

ExSIS: Extended Sure Independence Screening for Ultrahigh-dimensional Linear Models

Talal Ahmed and Waheed U. Bajwa

Abstract

Statistical inference can be computationally prohibitive in ultrahigh-dimensional linear models. Correlation-based variable screening, in which one leverages marginal correlations for removal of irrelevant variables from the model prior to statistical inference, can be used to overcome this challenge. Prior works on correlation-based variable screening either impose statistical priors on the linear model or assume specific post-screening inference methods. This paper first extends the analysis of correlation-based variable screening to arbitrary linear models and post-screening inference techniques. In particular, (i) it shows that a condition—termed the screening condition—is sufficient for successful correlation-based screening of linear models, and (ii) it provides insights into the dependence of marginal correlation-based screening on different problem parameters. Numerical experiments confirm that these insights are not mere artifacts of analysis; rather, they are reflective of the challenges associated with marginal correlation-based variable screening. Second, the paper explicitly derives the screening condition for arbitrary (random or deterministic) linear models and, in the process, it establishes that—under appropriate conditions—it is possible to reduce the dimension of an ultrahigh-dimensional, arbitrary linear model to almost the sample size even when the number of active variables scales almost linearly with the sample size. Third, it specializes the screening condition to sub-Gaussian linear models and contrasts the final results to those existing in the literature. This specialization formally validates the claim that the main result of this paper generalizes existing ones on correlation-based screening.

Keywords: Variable selection; Variable screening; Sure screening; Linear models; High-dimensional statistics; Sparse signal processing

I. INTRODUCTION

The ordinary linear model $y = X\beta + \text{noise}$, despite its apparent simplicity, has been the bedrock of signal processing, statistics, and machine learning for decades. The last decade, however, has witnessed a marked transformation of this model: instead of the classical low-dimensional setting in which the dimension, n , of y (henceforth, referred to as the sample size) exceeds the dimension, p , of β (henceforth, referred to as the number of features/predictors/variables), we are increasingly having to operate in the high-dimensional setting in which the number of variables far exceeds the sample size (i.e., $p \gg n$). While the high-dimensional setting should ordinarily lead to ill-posed problems, the *principle of parsimony*—which states that only a small number of variables typically

This work is supported in part by the National Science Foundation under awards CCF-1525276 and CCF-1453073, and by the Army Research Office under award W911NF-14-1-0295. Some of the results reported here were presented at IEEE Int. Workshop on Comput. Advances in Multi-Sensor Adaptive Process., 2017 [1]. The authors are with the Department of Electrical and Computer Engineering, Rutgers, The State University of New Jersey, 94 Brett Road, Piscataway, NJ 08854, USA. (Emails: talal.ahmed@rutgers.edu and waheed.bajwa@rutgers.edu).

affect the response y —has been employed in the literature to obtain unique solutions for myriad inference problems involving high-dimensional linear models. These inference problems range from sparse regression [2] to sparse representations [3] and sparse signal recovery [4], with guarantees provided in terms of various properties of X such as *spark* [3], *mutual coherence* [3], *null-space property* [5], [6] *restricted isometry property* [7], [8], etc.

Our focus in this paper is on *ultrahigh-dimensional* linear models, in which the number of variables can scale exponentially with the sample size: $\log p = \mathcal{O}(n^\alpha)$ for $\alpha \in (0, 1)$.¹ Such linear models are increasingly becoming common in application areas ranging from genomics [9]–[12] and proteomics [13]–[15] to sentiment analysis [16]–[18] and hyperspectral imaging [19]–[21]. While there exist a number of techniques in the literature—such as forward selection/matching pursuit, backward elimination [22], least absolute shrinkage and selection operator (LASSO) [2], elastic net [23], bridge regression [24], [25], adaptive LASSO [26], group LASSO [27], and Dantzig selector [28]—that can be employed for inference from high-dimensional linear models, all these techniques have super-linear (in the number of variables p) computational complexity. In the ultrahigh-dimensional setting, therefore, use of the aforementioned methods for statistical inference can easily become computationally prohibitive. Variable selection-based dimensionality reduction, commonly referred to as *variable screening*, has been put forth as a practical means of overcoming this *curse of dimensionality* [29]: since only a small number of (independent) variables actually contribute to the response (dependent variable) in the ultrahigh-dimensional setting, one can first—in principle—discard most of the variables (the screening step) and then carry out inference on a relatively low-dimensional linear model using any one of the sparsity-promoting techniques (the inference step). In this work, our focus is on the former step, i.e., the screening step. There are two main challenges that arise in the context of variable screening in ultrahigh-dimensional linear models. First, the screening algorithm should have low computational complexity (ideally, $\mathcal{O}(np)$). Second, the screening algorithm should be accompanied with mathematical guarantees that ensure the reduced linear model contains *all* relevant variables that affect the response. Our goal in this paper is to revisit one of the simplest screening algorithms, which uses marginal correlations between the variables $\{X_i\}_{i=1}^p$ and the response y for screening purposes [30], [31], and provide a comprehensive theoretical understanding of its screening performance for arbitrary (random or deterministic) ultrahigh-dimensional linear models.

A. Relationship to Prior Work

Researchers have long intuited that the (absolute) marginal correlation $|X_i^\top y|$ is a strong indicator of whether the i -th variable contributes to the response variable. Indeed, methods such as stepwise forward regression are based on this very intuition. It is only recently, however, that we have obtained a rigorous understanding of the role of marginal correlations in variable screening. One of the earliest screening works in this regard that is agnostic to the choice of the subsequent inference techniques is termed *sure independence screening* (SIS) [32]. SIS is based on simple thresholding of marginal correlations and satisfies the so-called *sure screening* property—which guarantees that all important variables survive the screening stage with high probability—for the case of normally distributed variables. An iterative variant of SIS, termed ISIS, is also discussed in [32], while [33] presents variants of SIS and

¹Recall Landau's *big- \mathcal{O}* notation: $f(n) = \mathcal{O}(g(n))$ if $\exists C > 0 : f(n) \leq Cg(n)$ and $f(n) = \Omega(g(n))$ if $g(n) = \mathcal{O}(f(n))$.

ISIS that can lead to reduced false selection rates of the screening stage. Extensions of SIS to generalized linear models are discussed in [33], [34], while its generalizations for semi-parametric (Cox) models and non-parametric models are presented in [35], [36] and [37], [38], respectively.

The marginal correlation $|X_i^\top y|$ can be considered an empirical measure of the Pearson correlation coefficient, which is a natural choice for discovering linear relations between the independent variables and the response. In order to perform ultrahigh-dimensional variable screening in the presence of non-linear relations between X_i 's and y and/or heavy-tailed variables, [39] and [40] have put forth screening using generalized (empirical) correlation and Kendall τ rank correlation, respectively.

The defining characteristics of the works referenced above is that they are agnostic to the inference technique that follows the screening stage. In recent years, screening methods have also been proposed for specific optimization-based inference techniques. To this end, [41] formulates a marginal correlations-based screening method, termed SAFE, for the LASSO problem. The goal in SAFE is to discard features that LASSO will not select, and [41] shows that it results in zero false selection rate. In [42], the so-called *strong rules* for variable screening in LASSO-type problems are proposed that are still based on marginal correlations and that result in discarding of far more variables than the SAFE method. The screening tests of [41], [42] for the LASSO problem are further improved in [43]–[45] by analyzing the dual of the LASSO problem. We refer the reader to [46] for an excellent review of these different screening tests for LASSO-type problems.

Notwithstanding these prior works, we have holes in our understanding of variable screening in ultrahigh-dimensional linear models. Works such as [41]–[45] necessitate the use of LASSO-type inference techniques after the screening stage. In addition, these works do not help us understand the relationship between the problem parameters and the dimensions of the reduced model. Stated differently, it is difficult to a priori quantify the computational savings associated with the screening tests proposed in [41]–[45]. Similar to [32], [33], [39], [40], and in contrast to [41]–[45], our focus in this paper is on screening that is agnostic to the post-screening inference technique. To this end, [39] lacks a rigorous theoretical understanding of variable screening using the generalized correlation. While [32], [33], [40] overcome this shortcoming of [39], these works have two major limitations. First, their results are derived under the assumption of statistical priors on the linear model (e.g., normally distributed X_i 's). In many applications, however, it can be a challenge to ascertain the distribution of the independent variables. Second, the analyses in [32], [33], [40] assume the variance of the response variable to be bounded by a constant; this assumption, in turn, imposes the condition $\|\beta\|_2 = \mathcal{O}(1)$. In contrast, defining $\beta_{\min} := \min_i |\beta_i|$, we establish in the sequel that the ratio $\frac{\beta_{\min}}{\|\beta\|_2}$ (and not $\|\beta\|_2$) directly influences the performance of marginal correlation-based screening procedures.

B. Our Contributions

Our focus in this paper is on marginal correlation-based screening of ultrahigh-dimensional linear models that is agnostic to the post-screening inference technique. To this end, we provide an extended analysis of the thresholding-based SIS procedure of [32]. The resulting screening procedure, which we term *extended sure independence screening* (ExSIS), provides new insights into marginal correlation-based screening of arbitrary (random

or deterministic) ultrahigh-dimensional linear models. Specifically, we first provide a simple, distribution-agnostic sufficient condition—termed the *screening condition*—for (marginal correlation-based) screening of linear models. This sufficient condition, which succinctly captures joint interactions among both the active and the inactive variables, is then leveraged to explicitly characterize the performance of ExSIS as a function of various problem parameters, including noise variance, the ratio $\frac{\beta_{\min}}{\|\beta\|_2}$, and model sparsity. The numerical experiments reported at the end of this paper confirm that the dependencies highlighted in this screening result are reflective of the actual challenges associated with marginal correlation-based screening and are not mere artifacts of our analysis.

Next, despite the theoretical usefulness of the screening condition, it cannot be explicitly verified in polynomial time for any given linear model. This is reminiscent of related conditions such as the *incoherence condition* [47], the *irrepresentable condition* [48], the *restricted isometry property* [8], and the *restricted eigenvalue condition* [49] studied in the literature on high-dimensional linear models. In order to overcome this limitation of the screening condition, we explicitly derive it for a family of linear models that corresponds to arbitrary (random or deterministic) linear models in which the (empirical) correlations between independent variables satisfy certain polynomial-time verifiable conditions. The ExSIS results for this family of linear models establish that, under appropriate conditions, it is possible to reduce the dimension of an ultrahigh-dimensional linear model to almost the sample size even when the number of active variables scales almost linearly with the sample size. This, to the best of our knowledge, is the first screening result that provides such explicit and optimistic guarantees *without* imposing a statistical prior on the distribution of the independent variables.

Finally, in order to highlight the generalizability of our results and their explicit relationship to existing results, we also derive the screening condition for a family of sub-Gaussian linear models in which the independent variables are independently drawn from (possibly different) sub-Gaussian distributions. We show that the ExSIS results for this family of linear models match with the SIS results derived in [32] for the case of normally distributed linear models. This validates the generality of the screening condition in relation to existing works.

C. Notation and Organization

The following notation is used throughout this paper. Lower-case letters are used to denote scalars and vectors, while upper-case letters are used to denote matrices. Given $a \in \mathbb{R}$, $\lceil a \rceil$ denotes the smallest integer greater than or equal to a . Given $q \in \mathbb{Z}_+$, we use $\llbracket q \rrbracket$ as a shorthand for $\{1, \dots, q\}$. Given a vector v , $\|v\|_p$ denotes its ℓ_p norm. Given a matrix A , A_j denotes its j -th column and $A_{i,j}$ denotes the entry in its i -th row and j -th column. Further, given a set $\mathcal{I} \subset \mathbb{Z}_+$, $A_{\mathcal{I}}$ (resp., $v_{\mathcal{I}}$) denotes a submatrix (resp., subvector) obtained by retaining columns of A (resp., entries of v) corresponding to the indices in \mathcal{I} . Finally, the superscript $(\cdot)^\top$ denotes the transpose operation.

The rest of this paper is organized as follows. We formulate the problem of marginal correlation-based screening in Sec. II. Next, in Sec. III, we define the screening condition and present one of our main results that establishes the screening condition as a sufficient condition for ExSIS. In Sec. IV, we derive the screening condition for arbitrary linear models based on the correlations between independent variables and discuss implications of the derived ExSIS results. In Sec. V, we derive the screening condition for sub-Gaussian linear models and discuss the

resulting ExSIS guarantees in relation to prior work. Finally, results of extensive numerical experiments on both synthetic and real data are reported in Sec. VI, while concluding remarks are presented in Sec. VII.

II. PROBLEM FORMULATION

Our focus in this paper is on the ultrahigh-dimensional ordinary linear model $y = X\beta + \eta$, where $y \in \mathbb{R}^n$, $X \in \mathbb{R}^{n \times p}$, and $\log p = \mathcal{O}(n^\alpha)$ for $\alpha \in (0, 1)$. In the literature, X is referred to as data/design/observation/sensing matrix with the rows of X corresponding to individual observations and the columns of X corresponding to individual features/predictors/variables, y is referred to as measurement/observation/response vector with individual responses given by $\{y_i\}_{i=1}^n$, β is referred to as the parameter vector, and η is referred to as modeling error or measurement/observation noise. Throughout this paper, we assume X has unit ℓ_2 -norm columns, $\beta \in \mathbb{R}^p$ is k -sparse with $k < n$ (i.e., $|\{i \in [[p]] : \beta_i \neq 0\}| = k < n$), and $\eta \in \mathbb{R}^p$ is a zero-mean Gaussian vector with (entry-wise) variance σ^2 and covariance $C_\eta = \sigma^2 I$. Here, η is taken to be Gaussian with covariance $\sigma^2 I$ for the sake of this exposition, but our analysis is trivially generalizable to other noise distributions and/or covariance matrices. Further, we make no a priori assumption on the distribution of X . Finally, we define $\mathcal{S} := \{i \in [[p]] : \beta_i \neq 0\}$ to be the set that indexes the non-zero components of β . Using this notation, the linear model can equivalently be expressed as

$$y = X\beta + \eta = X_{\mathcal{S}}\beta_{\mathcal{S}} + \eta. \quad (1)$$

Given (1), the goal of variable screening is to reduce the number of variables in the linear model from p (since $p \gg n$) to a moderate scale d (with $d \ll p$) using a fast and efficient method. Our focus here is in particular on screening methods that satisfy the so-called *sure screening* property [32]; specifically, a method is said to carry out sure screening if the d -dimensional model returned by it is guaranteed with high probability to retain all the columns of X that are indexed by \mathcal{S} . The motivation here is that once one obtains a moderate-dimensional model through sure screening of (1), one can use computationally intensive model selection, regression and estimation techniques on the d -dimensional model for reliable model selection (identification of \mathcal{S}), prediction (estimation of $X\beta$), and reconstruction (estimation of β), respectively.

In this paper, we study sure screening using marginal correlations between the response vector and the columns of X . The resulting screening procedure is outlined in Algorithm 1, which is based on the principle that the higher the correlation of a column of X with the response vector, the more likely it is that the said column contributes to the response vector (i.e., it is indexed by the set \mathcal{S}).

Algorithm 1: Marginal Correlation-based Independence Screening

- 1: **Input:** Design matrix $X \in \mathbb{R}^{n \times p}$, response vector $y \in \mathbb{R}^n$, and $d \in \mathbb{Z}_+$
 - 2: **Output:** $\widehat{\mathcal{S}}_d \subset [[p]]$ such that $|\widehat{\mathcal{S}}_d| = d$
 - 3: $w \leftarrow X^\top y$
 - 4: $\widehat{\mathcal{S}}_d \leftarrow \{i \in [[p]] : |w_i| \text{ is among the } d \text{ largest of all marginal correlations}\}$
-

The computational complexity of Algorithm 1 is only $\mathcal{O}(np)$ and its ability to screen ultrahigh-dimensional linear models has been investigated in recent years by a number of researchers [30], [31]. The fundamental difference

among these works stems from the manner in which the parameter d (the dimension of the screened model) is computed from (1). Our goal in this paper is to provide an extended understanding of the screening performance of Algorithm 1 for *arbitrary* (random or deterministic) design matrices. The term *sure independence screening* (SIS) was coined in [32] to refer to screening of ultrahigh-dimensional *sub-Gaussian* linear models using Algorithm 1.² In this vein, we refer to variable screening using Algorithm 1 and the analysis of this paper as *extended sure independence screening* (ExSIS). The main research challenge for ExSIS is specification of d for arbitrary matrices such that $\mathcal{S} \subset \widehat{\mathcal{S}}_d$ with high probability. Note that there is an inherent trade-off in addressing this challenge: the higher the value of d , the more likely is $\widehat{\mathcal{S}}_d$ to satisfy the sure screening property; however, the smaller the value of d , the lower the computational cost of performing model selection, regression, estimation, etc., on the d -dimensional problem. This leads us to the following research questions for ExSIS: (i) What are the conditions on X under which $\mathcal{S} \subset \widehat{\mathcal{S}}_d$? (ii) How small can d be for arbitrary matrices such that $\mathcal{S} \subset \widehat{\mathcal{S}}_d$? (iii) What are the constraints on the sparsity parameter k under which $\mathcal{S} \subset \widehat{\mathcal{S}}_d$? Note that there is also an interplay between the sparsity level k and the allowable value of d for sure screening: the lower the sparsity level, the easier it should be to screen a larger number of columns of X . Thus, an understanding of ExSIS also requires characterization of this relationship between k and d for marginal correlation-based screening. In the sequel, we not only address the aforementioned questions for ExSIS, but also characterize this relationship.

III. SUFFICIENT CONDITIONS FOR SURE SCREENING

In this section, we derive the most general sufficient conditions for ExSIS of ultrahigh-dimensional linear models. The results reported in this section provide important insights into the workings of ExSIS *without* imposing any statistical priors on X and β . We begin with a definition of the *screening condition* for the design matrix X .

Definition 1 ((k, b) -Screening Condition). *Fix an arbitrary $\beta \in \mathbb{R}^p$ that is k -sparse. The (normalized) matrix X satisfies the (k, b) -screening condition if there exists $0 < b(n, p) < \frac{1}{\sqrt{k}}$ such that the following hold:*

$$\max_{i \in \mathcal{S}} \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| \leq b(n, p) \|\beta\|_2, \quad (\text{SC-1})$$

$$\max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| \leq b(n, p) \|\beta\|_2. \quad (\text{SC-2})$$

The screening condition is a statement about the collinearity of the independent variables in the design matrix. The parameter $b(n, p)$ in the screening condition captures the similarity between (i) the columns of $X_{\mathcal{S}}$, and (ii) the columns of $X_{\mathcal{S}}$ and $X_{\mathcal{S}^c}$; the smaller the parameter $b(n, p)$ is, the less similar the columns are. Furthermore, since $k < (b(n, p))^{-2}$ in the screening condition, the parameter $b(n, p)$ reflects constraints on the sparsity parameter k .

We now present one of our main screening results for arbitrary design matrices, which highlights the significance of the screening condition and the role of the parameter $b(n, p)$ within ExSIS.

²The analysis in [32] requires a certain concentration property that has only been shown in [32] to hold for Gaussian matrices. However, this property can also be proved for sub-Gaussian matrices by appealing to matrix concentration inequalities; see, e.g., [50].

Theorem 1 (Sufficient Conditions for ExSIS). *Let $y = X\beta + \eta$ with β a k -sparse vector and the entries of η independently distributed as $\mathcal{N}(0, \sigma^2)$. Define $\beta_{\min} := \min_{i \in \mathcal{S}} |\beta_i|$ and $\tilde{\eta} := X^\top \eta$, and let \mathcal{G}_η be the event $\{\|\tilde{\eta}\|_\infty \leq 2\sqrt{\sigma^2 \log p}\}$. Suppose X satisfies the screening condition and assume $\frac{\beta_{\min}}{\|\beta\|_2} > 2b(n, p) + 4\frac{\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$. Then, conditioned on \mathcal{G}_η , Algorithm 1 satisfies $\mathcal{S} \subset \hat{\mathcal{S}}_d$ as long as $d \geq \left\lceil \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2b(n, p) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil$.*

We refer the reader to Sec. III-B for a proof of this theorem.

A. Discussion

Theorem 1 highlights the dependence of ExSIS on the observation noise, the ratio $\frac{\beta_{\min}}{\|\beta\|_2}$, the parameter $b(n, p)$, and model sparsity. We first comment on the relationship between ExSIS and observation noise η . Notice that the statement of Theorem 1 is dependent upon the event \mathcal{G}_η . However, for any $\epsilon > 0$, we have (see, e.g., [51, Lemma 6])

$$\Pr(\|\tilde{\eta}\|_\infty \geq \sigma\epsilon) < \frac{4p}{\epsilon\sqrt{2\pi}} \exp\left(-\frac{\epsilon^2}{2}\right). \quad (2)$$

Therefore, substituting $\epsilon = 2\sqrt{\log p}$ in (2), we obtain

$$\Pr(\mathcal{G}_\eta) \geq 1 - 2(p\sqrt{2\pi \log p})^{-1}. \quad (3)$$

Thus, Algorithm 1 possesses the sure screening property in the case of the observation noise η distributed as $\mathcal{N}(0, \sigma^2 I)$. We further note from the statement of Theorem 1 that the higher the *signal-to-noise ratio* (SNR), defined here as $\text{SNR} := \frac{\|\beta\|_2}{\sigma}$, the more Algorithm 1 can screen irrelevant/inactive variables. It is also worth noting here trivial generalizations of Theorem 1 for other noise distributions. In the case of η distributed as $\mathcal{N}(0, C_\eta)$, Theorem 1 has σ^2 replaced by the largest eigenvalue of the covariance matrix C_η . In the case of η following a non-Gaussian distribution, Theorem 1 has $2\sqrt{\sigma^2 \log p}$ replaced by distribution-specific upper bound on $\|X^\top \eta\|_\infty$ that holds with high probability.

In addition to the noise distribution, the performance of ExSIS also seems to be impacted by the *minimum-to-signal ratio* (MSR), defined here as $\text{MSR} := \frac{\beta_{\min}}{\|\beta\|_2} \in (0, \frac{1}{\sqrt{k}}]$. Specifically, the higher the MSR, the more Algorithm 1 can screen inactive variables. Stated differently, the independent variable with the weakest contribution to the response determines the size of the screened model. Finally, the parameter $b(n, p)$ in the screening condition also plays a central role in characterization of the performance of ExSIS. First, the smaller the parameter $b(n, p)$, the more Algorithm 1 can screen inactive variables. Second, the smaller the parameter $b(n, p)$, the more independent variables can be active in the original model; indeed, we have from the screening condition that $k < (b(n, p))^{-2}$. Third, the smaller the parameter $b(n, p)$, the lower the smallest allowable value of MSR; indeed, we have from the theorem statement that $\text{MSR} > 2b(n, p) + 4\frac{\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$.

It is evident from the preceding discussion that the screening condition (equivalently, the parameter $b(n, p)$) is one of the most important factors that helps understand the workings of ExSIS and helps quantify its performance. Unfortunately, the usefulness of this knowledge is limited in the sense that the screening condition cannot be utilized in practice. Specifically, the screening condition is defined in terms of the set \mathcal{S} , which is of course unknown. We overcome this limitation of Theorem 1 by implicitly deriving the screening condition for a class of arbitrary (random or deterministic) design matrices in Sec. IV and sub-Gaussian design matrices in Sec. V.

B. Proof of Theorem 1

We first provide an outline of the proof of Theorem 1, which is followed by its formal proof. Define $p_0 := p$, $\widehat{\mathcal{S}}_{p_0} := \llbracket [p] \rrbracket$, and $t_1 := |\{j \in \widehat{\mathcal{S}}_{p_0} : |w_j| \geq \min_{i \in \mathcal{S}} |w_i|\}|$. Next, fix a positive integer $p_1 < p_0$ and define

$$\widehat{\mathcal{S}}_{p_1} := \{i \in \widehat{\mathcal{S}}_{p_0} : |w_i| \text{ is among the } p_1 \text{ largest of all marginal correlations}\}.$$

The idea is to first derive an initial upper bound on t_1 , denoted by \bar{t}_1 , and then choose $p_1 = \lceil \bar{t}_1 \rceil$; trivially, we have $\mathcal{S} \subset \widehat{\mathcal{S}}_{p_1} \subset \widehat{\mathcal{S}}_{p_0}$. As a result, we get

$$y = X\beta + \eta = X_{\widehat{\mathcal{S}}_{p_0}} \beta_{\widehat{\mathcal{S}}_{p_0}} + \eta = X_{\widehat{\mathcal{S}}_{p_1}} \beta_{\widehat{\mathcal{S}}_{p_1}} + \eta. \quad (4)$$

Note that while deriving \bar{t}_1 , we need to ensure $\bar{t}_1 < p_0$; this in turn imposes some conditions on X that also need to be specified. Next, we can repeat the aforementioned steps to obtain $\widehat{\mathcal{S}}_{p_2}$ from $\widehat{\mathcal{S}}_{p_1}$ for a fixed positive integer $p_2 < p_1 < p_0$. Specifically, define

$$\widehat{\mathcal{S}}_{p_2} := \{i \in \widehat{\mathcal{S}}_{p_1} : |w_i| \text{ is among the } p_2 \text{ largest of all marginal correlations}\}$$

and $t_2 := |\{j \in \widehat{\mathcal{S}}_{p_1} : |w_j| \geq \min_{i \in \widehat{\mathcal{S}}_{p_1}} |w_i|\}|$. We can then derive an upper bound on t_2 , denoted by \bar{t}_2 , and then choose $p_2 = \lceil \bar{t}_2 \rceil$; once again, we have $\mathcal{S} \subset \widehat{\mathcal{S}}_{p_2} \subset \widehat{\mathcal{S}}_{p_1} \subset \widehat{\mathcal{S}}_{p_0}$. Notice further that we do require $\bar{t}_2 < p_1$, which again will impose conditions on X .

In similar vein, we can keep on repeating this procedure to obtain a decreasing sequence of numbers $\{\bar{t}_j\}_{j=1}^i$ and sets $\widehat{\mathcal{S}}_{p_0} \supset \widehat{\mathcal{S}}_{p_1} \supset \widehat{\mathcal{S}}_{p_2} \supset \dots \supset \widehat{\mathcal{S}}_{p_i} \supset \mathcal{S}$ as long as $\bar{t}_i < p_{i-1}$, where $\{p_j := \lceil \bar{t}_j \rceil\}_{j=1}^i$ and $i \in \mathbb{Z}_+$. The complete proof of Theorem 1 follows from a careful combination of these (analytical) steps. In order for us to be able to do that, however, we need two lemmas. The first lemma provides an upper bound on $t_i = |\{j \in \widehat{\mathcal{S}}_{p_{i-1}} : |w_j| \geq \min_{i \in \mathcal{S}} |w_i|\}|$ for $i \in \mathbb{Z}_+$, denoted by \bar{t}_i . The second lemma provides conditions on the design matrix X such that $\bar{t}_i < p_{i-1}$. The proof of the theorem follows from repeated application of the two lemmas.

Lemma 1. Fix $i \in \mathbb{Z}_+$ and suppose $\mathcal{S} \subset \widehat{\mathcal{S}}_{p_{i-1}}$, where $|\widehat{\mathcal{S}}_{p_{i-1}}| =: p_{i-1}$ and $p_{i-1} \leq p$. Further, suppose the design matrix X satisfies the (k, b) -screening condition for the k -sparse vector β and the event \mathcal{G}_η holds true. Finally, define $t_i := |\{j \in \widehat{\mathcal{S}}_{p_{i-1}} : |w_j| \geq \min_{i \in \mathcal{S}} |w_i|\}|$. Under these conditions, we have

$$t_i \leq \frac{p_{i-1} b(n, p) \|\beta\|_2 + \|\beta\|_1 + 2p_{i-1} \sqrt{\sigma^2 \log p}}{\beta_{\min} - b(n, p) \|\beta\|_2 - 2\sqrt{\sigma^2 \log p}} =: \bar{t}_i. \quad (5)$$

The proof of this lemma is provided in Appendix A. The second lemma, whose proof is given in Appendix B, provides conditions on X under which the upper bound derived on t_i for $i \in \mathbb{Z}_+$, denoted by \bar{t}_i , is non-trivial.

Lemma 2. Fix $i \in \mathbb{Z}_+$. Suppose $p_{i-1} > \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2b(n, p) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}}$ and $\frac{\beta_{\min}}{\|\beta\|_2} > 2b(n, p) + \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$. Then, we have $0 < \bar{t}_i < p_{i-1}$.

We are now ready to present a complete technical proof of Theorem 1.

Proof. The idea is to use Lemma 1 and Lemma 2 *repeatedly* to screen columns of X . Note, however, that this is simply an analytical technique and we do not *actually* need to perform such an iterative procedure to specify d in Algorithm 1. To begin, recall that we have $p_0 := p$, $\widehat{\mathcal{S}}_{p_0} := [[p]]$,

$$\widehat{\mathcal{S}}_{p_1} := \{i \in \widehat{\mathcal{S}}_{p_0} : |w_i| \text{ is among the } p_1 \text{ largest of all marginal correlations}\},$$

and $p_1 = \lceil \bar{t}_1 \rceil$, where \bar{t}_1 is defined in (5). By Lemma 1 and Lemma 2, we have $\mathcal{S} \subset \widehat{\mathcal{S}}_{p_1}$ and $p_1 < p_0$, respectively.

Next, given $p_1 > \left\lceil \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2b(n,p) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil$, we can use Lemma 1 and Lemma 2 to obtain $\widehat{\mathcal{S}}_{p_2}$ from $\widehat{\mathcal{S}}_{p_1}$ in a similar fashion. Specifically, let

$$\widehat{\mathcal{S}}_{p_2} := \{i \in \widehat{\mathcal{S}}_{p_1} : |w_i| \text{ is among the } p_2 \text{ largest of all marginal correlations}\}$$

and $p_2 = \lceil \bar{t}_2 \rceil$, where \bar{t}_2 is defined in (5). Then, by Lemma 1 and Lemma 2, we have $\mathcal{S} \subset \widehat{\mathcal{S}}_{p_2}$ and $p_2 < p_1$, respectively.

Notice that we can keep on repeating this procedure to obtain sub-models $\widehat{\mathcal{S}}_{p_1}, \widehat{\mathcal{S}}_{p_2}, \dots, \widehat{\mathcal{S}}_{p_l}$ such that $p_l \leq \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2b(n,p) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}}$ and $p_{l-1} > \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2b(n,p) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}}$. By repeated applications of Lemma 1 and Lemma 2, we have $\mathcal{S} \subset \widehat{\mathcal{S}}_{p_l}$. Further, we are also guaranteed that $p_l \leq \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2b(n,p) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}}$. Thus, we can choose $d \geq \left\lceil \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2b(n,p) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil$ in Algorithm 1 in one shot and have $\mathcal{S} \subset \widehat{\mathcal{S}}_d$. \blacksquare

IV. SCREENING OF ARBITRARY DESIGN MATRICES

In this section, we specialize Theorem 1 for a family of arbitrary linear models in which the design matrices satisfy certain polynomial-time verifiable conditions. We achieve this by deriving verifiable sufficient conditions that guarantee the screening condition for any arbitrary (random or deterministic) design matrix. These sufficient conditions are presented in terms of two measures of similarity among the columns of a design matrix. These measures, termed *worst-case coherence* and *average coherence*, are defined as follows.

Definition 2 (Worst-case and Average Coherences). *Let X be an $n \times p$ matrix with unit ℓ_2 -norm columns. The worst-case coherence of X is denoted by μ and is defined as [52]:*

$$\mu := \max_{i,j:i \neq j} |X_i^\top X_j|.$$

On the other hand, the average coherence of X is denoted by ν and is defined as [51]:

$$\nu := \frac{1}{p-1} \max_i \left| \sum_{j:j \neq i} X_i^\top X_j \right|.$$

Notice that both the worst-case and the average coherences are readily computable in polynomial time. Heuristically, the worst-case coherence is an indirect measure of pairwise similarity among the columns of X : $\mu \in [0, 1]$ with $\mu \searrow 0$ as the columns of X become less similar and $\mu \nearrow 1$ as at least two columns of X become more similar. The average coherence, on the other hand, is an indirect measure of both the collective similarity among the columns of X and the spread of the columns of X within the unit sphere: $\nu \in [0, \mu]$ with $\nu \searrow 0$ as the columns of X become

more spread out in \mathbb{R}^n and $\nu \nearrow \mu$ as the columns of X become less spread out. We refer the reader to [53] for further discussion of these two measures as well as their values for commonly encountered matrices.

We are now ready to describe the main results of this section. The first result connects the screening condition to the worst-case coherence. We will see, however, that this result suffers from the so-called square-root bottleneck: ExSIS analysis based solely on the worst-case coherence can, at best, handle $k = \mathcal{O}(\sqrt{n})$ scaling of the sparsity parameter. The second result overcomes this bottleneck by connecting the screening condition to both worst-case and average coherences. The caveat here is that this result imposes a mild statistical prior on the set \mathcal{S} .

A. ExSIS and the Worst-case Coherence

We begin by relating the worst-case coherence of an arbitrary design matrix X with unit-norm columns to the screening condition.

Lemma 3 (Worst-case Coherence and the Screening Condition). *Let X be an $n \times p$ design matrix with unit-norm columns. Then, we have*

$$\begin{aligned} \max_{i \in \mathcal{S}} \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| &\leq \mu \sqrt{k} \|\beta\|_2, \text{ and} \\ \max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| &\leq \mu \sqrt{k} \|\beta\|_2. \end{aligned}$$

The proof of this lemma is provided in Appendix C. It follows from Lemma 3 that a design matrix satisfies the screening condition with parameter $b(n, p) = \mu \sqrt{k}$ as long as $\mu < k^{-1} \Leftrightarrow k < \mu^{-1}$. We now combine this implication of Lemma 3 with Theorem 1 to provide a result for ExSIS of arbitrary linear models.

Theorem 2. *Let $y = X\beta + \eta$ with β a k -sparse vector and the entries of η independently distributed as $\mathcal{N}(0, \sigma^2)$. Suppose $k < \mu^{-1}$ and $\frac{\beta_{\min}}{\|\beta\|_2} > 2\mu\sqrt{k} + 4\frac{\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$. Then, Algorithm 1 satisfies $\mathcal{S} \subset \widehat{\mathcal{S}}_d$ with probability exceeding $1 - 2(p\sqrt{2\pi \log p})^{-1}$ as long as $d \geq \left\lceil \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2\mu\sqrt{k} - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil$.*

The proof of this theorem follows directly from Lemma 3 and Theorem 1. Next, a straightforward corollary of Theorem 2 shows that ExSIS of arbitrary linear models can in fact be carried out without explicit knowledge of the sparsity parameter k .

Corollary 2.1. *Let $y = X\beta + \eta$ with β a k -sparse vector and the entries of η independently distributed as $\mathcal{N}(0, \sigma^2)$. Suppose $p \geq 2n$, $k < \mu^{-1}$, and $\frac{\beta_{\min}}{\|\beta\|_2} > 2c_1\mu\sqrt{k} + 4c_2\frac{\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$ for some $c_1, c_2 > 2$. Then, Algorithm 1 satisfies $\mathcal{S} \subset \widehat{\mathcal{S}}_d$ with probability exceeding $1 - 2(p\sqrt{2\pi \log p})^{-1}$ as long as $d \geq \lceil \sqrt{n} \rceil$.*

Proof. Under the assumption of $\frac{\beta_{\min}}{\|\beta\|_2} > 2c_1\mu\sqrt{k} + \frac{4c_2\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$, notice that

$$d \geq \left\lceil \frac{\sqrt{k}}{2(c_1 - 1)\mu\sqrt{k} + \frac{4(c_2 - 1)\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil \tag{6}$$

is a sufficient condition for $d \geq \left\lceil \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2\mu\sqrt{k} - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil$. Further, note that $d \geq \lceil (2\mu)^{-1} \rceil$ is a sufficient condition for (6). Next, since $p \geq 2n$, we also have $\mu^{-1} \leq \sqrt{2n}$ from the Welch bound on the worst-case coherence of design matrices [54]. Thus, $d \geq \lceil \sqrt{n} \rceil$ is a sufficient condition for $d \geq \left\lceil \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2\mu\sqrt{k} - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil$. ■

It is worth reflecting here on the fact that Corollary 2.1 allows the number of remaining variables in the screened model to scale as $d \geq \lceil \sqrt{n} \rceil$. While this means that the screened model can have as few as $\lceil \sqrt{n} \rceil$ variables, the result suffers from the so-called square-root bottleneck as it constraints sparsity to scale only as $\mathcal{O}(\sqrt{n})$: $k = \mathcal{O}(\mu^{-1}) \Rightarrow k = \mathcal{O}(\sqrt{n})$ because of the Welch bound. We overcome this sparsity limitation by adding the average coherence into the mix and imposing a uniform statistical prior on the true model \mathcal{S} in the next section.

B. ExSIS and the Coherence Property

In order to break the square-root bottleneck for ExSIS of arbitrary linear models, we first define the notion of the coherence property.

Definition 3 (The Coherence Property). *An $n \times p$ design matrix X with unit-norm columns is said to obey the coherence property if there exists a constant $c_\mu > 0$ such that $\mu < \frac{1}{c_\mu \sqrt{\log p}}$ and $\nu < \frac{\mu}{\sqrt{n}}$.*

Heuristically, the coherence property, which was first introduced in [51], requires the independent variables to be sufficiently (marginally and jointly) uncorrelated. Notice that, unlike many conditions in high-dimensional statistics (see, e.g., [8], [47]–[49]), the coherence property is explicitly certifiable in polynomial time for any given design matrix. We now establish that the coherence property implies the design matrix satisfies the screening condition with high probability, where the probability is with respect to uniform prior on the true model \mathcal{S} .

Lemma 4 (Coherence Property and the Screening Condition). *Let X be an $n \times p$ design matrix that satisfies the coherence property with $c_\mu > 10\sqrt{2}$, and suppose $p \geq \max\{2n, \exp(5)\}$ and $k \leq \frac{n}{\log p}$. Further, assume \mathcal{S} is drawn uniformly at random from k -subsets of $[[p]]$. Then, with probability exceeding $1 - 4p^{-1}$, we have*

$$\begin{aligned} \max_{i \in \mathcal{S}} \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| &\leq c_\mu \mu \sqrt{\log p} \|\beta\|_2, \quad \text{and} \\ \max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| &\leq c_\mu \mu \sqrt{\log p} \|\beta\|_2. \end{aligned}$$

The proof of this lemma is provided in Appendix D. Lemma 4 implies that a design matrix that obeys the coherence property also satisfies the screening condition for *most* models with $b(n, p) = c_\mu \mu \sqrt{\log p}$ as long as $\mu < c_\mu^{-1} (k \log p)^{-1/2} \Leftrightarrow k < \frac{\mu^{-2}}{c_\mu^2 \log p}$. Comparing this with Lemma 3 and the resulting constraint $k < \mu^{-1}$ for the screening condition to hold in the case of arbitrary design matrices, we see that—at the expense of uniform prior on the true model—the coherence property results in a better bound on the screening parameter as long as $\log p = \mathcal{O}(\mu^{-1})$. We can now utilize Lemma 4 along with Theorem 1 to provide an improved result for ExSIS of arbitrary linear models.

Theorem 3. Let $y = X\beta + \eta$ with β a k -sparse vector and the entries of η independently distributed as $\mathcal{N}(0, \sigma^2)$. Further, assume X satisfies the coherence property with $c_\mu > 10\sqrt{2}$ and \mathcal{S} is drawn uniformly at random from k -subsets of $[[p]]$. Finally, suppose $p \geq \max\{2n, \exp(5)\}$, $k < \frac{\mu^{-2}}{c_\mu^2 \log p}$, and $\frac{\beta_{\min}}{\|\beta\|_2} > 2c_\mu \mu \sqrt{\log p} + \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$. Then, Algorithm 1 satisfies $\mathcal{S} \subset \widehat{\mathcal{S}}_d$ with probability exceeding $1 - 6p^{-1}$ as long as $d \geq \left\lceil \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2c_\mu \mu \sqrt{\log p} - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil$.

The proof of this theorem is omitted here since it follows in a straightforward manner from Lemma 4, Theorem 1, and a union bound argument. Nonetheless, it is worth mentioning here that the $k \leq \frac{n}{\log p}$ bound in Lemma 4 is omitted in Theorem 3 since $k < \frac{\mu^{-2}}{c_\mu^2 \log p} \Rightarrow k \leq \frac{n}{\log p}$ because of the Welch bound. The final result of this section is a corollary of Theorem 3 that removes the dependence of d on knowledge of the problem parameters.

Corollary 3.1. Let $y = X\beta + \eta$ with β a k -sparse vector and the entries of η independently distributed as $\mathcal{N}(0, \sigma^2)$. Further, assume X satisfies the coherence property with $c_\mu > 10\sqrt{2}$ and \mathcal{S} is drawn uniformly at random from k -subsets of $[[p]]$. Finally, suppose $p \geq \max\{2n, \exp(5)\}$, $k < \frac{\mu^{-2}}{c_\mu^2 \log p}$, and $\frac{\beta_{\min}}{\|\beta\|_2} > 2c_\mu c_1 \mu \sqrt{\log p} + 4c_2 \frac{\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$ for some $c_1, c_2 > 2$. Then, Algorithm 1 satisfies $\mathcal{S} \subset \widehat{\mathcal{S}}_d$ with probability exceeding $1 - 6p^{-1}$ as long as $d \geq \left\lceil \frac{n}{\log p} \right\rceil$.

Proof. Since $\frac{\beta_{\min}}{\|\beta\|_2} > 2c_1 c_\mu \mu \sqrt{\log p} + \frac{4c_2 \sqrt{\sigma^2 \log p}}{\|\beta\|_2}$, we have that

$$d \geq \left\lceil \frac{\sqrt{k}}{2(c_1 - 1)c_\mu \mu \sqrt{\log p} + \frac{4(c_2 - 1)\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil \quad (7)$$

is a sufficient condition for $d \geq \left\lceil \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2c_\mu \mu \sqrt{\log p} - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil$. The claim now follows because $d \geq \left\lceil \frac{n}{\log p} \right\rceil$ is a sufficient condition for (7), owing to the facts that $p \geq 2n$ and the Welch bound imply $\mu^{-1} \leq \sqrt{2n}$ and $k < \frac{\mu^{-2}}{c_\mu^2 \log p} \Rightarrow k \leq \frac{n}{\log p}$. ■

C. Discussion

Both Theorem 2 and Theorem 3 shed light on the feasibility of marginal correlation-based screening of linear models *without* imposing a statistical prior on the design matrix. While Theorem 2 in this regard provides the least restrictive results, it does suffer from the square-root bottleneck: $k = \mathcal{O}(\mu^{-1}) \Rightarrow k = \mathcal{O}(\sqrt{n})$. Theorem 3, on the other hand, overcomes this bottleneck at the expense of uniform prior on the true model as long as $\log p = \mathcal{O}(\mu^{-1})$; in this case, the condition on the sparsity parameter becomes $k = \mathcal{O}(\mu^{-2}/\log p)$. Therefore, Theorem 3 allows for sparsity scaling as high as $k = \mathcal{O}(n/\log p)$ for design matrices with $\mu = \mathcal{O}(n^{-1/2})$; see [53] for existence of such matrices. In addition, Theorem 2 and Theorem 3 also differ from each other in terms of their respective constraints on $\frac{\beta_{\min}}{\|\beta\|_2}$ for feasibility of marginal correlation-based screening; the constraint in Theorem 3 is less restrictive than in Theorem 2 for $\log p = \mathcal{O}(\mu^{-1})$.

Given that the constraints on sparsity and $\frac{\beta_{\min}}{\|\beta\|_2}$ in Theorem 3 are less restrictive than in Theorem 2, it is instructive to specialize the results in Theorem 3 for screening of some specific family of linear models. For this purpose, we consider the case of Gaussian design matrices and reproduce bounds on their worst-case and average coherences from [53].

Lemma 5 ([53, Theorem 8]). *Let $V = [V_{i,j}]$ be an $n \times p$ matrix with the entries $\{V_{i,j}\}_{i,j=1}^{n,p}$ independently distributed as $\mathcal{N}(0, 1)$ and $60 \log p \leq n \leq \frac{p-1}{4 \log p}$. Suppose the design matrix X is obtained by normalizing the columns of V , i.e., $X = V \text{diag}(1/\|V_1\|_2, \dots, 1/\|V_p\|_2)$. Then, with probability exceeding $1 - 11p^{-1}$, we have*

$$\begin{aligned} \mu &\leq \frac{\sqrt{15 \log p}}{\sqrt{n} - \sqrt{12 \log p}}, \quad \text{and} \\ \nu &\leq \frac{\sqrt{15 \log p}}{n - \sqrt{12n \log p}}. \end{aligned}$$

It can be seen from this lemma that Gaussian design matrices satisfy the coherence property for $\log p = \mathcal{O}(n^{1/2})$. We can therefore use Lemma 5 to specialize Corollary 3.1 to Gaussian matrices and obtain the following result.

Corollary 3.2. *Let $V = [V_{i,j}]$ be an $n \times p$ matrix with the entries $\{V_{i,j}\}_{i,j=1}^{n,p}$ independently distributed as $\mathcal{N}(0, 1)$, the design matrix X be obtained by normalizing the columns of V , and $y = X\beta + \eta$ with β a k -sparse vector and the entries of η independently distributed as $\mathcal{N}(0, \sigma^2)$. Further, assume \mathcal{S} is drawn uniformly at random from k -subsets of $[p]$. Finally, suppose $p \geq \max\{2n, \exp(5)\}$, $n \geq 2$, $\log p < \frac{\sqrt{n}}{c_\mu \sqrt{15}}$, $k < \frac{n}{15c_\mu^2 (\log p)^2}$, and $\frac{\beta_{\min}}{\|\beta\|_2} > 2\sqrt{15}c_\mu c_1 \frac{\log p}{\sqrt{n}} + 4c_2 \frac{\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$ for some $c_1, c_2 > 2, c_\mu > 10\sqrt{2}$. Then, Algorithm 1 satisfies $\mathcal{S} \subset \hat{\mathcal{S}}_d$ with probability exceeding $1 - 17p^{-1}$ as long as $d \geq \left\lceil \frac{n}{\log p} \right\rceil$.*

From Corollary 3.2, we can conclude that screening of Gaussian linear models using Algorithm 1 can be carried out with $d \geq \left\lceil \frac{n}{\log p} \right\rceil$ as long as: (i) $\log p = \mathcal{O}(n^{1/2})$, (ii) $k = \mathcal{O}(n/(\log p)^2)$, and (iii) $\frac{\beta_{\min}}{\|\beta\|_2} = \Omega\left(\frac{\log p}{\sqrt{n}} + \frac{\sqrt{\sigma^2 \log p}}{\|\beta\|_2}\right)$. We now compare this specialization of Theorem 3 for Gaussian matrices with the known results of SIS in the existing literature. To this end, we focus on the results reported in [32], which is one of the most influential SIS works for the family of Gaussian linear models. Specifically, we assume $\sigma = 0$ for the sake of simplicity of argument, and explicitly compare Corollary 3.2 with [32, Theorem 1] for the case of Gaussian design matrices with independent entries. The SIS results reported in [32, Theorem 1] hold under four specific conditions. In particular, Condition 3 in [32] requires that: (i) the variance of the response variable is $\mathcal{O}(1)$, (ii) $\beta_{\min} \geq \frac{c_\kappa}{n^\kappa}$ for some $c_\kappa > 0, \kappa \geq 0$, and (iii) $\min_{i \in \mathcal{S}} |\text{cov}(\beta_i^{-1} Y, X_i)| \geq c_3$ for some $c_3 > 0$. Notice that the $\mathcal{O}(1)$ variance condition is equivalent to having $\|\beta\|_2 = \mathcal{O}(1)$. We therefore also impose the condition $\|\beta\|_2 = \mathcal{O}(1)$ in Corollary 3.2 for comparison purposes.

Comparing the two results, we see that [32] ensures sure screening with slightly less-restrictive constraints. In [32, Theorem 1], SIS requires $\beta_{\min} = \Omega(n^{-\kappa})$ for $\kappa < 1/2$ and $\log p = \mathcal{O}(n^\alpha)$ for some $\alpha \in (0, 1 - 2\kappa)$. In contrast, substituting $\beta_{\min} = \Omega(n^{-\kappa})$ and $\log p = \mathcal{O}(n^\alpha)$ in Corollary 3.2 leads to constraints of $\kappa < 1/2$ and $0 < \alpha < \frac{1}{2} - \kappa$. The impact of this small gap in α for the two results can be more concretely analyzed in terms of the respective sparsity constraints. Specifically, [32, Theorem 1] imposes the sparsity constraint $k = \mathcal{O}(n^{2\kappa})$ for the sure screening result to hold. Invoking, the condition $\log p = \mathcal{O}(n^\alpha)$ with $\alpha \in (0, 1 - 2\kappa)$ reduces this constraint to $k = \mathcal{O}(n^{1-\alpha}) = \mathcal{O}\left(\frac{n}{\log p}\right)$, which is slightly less restrictive than the sparsity constraint of $k = \mathcal{O}(n/(\log p)^2)$ imposed in Corollary 3.2.

The aforementioned slight gaps in constraints lead to the following questions: are these gaps because Corollary 3.2 is based on the coherence property, which is only a sufficient condition for the original condition; or are these gaps because the screening condition is fundamentally weaker than the conditions derived in [32, Theorem 1] for sure screening? In the next section, we show that the answer to the first question is in the affirmative by evaluating the screening condition explicitly for sub-Gaussian design matrices and then specializing Theorem 1 to obtain tight conditions for sure screening of such matrices.

V. SCREENING OF SUB-GAUSSIAN DESIGN MATRICES

In this section, we explicitly characterize the implications of Theorem 1 for ExSIS of the family of sub-Gaussian design matrices. As noted in Sec. III, this effort primarily involves establishing the screening condition for sub-Gaussian matrices and specifying the parameter $b(n, p)$ for such matrices. We begin by first recalling the definition of a sub-Gaussian random variable.

Definition 4. A zero-mean random variable \mathcal{X} is said to follow a sub-Gaussian distribution $\text{subG}(B)$ if there exists a sub-Gaussian parameter $B > 0$ such that $\mathbb{E}[\exp(\lambda\mathcal{X})] \leq \exp\left(\frac{B^2\lambda^2}{2}\right)$ for all $\lambda \in \mathbb{R}$.

Our focus in this section is on design matrices in which entries are first independently drawn from sub-Gaussian distributions and then the columns are normalized. In contrast to prior works, however, we do not require the (pre-normalized) entries to be identically distributed. Rather, we allow each independent variable to be distributed as a sub-Gaussian random variable with a different sub-Gaussian parameter. Thus, the ExSIS analysis of this section is applicable to design matrices in which different columns might have different sub-Gaussian distributions. It is also straightforward to extend our analysis to the case where all (and not just across column) entries of the design matrix are non-identically distributed; we do not focus on this extension in here for the sake of notational clarity.

A. Main Result

The ExSIS of linear models involving sub-Gaussian design matrices mainly requires establishing the screening condition and characterization of the parameter $b(n, p)$ for sub-Gaussian matrices. We accomplish this by individually deriving (SC-1) and (SC-2) in Definition 1 for sub-Gaussian design matrices in the following two lemmas.

Lemma 6. Let $V = [V_{i,j}]$ be an $n \times p$ matrix with the entries $\{V_{i,j}\}_{i,j=1}^{n,p}$ independently distributed as $\text{subG}(b_j)$ with variances $\mathbb{E}[V_{i,j}^2] = \sigma_j^2$. Suppose the design matrix X is obtained by normalizing the columns of V , i.e., $X = V \text{diag}(1/\|V_1\|_2, \dots, 1/\|V_p\|_2)$. Finally, fix an arbitrary $\beta \in \mathbb{R}^p$ that is k -sparse, define $\frac{\sigma_*}{b_*} := \min_{j \in \mathcal{S}} \frac{\sigma_j}{b_j}$, and let $\log p \leq \frac{n}{16} \left(\frac{\sigma_*}{4b_*}\right)^4$. Then, with probability exceeding $1 - \frac{2k^2}{p^2}$, we have

$$\max_{i \in \mathcal{S}} \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| \leq \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2.$$

Lemma 7. Let $V = [V_{i,j}]$ be an $n \times p$ matrix with the entries $\{V_{i,j}\}_{i,j=1}^{n,p}$ independently distributed as $\text{subG}(b_j)$ with variances $\mathbb{E}[V_{i,j}^2] = \sigma_j^2$. Suppose the design matrix X is obtained by normalizing the columns of V , i.e.,

$X = V \text{diag}(1/\|V_1\|_2, \dots, 1/\|V_p\|_2)$. Finally, fix an arbitrary $\beta \in \mathbb{R}^p$ that is k -sparse, define $\frac{\sigma_*}{b_*} := \min_{j \in \mathcal{S}} \frac{\sigma_j}{b_j}$, and let $\log p \leq \frac{n}{16} \left(\frac{\sigma_*}{4b_*}\right)^4$. Then, with probability exceeding $1 - \frac{2(k+1)(p-k)}{p^2}$, we have

$$\max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| \leq \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2.$$

The proofs of Lemma 6 and Lemma 7 are provided in Appendix E and Appendix F, respectively. It now follows from a simple union bound argument that the screening condition holds for sub-Gaussian design matrices with probability exceeding $1 - 2(k+1)p^{-1}$. In particular, we have from Lemma 6 and Lemma 7 that $b(n, p) = \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right)$ for sub-Gaussian matrices. We can now use this knowledge and Theorem 1 to provide the main result for ExSIS of ultrahigh-dimensional linear models involving sub-Gaussian design matrices.

Theorem 4 (ExSIS and Sub-Gaussian Matrices). *Let $V = [V_{i,j}]$ be an $n \times p$ matrix with the entries $\{V_{i,j}\}_{i,j=1}^{n,p}$ independently distributed as $\text{subG}(b_j)$ with variances $\mathbb{E}[V_{i,j}^2] = \sigma_j^2$. Suppose the design matrix X is obtained by normalizing the columns of V , i.e., $X = V \text{diag}(1/\|V_1\|_2, \dots, 1/\|V_p\|_2)$. Next, let $y = X\beta + \eta$ with β a k -sparse vector and the entries of η independently distributed as $\mathcal{N}(0, \sigma^2)$. Finally, define $\frac{\sigma_*}{b_*} := \min_{j \in \mathcal{S}} \frac{\sigma_j}{b_j}$ and $\beta_{\min} := \min_{i \in \mathcal{S}} |\beta_i|$, and let $\log p \leq \frac{n}{16} \left(\frac{\sigma_*}{4b_*}\right)^4$ and $\frac{\beta_{\min}}{\|\beta\|_2} > 2\sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) + 4\frac{\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$. Then Algorithm 1 guarantees $\mathcal{S} \subset \widehat{\mathcal{S}}_d$ with probability exceeding $1 - 2(k+2)p^{-1}$ as long as*

$$d \geq \left\lceil \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2\sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil.$$

Proof. Let \mathcal{G}_p be the event that the design matrix X satisfies the screening condition with parameter $b(n, p) = \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right)$. Further, let \mathcal{G}_η be the event as defined in Theorem 1. It then follows from Lemma 6, Lemma 7, (3), and the union bound that the event $\mathcal{G}_p \cap \mathcal{G}_\eta$ holds with probability exceeding $1 - 2(k+2)p^{-1}$. The advertised claim now follows directly from Theorem 1. \blacksquare

B. Discussion

Since Theorem 4 follows from Theorem 1, it shares many of the insights discussed in Sec. III-A. In particular, Theorem 4 allows for exponential scaling of the number of independent variables, $\log p \leq \frac{n}{16} \left(\frac{\sigma_*}{4b_*}\right)^4$, and dictates that the number of independent variables, d , retained after the screening stage be increased with an increase in the sparsity level and/or the number of independent variables, while it can be decreased with an increase in the SNR, MSR, and/or the number of samples. Notice that the lower bound on d in Theorem 4 does require knowledge of the sparsity level. However, this limitation can be overcome in a straightforward manner, as shown below.

Corollary 4.1. *Let $V = [V_{i,j}]$ be an $n \times p$ matrix with the entries $\{V_{i,j}\}_{i,j=1}^{n,p}$ independently distributed as $\text{subG}(b_j^2)$ with variances $\mathbb{E}[V_{i,j}^2] = \sigma_j^2$. Suppose the design matrix X is obtained by normalizing the columns of V , i.e., $X = V \text{diag}(1/\|V_1\|_2, \dots, 1/\|V_p\|_2)$. Next, let $y = X\beta + \eta$ with β a k -sparse vector and the entries of η independently*

distributed as $\mathcal{N}(0, \sigma^2)$. Further, define $\frac{\sigma_*}{b_*} := \min_{j \in \mathcal{S}} \frac{\sigma_j}{b_j}$ and $\beta_{\min} := \min_{i \in \mathcal{S}} |\beta_i|$. Finally, let $k \leq \frac{n}{\log p}$, $\log p \leq \frac{n}{16} \left(\frac{\sigma_*}{4b_*}\right)^4$, and $\frac{\beta_{\min}}{\|\beta\|_2} > 2c_1 \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) + 4c_2 \frac{\sqrt{\sigma^2 \log p}}{\|\beta\|_2}$ for some constants $c_1, c_2 > 2$. Then Algorithm 1 guarantees $\mathcal{S} \subset \widehat{\mathcal{S}}_d$ with probability exceeding $1 - 2(k + 2)p^{-1}$ as long as $d \geq \left\lceil \frac{n}{\log p} \right\rceil$.

Proof. Theorem 4 and the condition $\frac{\beta_{\min}}{\|\beta\|_2} > 2c_1 \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) + \frac{4c_2 \sqrt{\sigma^2 \log p}}{\|\beta\|_2}$ dictates

$$d \geq \left\lceil \frac{\sqrt{k}}{2(c_1 - 1) \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) + \frac{4(c_2 - 1) \sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \right\rceil \quad (8)$$

for sure screening of sub-Gaussian design matrices. The claim now follows by noting that $d \geq \left\lceil \frac{n}{\log p} \right\rceil$ is a sufficient condition for (8) since $k \leq \frac{n}{\log p}$ and $\sigma_* \leq b_*$ for sub-Gaussian random variables. ■

Following our discussion in Sec. IV-C, a few remarks are in order now concerning the comparison between Corollary 4.1 and [32, Theorem 1]. We compare the two results for the case of Gaussian matrices with independent entries by once again imposing $\sigma = 0$ and $\|\beta\|_2 = \mathcal{O}(1)$. Recall that [32, Theorem 1] imposes constraints $\beta_{\min} = \Omega(n^{-\kappa})$ for $\kappa < 1/2$ and $\log p = \mathcal{O}(n^\alpha)$ for some $\alpha \in (0, 1 - 2\kappa)$. In comparison, substituting $\beta_{\min} = \Omega(n^{-\kappa})$ and $\log p = \mathcal{O}(n^\alpha)$ in Corollary 4.1 results in identical constraints of $\kappa < 1/2$ and $\alpha < 1 - 2\kappa$. Further, the sparsity constraint $k = \mathcal{O}(n/\log p)$ in [32, Theorem 1] also matches the sparsity constraint imposed in Corollary 4.1. To summarize, the ExSIS results derived in this section coincide with the ones in [32] for the case of Gaussian design matrices. However, our results are more general in the sense that they explicitly bring out the dependence of Algorithm 1 on the SNR and the MSR, which is something missing in [32].

VI. EXPERIMENTAL RESULTS

In this section, we present results from two synthetic-data experiments and one real-data experiment. In Section VI-A, we demonstrate that our insights on the ExSIS procedure are not mere artifacts of our analysis. In Section VI-B, we analyze the performance of various regularization-based screening procedures in comparison to the ExSIS procedure. Finally, in Section VI-C, we analyze the computational savings achieved from the use of ExSIS for screening of the feature space as part of sentiment analysis of IMDb movie reviews [55].

A. Oracle-based Screening

In order to ensure the insights offered by Theorem 1 are not mere artifacts of our analysis, we carry out numerical experiments to study the impact of relevant parameters on the screening performance of an *oracle* that has perfect knowledge of the minimum value of d required in Algorithm 1 to ensure $\mathcal{S} \subset \widehat{\mathcal{S}}_d$. In particular, we use these oracle-based experiments to verify the role of $b(n, p)$ and MSR in screening using Algorithm 1, as specified by Theorem 1. Since worst-case coherence is an indirect measure of pairwise similarity among the columns of X , and since $b(n, p)$ cannot be explicitly computed, we use μ as a surrogate for the value of $b(n, p)$ in our experiments.

The design matrix $X \in \mathbb{R}^{n \times p}$ in our experiments is generated such that it consists of independent and identically distributed Gaussian entries, followed by normalization of the columns of X . Among other parameters, $n = 500$, $p = 2000$, $k = 5$, and $\sigma = 0$ in the experiments. The entries of \mathcal{S} are chosen uniformly at random from $[[p]]$.


Furthermore, the non-zero entries in the parameter vector β are sampled from a uniform distribution $U[a, e]$; the value of a is set at 1 whereas $e \in [2, 10]$. Finally, the experiments comprise the use of an oracle to find the minimum possible value of d that can be used in Algorithm 1 while ensuring $\mathcal{S} \subset \widehat{\mathcal{S}}_d$. We refer to this minimum value of d as the *minimum model size* (MMS), and we use median of MMS over 400 runs of the experiment as a metric of difficulty of screening.

To analyze the impact of increasing μ (equivalently, $b(n, p)$) and MSR on screening using Algorithm 1, the numerical experiments are repeated for various values of μ and MSR. In particular, the worst-case coherence of X is varied by scaling its largest singular value, followed by normalization of the columns of X , while the MSR is increased by decreasing the value of e . In Fig. 1(a), we plot the median MMS against μ for different MSR values. The experimental results of the oracle performance offer two interesting insights. First, the median MMS increases with μ ; this shows that any analysis for screening using Algorithm 1 needs to account for the similarity between the columns of X . This relationship is captured by the parameter $b(n, p)$ in Theorem 1. Second, the difficulty of screening for an oracle increases with decreasing MSR values. This relationship is also reflected in Theorem 1: as $\|\beta\|_2$ increases for a fixed e , MSR decreases and the median MMS increases.

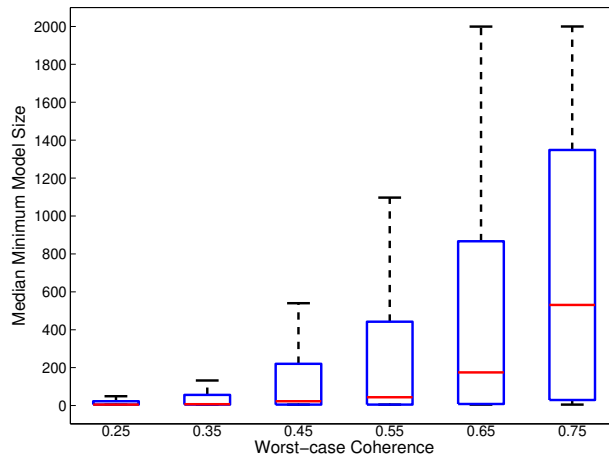
More interestingly, if we focus on the plot in Fig. 1(a) for $e = 10$, and we plot the relationship between μ and median MMS along with the interquartile range of MMS for each value of μ , it can be seen that there are instances when the oracle has to select all 2000 predictors to ensure $\mathcal{S} \subset \widehat{\mathcal{S}}_d$ (see boxplot for $\mu = 0.65$ and 0.75). In other words, no screening can be performed at all in these cases. This phenomenon is also reflected in Theorem 1: when $b(n, p)$ becomes too large, the condition imposed on MSR is no longer true and our analysis cannot be used for screening using Algorithm 1. Thus, this experiment shows that the condition imposed on MSR in Theorem 1 is not a mere artifact of our analysis: as μ (equivalently, $b(n, p)$) becomes large, one may not be able to perform screening at all while ensuring $\mathcal{S} \subset \widehat{\mathcal{S}}_d$.

B. Comparison with Screening Procedures for LASSO-type Methods

In this section, we use Gaussian data to compare the performance of ExSIS to that of screening procedures for LASSO-type methods. The design matrix $X \in \mathbb{R}^{n \times p}$ contains entries from a standard Gaussian distribution such that the pairwise correlation between the variables is ρ . In this experiment, we fix n at 200 while we consider two models with $p = 2000$ and $p = 5000$. For each value of p , we further consider two models with $\rho = 0.0$ and $\rho = 0.3$. Thus, in our experiments, we consider four setups with $(p, \rho) = (2000, 0.0)$, $(2000, 0.3)$, $(5000, 0.0)$ and $(5000, 0.3)$ to analyze the impact of dimensionality and pairwise correlation on performance of the screening procedures for LASSO-type methods in relation to ExSIS. For each of these four different setups, the model size is set at $|\mathcal{S}| = 5$, and the locations of the non-zero coefficients in the parameter vector β are chosen such that \mathcal{S} is a uniformly at random subset of $[[p]]$. The values of the non-zero coefficients in the parameter vector β are generated from $|z| + 2$ where z is distributed as a standard Gaussian random variable. Furthermore, the noise samples are generated from a standard Gaussian distribution, and the response vector y is generated using (1). Finally, the response vector y and the columns of the design matrix X are normalized to have unit norm.



(a)



(b)

Fig. 1: (a) For various values of worst-case coherence of the design matrix, the output of the oracle is computed. The relation between worst-case coherence and minimum model size is analyzed for various MSR values. (b) Boxplot of minimum model size for various values of worst-case coherence. The boxplot corresponds to the plot in Fig. 1(a) for $e = 10$. As the worst-case coherence becomes large, there are instances when oracle has to select all 2000 predictors to ensure $\mathcal{S} \subset \widehat{\mathcal{S}}_d$ (see boxplot for $\mu = 0.65$ and 0.75).

To analyze the performance of screening procedures for the LASSO method [2], the columns of X are screened using SAFE method [41] and strong rules [42] for LASSO. Recall that the LASSO problem can be expressed as

$$\hat{\beta} = \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1.$$

For each of these screening methods, we perform screening of X over a set of 200 values of the regularization parameter λ that are chosen uniformly from a linear scale. We compare the screening performance of SAFE method and strong rules for LASSO with the ExSIS method where $d = 2n$. Note that our selection of the value of d has some




Fig. 2: Gaussian data matrices are screened using LASSO-based screening procedures for various model parameters. In (a) and (c), $p = 2000$; whereas, in (b) and (d), $p = 5000$. For each value of p , the screening experiment is repeated for $\rho = 0.0$ and $\rho = 0.3$. In each experiment, the model size after screening and the corresponding detection rate is evaluated for different values of the regularization parameter λ . The shown results are median over 100 random draws of the data matrix X /parameter vector β /noise vector η in (1).

slack over the suggested value of d from Corollary 4.1 because the conditions on $\log p$ and $\frac{\beta_{\min}}{\|\beta\|_2}$ in Corollary 4.1 don't hold true in this experiment.³ To compare the performance of these various screening methods, we use two metrics: (i) the model size (number of variables) after screening, which is defined as $|\hat{S}_d|$, and (ii) the detection rate, which is defined as $\frac{|S \cap \hat{S}_d|}{|\hat{S}_d|}$. Using these metrics of performance, Fig. 2 shows the results of our simulations as median over 100 draws of the random design matrix X /parameter vector β /noise vector η in (1) for each of the four setups that we consider in this section.

Next, the design matrix X is also generated and screened using SAFE method [41] and strong rules [42] for elastic net [23] as explained before. Recall that the elastic net problem can be expressed as

$$\hat{\beta} = \arg \min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda_1 \|\beta\|_1 + \frac{1}{2} \lambda_2 \|\beta\|_2.$$

³In order for stated conditions to hold, we need significantly larger n (and p); however, running LASSO on such large problems has high computational needs.




Fig. 3: Gaussian data matrices are screened using elastic net-based screening procedures for various model parameters. In (a) and (c), $p = 2000$; whereas, in (b) and (d), $p = 5000$. For each value of p , the screening experiment is repeated for $\rho = 0.0$ and $\rho = 0.3$. In each experiment, the model size after screening and the corresponding detection rate is evaluated for different values of the regularization parameter λ . The shown results are median over 100 random draws of the data matrix X /parameter vector β /noise vector η in (1).

In our simulations, we use the parametrization $(\lambda_1, \lambda_2) = (\alpha\lambda, (1 - \alpha)\lambda)$ and we let $\alpha = 0.5$. Fig. 3 shows the screening performance of SAFE method for elastic net, strong rules for elastic net and ExSIS over 100 draws of the random design matrix X for each of the four setups, as described before. In both Figs. 2 and 3, the largest value of λ for which median detection rate is 1.0 is labeled with an asterisk for each optimization-based screening procedure. In other words, for each screening procedure, only if λ is smaller than the value of λ labeled by an asterisk, the screening procedure maintains a median detection rate of 1.0. Notice also that if the chosen value of λ is too small, no variable is deleted by the screening procedure. Thus, as can be seen in both the figures, there is only a narrow range of values of λ over which the optimization-based screening procedures are able to get rid of variables while maintaining a detection rate of 1.0. Thus, from a practical point of view, it is not trivial to use SAFE method or strong rules for screening because there is no way of ensuring that the chosen value of λ is within the narrow range of values of λ for which significant screening can be performed while maintaining a detection rate of 1.0. In comparison, the ExSIS method does not depend on the parameter λ , and in our experiments, it could always be used for screening while maintaining a median detection rate of 1.0 (as shown in both Figs. 2 and 3).

Before we end this discussion, note that, even within the narrow range of values of λ for which SAFE method or strong rules can be used for screening while maintaining a detection rate of 1.0, there is an even narrower range of values of λ for which SAFE method or strong rules delete more variables than ExSIS.

C. Sentiment Analysis of IMDb Movie Reviews and ExSIS

In high-dimensional classification, it has been shown that the presence of irrelevant variables increases the difficulty of classification, and the classification error tends to increase with dimensionality of the data model [56], [57]. Variable selection becomes important in high-dimensional classification as it can be used to discard the subset of irrelevant variables and reduce the dimensionality of the data model. Once the variable selection step is performed, classification can be performed based on the subset of relevant variables. In this section, we consider the problem of classifying IMDb movie reviews with positive or negative sentiments. In particular, we use variable selection to (i) reduce dimensionality of the data model, and (ii) learn a linear data model for classification (as explained later). To build and test our classification model, we make use of the IMDb movie reviews dataset [55], with the response being either a 1 (positive review) or a 0 (negative review), and we extract features using the *term frequency-inverse document frequency method* [58].

To increase the reliability of our results, the original dataset of 25K reviews is first randomly divided into five bins for five independent trials, with each bin further divided into 3K train and 2K test reviews. In each bin, before we use the 3K reviews for fitting a linear model on the feature space, we perform a preprocessing step to get rid of the features (words) that are highly correlated. Note that, when we refer to learning/fitting a linear model, we mean to estimate the vector β in (1). For learning the linear data model, we use LASSO as well as elastic net. For tuning the regularization parameter in LASSO as well as elastic net for each bin, we perform a five-fold cross validation experiment and choose the value of the regularization parameter that minimizes the mean square error on the training dataset. To evaluate the predictive power of the linear model, we use the notion of test *true positive* (TP) rate, which is the percentage of the remaining 2K test movie reviews that are correctly classified by the model. For classification of the test reviews, we use the trained linear model to estimate the response for each test review. If the estimated response is less than 0.5 for a test review, the test review is assigned a negative sentiment and vice versa. The above procedure is repeated for each of the five bins of data and the average prediction accuracy is reported in Table I. The average model size before variable selection in the five runs of the experiment is 21,345.

We also repeat the aforementioned experiment procedure but with a slight variation. For each of the five bins of data, we use Algorithm 1 to decrease dimensionality of the data model before performing the variable selection step. The objective of this variation in the experiment is to analyze the decrease in computational time and any change in the prediction accuracy when the variable selection step is preceded with a screening step using Algorithm 1. To choose the value of d in Algorithm 1, we verify that the training data matrix for each fold of data satisfies the coherence property, and then we choose $d = 2n$ where $n = 3000$. Note that the chosen value of d has some slack over the suggested value of d in Corollary 3.1 because the condition on $\frac{\beta_{\min}}{\|\beta\|_2}$ in Corollary 3.1 does not hold true in this experiment. After the screening step, we use LASSO and elastic net to learn and test a classification model as explained before. The results are reported in Table I.

TABLE I: Test true positive (TP) rates and computational times for experiments on the IMDb dataset, averaged over the five folds of the IMDb dataset. The standard deviation is also reported in parenthesis.

Training method	Test TP rate	Training time
LASSO	83.01 (0.77)	388.35 (26.66)
ExSIS-LASSO	82.23 (0.83)	177.43 (16.85)
Elastic net	84.35 (0.89)	272.46 (15.76)
ExSIS-Elastic net	82.06 (0.94)	111.20 (4.65)

Thus, we use LASSO and elastic net, both with and without screening, to train and test a linear model for classification of movie reviews. For each of these four cases, Table I reports the test TP rate, which is the percentage of correctly classified reviews within the test reviews. The computational time needed for learning the linear model is also reported as an average over the five folds of data. It can be seen from the table that Algorithm 1 reduces the training time by a factor of more than two, while there is only a small decrease in predictive power of the trained model.

VII. CONCLUSION

In this paper, we studied marginal correlation-based screening for ultrahigh-dimensional linear models. In our analysis, we provided sufficient conditions for sure screening of arbitrary linear models, which gave us novel insights on the screening procedure, while our simulation results verified that these insights are reflective of the challenges associated with marginal correlation-based screening. We also provided polynomial-time verifiable conditions for arbitrary (random or deterministic) matrices under which the dimension of the model can be reduced to almost the sample size. Furthermore, we specialized our sufficient conditions for sub-Gaussian linear models and demonstrated that our analysis coincides with the existing results for screening of such models. In our experiments with real-world data, we demonstrated the computational savings that can be achieved through ExSIS in high dimensional variable selection.

APPENDIX A

PROOF OF LEMMA 1

We begin by defining

$$w^{(i)} := X_{\widehat{\mathcal{S}}_{p_{i-1}}}^\top y = X_{\widehat{\mathcal{S}}_{p_{i-1}}}^\top X_S \beta_S + X_{\widehat{\mathcal{S}}_{p_{i-1}}}^\top \eta =: \xi^{(i)} + \tilde{\eta}^{(i)} \quad (9)$$

where $w^{(i)} \in \mathbb{R}^{|\widehat{\mathcal{S}}_{p_{i-1}}|}$ measures the correlation of the observation vector y with each column of $X_{\widehat{\mathcal{S}}_{p_{i-1}}}$. To derive an upper bound on t_i , we derive upper and lower bounds on $\sum_{j_1=1}^{p_{i-1}} |w_{j_1}^{(i)}|$. A simple upper bound on $\sum_{j_1=1}^{p_{i-1}} |w_{j_1}^{(i)}|$ is:

$$\sum_{j_1=1}^{p_{i-1}} |w_{j_1}^{(i)}| = \sum_{j_1=1}^{p_{i-1}} |\xi_{j_1}^{(i)} + \tilde{\eta}_{j_1}^{(i)}| \leq \sum_{j_1=1}^{p_{i-1}} (|\xi_{j_1}^{(i)}| + |\tilde{\eta}_{j_1}^{(i)}|) = \|\xi^{(i)}\|_1 + \|\tilde{\eta}^{(i)}\|_1. \quad (10)$$

Next, we recall that $w = X^\top y$ and we further define ξ and $\tilde{\eta}$ such that

$$w = X^\top y = X^\top X_S \beta_S + X^\top \eta =: \xi + \tilde{\eta}. \quad (11)$$

Now define $\mathcal{T}_i := \{j_1 \in \widehat{\mathcal{S}}_{p_{i-1}} : |w_{j_1}| \geq \min_{j_2 \in \mathcal{S}} |w_{j_2}|\}$. Then, a simple lower bound on $\sum_{j_1=1}^{p_{i-1}} |w_{j_1}^{(i)}|$ is:

$$\begin{aligned} \sum_{j_1=1}^{p_{i-1}} |w_{j_1}^{(i)}| &= \sum_{j_1 \in \widehat{\mathcal{S}}_{p_{i-1}}} |w_{j_1}| \geq \sum_{j_1 \in \mathcal{T}_i} |w_{j_1}| \\ &\geq \sum_{j_1 \in \mathcal{T}_i} \min_{j_2 \in \mathcal{T}_i} |w_{j_2}| \stackrel{(a)}{\geq} \sum_{j_1 \in \mathcal{T}_i} \min_{j_2 \in \mathcal{S}} |w_{j_2}| \\ &= t_i(\min_{j_2 \in \mathcal{S}} |w_{j_2}|) = t_i(\min_{j_2 \in \mathcal{S}} |\xi_{j_2} + \tilde{\eta}_{j_2}|) \\ &\geq t_i(\min_{j_2 \in \mathcal{S}} |\xi_{j_2}| - \max_{j_2 \in \mathcal{S}} |\tilde{\eta}_{j_2}|), \end{aligned} \quad (12)$$

where (a) follows from definition of \mathcal{T}_i . Combining (10) with (12), we get

$$t_i \leq \frac{\|\xi^{(i)}\|_1 + \|\tilde{\eta}^{(i)}\|_1}{\min_{j_2 \in \mathcal{S}} |\xi_{j_2}| - \max_{j_2 \in \mathcal{S}} |\tilde{\eta}_{j_2}|}. \quad (13)$$

We next bound $\|\xi^{(i)}\|_1$, $\max_{j_2 \in \mathcal{S}} |\tilde{\eta}_{j_2}|$, $\|\tilde{\eta}^{(i)}\|_1$ and $\min_{j_2 \in \mathcal{S}} |\xi_{j_2}|$ separately. First, we derive an upper bound on $\|\xi^{(i)}\|_1$:

$$\begin{aligned} \|\xi^{(i)}\|_1 &= \sum_{j_1 \in \widehat{\mathcal{S}}_{p_{i-1}}} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| \\ &\stackrel{(b)}{=} \sum_{j_1 \in \mathcal{S}} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| + \sum_{j_1 \in \widehat{\mathcal{S}}_{p_{i-1}} \setminus \mathcal{S}} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| \\ &\stackrel{(c)}{\leq} \sum_{j_1 \in \mathcal{S}} \left| \sum_{\substack{j_2 \in \mathcal{S} \\ j_2 \neq j_1}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| + \sum_{j_1 \in \mathcal{S}} |\beta_{j_1}| + \sum_{j_1 \in \widehat{\mathcal{S}}_{p_{i-1}} \setminus \mathcal{S}} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| \\ &\leq k \max_{\substack{j_1 \in \mathcal{S} \\ j_2 \in \mathcal{S} \\ j_2 \neq j_1}} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| + (p_{i-1} - k) \max_{j_1 \in \widehat{\mathcal{S}}_{p_{i-1}} \setminus \mathcal{S}} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| + \|\beta\|_1 \\ &\leq k \max_{\substack{j_1 \in \mathcal{S} \\ j_2 \in \mathcal{S} \\ j_2 \neq j_1}} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| + (p_{i-1} - k) \max_{j_1 \in \mathcal{S}^c} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| + \|\beta\|_1, \end{aligned} \quad (14)$$

where (b) follows since $\mathcal{S} \subset \widehat{\mathcal{S}}_{p_{i-1}}$ and (c) follows from the triangle inequality and the fact that the columns of X are unit norm. Next, we have

$$\max_{j_2 \in \mathcal{S}} |\tilde{\eta}_{j_2}| \leq \|\tilde{\eta}\|_\infty \leq 2\sqrt{\sigma^2 \log p}, \quad (15)$$

where the last inequality follows from conditioning on \mathcal{G}_η . Similarly, we have

$$\begin{aligned} \|\tilde{\eta}^{(i)}\|_1 &= \sum_{j_1 \in \widehat{\mathcal{S}}_{p_{i-1}}} |X_{j_1}^\top \eta| \leq \sum_{j_1 \in \widehat{\mathcal{S}}_{p_{i-1}}} \max_{j_2 \in \widehat{\mathcal{S}}_{p_{i-1}}} |X_{j_2}^\top \eta| \\ &= p_{i-1} \left(\max_{j_2 \in \widehat{\mathcal{S}}_{p_{i-1}}} |X_{j_2}^\top \eta| \right) \leq 2p_{i-1} \sqrt{\sigma^2 \log p} \end{aligned} \quad (16)$$

where the last inequality, again, follows from \mathcal{G}_η . Last, we lower bound $\min_{j_1 \in \mathcal{S}} |\xi_{j_1}|$ as follows:

$$\begin{aligned} \min_{j_1 \in \mathcal{S}} |\xi_{j_1}| &= \min_{j_1 \in \mathcal{S}} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| \\ &\stackrel{(d)}{=} \min_{j_1 \in \mathcal{S}} \left| \sum_{\substack{j_2 \in \mathcal{S} \\ j_2 \neq j_1}} X_{j_1}^\top X_{j_2} \beta_{j_2} + \beta_{j_1} \right| \\ &\geq \min_{j_1 \in \mathcal{S}} |\beta_{j_1}| - \max_{\substack{j_2 \in \mathcal{S} \\ j_2 \neq j_1}} \left| \sum_{j_1 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| = \beta_{\min} - \max_{\substack{j_2 \in \mathcal{S} \\ j_2 \neq j_1}} \left| \sum_{j_1 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| \end{aligned} \quad (17)$$

where (d) follows because the columns of X are unit norm. Combining (14), (15), (16), (17) with (13), we obtain

$$t_i \leq \frac{k \max_{j_1 \in \mathcal{S}} \left| \sum_{\substack{j_2 \in \mathcal{S} \\ j_2 \neq j_1}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| + (p_{i-1} - k) \max_{j_1 \in \mathcal{S}^c} \left| \sum_{j_2 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| + \|\beta\|_1 + 2p_{i-1} \sqrt{\sigma^2 \log p}}{\beta_{\min} - \max_{\substack{j_1 \in \mathcal{S} \\ j_2 \neq j_1}} \left| \sum_{j_1 \in \mathcal{S}} X_{j_1}^\top X_{j_2} \beta_{j_2} \right| - 2\sqrt{\sigma^2 \log p}}. \quad (18)$$

Assuming the (k, b) -screening condition for the matrix X holds, we finally obtain

$$\begin{aligned} t_i &\leq \frac{kb(n, p) + (p_{i-1} - k)b(n, p) + \|\beta\|_1 + 2p_{i-1} \sqrt{\sigma^2 \log p}}{\beta_{\min} - b(n, p) - 2\sqrt{\sigma^2 \log p}} \\ &= \frac{p_{i-1}b(n, p) + \|\beta\|_1 + 2p_{i-1} \sqrt{\sigma^2 \log p}}{\beta_{\min} - b(n, p) - 2\sqrt{\sigma^2 \log p}}. \end{aligned} \quad (19)$$

This completes the proof of the lemma. \blacksquare

APPENDIX B

PROOF OF LEMMA 2

For $\bar{t}_i < p_{i-1}$, we need

$$\begin{aligned} \frac{p_{i-1}b(n, p)\|\beta\|_2 + \|\beta\|_1 + 2p_{i-1}\sqrt{\sigma^2 \log p}}{\beta_{\min} - b(n, p)\|\beta\|_2 - 2\sqrt{\sigma^2 \log p}} &< p_{i-1} \\ \Leftrightarrow p_{i-1} &> \frac{\frac{\|\beta\|_1}{\|\beta\|_2}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2b(n, p) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}}. \end{aligned} \quad (20)$$

Since $\|\beta\|_1 \leq \sqrt{k}\|\beta\|_2$, we have

$$p_{i-1} > \frac{\sqrt{k}}{\frac{\beta_{\min}}{\|\beta\|_2} - 2b(n, p) - \frac{4\sqrt{\sigma^2 \log p}}{\|\beta\|_2}} \quad (21)$$

as a sufficient condition for (20). Thus, (21) is a sufficient condition for $\bar{t}_i < p_{i-1}$. \blacksquare

APPENDIX C

PROOF OF LEMMA 3

Notice that $\max_{\substack{i \in \mathcal{S} \\ j \in \mathcal{S} \\ j \neq i}} |X_i^\top X_j \beta_j| \leq \max_{\substack{i \in \mathcal{S} \\ j \in \mathcal{S} \\ j \neq i}} |X_i^\top X_j \beta_j| \leq \max_{\substack{i \in \mathcal{S} \\ j \in \mathcal{S} \\ j \neq i}} |X_i^\top X_j| \|\beta_j\|$. Further, we have

$$\max_{\substack{i \in \mathcal{S} \\ j \in \mathcal{S} \\ j \neq i}} |X_i^\top X_j| \|\beta_j\| \leq \max_{\substack{i \in \mathcal{S} \\ j \in \mathcal{S} \\ j \neq i}} \mu \|\beta_j\| \leq \mu \|\beta\|_1 \leq \mu \sqrt{k} \|\beta\|_2. \quad (22)$$

An identical argument also establishes that $\max_{i \in \mathcal{S}^c} |X_i^\top X_j \beta_j| \leq \mu \|\beta\|_1 \leq \mu \sqrt{k} \|\beta\|_2$. \blacksquare

APPENDIX D
PROOF OF LEMMA 4

The proof of Lemma 4 relies on the following two lemmas, which are formally proved in [51].

Lemma 8 ([51, Lemma 3]). *Assume \mathcal{S} is drawn uniformly at random from k -subsets of $[[p]]$ and choose a parameter $a \geq 1$. Let $\epsilon \in [0, 1)$ and $k \leq \min\{\epsilon^2\nu^{-2}, (1+a)^{-1}p\}$. Then, with probability exceeding $1 - 4k \exp(-\frac{(\epsilon - \sqrt{k\nu})^2}{16(2+a^{-1})^2\mu^2})$, we have*

$$\max_{i \in \mathcal{S}} \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| \leq \epsilon \|\beta\|_2.$$

Lemma 9 ([51, Lemma 4]). *Assume \mathcal{S} is drawn uniformly at random from k -subsets of $[[p]]$ and choose a parameter $a \geq 1$. Let $\epsilon \in [0, 1)$ and $k \leq \min\{\epsilon^2\nu^{-2}, (1+a)^{-1}p\}$. Then, with probability exceeding $1 - 4(p - k) \exp(-\frac{(\epsilon - \sqrt{k\nu})^2}{8(2+a^{-1})^2\mu^2})$, we have*

$$\max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| \leq \epsilon \|\beta\|_2.$$

Using a simple union bound with Lemma 8 and Lemma 9, we have

$$\begin{aligned} \max_{i \in \mathcal{S}} \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| &\leq \epsilon \|\beta\|_2, \text{ and} \\ \max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| &\leq \epsilon \|\beta\|_2 \end{aligned}$$

with probability exceeding $1 - \delta$ where $\delta = 4p \exp(-\frac{(\epsilon - \sqrt{k\nu})^2}{16(2+a^{-1})^2\mu^2})$. Fix $\epsilon = c_\mu \mu \sqrt{\log p}$ and $a = 9$. Then, the claim is that $\delta \leq 4p^{-1}$. Next, we will prove our claim. Before we can fix $\epsilon = c_\mu \mu \sqrt{\log p}$ and $a = 9$, we need to ensure that the chosen values of a , ϵ and the allowed values of k satisfy the assumptions in Lemma 8 and Lemma 9. First, note that $\epsilon < 1$ because of $\mu < \frac{1}{c_\mu \sqrt{\log p}}$. Second, $k \leq \frac{p}{1+a}$ because of $a = 9$, $p \geq \max\{2n, \exp(5)\}$ and $k \leq \frac{n}{\log p}$. Third, and last, $k \leq \frac{\epsilon^2}{(9\nu)^2}$ because of $\nu < \frac{\mu}{\sqrt{n}}$, $k \leq \frac{n}{\log p}$, $p \geq 2n$ and $c_\mu > 10\sqrt{2}$. Finally, we use $\sqrt{k\nu} \leq \frac{\epsilon}{9}$ with $a = 9$, $c_\mu > 10\sqrt{2}$ and $\epsilon = c_\mu \mu \sqrt{\log p}$ to obtain

$$\exp\left(-\frac{(\epsilon - \sqrt{k\nu})^2}{16(2+a^{-1})^2\mu^2}\right) \leq p^{-2}$$

and thus $\delta \leq 4p^{-1}$. ■

APPENDIX E
PROOF OF LEMMA 6

Since $X_i := \frac{V_i}{\|V_i\|_2}$, we can write

$$\max_{i \in \mathcal{S}} \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| = \max_{i \in \mathcal{S}} \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} \frac{V_i^\top}{\|V_i\|_2} \frac{V_j}{\|V_j\|_2} \beta_j \right|.$$

We next fix an $i' \in \mathcal{S}$ and derive a probabilistic bound on $|\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} X_{i'}^\top X_j \beta_j|$. This involves deriving both upper and lower probabilistic bounds on $\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} X_{i'}^\top X_j \beta_j$ below in Step 1 and Step 2, respectively.

Step 1 (Upper Bound): To provide an upper probabilistic bound on $\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} X_{i'}^\top X_j \beta_j$, we first establish that $\|V_j\|_2 \geq$

$\sqrt{\frac{n\sigma_j^2}{2}}$ for each $j \in \mathcal{S}$ with high probability in Step 1a and then we derive an upper probabilistic bound on $\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \frac{V_{i'}^\top}{\|V_{i'}\|_2} \frac{\sqrt{2}V_j}{\sqrt{n\sigma_j^2}} \beta_j$ in Step 1b. We then combine these two steps in Step 1c for the final upper probabilistic bound on $\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} X_{i'}^\top X_j \beta_j$.

Step 1a: Note that

$$\Pr \left[\|V_j\|_2 < \sqrt{\frac{n\sigma_j^2}{2}} \right] \leq \exp \left(-\frac{n}{8} \left(\frac{\sigma_j}{4b_j} \right)^4 \right) \quad (23)$$

for any $j \in \mathcal{S}$ [59, eq. (2.20)]. Next, let $\mathcal{G}_{u,a}$ be the event that $\|V_j\|_2 \geq \sqrt{\frac{n\sigma_j^2}{2}}$ for all $j \in \mathcal{S} \setminus \{i'\}$. Then

$$\begin{aligned} \Pr[\mathcal{G}_{u,a}^c] &= \Pr \left[\bigcup_{\substack{j \in \mathcal{S} \\ j \neq i'}} \left\{ \|V_j\|_2 < \sqrt{\frac{n\sigma_j^2}{2}} \right\} \right] \leq \sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \Pr \left[\|V_j\|_2 < \sqrt{\frac{n\sigma_j^2}{2}} \right] \leq \sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \exp \left(-\frac{n}{8} \left(\frac{\sigma_j}{4b_j} \right)^4 \right) \\ &\leq \sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \exp \left(-\frac{n}{8} \left(\frac{\sigma_*}{4b_*} \right)^4 \right) = (k-1) \exp \left(-\frac{n}{8} \left(\frac{\sigma_*}{4b_*} \right)^4 \right). \end{aligned} \quad (24)$$

Step 1b: Define

$$Y_{i'} := \sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \sqrt{\frac{2}{n\sigma_j^2}} \frac{V_{i'}^\top}{\|V_{i'}\|_2} V_j \beta_j,$$

and let $\mathcal{G}_{u,b}$ be the event that $Y_{i'} \leq \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*} \right) \|\beta\|_2$. Then, the claim is that $\Pr(\mathcal{G}_{u,b}) \geq 1 - \frac{1}{p^2}$. In order to prove this claim, let us define another event $\mathcal{G}'_v := \{V_i' = v_{i'}\}$. Then, defining $u_{i'} := \frac{v_{i'}}{\|v_{i'}\|_2}$, we have

$$\begin{aligned} M_{Y_{i'}}(\lambda | \mathcal{G}'_v) &:= \mathbb{E}[\exp(\lambda Y_{i'}) | \mathcal{G}'_v] = \mathbb{E} \left[\exp \left(\lambda \sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \sqrt{\frac{2}{n\sigma_j^2}} \frac{V_{i'}^\top}{\|V_{i'}\|_2} V_j \beta_j \right) | \mathcal{G}'_v \right] = \mathbb{E} \left[\exp \left(\lambda \sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \sqrt{\frac{2}{n\sigma_j^2}} u_{i'}^\top V_j \beta_j \right) \right] \\ &= \mathbb{E} \left[\prod_{\substack{j \in \mathcal{S} \\ j \neq i'}} \exp \left(\lambda \sqrt{\frac{2}{n\sigma_j^2}} u_{i'}^\top V_j \beta_j \right) \right] = \mathbb{E} \left[\prod_{\substack{j \in \mathcal{S} \\ j \neq i'}} \exp \left(\lambda \sqrt{\frac{2}{n\sigma_j^2}} \sum_{l=1}^n u_{i',l} V_{j,l} \beta_j \right) \right] \\ &= \mathbb{E} \left[\prod_{\substack{j \in \mathcal{S} \\ j \neq i'}} \prod_{l=1}^n \exp \left(\lambda \sqrt{\frac{2}{n\sigma_j^2}} u_{i',l} V_{j,l} \beta_j \right) \right] = \prod_{\substack{j \in \mathcal{S} \\ j \neq i'}} \prod_{l=1}^n \mathbb{E} \left[\exp \left(\lambda \sqrt{\frac{2}{n\sigma_j^2}} u_{i',l} V_{j,l} \beta_j \right) \right] \\ &\leq \prod_{\substack{j \in \mathcal{S} \\ j \neq i'}} \prod_{l=1}^n \exp \left(\frac{\lambda^2}{n\sigma_j^2} u_{i',l}^2 b_j^2 \beta_j^2 \right) = \prod_{\substack{j \in \mathcal{S} \\ j \neq i'}} \exp \left(\frac{\lambda^2}{n\sigma_j^2} b_j^2 \beta_j^2 \sum_{l=1}^n u_{i',l}^2 \right) = \prod_{\substack{j \in \mathcal{S} \\ j \neq i'}} \exp \left(\frac{\lambda^2}{n\sigma_j^2} b_j^2 \beta_j^2 \right) \\ &= \exp \left(\frac{\lambda^2}{n} \sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \frac{b_j^2 \beta_j^2}{\sigma_j^2} \right) \leq \exp \left(\frac{\lambda^2}{n} \left(\frac{b_*}{\sigma_*} \right)^2 \sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \beta_j^2 \right) \leq \exp \left(\frac{\lambda^2}{n} \left(\frac{b_*}{\sigma_*} \right)^2 \|\beta\|_2^2 \right). \end{aligned} \quad (25)$$

By the Chernoff bound on $Y_{i'}$, we next obtain

$$\begin{aligned} \Pr(Y_{i'} > a | \mathcal{G}'_v) &\leq \min_{\lambda > 0} \exp(-\lambda a) M_{Y_{i'}}(\lambda | \mathcal{G}'_v) \\ &\leq \min_{\lambda > 0} \exp(-\lambda a) \exp\left(\frac{\lambda^2}{n} \left(\frac{b_*}{\sigma_*}\right)^2 \|\beta\|_2^2\right) \\ &= \exp\left(-\frac{1}{4} \left(\frac{\sigma_*}{b_*}\right)^2 \frac{a^2 n}{\|\beta\|_2^2}\right). \end{aligned} \quad (26)$$

Substituting $a = \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2$ in (26), we obtain

$$\Pr\left(Y_{i'} > \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2 \mid \mathcal{G}'_v\right) \leq \frac{1}{p^2}. \quad (27)$$

Finally, by the law of total probability, we obtain

$$\begin{aligned} \Pr\left(Y_{i'} > \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2\right) &= \mathbb{E}_{V_{i'}} \left[\Pr\left(Y_{i'} > \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2 \mid \mathcal{G}'_v\right) \right] \\ &\leq \mathbb{E}_{V_{i'}} \left[\frac{1}{p^2} \right] = \frac{1}{p^2}. \end{aligned} \quad (28)$$

Thus, the event $\mathcal{G}_{u,b}$ holds with probability exceeding $1 - \frac{1}{p^2}$.

Step 1c: Conditioning on $\mathcal{G}_{u,a} \cap \mathcal{G}_{u,b}$, we have from (24) and (28) that

$$\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} X_{i'}^\top X_j \beta_j \leq \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2. \quad (29)$$

Further, note that

$$\begin{aligned} \Pr(\mathcal{G}_{u,a} \cap \mathcal{G}_{u,b}) &\geq 1 - \Pr(\mathcal{G}_{u,a}^c) - \Pr(\mathcal{G}_{u,b}^c) \\ &\geq 1 - (k-1) \exp\left(-\frac{n}{8} \left(\frac{\sigma_*}{4b_*}\right)^4\right) - \frac{1}{p^2} \\ &\stackrel{(a)}{\geq} 1 - \frac{(k-1)}{p^2} - \frac{1}{p^2} = 1 - \frac{k}{p^2}, \end{aligned}$$

where (a) follows since $\log p \leq \frac{n}{16} \left(\frac{\sigma_*}{4b_*}\right)^4$. Thus, (29) holds with probability exceeding $1 - \frac{k}{p^2}$.

Step 2 (Lower Bound): Our claim in this step is that $\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} X_{i'}^\top X_j \beta_j \geq -\sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2$ with probability

exceeding $1 - \frac{k}{p^2}$. To establish this claim, notice that

$$\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} X_{i'}^\top X_j \beta_j \geq -\sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2 \iff -\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} X_{i'}^\top X_j \beta_j \leq \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2. \quad (30)$$

Further, we have $-\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} X_{i'}^\top X_j \beta_j \equiv -\sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \frac{V_{i'}^\top}{\|V_{i'}\|_2} \frac{V_j}{\|V_j\|_2} \beta_j = \sum_{\substack{j \in \mathcal{S} \\ j \neq i'}} \frac{V_{i'}^\top}{\|V_{i'}\|_2} \frac{\tilde{V}_j}{\|V_j\|_2} \beta_j$, where $\tilde{V}_j := -V_j$ is still distributed as V_j because of the symmetry of sub-Gaussian distributions. The claim now follows from a repetition of the analysis carried out in Step 1.

Final Step: Step 1 and Step 2, along with the union bound, imply that

$$\left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_{i'}^\top X_j \beta_j \right| \leq \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*}\right) \|\beta\|_2$$

with probability exceeding $1 - \frac{2k}{p^2}$. Next, notice that

$$\begin{aligned}
& \Pr \left[\max_{i \in \mathcal{S}} \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| \leq \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*} \right) \|\beta\|_2 \right] \\
&= 1 - \Pr \left[\bigcup_{i \in \mathcal{S}} \left\{ \left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| > \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*} \right) \|\beta\|_2 \right\} \right] \\
&\geq 1 - \sum_{i \in \mathcal{S}} \Pr \left[\left| \sum_{\substack{j \in \mathcal{S} \\ j \neq i}} X_i^\top X_j \beta_j \right| > \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*} \right) \|\beta\|_2 \right] \\
&\geq 1 - \sum_{i \in \mathcal{S}} \frac{2k}{p^2} = 1 - \frac{2k^2}{p^2}.
\end{aligned} \tag{31}$$

This complete the proof of the lemma. ■

APPENDIX F

PROOF OF LEMMA 7

Once again, notice that

$$\max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| = \max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} \frac{V_i^\top}{\|V_i\|_2} \frac{V_j}{\|V_j\|_2} \beta_j \right|.$$

We next fix an $i' \in \mathcal{S}^c$. Similar to the proof of Lemma 6, the plan is to first derive a probabilistic bound on $\left| \sum_{j \in \mathcal{S}} X_{i'}^\top X_j \beta_j \right|$ and then use the union bound to provide a probabilistic bound on $\max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right|$. Using steps similar to the ones in the proof of Lemma 6, it is straightforward to establish that

$$\left| \sum_{j \in \mathcal{S}} X_{i'}^\top X_j \beta_j \right| \leq \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*} \right) \|\beta\|_2$$

with probability exceeding $1 - \frac{2(k+1)}{p^2}$. The union bound finally gives us

$$\begin{aligned}
& \Pr \left[\max_{i \in \mathcal{S}^c} \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| \leq \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*} \right) \|\beta\|_2 \right] \\
&= 1 - \Pr \left[\bigcup_{i \in \mathcal{S}^c} \left\{ \left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| > \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*} \right) \|\beta\|_2 \right\} \right] \\
&\geq 1 - \sum_{i \in \mathcal{S}^c} \Pr \left[\left| \sum_{j \in \mathcal{S}} X_i^\top X_j \beta_j \right| > \sqrt{\frac{8 \log p}{n}} \left(\frac{b_*}{\sigma_*} \right) \|\beta\|_2 \right] \\
&\geq 1 - \sum_{i \in \mathcal{S}^c} \frac{2(k+1)}{p^2} = 1 - \frac{2(k+1)(p-k)}{p^2}.
\end{aligned} \tag{32}$$

This completes the proof of the lemma. ■

REFERENCES

- [1] T. Ahmed and W. U. Bajwa, "Correlation-based ultrahigh-dimensional variable screening," in *Proc. IEEE Int. Workshop Comput. Advances Multi-Sensor Adaptive Process. (CAMSAP)*, 2017.
- [2] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Roy. Statist. Soc. B.*, vol. 58, no. 1, pp. 267–288, 1996.

- [3] D. L. Donoho and M. Elad, “Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization,” *Proc. Natl. Acad. Sci.*, vol. 100, no. 5, pp. 2197–2202, 2003.
- [4] E. J. Candès and M. B. Wakin, “An introduction to compressive sampling,” *IEEE Signal Processing Mag.*, vol. 25, no. 2, pp. 21–30, 2008.
- [5] R. Gribonval and M. Nielsen, “Sparse representations in unions of bases,” *IEEE Trans. Inform. Theory*, vol. 49, no. 12, pp. 3320–3325, 2003.
- [6] A. Cohen, W. Dahmen, and R. DeVore, “Compressed sensing and best k -term approximation,” *J. Amer. Math. Soc.*, vol. 22, no. 1, pp. 211–231, 2009.
- [7] E. J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [8] E. J. Candès, “The restricted isometry property and its implications for compressed sensing,” *Comptes Rendus Mathématique*, vol. 346, no. 9–10, pp. 589–592, 2008.
- [9] T. R. Golub et al., “Molecular classification of cancer: Class discovery and class prediction by gene expression monitoring,” *Science*, vol. 286, no. 5439, pp. 531–537, 1999.
- [10] X. Huang and W. Pan, “Linear regression and two-class classification with gene expression data,” *Bioinformatics*, vol. 19, no. 16, pp. 2072–2078, 2003.
- [11] T. Helleputte and P. Dupont, “Partially supervised feature selection with regularized linear models,” in *Proc. 26th Int’l Conf. Machine Learning*, 2009, pp. 409–416.
- [12] F. C. Stingo, Y. A. Chen, M. G. Tadesse, and M. Vannucci, “Incorporating biological information into linear models: A Bayesian approach to the selection of pathways and genes,” *Ann. Appl. Stat.*, vol. 5, no. 3, 2011.
- [13] Y. Yasui et al., “A data-analytic strategy for protein biomarker discovery: Profiling of high-dimensional proteomic data for cancer detection,” *Biostatistics*, vol. 4, no. 3, pp. 449–463, 2003.
- [14] R. Clarke et al., “The properties of high-dimensional data spaces: Implications for exploring gene and protein expression data,” *Nature Rev. Cancer*, vol. 8, no. 1, pp. 37–49, 2008.
- [15] A. I. Nesvizhskii, “A survey of computational methods and error rate estimation procedures for peptide and protein identification in shotgun proteomics,” *J. Proteomics*, vol. 73, no. 11, pp. 2092–2123, 2010.
- [16] B. Pang and L. Lee, “Opinion mining and sentiment analysis,” *Foundations and Trends in Information Retrieval*, vol. 2, no. 1–2, pp. 1–135, 2008.
- [17] B. Liu, “Sentiment analysis and opinion mining,” *Synthesis Lectures on Human Language Technologies*, vol. 5, no. 1, pp. 1–167, 2012.
- [18] R. Feldman, “Techniques and applications for sentiment analysis,” *Commun. ACM.*, vol. 56, no. 4, pp. 82–89, 2013.
- [19] D. Landgrebe, “Hyperspectral image data analysis,” *IEEE Signal Processing Mag.*, vol. 19, no. 1, pp. 17–28, 2002.
- [20] A. Plaza et al., “Recent advances in techniques for hyperspectral image processing,” *Remote Sens. Environ.*, vol. 113, pp. S110–S122, 2009.
- [21] J. Bioucas-Dias et al., “Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches,” *IEEE J. Select. Topics Appl. Earth Observ. Remote Sensing*, vol. 5, no. 2, pp. 354–379, 2012.
- [22] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning*. Springer Texts in Statistics, 2013.
- [23] H. Zou and T. Hastie, “Regularization and variable selection via the elastic net,” *J. Roy. Statist. Soc.*, vol. 67, no. 2, pp. 301–320, 2005.
- [24] W. J. Fu, “Penalized regressions: The bridge versus the lasso,” *J. Computat. Graph. Statist.*, vol. 7, no. 3, pp. 397–416, 1998.
- [25] J. Huang, J. L. Horowitz, and S. Ma, “Asymptotic properties of bridge estimators in sparse high-dimensional regression models,” *Ann. Stat.*, vol. 36, no. 2, pp. 587–613, 2008.
- [26] H. Zou, “The adaptive lasso and its oracle properties,” *J. Am. Statist. Ass.*, vol. 101, no. 476, pp. 1418–1429, 2006.
- [27] M. Yuan and Y. Lin, “Model selection and estimation in regression with grouped variables,” *J. Royal Statistical Soc. B*, vol. 68, no. 1, pp. 49–67, 2006.
- [28] E. Candès and T. Tao, “The Dantzig selector: Statistical estimation when p is much larger than n ,” *Ann. Stat.*, vol. 35, no. 6, pp. 2313–2351, 2007.
- [29] D. L. Donoho, “High-dimensional data analysis: The curses and blessings of dimensionality,” *AMS Math Challenges Lecture*, vol. 1, p. 32, 2000.
- [30] —, “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, 2006.

- [31] C. R. Genovese, J. Jin, L. Wasserman, and Z. Yao, “A comparison of the lasso and marginal regression,” *J. Machine Learning Res.*, vol. 13, pp. 2107–2143, 2012.
- [32] J. Fan and J. Lv, “Sure independence screening for ultrahigh dimensional feature space,” *J. Roy. Statist. Soc. B.*, vol. 70, no. 5, pp. 849–911, 2008.
- [33] J. Fan, R. Samworth, and Y. Wu, “Ultrahigh dimensional feature selection: Beyond the linear model,” *J. Machine Learning Res.*, vol. 10, pp. 2013–2038, 2009.
- [34] J. Fan and R. Song, “Sure independence screening in generalized linear models with NP-dimensionality,” *Ann. Stat.*, vol. 38, no. 6, pp. 3567–3604, 2010.
- [35] J. Fan, Y. Feng, and R. Song, “Nonparametric independence screening in sparse ultra-high dimensional additive models,” *J. Am. Statist. Ass.*, vol. 106, no. 494, pp. 544–557, 2011.
- [36] J. Fan, Y. Ma, and W. Dai, “Nonparametric independence screening in sparse ultra-high dimensional varying coefficient models,” *J. Am. Statist. Ass.*, vol. 109, no. 507, pp. 1270–1284, 2014.
- [37] J. Fan, Y. Feng, and Y. Wu, “High-dimensional variable selection for Cox’s proportional hazards model,” in *Borrowing Strength: Theory Powering Applications—A Festschrift for Lawrence D. Brown*. Institute of Mathematical Statistics, 2010, pp. 70–86.
- [38] S. D. Zhao and Y. Li, “Principled sure independence screening for Cox models with ultra-high-dimensional covariates,” *J. Multivariate Anal.*, vol. 105, no. 1, pp. 397–411, 2012.
- [39] P. Hall and H. Miller, “Using generalized correlation to effect variable selection in very high dimensional problems,” *J. Computat Graph. Statist.*, vol. 18, no. 3, pp. 533–550, 2009.
- [40] G. Li, H. Peng, J. Zhang, and L. Zhu, “Robust rank correlation based screening,” *Ann. Stat.*, vol. 40, no. 3, pp. 1846–1877, 2012.
- [41] L. E. Ghaoui, V. Viallon, and T. Rabbani, “Safe feature elimination for the lasso and sparse supervised learning problems,” *arXiv preprint arXiv:1009.4219*, 2010.
- [42] R. Tibshirani et al., “Strong rules for discarding predictors in lasso-type problems,” *J. Roy. Statist. Soc. B.*, vol. 74, no. 2, pp. 245–266, 2012.
- [43] Z. J. Xiang and P. J. Ramadge, “Fast lasso screening tests based on correlations,” in *Proc. IEEE Int. Conf. on Acoustics Speech and Sig. Proc. (ICASSP)*, 2012, pp. 2137–2140.
- [44] L. Dai and K. Pelckmans, “An ellipsoid based, two-stage screening test for BPDN,” in *Proc. 20th Eur. Sig. Proc. Conf.*, 2012, pp. 654–658.
- [45] J. Wang, P. Wonka, and J. Ye, “Lasso screening rules via dual polytope projection,” *J. Mach. Learn. Res.*, vol. 16, pp. 1063–1101, 2015.
- [46] Z. J. Xiang, Y. Wang, and P. J. Ramadge, “Screening tests for lasso problems,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 5, pp. 1008–1027, 2017.
- [47] M. J. Wainwright, “Sharp thresholds for high-dimensional and noisy sparsity recovery using-constrained quadratic programming (lasso),” *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2183–2202, 2009.
- [48] P. Zhao and B. Yu, “On model selection consistency of lasso,” *J. Machine Learning Res.*, vol. 7, pp. 2541–2563, 2006.
- [49] G. Raskutti, M. J. Wainwright, and B. Yu, “Restricted eigenvalue properties for correlated Gaussian designs,” *J. Machine Learning Res.*, vol. 11, pp. 2241–2259, 2010.
- [50] M. Rudelson and R. Vershynin, “Smallest singular value of a random rectangular matrix,” *Commun Pure Appl. Math.*, vol. 62, no. 12, pp. 1707–1739, 2009.
- [51] W. U. Bajwa, R. Calderbank, and S. Jafarpour, “Why Gabor frames? Two fundamental measures of coherence and their role in model selection,” *J. Commun. Netw.*, vol. 12, no. 4, pp. 289–307, 2010.
- [52] G. Davis, S. Mallat, and M. Avellaneda, “Adaptive greedy approximations,” *J. Construct. Approx.*, vol. 13, no. 1, pp. 57–98, 1997.
- [53] W. U. Bajwa, A. Calderbank, and D. G. Mixon, “Two are better than one: Fundamental parameters of frame coherence,” *Appl. Comput. Harmon. Anal.*, vol. 33, pp. 58–78, 2012.
- [54] L. Welch, “Lower bounds on the maximum cross correlation of signals,” *IEEE Trans. Inform. Theory*, vol. 20, no. 3, pp. 397–399, 1974.
- [55] A. L. Maas et al., “Learning word vectors for sentiment analysis,” in *Proc. 49th Ann. Meeting of the Assoc. for Computational Linguistics: Human Language Technologies*, June 2011, pp. 142–150.
- [56] J. Fan and Y. Fan, “High dimensional classification using features annealed independence rules,” *Ann. Stat.*, vol. 36, no. 6, p. 2605, 2008.
- [57] J. Fan, F. Han, and H. Liu, “Challenges of big data analysis,” *Nat. Sci. Rev.*, vol. 1, no. 2, pp. 293–314, 2014.
- [58] C. D. Manning, P. Raghavan, and H. Schütze, *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- [59] J. Wainwright, “High-dimensional statistics: A non-asymptotic viewpoint,” *In preparation. University of California, Berkeley*, 2015. [Online]. Available: https://www.stat.berkeley.edu/~mjwain/stat210b/Chap2_TailBounds_Jan22_2015.pdf