

Sparse Randomized Kaczmarz for Support Recovery of Jointly Sparse Corrupted Multiple Measurement Vectors

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

Abstract While single measurement vector (SMV) models have been widely studied in signal processing, there is a surging interest in addressing the multiple measurement vectors (MMV) problem. In the MMV setting, more than one measurement vector is available and the multiple signals to be recovered share some commonalities such as a common support. Applications in which MMV is a naturally occurring phenomenon include online streaming, medical imaging, and video recovery. This work presents a stochastic iterative algorithm for the support recovery of jointly sparse corrupted MMV. We present a variant of the Sparse Randomized

Natalie Durgin
Spiceworks
Austin, TX, 78746.

Rachel Grotheer
Goucher College
Baltimore, MD 21204.
e-mail: rachel.grotheer@goucher.edu

Chenxi Huang
Yale University
New Haven, CT 06511.

Shuang Li
Colorado School of Mines
Golden, CO 80401.

Anna Ma
Claremont Graduate University
Claremont, CA 91711.
e-mail: anna.ma@cgu.edu

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

Jing Qin
Montana State University
Bozeman, MT 59717.

Kaczmarz algorithm for corrupted MMV and compare our proposed method with an existing Kaczmarz type algorithm for MMV problems. We also showcase the usefulness of our approach in the online (streaming) setting and provide empirical evidence that suggests the robustness of the proposed method to the number of corruptions and the distribution from which the corruptions are drawn.

1 Introduction

In recent years, there has been a drastic increase in the amount of available data. This so-called “data deluge” has created a demand for fast, iterative algorithms that can be used to process large-scale data. Stochastic iterative algorithms, such as the Randomized Kaczmarz or Stochastic Gradient Descent algorithms, have become an increasingly popular option for processing large-scale data [3, 10]. These methods recover a signal $\mathbf{X} \in \mathbb{R}^n$ given a vector of measurements $\mathbf{Y} \in \mathbb{R}^m$ and a measurement matrix $\Phi \in \mathbb{R}^{m \times n}$ such that:

$$\mathbf{Y} = \Phi \mathbf{X}, \quad (1)$$

without accessing the full measurement matrix in a single iteration. We refer to (1) as a single measurement vector (SMV) model. In the multiple measurement vector (MMV) setting, one may have thousands of measurement vectors $\mathbf{Y}_{(:,j)}$ pouring in over time. Each measurement vector corresponds to a signal $\mathbf{X}_{(:,j)}$ where signals typically share a common property such as sparsity, smoothness, etc. For simplicity, let $\mathbf{Y} = [\mathbf{Y}_{(:,1)} \cdots \mathbf{Y}_{(:,J)}] \in \mathbb{R}^{m \times J}$ and $\mathbf{X} = [\mathbf{X}_{(:,1)} \cdots \mathbf{X}_{(:,J)}] \in \mathbb{R}^{n \times J}$. Since high-dimensional data is typically sparse in nature, a commonality of particular interest is *joint sparsity*, or when signals share the same support. The support of a vector \mathbf{v} is defined to be the set of indices for which \mathbf{v} is nonzero, i.e. $\text{supp}(\mathbf{v}) = \{i : \mathbf{v}_i \neq 0\}$.

Many algorithms have been developed for the MMV setting, especially in applications such as line spectral estimation [14, 20] and modal analysis [13]. The authors in these works extend the previous SMV-based algorithms as well as theoretical analysis in [4, 8, 19] to the MMV case. The theoretical bound in [13] also indicates that MMV settings could make the problem of compressed signal recovery much easier than in the SMV setting. In particular, the number of measurements needed for perfect recovery in each signal decreases as the number of signals increases, reducing the sample complexity per signal.

As a motivating example, consider diffuse optical tomography (DOT) where the goal is to find small areas of high contrast corresponding to the location of cancerous cells [2]. Since cancerous cells have a much larger absorption coefficient than healthy cells, a two-dimensional medical image can be interpreted as a sparse signal where each entry of the signal represents the absorption coefficient of a given pixel and the nonzero entries correspond to tumor locations. In a hyperspectral DOT setting, hundreds of different wavelengths are used to acquire a variety of images of the same tissue, allowing practitioners to obtain a more accurate location of tumors [11]. The hyperspectral imaging process results in a jointly sparse MMV, where each

wavelength produces a different image (or signal), and the joint support across all images indicates the locations of cancerous cells.

Signals may share a common support but it is improbable for them to be perfectly accurate. Since sensing mechanisms are not impervious to error, signals can contain corruptions. Other sources of corruption in signal processing include spikes in power supply, defective hardware, and adversarial agents [12]. Going back to the hyperspectral imaging example, “corruptions” in each signal may be caused by non-cancerous cells that absorb more light at a given wavelength than their neighbors. For example, if a cell contains an anomalous amount of melanin, then it absorbs more light at shorter wavelengths in the visible spectrum (i.e., violet or blue light) compared to a typical noncancerous cell [5, 15]. This produces a large nonzero absorption coefficient in the location of a healthy cell, i.e., a corruption. Corrupt entries erroneously indicate the presence of cancerous cells in a location with healthy cells.

Corruptions cause support recovery algorithms such as the MMV Sparse Randomized Kaczmarz (MMV-SRK) algorithm, which we describe in detail in Section 2, to fail due to the algorithmic dependence on the row norms of the signal approximation to estimate the support [1]. Thus, large corruptions in a signal with comparatively small entries may erroneously be included in the support estimate given by these algorithms. In the corrupt MMV setting, the availability of multiple measurement vectors becomes vital to the estimate of the true support. Clearly, if only a single measurement vector is available, there would be no way to distinguish a corrupt nonzero entry without any additional assumptions on the signal or corruption. Corrupt measurement signals have been studied in the context of the SMV model. In [18] and [12], additive noise in the measurement scheme is assumed to be sparse. Both works focus on the compressive sensing setting where $m \ll n$.

The primary objective of this work is to design an algorithm for recovering the support of jointly sparse, corrupt signals in the large-scale setting. We propose a new online algorithm called Sparse Randomized Kaczmarz for Corrupted MMV (cMMV-SRK) for support recovery. Note that the proposed algorithm can recover the signals with high accuracy based on our experiments, but we mainly focus on support recovery in this work. Our experiments show that the proposed algorithm outperforms the previously proposed Kaczmarz type algorithm in recovering the joint support from MMV when the signals are corrupted.

1.1 Problem Formulation

Given a set of linear measurements $\mathbf{Y}_{(:,j)} \in \mathbb{R}^m$ and a measurement matrix $\Phi \in \mathbb{R}^{m \times n}$ such that:

$$\mathbf{Y}_{(:,j)} = \Phi \mathbf{X}_{(:,j)} \quad \text{for } j = 1, \dots, J, \quad (2)$$

with $\mathbf{X}_{(:,j)} \in \mathbb{R}^n$. We assume that the data is large-scale, meaning we cannot access all of Φ at once (m and/or n are too large) and must only operate on one row of Φ at a time. We allow the system to be overdetermined ($m \gg n$) or underdetermined

($m \ll n$) and assume the $\mathbf{X}_{(:,j)}$'s are jointly sparse such that $\text{supp}(\mathbf{X}_{(:,j)}) = \mathcal{S}$ and $|\mathcal{S}| = k$. For an n -dimensional vector $\mathbf{X}_{(:,j)}$, let $\mathbf{X}_{(:,j)}|_s$ return $\mathbf{X}_{(:,j)}$ with zeros in the $n - s$ smallest (in magnitude) entries. We also assume that each column of \mathbf{X} contains one or more *corruptions* where a corruption is a non-zero entry occurring outside the joint support. In other words, instead of $\text{supp}(\mathbf{X}_{(:,j)}) \subset \mathcal{S}$, the joint support set, the support of $\mathbf{X}_{(:,j)}$ is:

$$\text{supp}(\mathbf{X}_{(:,j)}) = \mathcal{S} \cup \mathcal{C}_j, \quad \mathcal{C}_j \subset \{1, \dots, n\},$$

where \mathcal{C}_j is the ‘‘corrupt index set’’. Note that the \mathcal{C}_j are not necessarily the same for every j . In this work, our goal is to recover the joint support \mathcal{S} from the given linear measurements \mathbf{Y} .

The remainder of this manuscript is organized in the following way. Section 2 discusses the Sparse Randomized Kaczmarz method and the MMV-SRK algorithm. Section 3 provides a discussion on how corruptions can negatively impact the performance of MMV-SRK. Section 3 also presents our method, cMMV-SRK, a variant of SRK which works in the corrupted signal setting. Numerical experiments using this method are presented in Section 4 and we conclude with a summary of our contributions and future directions in Section 5.

2 Related and Existing Work

2.1 Sparse Randomized Kaczmarz

In this work, we utilize the Sparse Randomized Kaczmarz (SRK) algorithm to recover the support of each column of \mathbf{X} . The original Kaczmarz algorithm was first introduced in the early 1900s by Kaczmarz himself and was revitalized as the Algebraic Reconstruction Technique in the early 1980s [10, 6]. The Randomized Kaczmarz algorithm (RK) discussed throughout the paper was first introduced by Strohmer and Vershynin and enjoys an expected linear convergence rate to the solution of a consistent linear system [17]. The SRK algorithm is another variant designed specifically for overdetermined systems with sparse solutions. SRK has been empirically shown to solve underdetermined systems with sparse solutions as well [16].

Algorithm 1 outlines the SRK algorithm. Note that ties are broken lexicographically in all algorithms presented in this work. The estimated support size \hat{k} is a parameter of the algorithm and is typically chosen to be larger than the true support size k . In this variant, the algorithm runs for a specified number of iterations (up to τ). However, any stopping criteria one would use for an iterative algorithm, e.g., terminating after the residual meets a certain criteria, after the updates become small, etc., can be used. Algorithm 1 also differs from the original presented algorithm by [16] in that at every iteration the support estimate has size \hat{k} instead of starting with

Algorithm 1 Sparse Randomized Kaczmarz

```

1: procedure SRK( $\Phi \in \mathbb{R}^{m \times n}$ ,  $\mathbf{Y} \in \mathbb{R}^m$ ,  $\hat{k}$ ,  $\tau$ )
2:   Initialize  $\mathbf{X}^0 = \mathbf{0}^{n \times 1}$ 
3:   for  $t = 1, \dots, \tau$  do
4:     Choose row  $\Phi_{(i, \cdot)}$  with probability  $\frac{\|\Phi_{(i, \cdot)}\|_2^2}{\|\Phi\|_F^2}$ 
5:     Set support estimate  $\mathcal{S}^t = \text{supp}(\mathbf{X}^{t-1}|_{\hat{k}})$ 
6:     Set row weights  $\triangleright \mathbf{w} \in \mathbb{R}^n$ 

$$w_l = \begin{cases} 1 & : l \in \mathcal{S}^t \\ \frac{1}{\sqrt{t}} & : l \in \mathcal{S}^{t^c} \end{cases}$$

7:     Set  $\mathbf{a} = \mathbf{w} \cdot \Phi_{(i, \cdot)}$   $\triangleright \mathcal{S}^{t^c}$  is the complement set of  $\mathcal{S}^t$ 
8:     Update  $\mathbf{X}^t = \mathbf{X}^{t-1} + \frac{\mathbf{Y}_i - \mathbf{a}\mathbf{X}^{t-1}}{\|\mathbf{a}\|_2} \mathbf{a}^T$   $\triangleright \mathbf{a} \in \mathbb{R}^n$  is the weighted row of  $\Phi$ 
9:   end for
10:  return  $\mathbf{X}^\tau$ 
11: end procedure

```

n and shrinking the size to \hat{k} . We find that these modifications do not significantly affect the behavior of SRK.

Algorithm 1 has been shown empirically to find the solution to overdetermined, consistent (i.e., a solution exists) linear systems but there are no theoretical results supporting this. One can make a few observations about the behavior of SRK for support recovery. Concerning the support size estimate \hat{k} , it is clear that if $\hat{k} < k$ then the probability that the true support is contained in the support of the approximation is 0, i.e., $\mathbb{P}(\mathcal{S} \subset \text{supp}(\mathbf{X}^\tau)) = 0$. Additionally, if $\hat{k} = n$, then $\mathbb{P}(\mathcal{S} \subset \text{supp}(\mathbf{X}^\tau)) = 1$. In regards to the choice of weighting, as $t \rightarrow \infty$, $\frac{1}{\sqrt{t}} \rightarrow 0$ so that row elements inside the support estimate contribute to the approximation the most. If one has a weighting function that decreases too rapidly, the true support may not be captured in \mathcal{S}^t causing the algorithm to fail.

Although Algorithm 1 and the following algorithms require the Frobenius norm of the matrix, $\|\Phi\|_F^2$, for row selection, practically speaking row selections can be done uniformly at random to avoid using the full measurement matrix in a single iteration. Indeed, it is advantageous to select the rows at random to avoid introducing bias from rows with larger norms.

2.2 SRK for MMV

In this section we present an SRK-based approach to the MMV setting proposed by [1]. Because we are assuming joint sparsity in the MMV model, the estimated support of a signal reveals information about the support of all signals. The authors of [1] present Algorithm 2 to leverage this idea. There are a few key aspects to note

about this version of the SRK algorithm. First, the algorithm runs one iteration of SRK for every signal in the MMV model then updates the support estimate based on the row norms of the estimate \mathbf{X}^t . Due to this, the algorithm does not lend itself well to being extended for an online variant which only receives a small number (possibly 1) of signals at a time. Second, the algorithm uses the same selected row for each signal. It has been well observed that a random selection scheme reduces the possibility of a poor choice of row ordering and it may be advantageous to allow each signal to be projected onto a different randomly selected row [7, 9].

Algorithm 2 Sparse Randomized Kaczmarz for MMV

```

1: procedure MMV-SRK( $\Phi \in \mathbb{R}^{m \times n}$ ,  $\mathbf{Y} \in \mathbb{R}^{m \times J}$ ,  $\hat{k}$ ,  $\tau$ )
2:   Initialize  $\mathbf{X}^0 = \mathbf{0}^{n \times J}$ 
3:   for  $t = 1, \dots, \tau$  do
4:     Choose row  $\Phi_{(i, \cdot)}$  with probability  $\frac{\|\Phi_{(i, \cdot)}\|_2^2}{\|\Phi\|_F^2}$ 
5:     Set support estimate  $\mathcal{S}^t$ :  $\hat{k}$  indices with largest row norm of  $\mathbf{X}^{t-1}$ 
6:     Set row weights  $\triangleright \mathbf{w} \in \mathbb{R}^n$ 
           
$$w_l = \begin{cases} 1 & : l \in \mathcal{S}^t \\ \frac{1}{\sqrt{t}} & : l \in \mathcal{S}^{t^c} \end{cases}$$

            $\triangleright \mathcal{S}^{t^c}$  is the complement set of  $\mathcal{S}^t$ 
            $\triangleright \mathbf{a}$  is the weighted row of  $\Phi$ 
7:     Set  $\mathbf{a} = \mathbf{w} \cdot \Phi_{(i, \cdot)}$ 
8:     for  $j = 1, \dots, J$  do
9:       Update  $\mathbf{X}_{(\cdot, j)}^t = \mathbf{X}_{(\cdot, j)}^{t-1} + \frac{\mathbf{Y}_{(i, j)} - \mathbf{a} \mathbf{X}_{(\cdot, j)}^{t-1}}{\|\mathbf{a}\|_2} \mathbf{a}^T$ 
10:    end for
11:    Update  $\mathbf{X}^t = [\mathbf{X}_{(\cdot, 1)}^t | \dots | \mathbf{X}_{(\cdot, J)}^t]$ 
12:  end for
13:  return  $\mathbf{X}^t$ 
14: end procedure

```

3 Main Results

3.1 Corrupted MMV

To review, we are interested in constructing an algorithm that recovers the support of jointly sparse corrupted high-dimensional MMV, that is, where we can only access one row of the measurement matrix Φ at a time. To this end, we propose Algorithm 3, which we refer to as cMMV-SRK. We first note that the base of this algorithm is the SRK algorithm (Algorithm 1), which is an effective algorithm for large-scale problems due to its low memory footprint, requiring only one row of the measurement matrix to be used at a time. The proposed method, cMMV-SRK,

is adapted to the MMV setting using the intuition that the individual signals give us information about the common support between all signals. We keep track of a bin or tally vector \mathbf{b} that estimates the true support of the signals. In particular, we use the nonzeros in \mathbf{b} to indicate the estimated joint support. This binning process allows the algorithm to be robust in the face of corruptions in the signal, as the corruptions will receive a low number of tallies compared to the entries in the true support because the corruptions occur in random positions for every signal. Note that in the corrupted MMV case, we expect Algorithm 2 to fail as the support estimate step relies on the ℓ_2 -norm of the rows to be large if an index is in the support and small otherwise. The corruptions may be so large that a single corruption in a row could lead to mis-identification of the corrupt entry being in the joint support. Algorithm 3 can also account for signals being processed one at a time, as in an online or “streaming” setting.

For each signal, let $\tilde{\tau}_j$ be the number of SRK projections performed on $\mathbf{X}_{(:,j)}$ and let $\tilde{\tau} = [\tilde{\tau}_1 \cdots \tilde{\tau}_J]$. In the online setting, one can imagine that the amount of time before the next signal is acquired may vary due to, for example, stalls in the measurement process. The varying amount of time that the system has to process each signal is one of the major challenges of support recovery in the online setting. To improve the joint support estimate when $\tilde{\tau}_j$ varies, we weight the binning based on $\tau = \max_t \tilde{\tau}_t$. In other words, we let $\mathbf{b}_q = \mathbf{b}_q + \frac{\tilde{\tau}_j}{\tau}$ where \mathbf{b}_q is the q -th entry of \mathbf{b} and $\tilde{\tau}_j$ is the maximum number of inner iterations of SRK for the j^{th} signal. This reweighting scheme places a larger importance on support estimates which have had more time (iterations) to improve. In the online setting where the $\tilde{\tau}_j$ s are not known *a priori*, τ can be set manually. We adopt the following notation for cMMV-SRK: S_j^t is the estimated support at the t^{th} SRK iteration for $\mathbf{X}_{(:,j)}$, and \mathcal{S}^t denotes the joint support estimate.

If the number of inner iterations $\tilde{\tau}_j$ is large enough, the support estimate should be such that it contains the joint support (along with the corruption index). Because we are tallying the support estimate after every τ iterations, it is clear that the entries in the joint support will have an overwhelming number of tallies compared to all other entries. The experimental results in the next section support these claims and we leave the analytical study of these algorithms for future work.

4 Experiments

In this section, we compare Algorithm 2 and Algorithm 3 under a variety of settings, specifically comparing the robustness of each algorithm in the presence of corruptions in the signal. To test this robustness, we vary the number of corrupt entries, the distribution from which the corruptions are drawn, the number of columns in \mathbf{X} , and the number of projection computations $\tilde{\tau}_j$ made for each signal. In what follows, we will refer to τ as the number of SRK iterations. These experiments are summarized in Table 1. In all experiments, the results are averaged over 40 trials and the nonzero entries of \mathbf{X} are drawn independently and identically distributed (i.i.d.)

Algorithm 3 Sparse Randomized Kaczmarz for Corrupted MMV

```

1: procedure cMMV-SRK(  $\Phi \in \mathbb{R}^{m \times n}$ ,  $\mathbf{Y} \in \mathbb{R}^{m \times J}$ ,  $\hat{k}$ ,  $\tilde{\tau}$ )
2:   Initialize  $\mathbf{X}^0 = \mathbf{0}^{n \times J}$ ,  $\mathbf{b} = \mathbf{0}^{n \times 1}$ ,  $\tau = \max_j \tilde{\tau}_j$ 
3:   for  $j = 1, \dots, J$  do
4:     for  $t = 1, \dots, \tilde{\tau}_j$  do
5:       Choose row  $\Phi_{(i, \cdot)}$  with probability  $\frac{\|\Phi_{(i, \cdot)}\|_2^2}{\|\Phi\|_F^2}$ 
6:       Set support estimate  $S_j^t = \text{supp}(\mathbf{X}_{(\cdot, j)}^{t-1}|_{\hat{k}})$ 
7:       Set row weights  $\triangleright \mathbf{w} \in \mathbb{R}^n$ 
           
$$w_l = \begin{cases} 1 & : l \in S^t \\ \frac{1}{\sqrt{t}} & : l \in S^{t^c} \end{cases}$$

            $\triangleright S^{t^c}$  is the complement set of  $S^t$   

 $\triangleright \mathbf{a}$  is the weighted row of  $\Phi$ 
8:       Set  $\mathbf{a} = \mathbf{w} \cdot \Phi_{(i, \cdot)}$ 
9:       Update  $\mathbf{X}_{(\cdot, j)}^t = \mathbf{X}_{(\cdot, j)}^{t-1} + \frac{\mathbf{Y}_{(i, j)} - \mathbf{a}\mathbf{X}_{(\cdot, j)}^{t-1}}{\|\mathbf{a}\|_2^2} \mathbf{a}^T$ 
10:    end for
11:    If  $q \in S_j^t$  then  $\mathbf{b}_q = \mathbf{b}_q + \frac{\tilde{\tau}_j}{\tau}$ 
12:    Set initial support estimate for next signal  $\mathcal{S}^j = \text{supp}(\mathbf{b}|_{\hat{k}})$ 
13:  end for
14:  return Joint support estimate  $\mathcal{S}^J$ 
15: end procedure

```

from $\mathcal{N}(0, 1)$. On the horizontal axis we plot the number of “iterations” where a single iteration is defined by a projection. In other words, the horizontal axis represents every time Step 9 is performed in Algorithm 3 and Algorithm 2.

Table 1: Summary of Experiment Parameters. This table provides a summary of experiment parameters and figure references for each experiment.

Figures	Φ entries	Distribution of Corruptions	Number of Corruptions	$\tilde{\tau}_j$	J
Figure 1a	Gaussian	$\mathcal{N}(7, 1)$	1	40	300
Figure 1b	Uniform	$\mathcal{N}(7, 1)$	1	80	600
Figure 2a	Gaussian	$\mathcal{N}(0, 1)$	1	40	300
Figure 2b	Uniform	$\mathcal{N}(0, 1)$	1	80	600
Figure 3	Gaussian	$\mathcal{N}(7, 1)$	Varies	50	300
Figure 4a	Gaussian	$\mathcal{N}(7, 1)$	Varies	Varies	800

Figure 1 compares Algorithm 2 and Algorithm 3 with $m = 1000$, $n = 100$, and $k = 10$. The support size estimate is $\hat{k} = 1.5k$. To create \mathbf{X} , we uniformly at random select k of n indexes to be the joint support \mathcal{S} and set $\mathbf{Y} = \Phi\mathbf{X}$. The corrupt entries are drawn uniformly at random from $\{1, \dots, N\} \setminus \mathcal{S}$. To start off, each signal has one corrupt entry. We choose corruptions i.i.d. from $\mathcal{N}(7, 1)$ to simulate corruptions being large spikes in the signal (possibly caused by system malfunction). The maximum number of SRK iterations for each signal is $\tau = 300$. In Figure 1a, we

create $\Phi \in \mathbb{R}^{m \times n} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$ and $J = 300$. In Figure 1b, the entries of $\Phi \in \mathbb{R}^{m \times n}$ are drawn from a uniform distribution ranging from 0 to 1. We note that, in both cases, Algorithm 3 is able to recover the full support after a sufficient number of iterations, whereas Algorithm 2 is only able to recover at most about 20% of the support, regardless of the number of iterations. Since Algorithm 2 relies on row norms to estimate the joint support, it is to be expected that the relatively large value of the corruption would cause it to often be erroneously chosen to be part of the joint support estimate. As a result, this experiment highlights the advantage of the binning in Algorithm 3 in the presence of a single corruption with large magnitude.




Fig. 1: Comparing cMMV-SRK and MMV-SRK for support recovery when there is a single corrupt entry per signal whose magnitude is drawn from $\mathcal{N}(7, 1)$.

In Figure 2, we experiment further with the magnitude of the corruption. Here we have that $m = 1000$, $n = 100$, and $k = 10$ but instead of the corrupt entries being drawn from a normal distribution with mean 7 and standard deviation 1, it is drawn from a standard normal distribution, as are the entries in the support. This allows us to test the robustness of our method to the choice of distribution. Note that Algorithm 2 is able to find an increasingly accurate approximation for the support in this case, and will be able to recover the full support after a sufficiently large number of iterations. Because the magnitudes of the corruptions are smaller, the algorithm still has a chance of detecting the correct support using row norms to estimate \mathcal{S} . However, Algorithm 3 is able to obtain an accurate support estimate with a small number of projections compared to Algorithm 2.

In the next two experiments, we test the robustness of our proposed algorithm against multiple corruptions. In Figure 3, we allow for each signal to have multiple corruptions. For each signal, i.e. column of \mathbf{X} , we uniformly at random select an integer from 1-3 to be the number of corruptions. The value of the corruptions are drawn i.i.d. from $\mathcal{N}(7, 1)$ and an example of the resulting matrix can be seen in Figure 3a. The performance of the methods can be seen in Figure 3b. We note that the results of this experiment are very similar to those of the experiment in Figure 1 since the corruptions are drawn from the same distribution. Again, due to the use of




Fig. 2: Comparing cMMV-SRK and MMV-SRK for support recovery when there is a single corrupt entry per signal whose magnitude is drawn from $\mathcal{N}(0, 1)$.

row norms, in the presence of multiple corruptions Algorithm 2 gives a less accurate joint support estimate, recovering no more than about 15% of the support.




Fig. 3: Investigating the robustness of cMMV-SRK and MMV-SRK when a random number of (multiple) corruptions are introduced. Here, a signal can have between 1 and 3 corruptions. The corrupts magnitude of the corruptions are drawn from $\mathcal{N}(7, 1)$.

Figure 4a shows the performance results for Algorithm 3 in a simulated online setting. Instead of allowing the algorithm to loop for a fixed number of projections for each signal, for approximately 90% of the signals $\tilde{\tau}_j \in [5, 15]$ and for the remaining signals $\tilde{\tau}_j \in [95, 100]$. The purpose of this is to simulate a variation in the amount of time that a system has to work with a signal. The longer runs represent stalls in the online setting. For each signal, we first draw a random Bernoulli variable z with probability of success $p = 0.1$. If $z = 1$, then we choose an integer in $[95, 100]$ uniformly at random. If $z = 0$, then an integer in $[5, 15]$ is chosen uniformly at random.

Algorithm 2 cannot be investigated under this setting as the support estimate relies on processing all signals in every iteration. We note that with respect to other parameters, the only difference between this experiment and that in Figure 3 is the size of J . We choose J to be large enough such that the maximal number of projections made is 15000 (as in Figure 3).

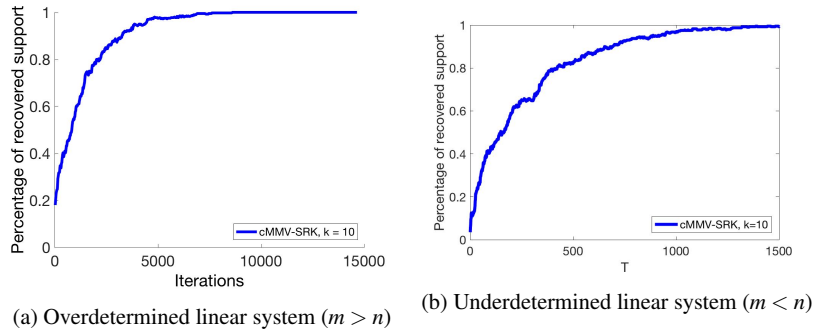


Fig. 4: Investigating the robustness of cMMV-SRK in a simulated online setting with a random number of (multiple) corruptions. Here, a signal can have between 1 and 3 corruptions whose magnitudes are drawn from $\mathcal{N}(7, 1)$ and we consider the over and under determined linear system settings.

The following experiment is motivated by compressed sensing and utilizes an underdetermined linear systems as opposed to an overdetermined system. We repeat the parameters as in Figure 4a with the exception of the measurement matrix, which has $m = 100$ rows and $n = 500$ columns, and the total number of signals $J = 1500$. The results can be found in Figure 4b. In the underdetermined case, the proposed algorithm is still successful in recovering the support of the signal.

Finally, we test the robustness of our algorithm on the hyperspectral diffuse optical tomography motivating problem discussed in the introduction. For this experiment, we simulate absorption coefficient values for a two-dimensional circular sample of tissue of radius 25 centimeters, centered at the origin, with a circular tumor of radius 5 centimeters centered at the point $(-15, -10)$. See Figure 5a. Each signal thus represents a reconstruction of the absorption coefficient value at each point in a mesh of size 541 over the sample area. The number of measurements corresponds to the number of source-detector pairs in the imaging process. We use a random Gaussian measurement matrix with $m = 248$ and $n = 541$, and have $J = 200$ total signals, each corresponding to a different wavelength at which the tissue was imaged. As in previous experiments, 1 to 3 corruptions for each signal are drawn from a normal distribution with mean equal to the average value of the absorption coefficient for the cancerous cells at each wavelength, and standard deviation a quarter of the distance between that value and the value of the absorption coefficient for the healthy cells. The online setting is not used for this experiment. The results can be found

in Figure 6. We see that the proposed algorithm is still successful in recovering the support of the signal.

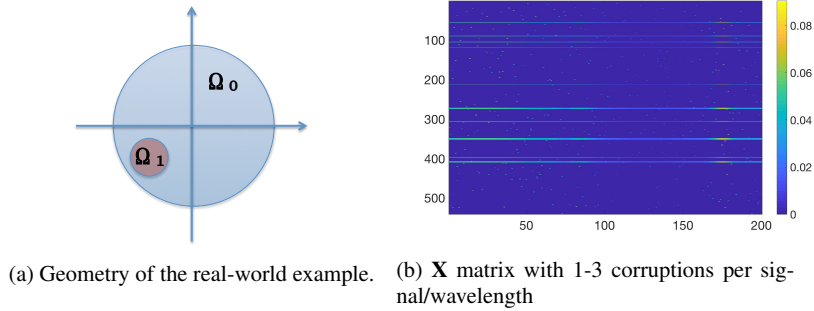


Fig. 5: Testing the robustness of cMMV-SRK in a real-world setting by using a simulated hyperspectral diffuse optical tomography (hyDOT) model. Healthy cells are in Ω_0 while cancerous cells are located in Ω_1 .

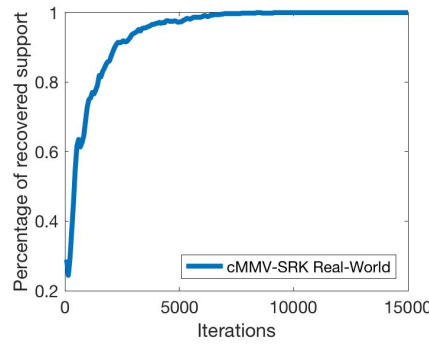


Fig. 6: Investigating the robustness of cMMV-SRK in a simulated real-world setting (hyperspectral diffuse optical tomography) when random multiple corruptions occur in each signal. In this setting, the measurement matrix is underdetermined ($m = 248$, $n = 541$).

The experiments shown in this section highlights the usefulness of Algorithm 3 for support recovery of jointly sparse MMVs, especially in the presence of (even quite large) corruptions. In each comparison between Algorithm 2 and Algorithm 3, our proposed method outperforms the previously proposed method for support recovery. Additionally, the proposed method lends itself to the online setting.

5 Conclusion

In this work, we construct an algorithm for the support recovery of corrupted jointly sparse MMV. Our empirical results demonstrate that the proposed algorithm outperforms the previously proposed method, cMMV-SRK, for support recovery of jointly sparse MMV, specifically when corruptions are present in the signal. Furthermore, empirical evidence indicates that our method is robust to the magnitude of corruptions and the number of corruptions. This improvement is due to the fact that the support estimate in Algorithm 2, as many other signal recovery approaches for the jointly sparse MMV problem, depends on the row norms of the signals, which in this case would be dominated by corruptions. In comparison, the estimate for Algorithm 3 only depends on the number of times an index appears in the support estimate of a signal. Lastly, our method lends itself well into the online setting when measurement vectors are streaming in continuously. We leave the analytical study of our method for future work.

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.

References

- [1] Aggarwal, H.K., Majumdar, A.: Extension of sparse randomized kaczmarz algorithm for multiple measurement vectors. pp. 1014–1019. IEEE In. Conf. Pat. Recog. (2014)
- [2] Arridge, S.R., Schotland, J.C.: Optical tomography: forward and inverse problems. *Inverse Probl.* **25**(12), 123,010 (2009)
- [3] Bottou, L.: Large-scale machine learning with stochastic gradient descent. In: Proc. of COMPSTAT, pp. 177–186. Springer (2010)
- [4] Candès, E.J., Fernandez-Granda, C.: Towards a mathematical theory of super-resolution. *Commun. Pur. Appl. Math.* **67**(6), 906–956 (2014)
- [5] Fodor, L., Elman, M., Ullmann, Y.: *Aesthetic applications of intense pulsed light*. Springer-Verlag London (2011)
- [6] Gordon, R., Bender, R., Herman, G.T.: Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography. *J. Theoret. Biol.* **29**, 471–481 (1970)
- [7] Hamaker, C., Solmon, D.C.: The angles between the null spaces of X-rays. *J. Math. Anal. Appl.* **62**(1), 1–23 (1978)
- [8] Heckel, R., Soltanolkotabi, M.: Generalized line spectral estimation via convex optimization. *IEEE T. Inform. Theory* **64**(6), 4001–4023 (2018)

- [9] Herman, G., Meyer, L.: Algebraic reconstruction techniques can be made computationally efficient. *IEEE T. Medical Imaging* **12**(3), 600–609 (1993)
- [10] Kaczmarz, S.: Angenäherte auflösung von systemen linearer gleichungen. *Bull. Int. Acad. Polon. Sci. Lett. Ser. A* pp. 335–357 (1937)
- [11] Larusson, F., Fantini, S., Miller, E.L.: Hyperspectral image reconstruction for diffuse optical tomography. *Biomed. Opt. Express* **2**(4), 946–965 (2011)
- [12] Laska, J.N., Davenport, M.A., Baraniuk, R.G.: Exact signal recovery from sparsely corrupted measurements through the pursuit of justice. In: *Asilomar Conf. Signals, Sys., Comput*, pp. 1556–1560. IEEE (2009)
- [13] Li, S., Yang, D., Tang, G., Wakin, M.B.: Atomic norm minimization for modal analysis from random and compressed samples. *IEEE T. Signal Proces.* **66**(7), 1817–1831 (2018)
- [14] Li, Y., Chi, Y.: Off-the-grid line spectrum denoising and estimation with multiple measurement vectors. *IEEE T. Signal Proces.* **64**(5), 1257–1269 (2016)
- [15] Lu, G., Fei, B.: Medical hyperspectral imaging: a review. *J. Biomed. Opt.* **19**(1), 010,901 (2014)
- [16] Mansour, H., Yilmaz, O.: A fast randomized kaczmarz algorithm for sparse solutions of consistent linear systems. *arXiv preprint arXiv:1305.3803* (2013)
- [17] Strohmer, T., Vershynin, R.: Comments on the randomized kaczmarz method. *J. Fourier Anal. Appl.* **15**(4), 437–440 (2009)
- [18] Studer, C., Kuppinger, P., Pope, G., Bolcskei, H.: Recovery of sparsely corrupted signals. *IEEE T. Inform. Theory* **58**(5), 3115–3130 (2012)
- [19] Tang, G., Bhaskar, B.N., Shah, P., Recht, B.: Compressed sensing off the grid. *IEEE T. Inform. Theory* **59**(11), 7465–7490 (2013)
- [20] Yang, Z., Xie, L.: Exact joint sparse frequency recovery via optimization methods. *IEEE T. Signal Proces.* **64**(19), 5145–5157 (2014)