] N. Meinshausen and P. Bühlmann, "High-dimensional graphs and variable selection with the lasso," *Ann. Statist.*, vol. 34, no. 3, pp. 1436–1462, June 2006.
- [3] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Scientific Comput.*, vol. 20, no. 1, pp. 33–61, Jan. 1998.
- [4] G. Davis, S. Mallat, and M. Avellaneda, "Adaptive greedy approximations," in *Constructive Approximation*. New York, NY: Springer, Mar. 1997, vol. 13, no. 1, pp. 57–98.
- [5] W. U. Bajwa, R. Calderbank, and S. Jafarpour, "Model selection: Two fundamental measures of coherence and their algorithmic significance," in *Proc. ISIT*, Austin, TX, June 2010, pp. 1568–1572.
- [6] J. A. Tropp, "Norms of random submatrices and sparse approximation," in *C. R. Acad. Sci., Ser. I*, Paris, 2008, vol. 346, pp. 1271–1274.
- [7] W. Alltop, "Complex sequences with low periodic correlations," *IEEE Trans. Inf. Theory*, vol. 26, no. 3, pp. 350–354, May 1980.
- [8] T. Strohmer and R. Heath, "Grassmanian frames with applications to coding and communication," *Appl. Comput. Harmon. Anal.*, vol. 14, no. 3, pp. 257–275, May 2003.
- [9] D. L. Donoho, "For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution," *Commun. Pure Appl. Math.*, vol. 59, no. 6, pp. 797–829, June 2006.
- [10] C. Genovesi, J. Jin, and L. Wasserman. Revisiting marginal regression. submitted. [Online]. Available: arXiv:0911.4080v1
- [11] C. L. Mallows, "Some comments on C_p ," *Technometrics*, vol. 15, no. 4, pp. 661–675, Nov. 1973.
- [12] H. Akaike, "A new look at the statistical model identification," *IEEE Trans. Autom. Control*, vol. 19, no. 6, pp. 716–723, Dec. 1974.
- [13] P. Massart, "A non-asymptotic theory for model selection," in *European Congress of Mathematics*, A. Laptev, Ed., Stockholm, Sweden, 2005, pp. 309–323.
- [14] G. Schwarz, "Estimating the dimension of a model," *Ann. Statist.*, vol. 6, no. 2, pp. 461–464, 1978.
- [15] D. P. Foster and E. I. George, "The risk inflation criterion for multiple regression," *Ann. Statist.*, vol. 22, no. 4, pp. 1947–1975, 1994.
- [16] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, no. 2, pp. 227–234, Apr. 1995.
- [17] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Roy. Statist. Soc. Ser. B*, vol. 58, no. 1, pp. 267–288, 1996.
- [18] P. Zhao and B. Yu, "On model selection consistency of lasso," *J. Mach. Learning Res.*, vol. 7, pp. 2541–2563, 2006.
- [19] M. J. Wainwright, "Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso)," *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2183–2202, May 2009.
- [20] E. J. Candès and Y. Plan, "Near-ideal model selection by ℓ_1 minimization," *Ann. Statist.*, vol. 37, no. 5A, pp. 2145–2177, Oct. 2009.

- [21] V. Saligrama and M. Zhao, "Thresholded basis pursuit: An LP algorithm for achieving optimal support recovery for sparse and approximately sparse signals from noisy random measurements," submitted to *IEEE Trans. Inf. Theory*. [Online]. Available: arXiv:0809.4883v3
- [22] K. Schnass and P. Vandergheynst, "Average performance analysis for thresholding," *IEEE Signal Process. Lett.*, vol. 14, no. 11, pp. 828–831, Nov. 2007.
- [23] A. K. Fletcher, S. Rangan, and V. K. Goyal, "Necessary and sufficient conditions for sparsity pattern recovery," *IEEE Trans. Inf. Theory*, vol. 55, no. 12, pp. 5758–5772, Dec. 2009.
- [24] G. Reeves and M. Gastpar, "A note on optimal support recovery in compressed sensing," in *Proc. 43rd Asilomar Conf. Signals, Syst. Comput.*, Pacific Grove, CA, Nov. 2009.
- [25] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, 1994.
- [26] "Special issue on compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, Mar. 2008.
- [27] E. J. Candès and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n ," *Ann. Statist.*, vol. 35, no. 6, pp. 2313–2351, Dec. 2007.
- [28] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [29] W. Dai and O. Milenkovic, "Subspace pursuit for compressive sensing signal reconstruction," *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2230–2249, May 2009.
- [30] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comput. Harmon. Anal.*, vol. 26, no. 3, pp. 301–321, 2009.
- [31] T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *Appl. Comput. Harmon. Anal.*, vol. 27, no. 3, pp. 265–274, 2009.
- [32] A. C. Gilbert, M. J. Strauss, J. A. Tropp, and R. Vershynin, "One sketch for all: Fast algorithms for compressed sensing," in *Proc. STOC*, San Diego, CA, June 2007, pp. 237–246.
- [33] A. C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. Strauss, "Near-optimal sparse Fourier representations via sampling," in *Proc. STOC*, Montreal, Canada, May 2002, pp. 152–161.
- [34] M. A. Iwen, "A deterministic sub-linear time sparse Fourier algorithm via non-adaptive compressed sensing methods," in *Proc. SODA*, San Francisco, CA, Jan. 2008, pp. 20–29.
- [35] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," in *C. R. Acad. Sci., Ser. I*, Paris, 2008, vol. 346, pp. 589–592.
- [36] L. Welch, "Lower bounds on the maximum cross correlation of signals," *IEEE Trans. Inf. Theory*, vol. 20, no. 3, pp. 397–399, May 1974.
- [37] G. E. Pfander, H. Rauhut, and J. Tanner, "Identification of matrices having a sparse representation," *IEEE Trans. Signal Process.*, vol. 56, no. 11, pp. 5376–5388, Nov. 2008.
- [38] M. A. Herman and T. Strohmer, "High-resolution radar via compressed sensing," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2275–2284, June 2009.
- [39] J. A. Tropp, "On the conditioning of random subdictionaries," *Appl. Comput. Harmon. Anal.*, vol. 25, pp. 1–24, 2008.
- [40] M. J. Wainwright, "Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting," *IEEE Trans. Inf. Theory*, vol. 55, no. 12, pp. 5728–5741, Dec. 2009.
- [41] M. Akcakaya and V. Tarokh, "Shannon-theoretic limits on noisy compressive sampling," *IEEE Trans. Inf. Theory*, vol. 56, no. 1, pp. 492–504, Jan. 2010.
- [42] D. V. Sarwate, "Meeting the Welch bound with equality," in *Proc. SETA*, 1998, pp. 79–102.
- [43] M. Rudelson and R. Vershynin, "Non-asymptotic theory of random matrices: Extreme singular values," to appear in *Proc. Int. Congr. of Mathematicians*, Hyderabad, India, Aug. 2010.
- [44] J. Lawrence, G. E. Pfander, and D. Walnut, "Linear independence of Gabor systems in finite dimensional vector spaces," *J. Fourier Anal. Appl.*, vol. 11, no. 6, pp. 715–726, Dec. 2005.
- [45] S. A. Geršgorin, "Über die abgrenzung der eigenwerte einer matrix," *Izv. Akad. Nauk SSSR Ser. Fiz.-Mat.*, vol. 6, pp. 749–754, 1931.
- [46] R. A. Devore, "Nonlinear approximation," in *Acta Numerica*, A. Iserles, Ed. Cambridge, U.K.: Cambridge University Press, 1998, vol. 7, pp. 51–150.
- [47] W. U. Bajwa, J. Haupt, G. Raz, and R. Nowak, "Compressed channel sensing," in *Proc. CISS*, Princeton, NJ, Mar. 2008, pp. 5–10.
- [48] J. Haupt, W. U. Bajwa, G. Raz, and R. Nowak, "Toeplitz compressed sensing matrices with applications to sparse channel estimation," to appear in *IEEE Trans. Inf. Theory*, 2010.
- [49] O. Christensen, *Frames and Bases*. Boston, MA: Birkhäuser, 2008.
- [50] W. U. Bajwa, A. M. Sayeed, and R. Nowak, "Learning sparse doubly-selective channels," in *Proc. 45th Annu. Allerton Conf. Commun., Control, and Computing*, Monticello, IL, Sept. 2008, pp. 575–582.
- [51] C. McDiarmid, "On the method of bounded differences," in *Surveys in Combinatorics*, J. Siemons, Ed. Cambridge University Press, 1989, pp. 148–188.
- [52] R. Motwani and P. Raghavan, *Randomized Algorithms*. New York, NY: Cambridge University Press, 1995.
- [53] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. Signal Process.*, vol. 57, no. 7, pp. 2479–2493, Jul. 2009.
- [54] K. Azuma, "Weighted sums of certain dependent random variables," *Tohoku Math. J.*, vol. 19, no. 3, pp. 357–367, 1967.
- [55] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*. Upper Saddle River, NJ: Prentice Hall, 1998.



Waheed U. Bajwa received the B.E. (with Honors) degree in electrical engineering from the National University of Sciences and Technology, Islamabad, Pakistan in 2001, and the M.S. and Ph.D. degrees in electrical engineering from the University of Wisconsin-Madison, Madison, WI in 2005 and 2009, respectively. He is currently a Postdoctoral Research Associate in the Program in Applied and Computational Mathematics at Princeton University. He was affiliated with Communications Enabling Technologies, Islamabad, Pakistan—the research arm of Avaz Networks Inc., Irvine, CA (now Quartics LLC)—from 2000–2003, with the Center for Advanced Research in Engineering, Islamabad, Pakistan during 2003, and with the RF and Photonics Lab of GE Global Research, Niskayuna, NY during the summer of 2006. His research interests include high-dimensional inference, statistical signal processing, wireless communications, learning theory, and applications in biological sciences, cyber-physical systems, radar and image processing, and cognitive/ad-hoc networks. He received the Best in Academics Gold Medal and President's Gold Medal in Electrical Engineering from the National University of Sciences and Technology (NUST) in 2001, and the Morgridge Distinguished Graduate Fellowship from the University of Wisconsin-Madison in 2003. He was Junior NUST Student of the Year (2000), Wisconsin Union Poker Series Champion (Spring 2008), and President of the University of Wisconsin-Madison chapter of Golden Key International Honor Society (2009). He is currently a Member of the IEEE, Pakistan Engineering Council, and Golden Key International Honor Society.



Robert Calderbank received the B.Sc. degree in 1975 from Warwick University, England, the M.Sc. degree in 1976 from Oxford University, England, and the Ph.D. degree in 1980 from the California Institute of Technology, all in mathematics. He is a Professor of Electrical Engineering and Mathematics at Princeton University where he directs the Program in Applied and Computational Mathematics. He joined Bell Telephone Laboratories as a Member of Technical Staff in 1980 and retired from AT&T in 2003 as Vice President for Research. He has made significant contributions to a wide range of research areas, from algebraic coding theory and quantum computing to wireless communication and active sensing. He served as Editor in Chief of the IEEE TRANSACTIONS ON INFORMATION THEORY from 1995 to 1998 and as Associate Editor for Coding Techniques from 1986 to 1989. He was a Member of the Board of Governors of the IEEE Information Theory Society from 1991 to 1996 and began a second term in 2006. He was honored by the IEEE Information Theory Prize Paper Award in 1995 for his work on the Z_4 linearity of Kerdock and Preparata Codes (joint with A.R. Hammons Jr., P.V. Kumar, N.J.A. Sloane, and P. Sole), and again in 1999 for the invention of space-time codes (joint with V. Tarokh and N. Seshadri). He received the 2006 IEEE Donald G. Fink Prize Paper Award and the IEEE Millennium Medal, and was elected to the US National Academy of Engineering in 2005.



Sina Jafarpour received his B.Sc. degree in Computer Engineering with great distinction from Sharif University of Technology in July 2007, and his M.A. degree in Computer Science from Princeton University in January 2009. He is currently a Ph.D. candidate in the Department of Computer Science, Princeton University, where he has been working towards the Ph.D. degree as a Member of the Machine Learning and Theoretical Computer Science groups in the Department of Computer Science. His research interests include Compressive Sampling, Machine Learning,

Coding Theory, and Signal Processing. He received the Princeton University Fellowship award in July 2007. In October 2008, his research on Compressive Learning and Detection received one of the five best student paper awards from Google and the New York Academy of Sciences, and in November 2009, his research on Expander-Based Compressed Sensing in the presence of Poisson noise received the best student paper award of the Asilomar Conference on Signals, Systems, and Computers.