

Compressed Anomaly Detection with Multiple Mixed Observations

Natalie Durgin, Rachel Grotheer, Chenxi Huang, Shuang Li, Anna Ma,
Deanna Needell and Jing Qin

Abstract We consider a collection of independent random variables that are identically distributed, except for a small subset which follows a different, anomalous distribution. We study the problem of detecting which random variables in the collection are governed by the anomalous distribution. Recent work proposes to solve this problem by conducting hypothesis tests based on mixed observations (e.g. linear combinations) of the random variables. Recognizing the connection between taking mixed observations and compressed sensing, we view the problem as recovering the “support” (index set) of the anomalous random variables from multiple mea-

Natalie Durgin
Spiceworks
Austin, TX, 78746.
e-mail: njdurgin@gmail.com

Rachel Grotheer
Goucher College
Baltimore, MD 21204.

Chenxi Huang
Yale University
New Haven, CT 06511.
e-mail: chenxi.huang@yale.edu

Shuang Li
Colorado School of Mines
Golden, CO 80401.

Anna Ma
Claremont Graduate University
Claremont, CA 91711.

Deanna Needell
University of California, Los Angeles
Los Angeles, CA 90095.

Jing Qin
Montana State University
Bozeman, MT 59717.

surement vectors (MMVs). Many algorithms have been developed for recovering jointly sparse signals and their support from MMVs. We establish the theoretical and empirical effectiveness of these algorithms in detecting anomalies. We also extend the LASSO algorithm to an MMV version for our purpose. Further, we perform experiments on synthetic data, consisting of samples from the random variables, to explore the trade-off between the number of mixed observations per sample and the number of samples required to detect anomalies.

1 Introduction

The problem of anomaly detection has been the focus of interest in many fields of science and engineering including network tomography, cognitive radio, and radar [35, 39, 26, 4]. In this paper, we study the problem of identifying a small number of anomalously distributed random variables within a much larger collection of independent and otherwise identically distributed random variables. We call the random variables following the anomalous distribution *anomalous random variables*. A conventional approach to detect these anomalous random variables is to sample from each random variable individually and then apply hypothesis testing techniques [28, 29, 30].

A recent paper [13] proposes to perform hypothesis testing on mixed observations (e.g. linear combinations) of random variables instead of on samples from individual random variables. They call this technique *compressed hypothesis testing*. Such an approach is motivated by the recent development of *compressed sensing* [7, 15, 22, 20], a signal processing paradigm that shows a small number of random linear measurements of a signal is sufficient for accurate reconstruction. Now a large body of work in this area shows that optimization-based [16, 9, 15, 8, 36] and iterative [37, 32, 6] methods can reconstruct the signal accurately and efficiently when the samples are taken via a sensing matrix satisfying certain incoherence properties [9, 10]. Compressed sensing is also studied in a Bayesian framework, where signals are assumed to obey some prior distribution [24, 40, 2].

The results presented in [13] show that the “mixed” measurement approach achieves better detection accuracy from fewer samples when compared to the conventional “un-mixed” approach. However, compressed hypothesis testing requires that the distributions of the random variables are known *a priori*, which may not be available in practice. Further, as the authors pointed out, their proposed approach requires conducting a large number of hypothesis tests, especially when the number of random variables in the collection is large, rendering such an approach computationally prohibitive. Two efficient algorithms are proposed as alternatives in [13], but no analytical study of their performance is provided.

We propose new methods for detecting anomalous random variables that require minimal knowledge of the distributions, are computationally efficient, and whose performance is easy to characterize. We begin by generalizing the compressed hypothesis testing method and posing our problem as a *multiple measurement vector*

(MMV) problem [23, 3, 17, 18, 11, 14, 31, 5]. In the MMV compressed sensing setting, a collection of signals are recovered simultaneously, under the assumption that they have some commonalities, such as sharing the same support. A related vein of work involves signals that are smoothly varying, where the support may not be consistent but changes slowly over time [1, 21, 34]. While the compressed hypothesis testing in [13] is certainly motivated by compressed sensing techniques, the authors do not formally frame the anomaly detection problem in the compressed sensing setting. Also, they do not focus on compressed sensing algorithms that might eliminate the need for prior knowledge of the distributions, and might lead to more efficient detection for large collections of random variables.

In the following, we view the collection of random variables as a random vector and aim to identify the indices of the anomalous random variables within the random vector. We also draw an analogy between the collection of independent samples from the random vector and an ensemble of signals where in practice these signals often become available over time. More specifically, we consider a random vector, $X = (X_1, \dots, X_N)$, where the X_n 's are independent random variables. We assume that each X_n follows one of the two distributions, $\mathcal{D}_1, \mathcal{D}_2$. We call \mathcal{D}_1 the *prevalent distribution*, and \mathcal{D}_2 the *anomalous distribution*. We let $\mathfrak{N} = \{n \in \mathbb{N} : 1 \leq n \leq N\}$ denote the index set of the random variables, X_n , and let \mathfrak{K} denote the index set of the K random variables that follow the anomalous distribution. Let $x_{(\cdot,t)} \in \mathbb{R}^N$ denote the independent realization of the random vector at time t . At each time-step t , we obtain M mixed observations by applying the sensing matrix $\phi_t \in \mathbb{R}^{M \times N}$,

$$y_t = \phi_t x_{(\cdot,t)}, \quad 1 \leq t \leq T,$$

with $y_t \in \mathbb{R}^M$. Thus the goal of the anomaly detection problem in this setting is to recover the index set \mathfrak{K} from the MMVs $y_t, t = 1, \dots, T$.

The signals $x_{(\cdot,t)}$ in our formulation are not necessarily sparse and may have different supports since they are samples from a random vector and are changing over time. Nevertheless, there is still a close connection between our formulation and that for recovering the common sparse support of a collection of signals from MMVs. The index set of the anomalous random variables, which corresponds to the index set of the anomalies (realizations of anomalous random variables) in the signals $x_{(\cdot,t)}$, is shared by all signals. This index set can thus be viewed as the common ‘‘support’’ of the anomalies in the signals, which motivates us to consider the applicability of many MMV algorithms designed for signal reconstruction. Further, the analytical studies of many of these algorithms are readily available. We therefore investigate which of these MMV algorithms can be applied or adapted to the anomaly detection problem under consideration and analyze their performance in detection accuracy in theory and through numerical experiments. We focus on algorithms presented in [3].

1.1 Contributions

In this paper, by extending the definitions of two so-called *joint sparsity models* (JSMs) from [3], we introduce two new signal models, JSM-2R and JSM-3R, for the problem of anomaly detection. For JSM-2R and JSM-3R signals, we adapt several MMV signal reconstruction algorithms to anomaly detection. Additionally, we develop a new algorithm for the JSM-2R model that extends the *Least Absolute Shrinkage and Selection Operator* (LASSO) algorithm [12] to the MMV framework. We show theoretically and numerically that these algorithms accurately detect the anomalous random variables. We also provide numerical results which demonstrate the trade-off between the number of time-steps, and the number of mixed observations per time-step needed to detect the anomalies.

1.2 Organization

In Section 2, we introduce the models JSM-2R, JSM-3R and the four algorithms we have repurposed from MMV signal recovery into MMV anomaly detection, as well as our new LASSO algorithm. We also provide theoretical guarantees in this section. In Section 3, we explore the performance of these algorithms by conducting numerical experiments for some strategic choices of the parameters involved. Finally, we conclude in Section 4. To help keep track of notation, we provide a handy reference table in the appendix. We adopt the convention that random variables will be upper case and their realizations will be lower case. All matrix entries will have two, subscripted indices. The first index will indicate the row position, the second will indicate the column position.

2 Method

In this section, we introduce two new signal models for the anomaly detection problem and describe five algorithms for detecting anomalous random variables under these signal models. We also provide theoretical guarantees for the algorithms.

Recall that we consider the problem of detecting K anomalous random variables from a collection of N random variables where $K \ll N$. The anomalous random variables have a different probability distribution from that of the remaining $N - K$ random variables. We seek to identify the K anomalous random variables, from T independent realizations of the N random variables. To emphasize our framing of this random variable problem as a compressed sensing problem, we refer to the independent realizations as signals. These T signals have an important commonality: they share the same indices of anomalous entries (realizations of anomalous random variables).

Commonality among signals has already been explored in the field of distributed compressed sensing for recovering signals that have specific correlation among them. Three *joint sparsity models* (JSMs) were introduced in [3] to characterize different correlation structures. To utilize the commonality of the signals for anomaly detection, we propose two new signal models that are motivated by two of the JSMs defined in [3], namely, JSM-2 and JSM-3. Since the signals under consideration are realizations of random variables, we term the new models JSM-2R and JSM-3R, respectively, where the appended “R” indicates the “random variable” version of the existing JSMs.

Before we define the new models, we first briefly describe JSM-2 and JSM-3. The JSM-2 signals are jointly sparse signals that share the same support (the indices of non-zero entries). The JSM-3 signals consist of two components: a non-sparse “common component” shared by all signals and a sparse “innovation component” that is different for each signal. But the innovation components of the JSM-3 signals share the same support. We next extend these definitions to the signals in the anomaly detection setting. The new JSM-2R and JSM-3R models are defined as follows:

Definition 1 (JSM-2R and JSM-3R). Let the random variable $X_n \sim \mathcal{D}_1$ if $n \notin \mathfrak{K}$ and $X_n \sim \mathcal{D}_2$ if $n \in \mathfrak{K}$ where \mathfrak{K} is the set of the anomalous indices. For a signal ensemble $x \in \mathbb{R}^{N \times T}$ where each of its entries $x_{(n,t)}$ denotes the realization of X_n at time t ,

1. x is a JSM-2R signal ensemble when: $|x_{(n,t)}|$ is small if $n \notin \mathfrak{K}$ and $|x_{(n,t)}|$ is large if $n \in \mathfrak{K}$;
2. x is a JSM-3R signal ensemble when: $x_{(n,t)} = x_n^C + x_{(n,t)}^I$ such that $|x_{(n,t)}^I|$ is small if $n \notin \mathfrak{K}$ and $|x_{(n,t)}^I|$ is large if $n \in \mathfrak{K}$. x_n^C is a common component shared by all t , and $x_{(n,t)}^I$ is an innovation component that is different for different t .

The JSM-2R signal model assumes a small amplitude for variables generated from the prevalent distribution and a large amplitude for variables generated from the anomalous distribution. Such a model characterizes a scenario where anomalies exhibit large spikes. This model relates to a sparse signal model where the support of the sparse signal corresponds to the set of indices of the anomalous random variables. In fact, when $\mathcal{D}_1 = \mathcal{N}(0, \sigma^2)$ and $\mathcal{D}_2 = \mathcal{N}(\mu, \sigma^2)$ with $\mu \neq 0$, the JSM-2R signal is a sparse signal with additive Gaussian noise. An example of anomalies following the JSM-2R model is a network where some of the sensors completely malfunction and produce signals with vastly different amplitudes from the rest of the sensors.

Different from the JSM-2R signals, the JSM-3R signal model introduced above does not have constraints on the amplitude of the signal entries $x_{(n,t)}$. Rather the signals at different time-steps are assumed to share an unknown common component x_n^C while having a different innovation component $x_{(n,t)}^I$ for signals at different time-steps. Of note, the common component x_n^C from the prevalent distribution may or may not be the same as that from the anomalous distribution. Further, the innovation component $x_{(n,t)}^I$ is assumed to follow the JSM-2R signal model. Such a model

characterizes a scenario where there exists a background signal that does not change over time and the anomalies exhibit large spikes on top of the background signal. Because of the common component, the JSM-3R signals no longer correspond to a sparse signal model. The JSM-3R model has applications in geophysical monitoring where a constant background signal is present and anomalies appear as large spikes of erratic behavior. Figure 1 provides a visual illustration of the model nuances.

2.1 Algorithms

We briefly describe the five algorithms we study in this paper, among which three are for JSM-2R signals and two are for JSM-3R signals. Two of the algorithms for JSM-2R signals were originally proposed for recovering JSM-2 signals, including the *one-step greedy algorithm* (OSGA) and the *multiple measurement vector simultaneous orthogonal matching pursuit* (MMV-SOMP) algorithm. We further propose a new MMV version of the LASSO algorithm for detecting anomalies for JSM-2R signals and investigate its performance via numerical experiments. The two algorithms for JSM-3R were also proposed in [3] for recovering JSM-3 signals, including the *Transpose Estimation of Common Component* (TECC) algorithm and the *Alternating Common and Innovation Estimation* (ACIE) algorithm.

For each of the presented algorithms, the goal is to identify the indices of the anomalous random variables from the mixed measurements $y_t = \phi_t x_{(\cdot,t)}$ for $t = 1, 2, \dots, T$. The number of anomalies K is assumed to be known *a priori*. We first describe the three algorithms that we applied to solve the anomaly detection problem for JSM-2R signals.

2.1.1 OSGA

The OSGA is a non-iterative greedy algorithm introduced in [3] to recover the support of JSM-2 signals based on inner products of the measurement and columns of the sensing matrix (Algorithm 1). We show in Theorem 1 that with some condition on the prevalent and anomalous distributions, the OSGA is able to recover the anomaly indices under the JSM-2R model, using a small number of measurements per time-step. Although the OSGA is shown to work asymptotically, it may not perform well when only a small number of time-steps are available. Empirical evidence has confirmed this conjecture when the OSGA is used to reconstruct JSM-2 signals [3]. Thus we further consider approaches like matching pursuit [27, 33] for our problem. Next, we describe the MMV version of orthogonal matching pursuit algorithm proposed in [3].

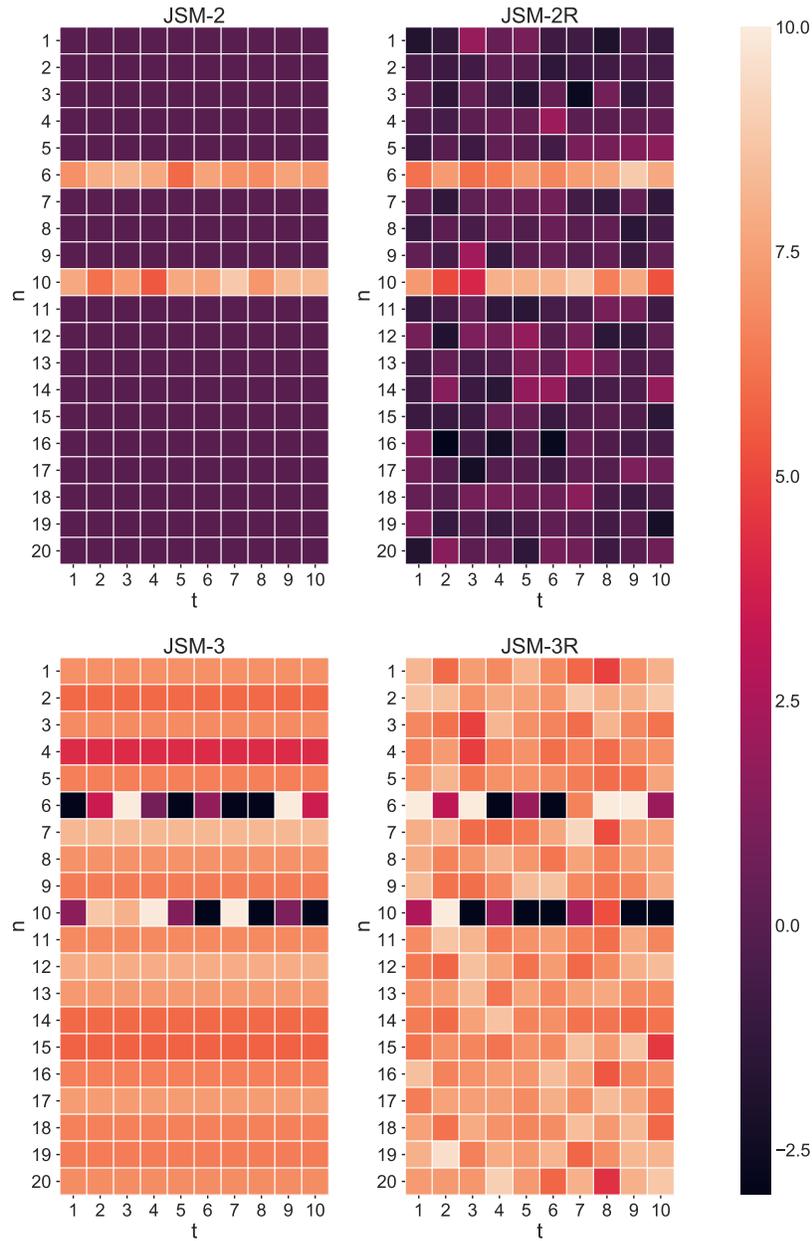


Fig. 1 Illustration of the existing joint sparsity models (JSM-2 and JSM-3) and our proposed models developed for anomaly detection (JSM-2R and JSM-3R). The distributions used to generate this example are the same as the ones used for the numerical experiments in Section 3, see Table 1. The index set of the anomalies is $\mathfrak{R} = \{6, 10\}$.

Algorithm 1 OSGA

```

1: Input:  $y_1, \dots, y_T, \phi_t, K$ .
2: Output:  $\hat{\mathcal{R}}$ .
3: for  $n = 1, 2, \dots, N$  do
4:   Compute  $\xi_n = \frac{1}{T} \sum_{t=1}^T \langle y_t, \phi_{t(\cdot, n)} \rangle^2$ 
5: end for
6: return  $\hat{\mathcal{R}} = \{n, \text{ for the } K \text{ largest } \xi_n\}$ 

```

2.1.2 MMV-SOMP

The MMV-SOMP algorithm is an iterative greedy pursuit algorithm for recovery of jointly sparse signals. SOMP was first proposed in [38] and was adapted to the MMV framework in [3]. Since our focus is not on signal recovery but on detecting anomalous entries, we adapt this algorithm for our JSM-2R signal model. The adapted algorithm is presented in Algorithm 2, which identifies the anomaly indices one at a time. In each iteration, the column index of the sensing matrices that accounts for the largest residual across signals of all time-steps is selected. Then the remaining columns of each sensing matrix (for each time-step) are orthogonalized. The algorithm stops after K iterations where K is the number of anomalous random variables. We show through numerical experiments in Section 3 that the adapted MMV-SOMP algorithm performs better than the OSGA for a small number of time-steps.

2.1.3 MMV-LASSO

The LASSO algorithm aims to find a sparse solution to the regression problem by constraining the L_1 norm of the solution [12]. The LASSO algorithm was also considered in [13] as an efficient algorithm for anomaly detection from mixed observations. However, the authors of [13] considered the LASSO algorithm when using only one measurement at each time-step. In this paper, we further extend the LASSO algorithm to a more general setting for MMV and term it the MMV-LASSO algorithm. The MMV-LASSO algorithm is described in Algorithm 3. The measurements $y_t \in \mathbb{R}^M$ up to T time-steps are concatenated vertically to become a vector $y \in \mathbb{R}^{(MT) \times 1}$; the sensing matrices $\phi_t \in \mathbb{R}^{M \times N}$ are also concatenated vertically to become $\phi \in \mathbb{R}^{(MT) \times N}$. The concatenated measurements and sensing matrices are then fed to the regular LASSO algorithm, where the anomaly indices are found by taking indices corresponding to the K largest amplitudes of the estimate. The LASSO problem, that is, Step 4 in Algorithm 3, can be tackled by various approaches [19, 25], which is out of scope of this paper.

We next describe two algorithms for anomaly detection for JSM-3R signals.

Algorithm 2 MMV-SOMP

-
- 1: **Input:** $y_1, \dots, y_T, \phi_t, K$.
 - 2: **Output:** $\widehat{\mathfrak{R}}$.
 - 3: **Initialize:** $\widehat{\mathfrak{R}} = \emptyset$, residual $r_t^0 = y_t$.
 - 4: **for** $k = 1, \dots, K$ **do**.
 - 5: Select

$$n_k = \underset{n}{\operatorname{arg\,max}} \sum_{t=1}^T \frac{|\langle r_t^{k-1}, \phi_{t(\cdot, n)} \rangle|}{\|\phi_{t(\cdot, n)}\|_2}.$$

- 6: Update $\widehat{\mathfrak{R}} = [\widehat{\mathfrak{R}}, n_k]$.
- 7: Orthogonalize selected basis vector against previously selected vectors for all $t, 1 \leq t \leq T$:

$$\begin{aligned} \gamma_t^0 &= \phi_{t(\cdot, k)}, & \text{if } k = 1, \\ \gamma_t^k &= \phi_{t(\cdot, n_k)} - \sum_{l=0}^{k-1} \frac{\langle \phi_{t(\cdot, n_k)}, \gamma_t^l \rangle}{\|\gamma_t^l\|_2^2} \gamma_t^l, & \text{if } k > 1. \end{aligned}$$

- 8: Update the residual for all $t, 1 \leq t \leq T$,

$$r_t^k = r_t^{k-1} - \frac{\langle r_t^{k-1}, \gamma_t^k \rangle}{\|\gamma_t^k\|_2^2} \gamma_t^k.$$

- 9: **end for**
 - 10: **return** $\widehat{\mathfrak{R}}$
-

Algorithm 3 MMV-LASSO

-
- 1: **Input:** $y_1, \dots, y_T, \phi_t, K$.
 - 2: **Output:** $\widehat{\mathfrak{R}}$.
 - 3: Let $y = [y_1^\top, \dots, y_T^\top]^\top$ and $\phi = [\phi_1^\top, \dots, \phi_T^\top]^\top$
 - 4: Solve

$$\hat{x} = \underset{x}{\operatorname{arg\,min}} \frac{1}{2} \|y - \phi x\|_2^2 + \lambda \|x\|_1$$

- 5: Let \hat{x}_n denote the n -th element of \hat{x}
 - 6: **return** $\widehat{\mathfrak{R}} = \{n, \text{ for the } K \text{ largest } |\hat{x}_n|\}$
-

2.1.4 TECC

The key difference between JSM-2R and JSM-3R signals is that JSM-3R signals share a common component that is unknown. Thus the two algorithms for the JSM-3R signals aim to first estimate the common component from the mixed measurement and subtract the contribution of this component from the measurement. The TECC algorithm was proposed in [3] for recovering JSM-3 signals. We also adapt the algorithm to focus only on detecting the anomalous indices of JSM-3R signals, and the adapted algorithm can be found in Algorithm 4. The first step of the TECC algorithm estimates the common component of the JSM-3R signals. Using this estimate, the contribution of the remaining innovation component to the measurement can be estimated. Then algorithms for JSM-2R signals can be applied to identify the

anomaly indices. We show in Theorem 2 that the TECC algorithm is able to identify the anomalous variables under some conditions on the prevalent and anomalous distributions. Similar to the OSGA, while Theorem 2 guarantees the success of the TECC algorithm in the asymptotic case as T goes to infinity, it may not perform well for a small T . Next we describe an alternative algorithm also proposed in [3] for cases with a small T .

Algorithm 4 TECC

- 1: **Input:** $y_1, \dots, y_T, \phi_t, K$.
 - 2: **Output:** $\hat{\mathfrak{K}}$.
 - 3: Let $y = [y_1^\top, \dots, y_T^\top]^\top$, and $\phi = [\phi_1^\top, \dots, \phi_T^\top]^\top$
 - 4: Calculate $\hat{x}^C = \frac{1}{TM} \phi^\top y$.
 - 5: Calculate $\hat{y}_t = y_t - \phi_t \hat{x}^C$.
 - 6: Estimate $\hat{\mathfrak{K}}$ from \hat{y}_t by Algorithm 1, 2 or 3
 - 7: **return** $\hat{\mathfrak{K}}$
-

2.1.5 ACIE

The ACIE algorithm is an extension of the TECC algorithm, also introduced in [3], based on the observation that the initial estimate of the common component may not be sufficiently accurate for subsequent steps. Instead of one-time estimation in the TECC algorithm, the ACIE algorithm iteratively refines the estimates of the common component and the innovation components. The ACIE algorithm can also be easily adapted for the JSM-3R signals for anomaly detection. In the ACIE algorithm described in Algorithm 5, we first obtain an initial estimate of the anomaly index set $\hat{\mathfrak{K}}$ using the TECC algorithm. Then for each iteration, we build a basis B_t for \mathbb{R}^M where M is the number of measurements at each time-step: $B_t = [\phi_{t, \hat{\mathfrak{K}}}, q_t]$, where $\phi_{t, \hat{\mathfrak{K}}}$ is the subset of the basis vectors in ϕ_t corresponding to the indices in $\hat{\mathfrak{K}}$ and q_t has orthonormal columns that spans the orthogonal complement of $\phi_{t, \hat{\mathfrak{K}}}$. Then we can project the measurements onto q_t to obtain the part of the measurement caused by signals not in $\hat{\mathfrak{K}}$:

$$\tilde{y}_t = q_t^\top y_t, \quad (1)$$

$$\tilde{\phi}_t = q_t^\top \phi_t. \quad (2)$$

Then \tilde{y}_t and $\tilde{\phi}_t$ are used to refine the estimate of the common component. After subtracting the contribution of this estimated common component, algorithms such as OSGA and MMV-SOMP described above can be applied to detect the anomalies.

Algorithm 5 ACIE

1: **Input:** $y_1, \dots, y_T, \phi_t, K, L$ (iteration counter).
2: **Output:** $\hat{\mathfrak{K}}$.
3: Let $y = [y_1^\top, \dots, y_T^\top]^\top$
4: Obtain an initial estimate of $\hat{\mathfrak{K}}$ from Algorithm 4
5: **for** $l = 1, 2, \dots, L$ **do**
6: Update \tilde{y}_t and $\tilde{\phi}_t$ according to Equations (1) and (2) for all $t, 1 \leq t \leq T$
7: Update $\tilde{x}^C = \tilde{\phi}^\dagger \tilde{y}$, where $\tilde{y} = [\tilde{y}_1^\top, \dots, \tilde{y}_T^\top]^\top$, $\tilde{\phi} = [\tilde{\phi}_1^\top, \dots, \tilde{\phi}_T^\top]$ and $\tilde{\phi}^\dagger = (\tilde{\phi}^\top \tilde{\phi})^{-1} \tilde{\phi}^\top$
8: **end for**
9: Calculate $\hat{y}_t = y_t - \phi_t \tilde{x}^C$
10: Estimate $\hat{\mathfrak{K}}$ from \hat{y}_t by Algorithm 1, 2 or 3
11: **return** $\hat{\mathfrak{K}}$

2.2 Theoretical Guarantees

In this section we show theoretically that Algorithm 1 and Algorithm 4 (coupled with Algorithm 1 in step 6) can detect anomalies for the JSM-2R and JSM-3R settings, respectively.

Recall that Algorithm 1 is designed for JSM-2R signals where variables generated from the prevalent distribution are much smaller in amplitude than those from the anomalous distribution. The following theorem shows that for JSM-2R signals, the OSGA is able to identify the indices of the anomalous variables asymptotically, with very few measurements at each time-step.

Theorem 1. [Adapted from [3] Theorem 8] *Let the $M \times N$ sensing matrix, ϕ_t , contain entries that are i.i.d $\sim \mathcal{N}(0, 1)$ at each time-step t . Suppose the random variables, X_n , are distributed with $\mathcal{D}_1 = \mathcal{N}(0, \sigma_1^2)$ if $n \notin \mathfrak{K}$ and $\mathcal{D}_2 = \mathcal{N}(\mu_2, \sigma_2^2)$ if $n \in \mathfrak{K}$. Assuming $\mu_2^2 + \sigma_2^2 > \sigma_1^2$, then with $M \geq 1$ measurements per time-step, OSGA recovers \mathfrak{K} with probability approaching one as $T \rightarrow \infty$.*

Before diving into the proof of Theorem 1, we first observe that the signals correspond to the JSM-2R signals: with a zero mean and a potentially small variance σ_1^2 for the prevalent distribution \mathcal{D}_1 , the signal entry $x_{(n,t)}$, $n \notin \mathfrak{K}$ (i.e. the realization of X_n at the time-step t) is expected to have small amplitude. In contrast, with a non-zero mean μ_2 and a similar or possibly larger variance σ_2^2 for the anomalous distribution \mathcal{D}_2 , the amplitude of $x_{(n,t)}$, $n \in \mathfrak{K}$ is expected to be much larger.

Proof. We assume, for convenience and without loss of generality, that the anomalous random variables are indexed by, $\mathfrak{K} = \{1, 2, \dots, K\}$, and the prevalent random variables are indexed by $\mathfrak{N} \setminus \mathfrak{K} = \{K+1, \dots, N\}$. Consider that the test statistic $\xi_n = \frac{1}{T} \sum_{t=1}^T \langle y_t, \phi_{t(\cdot, n)} \rangle^2$ is the sample mean of the random variable $\langle Y, \Phi_{(\cdot, n)} \rangle^2$, so by the Law of Large Numbers,

$$\lim_{T \rightarrow \infty} \xi_n = \mathbb{E}[\langle Y, \Phi_{(\cdot, n)} \rangle^2].$$

We select an arbitrary index n from each of the anomalous random variable index set and the prevalent random variable index set, and compute $\mathbb{E}[\langle Y, \Phi_{(\cdot, n)} \rangle^2]$ in each

case. As the final step, we compare the expected values of the two ξ_n and establish that they are distinguishable under very general conditions. Without loss of generality, we select $n = K + 1$ for the “prevalent case” and $n = 1$ for the “anomalous case”. Note that [3] refers to these cases respectively as the “bad statistics” and the “good statistics” in their setting. For them, “bad” reflects an incorrect estimate of the sparse support and “good” reflects a correct estimate of the sparse support.

Prevalent Case: Substituting ΦX for Y in $\langle Y, \Phi_{(\cdot, K+1)} \rangle$ and rearranging we obtain $\langle Y, \Phi_{(\cdot, K+1)} \rangle = \sum_{n=1}^N X_n \langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle$. We can then write,

$$\begin{aligned} \mathbb{E}[\langle Y, \Phi_{(\cdot, K+1)} \rangle^2] &= \mathbb{E} \left[\left(\sum_{n=1}^N X_n \langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle \right)^2 \right] \\ &= \mathbb{E} \left[\sum_{n=1}^N (X_n)^2 \langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle^2 \right] \\ &\quad + \mathbb{E} \left[\sum_{n=1}^N \sum_{\substack{l=1 \\ l \neq n}}^N X_n X_l \langle \Phi_{(\cdot, l)}, \Phi_{(\cdot, K+1)} \rangle \langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle \right] \\ &= \sum_{n=1}^N \mathbb{E}[(X_n)^2] \mathbb{E}[\langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle^2] \\ &\quad + \sum_{n=1}^N \sum_{\substack{l=1 \\ l \neq n}}^N \mathbb{E}[X_n] \mathbb{E}[X_l] \mathbb{E}[\langle \Phi_{(\cdot, l)}, \Phi_{(\cdot, K+1)} \rangle \langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle]. \end{aligned}$$

The last step follows from the independence of Φ and X and the independence of the X_n 's from each other. We claim that the cross-terms above sum to zero. To see this, we set $\Phi_{(\cdot, l)} = a$, $\Phi_{(\cdot, K+1)} = b$ and $\Phi_{(\cdot, n)} = c$, where the entries of the vectors a, b, c are all i.i.d. $\mathcal{N}(0, 1)$. We note that if $l, K + 1$, and n are mutually distinct, then a, b, c are mutually independent. In this case we have,

$$\begin{aligned} \mathbb{E}[\langle a, b \rangle \langle c, b \rangle] &= \mathbb{E}[a^\top b c^\top b] \\ &= \mathbb{E}[a^\top] \mathbb{E}[b c^\top b] \\ &= 0. \end{aligned}$$

Since the cross-terms assume $l \neq n$, we consider the cases when either $n = K + 1$ or $l = K + 1$. In the case where $n = K + 1$ we have,

$$\begin{aligned} \mathbb{E}[\langle a, b \rangle \langle b, b \rangle] &= \mathbb{E}[a^\top b b^\top b] \\ &= \mathbb{E}[a^\top] \mathbb{E}[b b^\top b] \\ &= 0. \end{aligned}$$

Similarly, in the case where $l = K + 1$ we have,

$$\begin{aligned}
\mathbb{E}[\langle b, b \rangle \langle c, b \rangle] &= \mathbb{E}[b^\top b c^\top b] \\
&= \mathbb{E}[c^\top b b^\top b] \\
&= \mathbb{E}[c^\top] \mathbb{E}[b b^\top b] \\
&= 0.
\end{aligned}$$

Thus, all cross-terms vanish so returning to our original goal we may claim,

$$\begin{aligned}
\mathbb{E}[\langle Y, \Phi_{(\cdot, K+1)} \rangle^2] &= \sum_{n=1}^N \mathbb{E}[(X_n)^2] \mathbb{E}[\langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle^2] \\
&= \sum_{n=1}^K \mathbb{E}[(X_n)^2] \mathbb{E}[\langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle^2] + \mathbb{E}[(X_{K+1})^2] \mathbb{E}[\|\Phi_{(\cdot, K+1)}\|^4] \\
&\quad + \sum_{n=K+2}^N \mathbb{E}[(X_n)^2] \mathbb{E}[\langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle^2].
\end{aligned}$$

Examining each expected value individually, we recall that for $n \in \{1, \dots, K\} = \mathfrak{R}$ the X_n were distributed with \mathcal{D}_2 and thus $\mathbb{E}[(X_n)^2] = \mathbb{E}[X_n]^2 + \text{Var}(X_n) = \mu_2^2 + \sigma_2^2$. Recalling that the rest of the X_n are distributed with \mathcal{D}_1 which has $\mu = 0$, we have that $\mathbb{E}[(X_n)^2] = \sigma_1^2$ in the subsequent cases. In [3] they establish that $\mathbb{E}[\|\Phi_{(\cdot, K+1)}\|^4] = M(M+2)$ and $\mathbb{E}[\langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, K+1)} \rangle^2] = M$, and we may use these results without further argument because we make the same assumptions about Φ . Finally, substituting the expected values we have just calculated, we have that as T grows large, the statistic ξ_n when $n \notin \mathfrak{R}$ converges to

$$\begin{aligned}
\mathbb{E}[\langle Y, \Phi_{(\cdot, K+1)} \rangle^2] &= K(\mu_2^2 + \sigma_2^2)M + \sigma_1^2 M(M+2) + (N-K-1)\sigma_1^2 M \\
&= M[K(\mu_2^2 + \sigma_2^2) + (M+1+N-K)\sigma_1^2]. \tag{3}
\end{aligned}$$

Anomalous Case: With $n = 1$, we proceed as in the previous case,

$$\begin{aligned}
\mathbb{E}[\langle Y, \Phi_{(\cdot, 1)} \rangle^2] &= \mathbb{E} \left[\left(\sum_{n=1}^N X_n \langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, 1)} \rangle \right)^2 \right] \\
&= \sum_{n=1}^N \mathbb{E}[(X_n)^2] \mathbb{E}[\langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, 1)} \rangle^2], \\
&= \mathbb{E}[(X_1)^2] \mathbb{E}[\|\Phi_{(\cdot, 1)}\|^4] + \sum_{n=2}^K \mathbb{E}[(X_n)^2] \mathbb{E}[\langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, 1)} \rangle^2] \\
&\quad + \sum_{n=K+1}^N \mathbb{E}[(X_n)^2] \mathbb{E}[\langle \Phi_{(\cdot, n)}, \Phi_{(\cdot, 1)} \rangle^2] \\
&= (\mu_2^2 + \sigma_2^2)M(M+2) + (K-2)(\mu_2^2 + \sigma_2^2)M + (N-K)\sigma_1^2 M \\
&= M[(M+1+K)(\mu_2^2 + \sigma_2^2) + (N-K)\sigma_1^2]. \tag{4}
\end{aligned}$$

Combining the results of (3) and (4), we have

$$\lim_{T \rightarrow \infty} \xi_n = \begin{cases} M[(M+1+K)(\mu_2^2 + \sigma_2^2) + (N-K)\sigma_1^2] & n \in \mathfrak{R} \\ M[K(\mu_2^2 + \sigma_2^2) + (M+1+N-K)\sigma_1^2] & n \notin \mathfrak{R}. \end{cases}$$

The difference in the two expectations is thus,

$$M(M+1)(\mu_2^2 + \sigma_2^2 - \sigma_1^2).$$

For any $M \geq 1$ and $\mu_2^2 + \sigma_2^2 > \sigma_1^2$, the expected value of ξ_n in the ‘‘anomalous case’’ is strictly larger than the expected value of ξ_n in the ‘‘prevalent case’’. Therefore, as T increases, OSGA can distinguish between the two expected values of ξ_n with overwhelming probability. \square

The next theorem shows that asymptotically, Algorithm 4 is able to detect anomalous variables with very few measurements at each time-step, for JSM-3R signals. Recall that JSM-3R signals have an unknown common component shared by signals at all time-steps, while each signal has a different innovation component that follows the JSM-2R model. The following theorem and proof assume that Algorithm 1 is implemented for step 6 of Algorithm 4. Once results like Theorem 1 exist for Algorithms 2 and 3, then any JSM-2R algorithm could be used in step 6, and Theorem 2 would still hold.

Theorem 2. [Adapted from [3] Theorem 10] *Let the $M \times N$ sensing matrix ϕ_t at each time-step t contain entries that are i.i.d. $\sim \mathcal{N}(0, 1)$. For random variables X_n that are distributed with $\mathcal{D}_1 = \mathcal{N}(\mu_1, \sigma_1^2)$ if $n \notin \mathfrak{R}$ and $\mathcal{D}_2 = \mathcal{N}(\mu_2, \sigma_2^2)$ if $n \in \mathfrak{R}$, if $\sigma_2^2 > \sigma_1^2$ and with $M \geq 1$, TECC algorithm (with OSGA) recovers \mathfrak{R} with probability approaching one as $T \rightarrow \infty$.*

We first note that the signals in Theorem 2 correspond to the JSM-3R signals: for $n \notin \mathfrak{R}$, the signal entries $x_{(n,t)}$ can be written as $x_{(n,t)} = \mu_1 + x_{n,t}^I$ where $x_{n,t}^I$ are i.i.d. $\sim \mathcal{N}(0, \sigma_1^2)$. With zero-mean and a potentially small variance, the amplitude of $x_{n,t}^I$, $n \notin \mathfrak{R}$ is expected to be small. For $n \in \mathfrak{R}$, the signal entries $x_{(n,t)}$ can be written as $x_{(n,t)} = \mu_2 + x_{n,t}^I$ where $x_{n,t}^I$ are i.i.d. $\sim \mathcal{N}(0, \sigma_2^2)$. With a larger variance σ_2^2 , the amplitude of $x_{n,t}^I$, $n \in \mathfrak{R}$ is expected to be much larger.

Proof. By the common component estimation from Algorithm 4, we have:

$$\begin{aligned} \widehat{x}^C &= \frac{1}{TM} \phi^\top y \\ &= \frac{1}{M} \frac{1}{T} \sum_{t=1}^T \phi_t^\top y_t \\ &= \frac{1}{M} \left(\frac{1}{T} \sum_{t=1}^T \phi_t^\top \phi_t x_{(\cdot,t)} \right). \end{aligned}$$

Note that this is $1/M$ times the sample mean of the random variable $\Phi^\top \Phi X$. Letting I_N denote the $N \times N$ identity matrix, we note that since Φ has independent $\mathcal{N}(0, 1)$ entries then $\mathbb{E}[\Phi^\top \Phi] = M I_N$. Since Φ is fully independent of X ,

$$\frac{1}{M} \mathbb{E}[\Phi^\top \Phi X] = \frac{1}{M} \mathbb{E}[\Phi^\top \Phi] \mathbb{E}[X] = I_N \mathbb{E}[X] = \mathbb{E}[X].$$

Invoking the Law of Large Numbers, we have

$$\lim_{T \rightarrow \infty} \hat{x}^C = \mathbb{E}[X].$$

Let $\hat{X} = X - \hat{x}^C$, then as $T \rightarrow \infty$, \hat{X}_n is distributed as $\mathcal{N}(0, \sigma_1^2)$ if $n \notin \mathfrak{K}$ and $\mathcal{N}(0, \sigma_2^2)$ if $n \in \mathfrak{K}$. Since $\hat{Y} = Y - \Phi \hat{x}^C = \Phi(X - \hat{x}^C) = \Phi \hat{X}$, it follows from Theorem 1 that with $M \geq 1$ and $\sigma_2^2 > \sigma_1^2$, the TECC with OSGA recovers \mathfrak{K} with probability approaching one as $T \rightarrow \infty$. \square

3 Experiments

In this section, we evaluate numerically the performance of Algorithms 1, 2, 3, 4 and 5 for anomaly detection. More specifically, we examine the success rate of determining the anomalous index set \mathfrak{K} from the signal matrix $x \in \mathbb{R}^{N \times T}$, whose columns are signals obtained at each time-step and share the same anomalous indices. The performance is assessed under various settings, by varying the number of anomalies, the number of columns in x (i.e. the time-steps) and the number of mixed measurements M at each time-step. Our focus is on the trade-off between the number of measurements M and the number of time-steps T required to identify \mathfrak{K} for varying numbers of anomalies.

In all experiments, the measurement matrices $\phi_t \in \mathbb{R}^{M \times N}$ comprise independent, $\mathcal{N}(0, 1)$ entries and the measurement vectors $y_t \in \mathbb{R}^M$ are calculated by $y_t = \phi_t x_{(:,t)}$ for $t = 1, \dots, T$. To obtain an estimate of an algorithm's recovery success rate with high confidence, instead of using a fixed number of random trials across the different parameter combinations, we adaptively determine the necessary number of trials with a Jeffreys interval i.e., a Bayesian two-tailed binomial proportion confidence interval. When the 95% confidence interval around the true success rate shrinks to a width smaller than 0.1, we report the current proportion of successes as the recovery accuracy for the algorithm. The signals $x_{(:,t)}$ are generated under the two models corresponding to the JSM-2R and JSM-3R signal definitions introduced in Section 2. Algorithms 1, 2 and 3 are applied to the JSM-2R signals while Algorithms 4 and 5 are applied to the JSM-3R signals.

The experiments are summarized in Table 1. The JSM-2R experiments assume a zero mean for the prevalent distribution and a much larger mean for the anomalous distribution while letting the variance be small. As shown in the previous section, signals generated from these distributions satisfy the definitions of JSM-2R. For

JSM-3R experiments, we explore two settings: First, the prevalent and anomalous distributions are assumed to have different means; second, the two distributions have the same mean. Recall from the previous section, we show that the means of the distributions are the common components for the JSM-3 signals generated from these distributions. Note that the algorithms for JSM-3R signals have no knowledge of the mean of the prevalent or anomalous distributions.

Table 1

Signal model	\mathcal{D}_1	\mathcal{D}_2	Algorithms
JSM-2R	$\mathcal{N}(0, 1)$	$\mathcal{N}(7, 1)$	OSGA, MMV-SOMP, MMV-LASSO
JSM-3R	$\mathcal{N}(7, 1)$	$\mathcal{N}(0, 10^2)$	TECC, ACIE
JSM-3R	$\mathcal{N}(7, 1)$	$\mathcal{N}(7, 10^2)$	TECC, ACIE

We chose the distributions in Table 1 for our numerical simulations to remain consistent with [13]. We observe, in the JSM-2R experiments, that the distributions $\mathcal{N}(0, 1)$ and $\mathcal{N}(7, 1)$ have their means separated by three standard deviations each, with one additional standard deviation in between for good measure. This ensures that the distributions are statistically distinct from each other. We have not explored how the detection accuracy is affected as we vary the proportion of overlap in the two distributions.

3.1 JSM-2R

We now present the results of recovering the anomalous index set for the JSM-2R signals. The signal length is fixed at $N = 100$ and results of $K = 1, 5,$ and 10 anomalies are presented. For each K value, K random variables follow the distribution $\mathcal{N}(7, 1)$ and the other $N - K$ random variables follow another distribution $\mathcal{N}(0, 1)$. The goal is to recover the index set \mathfrak{R} of these K random variables. Figure 2 shows the success rate of identifying \mathfrak{R} for the three K values using the OSGA. Each dot in the figure denotes the success rate for a specific M (number of measurements per time-step) and a specific T (number of time-steps) estimated from a number of trials, and the value is indicated by the color (see the colorbar). Both M and T take values from 1 to 100. Figures 3 and 4 plot the success rate for MMV-SOMP and MMV-LASSO algorithms respectively.

For all three algorithms, the success rate of anomaly identification increases as the number of measurements M increases and/or as the number of time-steps T increases. A 100% success of identification is obtained with a sufficiently large number of measurements and time-steps. There are some important differences in performance among the three algorithms.

Firstly, for the OSGA and the MMV-SOMP algorithms, with a sufficiently large number of time-steps, the minimum number of measurements at each time-step re-

quired for anomaly detection increases with the number of anomalies present. The MMV-LASSO performance is less affected by varying the number of anomalies than the performance of the other two algorithms. Secondly, Figures 2 and 3 reveal that MMV-SOMP requires fewer time-steps than the OSGA to reach 100% success for a given number of measurements. Thirdly, the MMV-LASSO algorithm requires significantly fewer measurements and time-steps for 100% success compared with OSGA and MMV-SOMP. Finally, there is asymmetry between the effect of increasing the number of measurements versus that of increasing the number of time-steps on the performance of OSGA and MMV-SOMP. For these two algorithms, increasing the number of measurements is more effective than increasing the number of time-steps for improving the performance. No obvious asymmetry of recovery performance is found for the MMV-LASSO algorithm. The near symmetry phenomenon of the MMV-LASSO is expected since doubling either M or T doubles the number of rows in the matrix ϕ in Algorithm 3, providing similar amounts of information for the algorithm.

For comparison with a benchmark, we note that in [13], the authors propose LASSO as an efficient algorithm to detect anomalies. The performance of their proposed method is shown as the first row, $M = 1$, in the phase diagrams of Figure 4. Here, we expand the application of LASSO by allowing for a trade-off between the number of measurements per time-step, M , and the number of time-steps, T , for which measurements are taken. Applications with an ability to store multiple measurements at each time-step, while seeking to minimize the time needed to accumulate data, might prefer to use the MMV-LASSO of Algorithm 3 to detect anomalies.

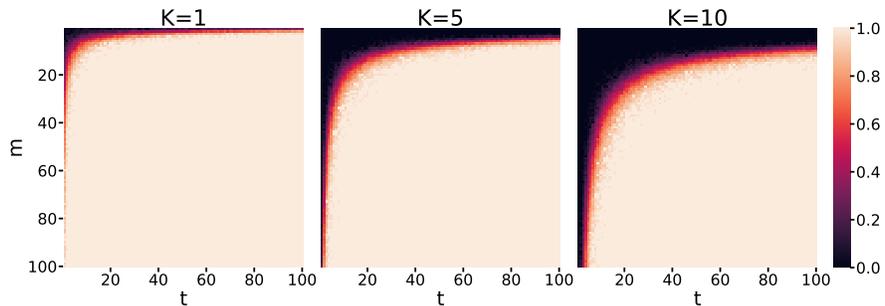


Fig. 2 The recovery phase transition for the OSGA with $K = 1, 5,$ and 10 anomalous random variables.

In these experiments, we have assumed that we know the number of anomalies, K . To explore the possibility of estimating the number of anomalies as we detect them, we consider the following experiments.

1. For OSGA, we calculate the test statistics, ξ_n , in Algorithm 1 for all $n = 1, 2, \dots, N$ and sort them in descending order; then determine whether the amplitudes of the ξ_n can be used to estimate K .

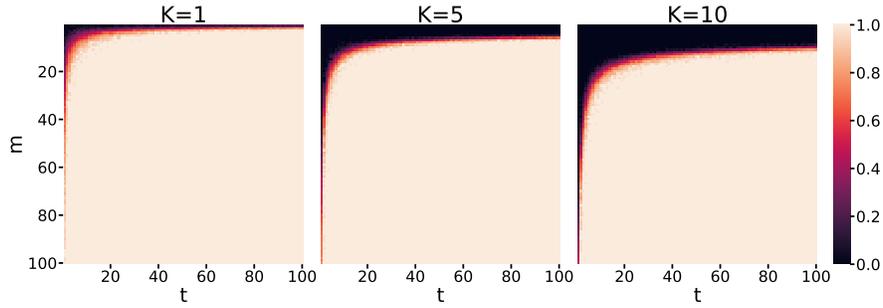


Fig. 3 The recovery phase transition for the MMV-SOMP algorithm with $K = 1, 5,$ and 10 anomalous random variables.

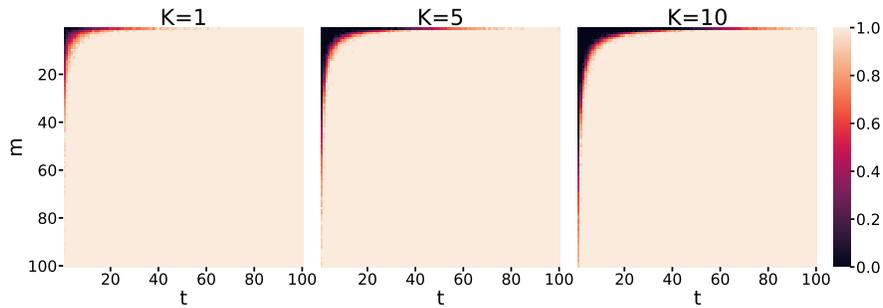


Fig. 4 The recovery phase transition for the MMV-LASSO algorithm with $K = 1, 5,$ and 10 anomalous random variables.

2. Similarly, for MMV-SOMP, we use the amplitude of $\sum_{t=1}^T \frac{|\langle r_t^{k-1}, \phi_{t(\cdot, n)} \rangle|}{\|\phi_{t(\cdot, n)}\|_2}$ in Algorithm 2 to determine the number of anomalies.
3. Lastly, for MMV-LASSO, we calculate the reconstructed signal $|\hat{x}|$ in Algorithm 3 and sort the entries in descending order, and determine K based on the amplitudes.

In each case, we fix $M = 50$ and $T = 50$ to ensure that recovery is possible if K is known (we can see this from the results in Figures 2, 3, 4). The results shown in Figure 5 demonstrate the potential of these methods to estimate K . Theoretical justification of these methods is left as future work.

3.2 JSM-3R

We next present the results of recovering the anomalous index set for the JSM-3R signals. Similar to JSM-2R signals, the length of the signal is set to $N = 100$ and the number of anomalies takes values of $K = 1, 5,$ and 10 . Unlike the JSM-2R signals, the $N - K$ random variables now follow the distribution $\mathcal{N}(7, 1)$ while the

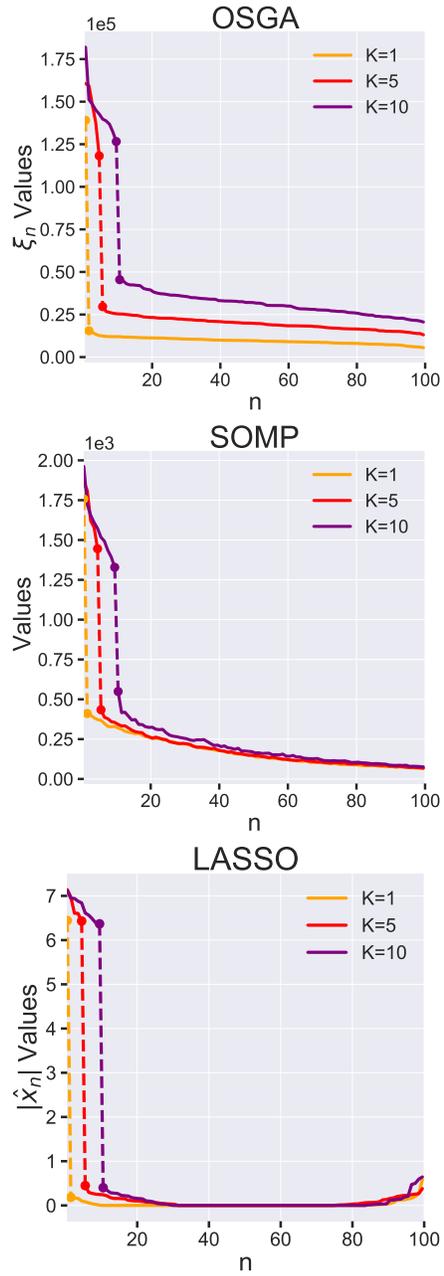


Fig. 5 Plots of the values from which indices are selected for $\hat{\mathbf{r}}$ in the JSM-2R algorithms. The dotted line denotes the drop between the top K values and the remaining $N - K$ values.

K anomalous random variables follow the distribution $\mathcal{N}(0, 10^2)$ or $\mathcal{N}(7, 10^2)$. In order to make the comparison fair between the two algorithms, we implement the OSGA for both step 6 of the TECC algorithm and step 10 of the ACIE algorithm. The number of iterations L in the ACIE algorithm is set to $L = 5$. The performance of the TECC and ACIE algorithms for varying numbers of measurements M and time-steps T when the anomalous distribution follows $\mathcal{N}(0, 10^2)$ is presented in Figures 6 and 7, where both M and T range from 1 to 100. The performance for the setting where the anomalous variables are distributed as $\mathcal{N}(7, 10^2)$ is similar to Figures 6 and 7 and is thus omitted.

With a sufficiently large number of measurements and time-steps, both algorithms are able to achieve 100% success in recovery of the anomalous index set. For a fixed number of time-steps, the minimum number of measurements required for identification increases as the number of anomalies increases for both algorithms. There is improvement in performance of the ACIE algorithm over the TECC algorithm. The ACIE algorithm requires fewer time-steps to reach 100% recovery success, for a given number of measurements; similarly, it requires fewer measurements for 100% recovery success with a given number of time-steps.

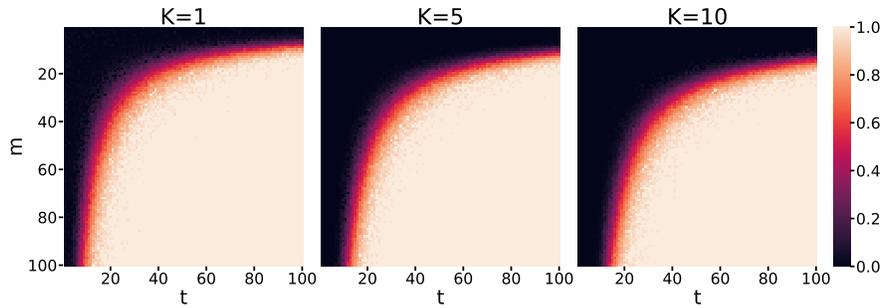


Fig. 6 The recovery phase transition for the TECC algorithm with $K = 1, 5,$ and 10 anomalous random variables. Here the prevalent distribution is $\mathcal{N}(7, 1)$ and the anomalous distribution is $\mathcal{N}(0, 10^2)$.

Thus far, we have assumed that the prevalent and anomalous distributions have very different variances, $\sigma_1 = 1$ and $\sigma_2 = 10$ in these experiments. To investigate the performance of these algorithms as the ratio of the variance changes, we experiment by setting $\sigma_2/\sigma_1 = 2, 5,$ and 10 , for $K = 1, 5,$ and 10 . Figure 8 shows the phase transition for the TECC algorithm as we vary the ratio of the variances, and Figure 9 shows the phase transition for the ACIE algorithm as we vary the ratio of the variances. In both cases, the algorithms are behaving as we might expect. The smaller the ratio between the variances, the more measurements and time-steps it takes to detect the anomalies.

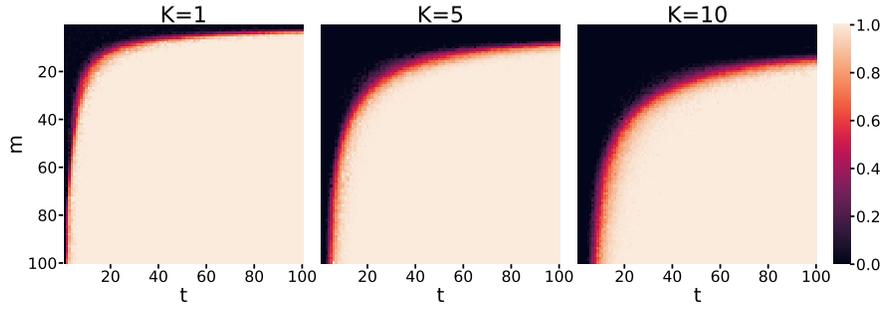


Fig. 7 The recovery phase transition for the ACIE algorithm with $K = 1, 5$, and 10 anomalous random variables. Here the prevalent distribution is $\mathcal{N}(7, 1)$ and the anomalous distribution is $\mathcal{N}(0, 10^2)$.

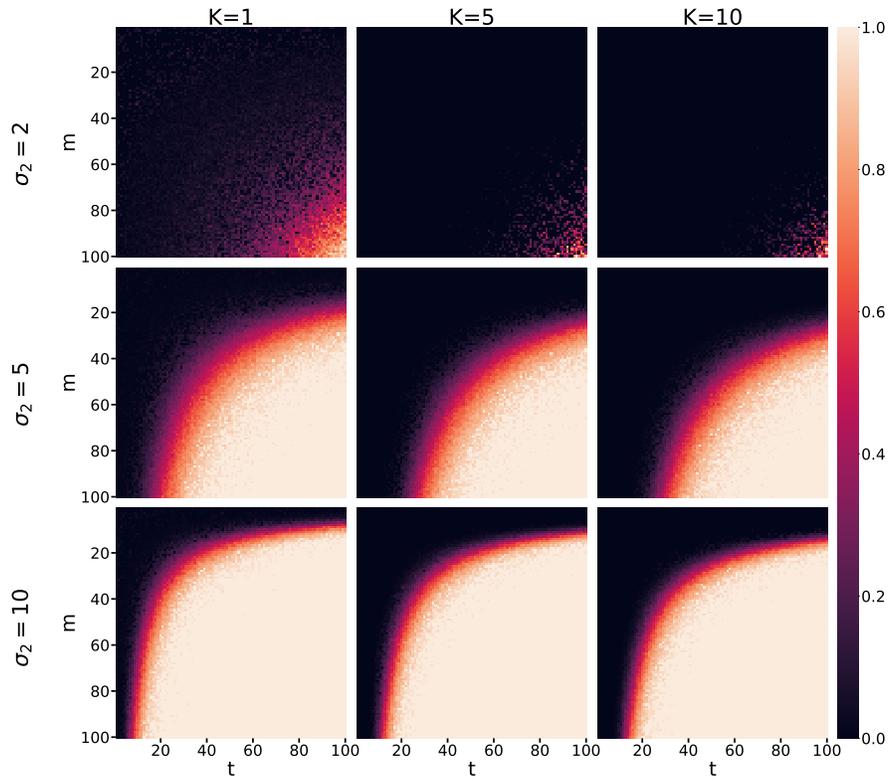


Fig. 8 The recovery phase transition for the TECC algorithm with $K = 1, 5$ and 10 anomalous random variables. Here the prevalent distribution is $\mathcal{N}(7, 1)$ and the anomalous distribution is $\mathcal{N}(0, \sigma_2^2)$, with $\sigma_2 = 2, 5$, and 10 shown.

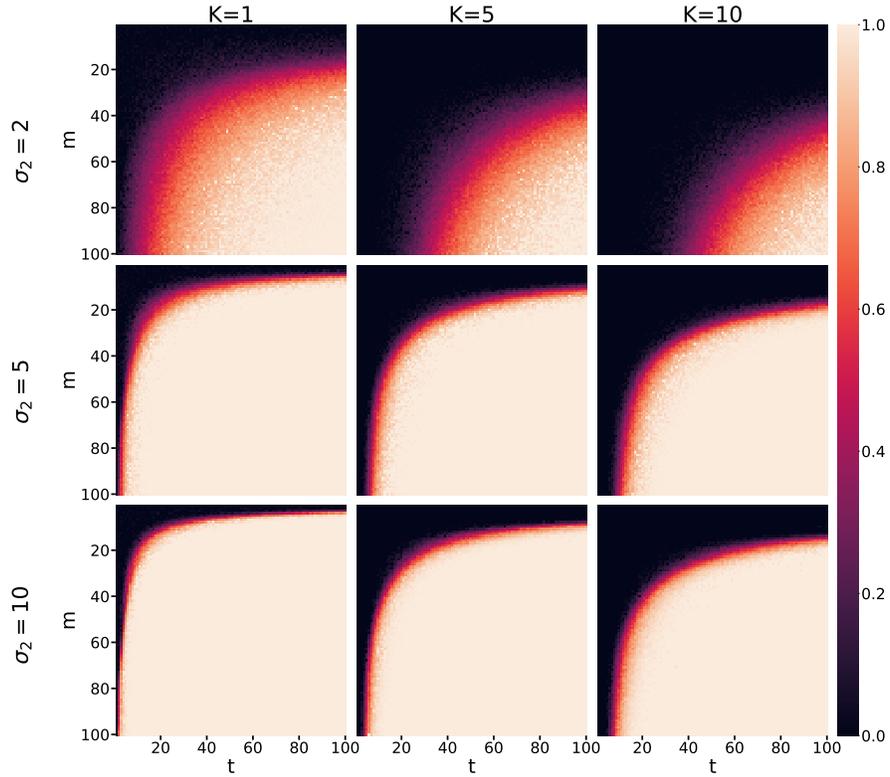


Fig. 9 The recovery phase transition for the ACIE algorithm with $K = 1, 5$ and 10 anomalous random variables. Here the prevalent distribution is $\mathcal{N}(7, 1)$ and the anomalous distribution is $\mathcal{N}(0, \sigma_2^2)$, with $\sigma_2 = 2, 5$, and 10 shown.

4 Conclusion

In this paper, we formally posed the problem of detecting anomalously distributed random variables as an MMV problem, by drawing an analogy between samples of the random variables and ensembles of signals. We further established two signal models characterizing possible correlation structures among signals that contain anomalous entries. Based on the new signal models, we showed through theoretical and numerical analysis that many of the MMV algorithms for sparse signal recovery can be adapted to the anomaly detection problem. For two of the proposed algorithms, we provided theoretical guarantees of anomaly detection in the asymptotic case. Our experimental results on synthetic data show good performance for signals conforming to either model, when a sufficiently large number of time-steps is available.

While these algorithms succeed in detecting anomalies, there is still room for optimizing performance. Currently these algorithms require storing the sensing ma-

trices at each time-step in memory. In future work, we would like to explore optimal ways to design sensing matrices to reduce the memory burden. Having provided asymptotic anomaly detection guarantees for two algorithms, we are further interested in providing such guarantees for all the algorithms presented. Additionally, we are interested in characterizing the performance bounds for each algorithm's finite sample case. Theorem 2 shows that only when the variances of the anomalous and prevalent distributions are distinct can the anomalies be detected by the algorithm. With additional information about the means of the distributions, the algorithms could be extended to identify the differences in means and detect anomalies even with identical variances. Finally, our current work assume that the signals follow Gaussian distributions. We are interested in expanding these algorithms to distributions which might not be distinguishable with the current approach. It would be interesting to investigate under what kinds of heavy-tailed distributions these algorithms start to fail.

Acknowledgements The initial research for this effort was conducted at the Research Collaboration Workshop for Women in Data Science and Mathematics, July 17-21 held at ICERM. Funding for the workshop was provided by ICERM, AWM and DIMACS (NSF grant CCF-1144502). SL was supported by NSF CAREER grant CCF-1149225. DN was partially supported by the Alfred P. Sloan Foundation, NSF CAREER #1348721, and NSF BIGDATA #1740325. JQ was supported by NSF DMS-1818374.

Appendix

Here we provide a summary of notation for reference.

N	Number of random variables
\mathfrak{N}	Set of random variables indices, $\{n \in \mathbb{N} : 1 \leq n \leq N\}$
K	Number of anomalous random variables
\mathfrak{K}	Set of anomalous random variable indices, $\mathfrak{K} \subset \mathfrak{N}$, $ \mathfrak{K} = K$
M	Number of measurements per time-step
m	Measurement index, $1 \leq m \leq M$
T	Number of time-steps measured
t	Time-step index, $1 \leq t \leq T$
\mathcal{D}_1	Prevalent distribution
\mathcal{D}_2	Anomalous distribution
X	Random vector comprising independent random variables X_1, \dots, X_N
x	$N \times T$ -dimensional matrix of independent realizations of X for all T time-steps
Φ	$M \times N$ -dimensional sensing matrix, i.i.d. $\sim \mathcal{N}(0, 1)$ entries
ϕ_t	$M \times N$ -dimensional realization of Φ at time t
ϕ	$(M \cdot T) \times N$ -dimensional vertical concatenation of the ϕ_t , $[\phi_1^\top, \dots, \phi_T^\top]^\top$
y_t	M -dimensional result of measuring the signal, $\phi_t \cdot x_{(\cdot, t)}$, at time t
y	$(M \cdot T)$ -dimensional vertical concatenation of the y_t , $[y_1^\top, \dots, y_T^\top]^\top$
Y	M -dimensional random vector defined by ΦX
JSM	Joint Sparsity Model, introduced in [3]
JSM-2	Signals are nonzero only on a common set of indices
JSM-3	Signals consist of common non-sparse component and a sparse innovation
JSM-2R	“Random variable” version of JSM-2
JSM-3R	“Random variable” version of JSM-3
OSGA	One-step greedy algorithm
MMV	Multiple Measurement Vector
MMV-LASSO	MMV Least Absolute Shrinkage and Selection Operator
MMV-SOMP	MMV Simultaneous Orthogonal Matching Pursuit
TECC	Transpose Estimation of Common Component
ACIE	Alternating Common and Innovation Estimation

We adopt the convention that random variables will be upper case and their realizations will be lower case. All matrix entries will have two, subscripted indices. The first index will indicate the row position, the second will indicate the column position. We will indicate row and column vectors by substituting \cdot for the respective index.

References

- [1] Angelosante, D., Giannakis, G.B., Grossi, E.: Compressed sensing of time-varying signals. In: 16th International Conference on Digital Signal Processing, pp. 1–8. Santorini-Hellas, Greece (2009)
- [2] Baron, D., Sarvotham, S., Baraniuk, R.G.: Bayesian compressive sensing via belief propagation. *IEEE Trans. Signal Process.* **58**(1), 269–280 (2010)
- [3] Baron, D., Wakin, M.B., Duarte, M.F., Sarvotham, S., Baraniuk, R.G.: Distributed compressed sensing (2005). Preprint available at https://www.ece.rice.edu/~shri/pub/DCS_TR.pdf
- [4] Basseville, M., Nikiforov, I.V.: Detection of Abrupt Changes: Theory and Application. Prentice-Hall, Inc., Upper Saddle River, NJ, USA (1993)
- [5] van den Berg, E., Friedlander, M.P.: Theoretical and empirical results for recovery from multiple measurements. *IEEE Trans. Inf. Theory* **56**(5), 2516–2527 (2010)
- [6] Blumensath, T., Davies, M.E.: Iterative hard thresholding for compressed sensing. *Appl. Comput. Harmon. Anal.* **27**(3), 265–274 (2009)
- [7] Candès, E.J.: Compressive sampling. In: Proceedings of the International Congress of Mathematicians, vol. 3, pp. 1433–1452. Madrid, Spain (2006)
- [8] Candès, E.J., Romberg, J., Tao, T.: Stable signal recovery from incomplete and inaccurate measurements. *Commun. Pure Appl. Math.* **59**(8), 1207–1223 (2006)
- [9] Candès, E.J., Tao, T.: Decoding by linear programming. *IEEE Trans. Inf. Theory* **51**(12), 4203–4215 (2005)
- [10] Candès, E.J., Tao, T.: Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inf. Theory* **52**(12), 5406–5425 (2006)
- [11] Chen, J., Huo, X.: Theoretical results on sparse representations of multiple-measurement vectors. *IEEE Trans. Signal Process.* **54**(12), 4634–4643 (2006)
- [12] Chen, S.S., Donoho, D.L., Saunders, M.A.: Atomic decomposition by basis pursuit. *SIAM Rev.* **43**(1), 129–159 (2001)
- [13] Cho, M., Xu, W., Lai, L.: Compressed Hypothesis Testing: To Mix or Not to Mix? arXiv preprint arXiv:1609.07528 (2016)
- [14] Cotter, S.F., Rao, B.D., Engan, K., Kreutz-Delgado, K.: Sparse solutions to linear inverse problems with multiple measurement vectors. *IEEE Trans. Signal Process.* **53**(7), 2477–2488 (2005)
- [15] Donoho, D.L.: Compressed sensing. *IEEE Trans. Inf. Theory* **52**(4), 1289–1306 (2006)
- [16] Donoho, D.L., Huo, X.: Uncertainty principles and ideal atomic decomposition. *IEEE Trans. Inf. Theory* **47**(7), 2845–2862 (2001)
- [17] Duarte, M.F., Wakin, M.B., Baron, D., Sarvotham, S., Baraniuk, R.G.: Measurement bounds for sparse signal ensembles via graphical models. *IEEE Trans. Inf. Theory* **59**(7), 4280–4289 (2013)

- [18] Duarte, M.R., Wakin, M.B., Baron, D., Baraniuk, R.G.: Universal distributed sensing via random projections. In: 5th International Conference on Information Processing in Sensor Networks, pp. 177–185. Nashville, TN (2006)
- [19] Efron, B., Hastie, T., Johnstone, I., Tibshirani, R.: Least angle regression. *Ann. Stat.* **32**(2), 407–499 (2004)
- [20] Eldar, Y.C., Kutyniok, G.: Compressed sensing: Theory and applications. Cambridge University Press (2012)
- [21] Filos, J., Karseras, E., Dai, W., Yan, S.: Tracking dynamic sparse signals with hierarchical kalman filters: A case study. In: 18th International Conference on Digital Signal Processing, pp. 1–6. Fira, Greece (2013)
- [22] Foucart, S., Rauhut, H.: A Mathematical Introduction to Compressive Sensing. Birkhäuser Basel (2013)
- [23] Haupt, J., Nowak, R.: Signal reconstruction from noisy random projections. *IEEE Trans. Inf. Theory* **52**(9), 4036–4048 (2006)
- [24] Ji, S., Xue, Y., Carin, L.: Bayesian compressive sensing. *IEEE Trans. Signal Process.* **56**(6), 2346–2356 (2008)
- [25] Keerthi, S.S., Shevade, S.: A fast tracking algorithm for generalized LARS/LASSO. *IEEE Trans. Neural Netw.* **18**(6), 1826–1830 (2007)
- [26] Lai, L., Fan, Y., Poor, H.V.: Quickest detection in cognitive radio: A sequential change detection framework. In: IEEE Global Telecommunications Conference, pp. 1–5. New Orleans, LO (2008)
- [27] Mallat, S.G., Zhang, Z.: Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Process.* **41**(12), 3397–3415 (1993)
- [28] Malloy, M., Nowak, R.: On the limits of sequential testing in high dimensions. In: 45th Asilomar Conference on Signals, Systems and Computers, pp. 1245–1249. Pacific Grove, CA (2011)
- [29] Malloy, M., Nowak, R.: Sequential analysis in high-dimensional multiple testing and sparse recovery. In: IEEE International Symposium on Information Theory, pp. 2661–2665. St. Petersburg, Russia (2011)
- [30] Malloy, M.L., Tang, G., Nowak, R.D.: Quickest search for a rare distribution. In: 46th Annual Conference on Information Sciences and Systems, pp. 1–6. Princeton, NJ (2012)
- [31] Mishali, M., Eldar, Y.C.: Reduce and boost: Recovering arbitrary sets of jointly sparse vectors. *IEEE Trans. Signal Process.* **56**(10), 4692–4702 (2008)
- [32] Needell, D., Vershynin, R.: Signal recovery from incomplete and inaccurate measurements via Regularized Orthogonal Matching Pursuit. *IEEE J. Sel. Topics Signal Process.* **4**(2), 310–316 (2010)
- [33] Pati, Y.C., Rezaifar, R., Krishnaprasad, P.S.: Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In: 27th Asilomar Conference on Signals, Systems and Computers, vol. 1, pp. 40–44. Pacific Grove, CA (1993)
- [34] Patterson, S., Eldar, Y.C., Keidar, I.: Distributed compressed sensing for static and time-varying networks. *IEEE Trans. Signal Process.* **62**(19), 4931–4946 (2014)

- [35] Poor, H.V., Hadjiliadis, O.: *Quickest Detection*. Cambridge University Press (2008)
- [36] Tropp, J.A.: Just relax: Convex programming methods for identifying sparse signals in noise. *IEEE Trans. Inf. Theory* **52**(3), 1030–1051 (2006)
- [37] Tropp, J.A., Gilbert, A.C.: Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Trans. Inf. Theory* **53**(12), 4655–4666 (2007)
- [38] Tropp, J.A., Gilbert, A.C., Strauss, M.J.: Simultaneous sparse approximation via greedy pursuit. In: *IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 5, pp. 721–724. Philadelphia, PA (2005)
- [39] Xia, Y., Tse, D.: Inference of link delay in communication networks. *IEEE J. Sel. Areas Commun.* **24**(12), 2235–2248 (2006)
- [40] Yu, G., Sapiro, G.: Statistical compressed sensing of Gaussian mixture models. *IEEE Trans. Signal Process.* **59**(12), 5842–5858 (2011)