Signal Recovery from Unlabeled Samples

Saeid Haghighatshoar, *Member, IEEE*, Giuseppe Caire, *Fellow, IEEE*

Abstract—In this paper, we study the recovery of a signal from a collection of unlabeled and possibly noisy measurements via a measurement matrix with random i.i.d. Gaussian components. We call the measurements unlabeled since their order is missing, namely, it is not known a priori which elements of the resulting measurements correspond to which row of the measurement matrix. We focus on the special case of ordered measurements, where only a subset of the measurements is kept and the order of the taken measurements is preserved. We identify a natural duality between this problem and the traditional Compressed Sensing, where we show that the unknown support (location of nonzero elements) of a sparse signal in Compressed Sensing corresponds in a natural way to the unknown location of the measurements kept in unlabeled sensing. While in Compressed Sensing it is possible to recover a sparse signal from an under-determined set of linear equations (less equations than the dimension of the signal), successful recovery in unlabeled sensing requires taking more samples than the dimension of the signal. We develop a low-complexity alternating minimization algorithm to recover the initial signal from the set of its unlabeled samples. We also study the behavior of the proposed algorithm for different signal dimensions and number of measurements both theoretically and empirically via numerical simulations. The results are a reminiscent of the phase-transition similar to that occurring in Compressed Sensing.

Index Terms—Unlabeled Sensing, Compressed Sensing, Alternating minimization algorithm.

I. INTRODUCTION

The recovery of a vector-valued signal $y \in \mathbb{R}^k$ from a set of linear and possibly noisy measurements $x = By + w$, with an $n \times k$ measurement matrix B, is the classical problem of linear regression in statistical inference and is arguably the most widely-studied problem in statistics, mathematics, and computer science. For $n \geq k$, the recovery of y corresponds to an over-determined set of linear equations, whose statistically optimal solution under the additive Gaussian noise w is given by the well-known method of *least-square*. For $n < k$, on the other hand, one deals with an under-determined set of linear equations, which is only solvable if some additional a priori information about the signal y is available. For example, the whole field of Compressed Sensing (CS) deals with the recovery of the signal y when it is sparse or more generally compressible, i.e., it has only a few significantly large coefficients when represented in a suitable bases [1–3].

Almost all the past research in linear regression mainly deals with exploiting the underlying signal structure, whereas it is generally assumed that the regression matrix B is fully known. In practice, the matrix B is implemented through a

measurement device, where due to physical limitations, there might be some uncertainty or mismatch between the intended matrix B and the one realized via measurement devices. This has resulted in a surge of interest in generalized linear regression problems in which the matrix B is mismatched [4, 5] or is known only up to some unknown transformation. In this paper, we are interested in a case frequently encountered in practical problems given by

$$x = SBy + w, \qquad (1)$$

where S is an unknown linear operator acting on the fully known measurement matrix B. It is assumed that although S is not a priori known, it belongs to a fully known set of linear transformations $\mathcal{T}$. An interesting special case of (1) arises when the set $\mathcal{T}$ consists of matrices S having 0-1 components with only a single 1 in each row. In such a case, each $S \in \mathcal{T}$ corresponds to sampling (selecting) some of elements of By, where the label of the measurements is not fully known due to unknown S (unlabeled sampling). Identifying S in (1), therefore, corresponds to associating the obtained measurements x to their corresponding linear equations via the matrix B. Once S is identified, (1) reduces to a linear regression problem whose solution can be obtained via standard techniques.

In [6, 7], a variant of this problem coined *unlabeled sensing* was studied when $\mathcal{T}$ is the set of all permutations of n measurements taken by B. It was shown that if the measurement matrix B is generated randomly, any arbitrary k-dim signal can be recovered from a set of $n = 2k$ noiseless unlabeled measurements, where this bound was shown to be tight. In [8], the authors studied a similar problem but rather than recovering the unknown signal y, they obtained a scaling law of the *signal-to-noise ratio* (SNR) required for identifying the unknown permutation matrix in $\mathcal{T}$. A practical scenario well-modeled by (1) is sampling in the presence of jitter [9], in which $\mathcal{T}$ consists of 0-1 matrices with 1s as their diagonal elements and with some off-diagonal 1s representing the location of the jittered samples. A similar problem arises in molecular channels where due to synchronization issues the correct label of measurements is available up to a jitter [10]. Reconstruction of the phase-space dynamics of a linear/nonlinear dynamical system from the jittered time-domain output of the system is another example of (1), where robust reconstruction algorithms are needed to eliminate the effect of jitter [11]. Unlabeled regression (1) also arises in noncooperative multi-target tracking, e.g., in radar, where the receiver only observes the unlabeled data associated to all the targets, thus, $\mathcal{T}$ consists of the set of all permutations corresponding to all possible data-associations. Once the observations are suitably associated to the targets, the location of the targets can be estimated/predicted via a standard linear

regression typically implemented as a Kalman Filter [12]. The data association, however, becomes quite challenging when a large number of targets are tracked simultaneously. For example, a naive approach requires an exhaustive search over a large number of permutations, which is formidable in practice. A quite similar scenario arises in robotics where a well-known classical problem is *simultaneous localization and mapping* (SLAM) where several robots measure their relative coordinates and recovery of the underlying geometry requires suitably permuting the data.

A different line of work well-modeled by (1) is the genome assembly problem from shotgun reads [13, 14] in which a vector (sequence) $y \in \{A, T, G, C\}^k$ of length k is assembled from an unknown permutation of its sub-vector measurements called "reads". Designing efficient recovery (assembly) algorithms is still an active research area (please refer to [14] and the refs. therein.). Communication over the classical noisy deletion channel is another example than can be modeled by (1), where B represents the linear encoding matrix with elements belonging to a finite field $\mathbb{F}_q$, for some prime number q, the vector $y \in \mathbb{F}_q^k$ denotes the k-dim vector containing k information symbols, $w \in \mathbb{F}_q^m$ is the additive noise of the channel, and where the set $\mathcal{T}$ consists of all selection operators that keep only $m \leq n$ out of n encoded symbols in By while preserving their order. In particular, in contrast with the erasure channel, where the location of erased symbols is known, in a deletion channel the location of deleted symbols is missing. This makes designing good encoding and decoding techniques as well as identifying the capacity quite challenging [15].

Contribution. Since satisfactory efficient algorithms to solve the general unlabeled sampling problem are yet unknown, in this paper we make progress towards this goal by addressing a relevant subproblem that we name Unlabeled Ordered Sampling (UOS). More specifically, we study a variant of (1) with a random measurement matrix B and with a $\mathcal{T}$ given by the set of all 0-1 *ordered sampling* matrices, where each $S \in \mathcal{T}$ selects only m out of n components of By for some $m \leq n$ while preserving their relative order. We discover a duality between this problem and the traditional Compressed sensing problem, where the unknown location of samples in the former corresponds in a natural way to the unknown location of nonzero coefficients of the signal in the latter. To the best of our knowledge, this is the first paper addressing the underlying duality connection between the unlabeled sensing in (1) and the traditional Compressed Sensing. Also, designing a low-complexity algorithm for recovering the desired signal from its unlabeled samples is generally considered to be a really challenging problem [6–9]. In particular, in almost all practical situations $\mathcal{T}$ consists of very large number of transformations, thus, a naive exhaustive search over $\mathcal{T}$ is formidable. For example, $\mathcal{T}$ in the unlabeled sensing problem studied in [6–8] consists of all possible permutation matrices, which is enormously large for all practical signal dimensions. In this paper, however, we are able to exploit the underlying ordered property of the unlabeled samples to design an efficient *Alternating Minimization* algorithm to recover the target signal y. We also analyze the noise stability of the proposed algorithm theoretically and empirically via numerical simulations.

Notation. We denote vectors with small letters (e.g., x), matrices with capital letters (e.g., X), and sets with capital calligraphic letters (e.g., $\mathcal{X}$). We denote by $\{\mathscr{X}_H : H \in \mathcal{H}\}$ a stochastic process consisting of random variables $\mathscr{X}_H$ labelled with the elements of the set $\mathcal{H}$. For integers $k_1, k_2 \in \mathbb{Z}$, we use the shorthand notation $[k_1 : k_2] = \{k_1, k_1 + 1, \ldots, k_2\}$, where the set is empty when $k_2 < k_1$, and use the simplified notation $[k_1]$ for $[1 : k_1]$. We denote the k-th component of a vector x by x_k and a sub-vector of x with indices in the range $k_1 : k_2$ by $x_{k_1 : k_2}$. We denote the element of a matrix X at row r and column c with $X_{r,c}$ and use an indexing notation similar to that for vectors for submatrices of X, namely, $X_{r,c}$, $X_{r,c_1 : c_2}$, $X_{r_1 : r_2, c}$, and $X_{r_1 : r_2, c_1 : c_2}$. We use $\mathsf{Tr}(.)$ for the trace operator and $\|X\|_{\mathsf{F}} = \sqrt{\mathsf{Tr}(XX^{\mathsf{H}})}$ for the Frobenius norm of a matrix. We denote a sequence of vectors and sequence of matrices by upper indices, e.g., $x^1, x^2, \cdots$ and $X^1, X^2, \cdots$. We denote a Gaussian distribution with a mean μ and a variance σ^2 by $\mathsf{N}(\mu, \sigma^2)$. We use $O(.)$ for the big-O notation, where for two sequences $\{a_n\}_{n=0}^{\infty}$ and $\{b_n\}_{n=0}^{\infty}$ we say $a_n = O(b_n)$ if there are positive numbers c_1, c_2 such that $a_n \leq c_1 b_n + c_2$ for $n \geq 0$.

II. STATEMENT OF THE PROBLEM

In this section, we first start from the more familiar Compressed Sensing (CS) problem. Then, we introduce the Unlabeled Ordered Sampling (UOS) and make a duality connection between the two.

A. Basic Setup

Let n and m be positive integers with $m \leq n$, and let $\mathcal{I} = \{i_1, \ldots, i_m\} \subset [n]$ be a subset of $[n]$ consisting of ordered elements of $[n]$ satisfying $i_l < i_{l+1}$. We define the *Lift-Up* (LU) operator associated to $\mathcal{I}$ as a linear map from $\mathbb{R}^m$ to $\mathbb{R}^n$ given by the $\{0, 1\}$-valued $n \times m$ tall matrix L with

$$L_{p,q} = \begin{cases} 1 & (p,q) = (i_l, l) \text{ for } l \in [m], \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

The operator L embeds m components of $x = (x_1, \ldots, x_m)^{\mathsf{T}}$ into the index set $\mathcal{I}$ in the n-dim vector Lx, while keeping their relative order, i.e., $(Lx)_{i_l} = x_l$ for $l \in [m]$, and fills the rest with 0. For example, for $n = 4$, $m = 3$, $x = (x_1, x_2, x_3)$, and $\mathcal{I} = \{1, 3, 4\}$, we have $Lx = (x_1, 0, x_2, x_3)$. For fixed m, n with $m \leq n$, we define the collection of all $\binom{n}{m}$ LU operators by $\mathcal{L}_{n,m}$.

B. Compressed Sensing

In CS [1–3], the goal is to recover a sparse or compressible signal by taking less measurements than the dimension of the signal. We call a signal $z \in \mathbb{R}^n$ m-sparse (m-compressible) if it has only m nonzero (significantly large) components. In this paper, for simplicity, we focus on sparse rather than compressible signals.

We fix m, n with $m \leq n$ as before. We define an instance of the CS problem for an m-sparse signal $z \in \mathbb{R}^n$ by the triple (x, L, A), where $x \in \mathbb{R}^m$ denotes the nonzero elements of z, where $L \in \mathcal{L}_{n,m}$ encodes the location of these nonzero
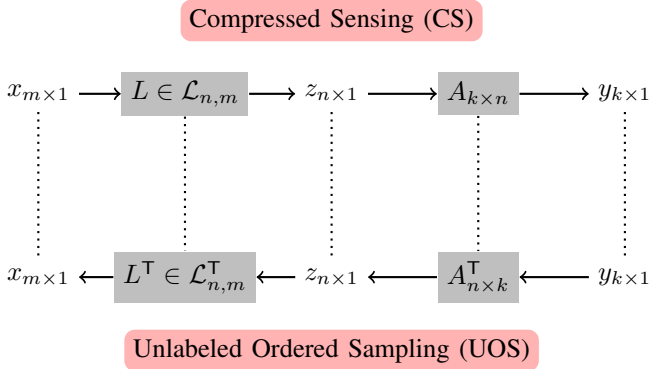
Fig. 1: Comparison between Compressed Sensing and Unlabeled Ordered Sensing.

elements, and where A is the $k \times n$ matrix whose rows correspond to k linear measurements. The n-dim m-sparse signal z is generated by embedding the m components of x via lifting matrix L by $z = Lx$, where it is seen that z, albeit being n-dim, has at most m nonzero components. This has been illustrated in Fig. 1. In CS, the sparse signal z is sampled via the measurement matrix A producing k linear measurements $y = Az$. The crucial idea in CS is that by exploiting the underlying sparsity, z can be recovered by taking less samples than its embedding dimension n. The practical regime of parameters for the CS is given by $m \leq k \leq n$.

C. Unlabeled Ordered Sampling

By changing the role of y and x in the CS problem, we obtain an instance of the dual problem which we call *Unlabeled Ordered Sampling* (UOS). This is also illustrated in Fig. 1. In the dual problem, a k-dim signal y is oversampled via the tall matrix A^T, which gives n measurements in $z = A^\mathsf{T} y$. The resulting over-complete measurements ($n \geq k$) are subsampled by the $m \times n$ matrix L^T, which selects only m out of n measurements in z and gives the final measurement vector x. Compared with the CS, where the support of the nonzero values in the signal z is unknown, in UOS the indices/labels of the final measurements are missing. However, it is not difficult to check that, due to the special structure of L^T, the relative order of the measurements is still preserved. We define the set of all such selection matrices by $\mathcal{S}_{m,n} := \{L^\mathsf{T} : L \in \mathcal{L}_{n,m}\}$. In contrast with the lift-up matrices in $\mathcal{L}_{n,m}$, which embed signals with a lower dimension m in a higher dimension n, the selection matrices in $\mathcal{S}_{m,n}$ reduce the dimensionality by selecting only m out of n components of their input signal (while keeping their order). The practical regime of parameters for the UOS is given by $k \leq m \leq n$.

III. RESTRICTED ISOMETRY PROPERTY

A. Basic Setup

For the rest of the paper, we focus on UOS problem. We consider a k-dim signal y and an n-dim vector of measurements $z = By$ taken via the $n \times k$ measurement matrix B, where $B = A^\mathsf{T}$ with A being the $k \times n$ matrix in the Compressed Sensing variant (see Fig. 1). This is illustrated

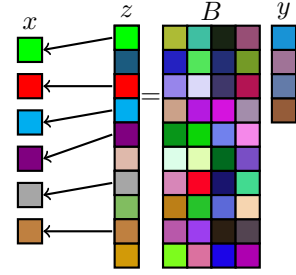in Fig. 2. An instance of UOS problem is defined by the

Fig. 2: Unlabeled Ordered Sampling.

triple (y, S, B), where the goal is to recover y from the noisy unlabeled measurements $x = SBy + w$, with w denoting the m-dim measurement noise, without explicit knowledge of the selection matrix $S \in \mathcal{S}_{m,n}$. This corresponds to a variant of unlabeled sensing problem in (1) with the set of possible transformations $\mathcal{T}$ given by $\mathcal{S}_{m,n}$. Using the vec notation, we can write

$$x = \text{vec}(SBy) + w = (y^\mathsf{T} \otimes S)\mathbb{b} + w, \tag{3}$$

where $\mathbb{b} = \text{vec}(B)$ denotes the vector obtained by stacking the columns of matrix B and where w denotes the measurement noise. Note that $\mathbb{b}$ establishes a linear map $H \mapsto H\mathbb{b}$ from the signal set $\mathcal{A} := \{y^\mathsf{T} \otimes S : y \in \mathbb{R}^k, S \in \mathcal{S}_{m,n}\}$ into $\mathbb{R}^m$. We will use both the notations SBy and $(y^\mathsf{T} \otimes S)\mathbb{b}$ interchangeably across the paper. Also, for a signal $H \in \mathcal{A}$, we define the measurement *Signal-to-Noise Ratio* (SNR) by $\text{snr} = \frac{\|H\mathbb{b}\|^2}{\|w\|^2}$.

B. Restricted Isometry Property on $\mathcal{A}$

Our goal is to recover the initial signal y. Since the signal set $\mathcal{A}$ is unbounded, a requirement is that at least $\|y\|$ should be feasibly recovered from the noiseless measurements $(y^\mathsf{T} \otimes S)\mathbb{b}$. A sufficient condition for this to hold is the *Restricted Isometry Property* (RIP) over $\mathcal{A}$, which resembles a similar property used in the traditional CS.

Definition *1 (RIP over $\mathcal{A}$):* Let B be the measurement matrix and let $\mathbb{b} = \text{vec}(B)$. The linear map $H \mapsto H\mathbb{b}$ induced by $\mathbb{b}$ satisfies the RIP with a constant $\epsilon \in (0, 1)$ over $\mathcal{A}$ if

$$(1 - \epsilon)\|H\|_\mathsf{F}^2 \leq \|H\mathbb{b}\|^2 \leq (1 + \epsilon)\|H\|_\mathsf{F}^2 \tag{4}$$

holds for all $H \in \mathcal{A}$, where $\|.\|_\mathsf{F}$ denotes the Frobenius norm of a matrix. $\diamond$

Note that since the signal set $\mathcal{A}$ is star-shaped, i.e., $\lambda \mathcal{A} \subset \mathcal{A}$ for any $\lambda \in \mathbb{R}$, without any loss of generality, we can investigate the RIP over the subset $\mathcal{A}^\circ := \{H \in \mathcal{A} : \frac{\|H\|_\mathsf{F}^2}{m} = 1\}$ consisting of all signals $y^\mathsf{T} \otimes S$ with $\|y\| = 1$.

In this section, we prove that under suitable conditions on m, n, k, we can obtain a matrix B satisfying the RIP condition in Definition 4 by sampling components of B i.i.d. from the Gaussian distribution. We need some notation first. Let $H \in \mathcal{A}$ and let us define $\mathscr{D}_H = \frac{\|H\mathbb{b}\|^2}{m}$, where now we assume that the components of $\mathbb{b}$ are i.i.d. $\mathsf{N}(0, 1)$. It is not difficult to check that, for such a $\mathbb{b}$, we have $\mathbb{E}[\mathscr{D}_H] = \frac{\|H\|_\mathsf{F}^2}{m}$. Thus, the RIP condition in (4) can be equivalently written as

$$\left| \mathscr{D}_H - \mathbb{E}[\mathscr{D}_H] \right| \leq \epsilon \mathbb{E}[\mathscr{D}_H], \tag{5}$$

which should hold for all $H \in \mathcal{A}^\circ$. Finding a suitable RIP constant $\epsilon \in (0,1)$ requires obtaining a concentration bound for the random variable $\mathcal{R}_{\mathcal{A}^\circ} := \sup_{H \in \mathcal{A}^\circ} |\mathscr{D}_H - \mathbb{E}[\mathscr{D}_H]|$. Note that for a Gaussian vector $\mathbb{b}$, as H ranges over $\mathcal{A}^\circ$, we obtain the stochastic process $\{\mathscr{X}_H : H \in \mathcal{A}^\circ\}$ where $\mathscr{X}_H = H\mathbb{b}$ is a Gaussian vector in $\mathbb{R}^m$ parametrized with $H \in \mathcal{A}^\circ$. The random variables $\mathcal{R}_{\mathcal{A}^\circ}$ denotes the maximum deviation of the square of ℓ_2 norm of this vector Gaussian process $\mathscr{D}_H := \frac{\|\mathscr{X}_H\|^2}{m}$ from its average $\mathbb{E}[\mathscr{D}_H]$ over the whole index set $\mathcal{A}^\circ$. Our next result indicates that for suitably selected parameters m, n, k and also ϵ, the concentration bound (5) holds simultaneously for all $H \in \mathcal{A}^\circ$ with a very high probability.

Proposition 1: Let B be a random matrix with i.i.d. $\mathsf{N}(0,1)$ components and let $\mathbb{b} = \text{vec}(B)$. Then, there is a constant $c > 0$ such that B satisfies the RIP condition on $\mathcal{A}$ with a probability larger than $1 - 2(1 + \frac{2}{\epsilon})^k \binom{n}{m} e^{-cm\epsilon^2}$. $\quad\square$

Proof: First, note that for a given $H = y^\mathsf{T} \otimes S$, from $H\mathbb{b} = SBy$ and $\mathbb{E}[\|H\mathbb{b}\|^2] = \|H\|_\mathsf{F}^2$, the concentration bound in (5) can be equivalently written as

$$(1 - \epsilon)m \leq \|SBy\|^2 \leq (1 + \epsilon)m, \tag{6}$$

where we used the fact $\|H\|_\mathsf{F}^2 = m\|y\|^2 = m$ for $H \in \mathcal{A}^\circ$. We first consider $H = y^\mathsf{T} \otimes S \in \mathcal{A}^\circ$ for a fixed S and y. Note that for such an S and y, the random vector SBy has a Gaussian distribution with i.i.d. $\mathsf{N}(0,1)$ components. Using the Gaussian concentration result, there is a $c > 0$ such that

$$\mathbb{P}\left[\left|\|SBy\|^2 - m\right| > \frac{\epsilon}{2}m\right] \leq 2e^{-cm\epsilon^2}. \tag{7}$$

We first generalize this concentration result to the case where y is an arbitrary vector in the unit sphere $\mathbb{S}^{k-1}$, which would imply that the operator norm of SB is concentrated around $\sqrt{m}$. To prove the result, we first consider a discrete $\frac{\epsilon}{2}$-net (grid) of minimal size N over $\mathbb{S}^{k-1}$ given by $\mathcal{N}_\epsilon = \{g^1, \ldots, g^N\}$ that satisfies $\sup_{y \in \mathbb{S}^{k-1}} \min_{g \in \mathcal{N}_\epsilon} \|g - y\| \leq \frac{\epsilon}{2}$. Consider a set of N spheres with centers belonging to the net $\mathcal{N}_\epsilon$ and each having a radius $\frac{\epsilon}{2}$. It is not difficult to see that all these spheres lie inside a sphere of radius $1 + \frac{\epsilon}{2}$ centered at the origin. Thus, using the volume inequality $N(\frac{\epsilon}{2})^k \text{vol}(\mathbb{B}_2^k) \leq (1 + \frac{\epsilon}{2})^k \text{vol}(\mathbb{B}_2^k)$ for the k-dim unit ball $\mathbb{B}_2^k$, we obtain that such a minimal net consists of at most $N \leq (1 + \frac{2}{\epsilon})^k$ points. Notice that the operator norm of SB can be well estimated using points in $\mathcal{N}_\epsilon$, i.e.,

$$\max_{y \in \mathbb{S}^{k-1}} \|SBy\| \leq (1 + \frac{\epsilon}{2}) \max_{g \in \mathcal{N}_\epsilon} \|SBg\|. \tag{8}$$

Using (8) and (7) and applying the union bound over the net $\mathcal{N}_\epsilon$, we obtain that

$$\mathbb{P}\left[\sup_{y \in \mathbb{S}^{k-1}} \left|\|SBy\|^2 - m\right| > m\epsilon\right]$$
$$\leq \mathbb{P}\left[\sup_{g \in \mathcal{N}_\epsilon} \left|\|SBg\|^2 - m\right| > m\frac{\epsilon}{2}\right]$$
$$\leq 2(1 + \frac{2}{\epsilon})^k e^{-cm\epsilon^2}. \tag{9}$$

Finally applying the union bound over all $\binom{n}{m}$ possible selection matrices $S \in \mathcal{S}_{m,n}$, we have

$$\mathbb{P}\left[\sup_{y \in \mathbb{S}^{k-1}, S} \left|\|SBy\|^2 - m\right| > m\epsilon\right] \leq 2(1 + \frac{2}{\epsilon})^k \binom{n}{m} e^{-cm\epsilon^2},$$

which implies the desired RIP bound on $\mathcal{A}$ as in (5) with a probability of at least $1 - 2(1 + \frac{2}{\epsilon})^k \binom{n}{m} e^{-cm\epsilon^2}$. $\quad\blacksquare$

C. Restricted Isometry Property on $\mathcal{A} - \mathcal{A}$

In terms of signal recovery, we need an RIP for $\mathcal{A} - \mathcal{A} := \{H - H' : H, H' \in \mathcal{A}\}$ denoting the Minkowski difference of $\mathcal{A}$. Our approach in this section follows from similar techniques proposed in [16].

Definition 2 (RIP over $\mathcal{A} - \mathcal{A}$): Let B be the measurement matrix and let $\mathbb{b} = \text{vec}(B)$. The linear map $H \mapsto H\mathbb{b}$ induced by $\mathbb{b}$ satisfies the RIP with a constant $\epsilon \in (0,1)$ over $\mathcal{A} - \mathcal{A}$ if

$$(1 - \epsilon)\|H - H'\|_\mathsf{F}^2 \leq \|(H - H')\mathbb{b}\|^2 \leq (1 + \epsilon)\|H - H'\|_\mathsf{F}^2$$

holds for all $H, H' \in \mathcal{A}$. $\quad\diamond$

Since $0 \in \mathcal{A}$, as we will see, this RIP on $\mathcal{A} - \mathcal{A}$ implies the RIP we developed for $\mathcal{A}$ in the previous section. Therefore, without loss of generality, we assume that $\|H\|_\mathsf{F}$ and as a result $\|y\|$ for any arbitrary H can be well estimated from $\|H\mathbb{b}\| = \|SBy\|$ and focus on the normalized signal set $\mathcal{A}^\circ$. Similar to Section III-B, we show that, under suitable conditions on m, n, k, a matrix B with i.i.d. $\mathsf{N}(0,1)$ components satisfies the RIP over $\mathcal{A} - \mathcal{A}$ with a very high probability.

Let B be an $n \times k$ matrix with i.i.d. $\mathsf{N}(0,1)$ components and let $\mathbb{b} = \text{vec}(B)$. For $H, H' \in \mathcal{A}^\circ$, we define $\mathscr{D}_{H,H'} = \frac{\|(H-H')\mathbb{b}\|^2}{m}$, where we can check that $\mathbb{E}[\mathscr{D}_{H,H'}] = \frac{\|H-H'\|_\mathsf{F}^2}{m}$. Therefore, we can write the RIP condition in Definition 2 as

$$\left|\mathscr{D}_{H,H'} - \mathbb{E}[\mathscr{D}_{H,H'}]\right| \leq \epsilon \mathbb{E}[\mathscr{D}_{H,H'}], \tag{10}$$

which should hold simultaneously for all $H, H' \in \mathcal{A}^\circ$. First, note that $d_{H,H'} = \sqrt{\mathbb{E}[\mathscr{D}_{H,H'}]}$ defines a metric over $\mathcal{A}^\circ$, where for $H = y \otimes S$ and $H' = y' \otimes S'$ belonging to $\mathcal{A}^\circ$

$$d_{H,H'} = \sqrt{\nu_{S,S'}\|y - y'\|^2 + (1 - \nu_{S,S'})(\|y\|^2 + \|y'\|^2)}, \tag{11}$$

where $\nu_{S,S'} = \frac{\text{Tr}(S^\mathsf{T}S')}{m}$ denotes the fraction of identical rows in the selection matrices S and S' (similarity metric). Note that $d_{H,H'} \leq 2$ for any $H, H' \in \mathcal{A}^\circ$, thus, $\mathcal{A}^\circ$ is a bounded set under this metric. This implies that we can obtain the desired RIP over $\mathcal{A}^\circ - \mathcal{A}^\circ$ by deriving a concentration bound for

$$\mathcal{R}_{\mathcal{A}^\circ - \mathcal{A}^\circ} := \sup_{H, H' \in \mathcal{A}^\circ} \left|\mathscr{D}_{H,H'} - \mathbb{E}[\mathscr{D}_{H,H'}]\right|, \tag{12}$$

where we obtain the following result.

Proposition 2: Let B be a random matrix with i.i.d. $\mathsf{N}(0,1)$ components and let $\mathbb{b} = \text{vec}(B)$. Then, there is a constant $c > 0$ such that B satisfies the RIP condition on $\mathcal{A} - \mathcal{A}$ with a probability larger than $1 - 4(1 + \frac{2}{\epsilon})^{2k} \binom{n}{m}^2 e^{-cm\epsilon^2}$. $\quad\square$

Proof: Proof in Appendix A. $\quad\blacksquare$

D. Guarantee for Signal Recovery

Using Proposition 2, we obtain a universal recovery guarantee for all the signals $H \in \mathcal{A}$. Let $x = H\mathbb{b} + w$ be the set of m linear and noisy measurements from H, where w denotes the measurement noise with $\|w\| = \frac{\|H\mathbb{b}\|_{\mathsf{F}}}{\sqrt{\mathsf{snr}}}$ with snr denoting the measurement SNR as before. We assume that m, n, k are selected such that $\mathbb{b}$ satisfies the RIP over $\mathcal{A} - \mathcal{A}$ with a small $\epsilon \in (0, 1)$ as in Proposition 2. From Proposition 1, this implies the RIP over $\mathcal{A}$ with some $\epsilon' \ll \epsilon$, thus, $\|H\mathbb{b}\|_{\mathsf{F}} \approx \|H\|_{\mathsf{F}}$ for all $H \in \mathcal{A}$ with a very high probability. We consider the following recovery algorithm

$$\widehat{H} = \arg\min_{H' \in \mathcal{A}} \|H'\mathbb{b} - x\|. \qquad (13)$$

Theorem 1: Let B be a $n \times k$ measurement matrix with i.i.d. $\mathsf{N}(0, 1)$ components and let $\mathbb{b} = \mathrm{vec}(B)$. Let H be an arbitrary signal in $\mathcal{A}$ and let $x = H\mathbb{b} + w$ be the set of m linear and noisy measurements from H. Let $\widehat{H}$ be an estimate of H obtained from (13). Then, $\|H - \widehat{H}\|_{\mathsf{F}} \leq \frac{2\|H\|_{\mathsf{F}}}{\sqrt{\mathsf{snr}(1-4\epsilon)}}$ with a probability at least $1 - 4(1 + \frac{2}{\epsilon})^{2k} \binom{n}{m}^2 e^{-cm\epsilon^2}$. $\qquad \square$

Proof: We first assume that $\mathbb{b}$ satisfies the RIP over $\mathcal{A} - \mathcal{A}$, which from Proposition 2 holds with a probability of at least $1 - 4(1 + \frac{2}{\epsilon})^{2k} \binom{n}{m}^2 e^{-cm\epsilon^2}$ if B has i.i.d. $\mathsf{N}(0, 1)$ components. Since $\widehat{H}$ is the solution of (13) and H itself is a feasible solution, we have that

$$\|\widehat{H}\mathbb{b} - x\| \leq \|H\mathbb{b} - x\| = \|w\|. \qquad (14)$$

Applying the triangle inequality, we obtain that $\|(\widehat{H} - H)\mathbb{b}\| \leq \|w\| + \|w\| = 2\|w\|$. Finally, applying the RIP condition implied by Proposition 2 to $\mathcal{A}^\circ$ scaled by $\|H\|_{\mathsf{F}}$, we have

$$\sqrt{1 - 4\epsilon} \frac{\|H - \widehat{H}\|_{\mathsf{F}}}{\|H\|_{\mathsf{F}}} \leq \frac{\|(H - \widehat{H})\mathbb{b}\|}{\|H\|_{\mathsf{F}}} \leq \frac{2\|w\|}{\|H\|_{\mathsf{F}}} \qquad (15)$$

$$\overset{(a)}{\simeq} \frac{2\|w\|}{\|H\mathbb{b}\|} = \frac{2}{\sqrt{\mathsf{snr}}}, \qquad (16)$$

where in (a) we used the RIP over $\mathcal{A}$, which holds with an $\epsilon' \ll \epsilon$. This implies the desired result $\|H - \widehat{H}\|_{\mathsf{F}} \leq \frac{2\|H\|_{\mathsf{F}}}{\sqrt{\mathsf{snr}(1-4\epsilon)}}$ and completes the proof. $\blacksquare$

Remark 1: The recovery guarantee provided by Theorem 1 for the solution $\widehat{H}$ of (13) continues to hold for any other solution $\widehat{H}$ that merely satisfies the feasibility condition

$$\widehat{H} \in \{H' : \|H'\mathbb{b} - x\| \leq \beta \|w\|\}, \qquad (17)$$

for any $\beta \geq 1$, where we obtain $\|H - \widehat{H}\|_{\mathsf{F}} \leq \frac{(\beta+1)\|H\|_{\mathsf{F}}}{\sqrt{\mathsf{snr}(1-4\epsilon)}}$. $\diamond$

For a suitably selected set of parameters n, m, k, and ϵ, Theorem 1 guarantees the stable universal recovery of any $H \in \mathcal{A}$ up to relative recovery precision ϵ with a high probability provided that $4(1 + \frac{2}{\epsilon})^{2k} \binom{n}{m}^2 e^{-cm\epsilon^2}$ is small. For an asymptotic regime where $m, n, k \to \infty$, this condition is satisfied provided that

$$\frac{k}{n} \log(1 + \frac{2}{\epsilon}) + \frac{1}{n} \log \binom{n}{m} - \frac{c\epsilon^2 m}{2n} < 0. \qquad (18)$$

In particular, if $\frac{m}{n} \to \rho$ and $\frac{k}{n} \to \delta$ in such a regime, the condition (18) takes the form

$$\delta \log(1 + \frac{2}{\epsilon}) + h(\theta) - \frac{c\epsilon^2}{2}(1 - \theta) < 0, \qquad (19)$$

where $\theta := 1 - \rho$ denotes the fractional sampling loss and where $h(\theta) = -\theta \log(\theta) - (1 - \theta) \log(1 - \theta)$ denotes the entropy function. It is seen that for a given precision ϵ, the recovery is successful for a sufficiently small θ and δ although it typically requires a large oversampling factor $\frac{n}{k} = \frac{1}{\delta}$. When only $o(n)$ number of samples are missing, i.e., $m = n - o(n)$, then an oversampling ratio of order $\frac{n}{k} \approx \frac{2 \log_2(1 + \frac{2}{\epsilon})}{c\epsilon^2}$ would be enough to compensate for the missing labels of the measurements.

Theorem 1 provides a universal recovery guarantee up to precision ϵ for any estimate $\widehat{H}$ obtained from (13) or (17). An implicit requirement, however, is to design a suitable algorithm to find such an estimate $\widehat{H}$. In the next section, we provide a low-complexity algorithm using alternating minimization. We also prove that the RIP over $\mathcal{A} - \mathcal{A}$ is a sufficient condition for the proposed algorithm to return a *good* estimate $\widehat{H}$. We also investigate via numerical simulations the feasible range of parameters m, n, k for which this algorithm is able to find such an estimate.

IV. RECOVERY ALGORITHM

Let $\check{y}$ be the desired signal and let $\check{x} = \check{S}B\check{y} + w$ be the vector of noisy unlabeled measurements. We define the following cost function for the recovery of the signal $\check{y}$ as in (13):

$$f(S, y) = \|\check{x} - SBy\|^2, \ S \in \mathcal{S}_{m,n}, \ y \in \mathbb{R}^k. \qquad (20)$$

Note that for a fixed B and random i.i.d. Gaussian noise w, the minimizer of $f(S, y)$ denoted by (S^*, y^*) gives the maximum likelihood (ML) estimate of $(\check{S}, \check{y})$ and consequently the ML estimate y^* of the desired signal $\check{y}$. Finding the ML estimate y^*, however, requires a joint optimization over (S, y). This seems to require searching over all $\binom{n}{m}$ possible $S \in \mathcal{S}$, which might be intractable to do for large n and m. Here, instead of joint search over S and y, we use an iterative alternating projection with respect to y and S individually to reduce the complexity. We define (S^t, y^t) as the estimate of $\check{S}$ and $\check{y}$ at iteration $t = 1, 2, \ldots$ during the algorithm, where we update them via the sequential projection operations

$$S^t \mapsto y^t = \arg\min_{y \in \mathcal{Y}} f(S^t, y), \qquad (21)$$

$$y^t \mapsto S^{t+1} = \arg\min_{S \in \mathcal{S}} f(S, y^t), \qquad (22)$$

where $\mathcal{Y} := \mathbb{R}^k$, where for simplicity we dropped the dependence on m, n in $\mathcal{S}_{m,n}$, and where we used the notation $S \mapsto y$ for $y = \arg\min_{y \in \mathcal{Y}} f(S, y)$ and $y \mapsto S$ for $S = \arg\min_{S \in \mathcal{S}} f(S, y)$ for the projection operators onto $\mathcal{Y}$ and $\mathcal{S}$. We initialize the algorithm with a random $S^1 \in \mathcal{S}$ and generate the sequential estimates $S^1 \mapsto y^1 \mapsto S^2 \mapsto y^2 \mapsto \cdots$.

A. Projection on $\mathcal{Y}$

For a fixed S^t in (21), finding y^t given by $S^t \mapsto y^t$ boils down to obtaining the least-square solution of an over-complete set of linear equations (via the linear matrix $S^t B$). The result is given by $y^t = (S^t B)^\dagger x$, where $\dagger$ denotes the pseudo-inverse operator (for the tall matrix $S^t B$).

B. Projection on $\mathcal{S}$

Let y^t be the optimal solution obtained from (21) and set $z^t = By^t$. Finding the optimal selection matrix S^{t+1} at (22) given by $y^t \mapsto S^{t+1}$ requires extracting a sub-vector of z^t of dimension m, while keeping the relative order of the components, that is closest to m-dim vector of measurements $\check{x}$ in ℓ_2 distance. We formulate this as a Dynamic Programming (DP) problem as follows. We define a 2D table of size $m \times n$ whose elements are labelled with $(r,c) \in [m] \times [n]$ and have the value $T_{r,c} \in \mathbb{R}_+ \cup \{\infty\}$ given by

$$
T_{r,c} = \begin{cases} \begin{array}{l} \text{minimum squared-distance between} \\ \text{the subvector } \check{x}_{1:r} \text{ and a subsequence} \\ \text{of } z_{1:c}^t \text{ of length } r, \end{array} & c \geq r, \\ \infty, & \text{otherwise.} \end{cases}
$$

We initialize the diagonal elements of the table with $T_{i,i} = \|\check{x}_{1:i} - z_{1:i}^t\|^2$ since there is only one way to match the first i elements of z^t with the i elements in $\check{x}_{1:i}$. We also initialize the elements in the first column of the table, i.e., $T_{1,j}$ for $j \in [n]$, with $T_{1,j} = \min_{j' \in [j]} |\check{x}_1 - z_{j'}^t|^2$ since the single element $\check{x}_1$ should be matched with the closest element in the sub-vector $z_{1:j}^t$ consisting of the first j elements of z^t.

To find the value of a typical element $T_{r,c}$ in the table, we need to match $\check{x}_{1:r}$ with a suitable subsequence of $z_{1:c}^t$ of length r. In the optimal matching, the last component $\check{x}_r$ is matched either with z_c^t or with $z_{c'}^t$ for some $c' < c$. Thus, $T_{r,c}$ can be computed from the already computed elements of the table as follows:

$$
T_{r,c} = \min\left\{T_{r-1,c-1} + |\check{x}_r - z_c^t|^2, T_{r,c-1}\right\}. \quad (23)
$$

Notice that with the already mentioned initialization and the recursion (23), we can complete the whole table by filling its d-th diagonal consisting of elements $\{T_{1,d}, T_{2,d+1}, T_{3,d+2}, \dots\}$ for $d = 2, \dots, n$ one at a time. Overall, filling the whole table requires $O(mn)$ operations.

After filling the whole table, we can find the indices of those m elements of z^t that are optimally matched to the elements of $\check{x}$ as follows. We start from the element $T_{m,n}$ located at the up-right corner of the table at location (m,n). Note that by definition, $i \to T_{m,i}$, for $i \in [n]$, is a decreasing sequence of i since by increasing i the subsequence $z_{1:i}^t$ becomes longer and provides more options to find a better subsequence thereof matching $\check{x}$. Thus, we can check that the index of the last element in z^t that is matched to the last element $\check{x}_m$ of $\check{x}$ in the optimal matching is given by

$$
i_m = \min\left\{i \in [n] : T_{m,i} = T_{m,n}\right\}. \quad (24)
$$

To find the next largest index i_{m-1}, we apply the same argument to the sub-table $T_{1:m-1,1:i_m-1}$ and its up-right
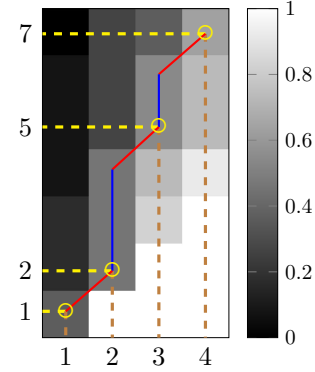

Fig. 3: Illustration of the Dynamic Programming table for matching two sequences of length 7 and 4 respectively. For the illustrated example, the indices of the matched elements in the larger vector are given by $\{1,2,5,7\}$.

corner element located at $(m-1, i_m - 1)$, where we obtain the following recursive formula for the remaining indices:

$$
i_{m-\gamma} = \min\left\{i \in [i_{m-\gamma+1} - 1] : T_{m-\gamma,i} = T_{m-\gamma,i_{m-\gamma+1}-1}\right\},
$$

where $\gamma \in [m-1]$. This has been illustrated for a vector z of dimension 7 and a vector x of dimension 4 in Fig. 3.

V. PERFORMANCE ANALYSIS

In this section, we analyze the performance of our proposed algorithm under the assumption that the measurement matrix B satisfies the RIP over $\mathcal{A} - \mathcal{A}$. As in Section IV, we denote the desired signal to be recovered by $\check{H} = \check{y}^\mathsf{T} \otimes \check{S}$, the noisy samples thereof by $\check{x} = \check{H}\mathbb{b} + w = \check{S}B\check{y} + w$, and the sequence generated by the alternating minimization algorithm by $S^1 \mapsto y^1 \mapsto S^2 \mapsto \dots$. We prove that for a suitable initialization for S^1, the algorithm recovers a good estimate of the desired signal. First, note that applying the RIP to (S,y) and the target signal $(\check{S}, \check{y})$ and using the triangle inequality, we have that

$$
g(S,y) = \|\check{x} - SBy\| \leq \sqrt{m(1+\epsilon)}\, d_{H,\check{H}} + \|w\|
$$
$$
=: h(S,y), \quad (25)
$$

where $g(S,y) := \sqrt{f(S,y)}$, where $H = y^\mathsf{T} \otimes S$, and where $d_{H,\check{H}}$ denotes the distance between H and $\check{H}$ given by

$$
d_{H,\check{H}} = \sqrt{\nu_{S,\check{S}}\|y - \check{y}\|^2 + (1 - \nu_{S,\check{S}})(\|y\|^2 + \|\check{y}\|^2)}.
$$

as in (11). It is worthwhile to remark here that although $h(S,y)$ is not computable–since it requires the knowledge of $\check{H}$ which is exactly what we are trying to obtain–the alternating minimization can be applied to $h(S,y)$ (equivalently, to $d_{\check{H},H}$), where we wish to minimize the distance by alternating minimization on y and S with a suitable initialization of S. This conceptual (non-computable) result serves to prove the desired result on the convergence of the actual (computable) algorithm. We first prove that with a good initialization S^1, the minimum of h under alternating minimization converges to $\|w\|$.

Proposition 3: Let $\check{H} = \check{y} \otimes \check{S}$ be as before and let $h(S,y)$ be given as in (25). Let S^1 be an initialization for S such

<div style="text-align:center">(a) Random initialization.</div>

<div style="text-align:center">(b) Genie-aided initialization.</div>
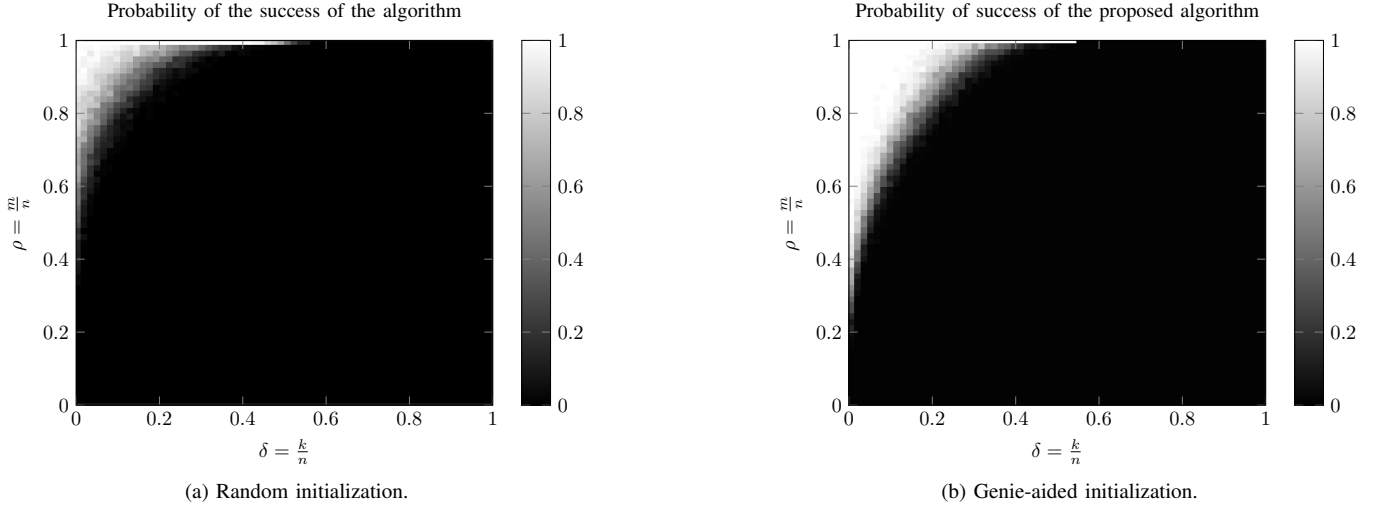
Fig. 4: Probability of success of the alternating minimization algorithm as a function of $\delta = \frac{k}{n}$ and $\rho = \frac{m}{n}$ with a random initialization S^1 for the selection matrix in (a) and a genie-aide initialization S^1 with $\nu_{S^1,\check{S}} = \nu_0 = 0.2$ in (b), where in the latter the remaining rows of S^1 (remaining 80% rows) are selected randomly.

that $\nu_{S^1,\check{S}} \geq \nu_0 > 0$. Let $y_h^t := \arg\min_{y'} h(S_h^t, y')$ and $S_h^{t+1} := \arg\min_{S'} h(S', y_h^t)$ be the sequence generated by the alternating minimization applied to h with the initialization $S_h^1 = S^1$. Then, $\lim_{t\to\infty} h(S_h^t, y_h^t) = \|w\|$.

Proof: We prove a stronger result that the convergence happens in 2 iterations. We can check that for the given initialization S^1, the optimal solution y_h^1 is given by $\nu_0 \check{y}$. In particular, replacing in $h(S, y)$, we have that

$$h(S, y_h^1) = \chi \sqrt{\|\check{y}\|^2 (1 + \nu_0^2) - 2\nu_0 \|\check{y}\|^2 \nu_{S,\check{S}}} + \|w\|, \quad (26)$$

where $\chi = \sqrt{m(1+\epsilon)}$. It is not difficult to check that the optimal solution $S_h^2 = \arg\min_S h(S, y_h^1)$ is achieved for $\nu_{S_h^2, \check{S}} = 1$, thus, $S_h^2 = \check{S}$. For S_h^2, we have $h(S_h^2, y) = \chi \|y - \check{y}\| + \|w\|$, thus, it results that $y_h^2 = \arg\min_{y'} h(S_h^2, y) = \check{y}$ and $h(S_h^2, y_h^2) = \|w\|$. We can also check that $S_h^t = \check{S}$ and $y_h^t = \check{y}$ for $t \geq 3$, thus, $h(S_h^t, y_h^t)$ converges to $\|w\|$ in only 2 iterations. This completes the proof. ∎

Theorem 2: Let $\check{H} = \check{y}^\mathsf{T} \otimes \check{S}$ and S^1 be as in Proposition 3. Let $S^1 \mapsto y^1 \mapsto S^2 \dots$ be the sequence generated by the alternating minimization applied to g. Then, $\lim_{t\to\infty} g(S^t, y^t) \leq \|w\|$.

Proof: From (25), we have that for a given S and y:

$$\min_{y'} g(S, y') \leq \min_{y'} h(S, y'), \quad (27)$$

$$\min_{S'} g(S', y) \leq \min_{S'} h(S', y). \quad (28)$$

This implies that, starting from a common initialization S^1, the minimum of g under alternating minimization algorithm is always less than the minimum of h. In particular, denoting by $\{y_h^t, S_h^t\}_{t=1}^\infty$ the sequence generated by alternating minimization applied to h with the initialization $S_h^1 = S^1$, we have that $g(S^t, y^t) \leq h(S_h^t, y_h^t)$. Taking the limit as t tends to infinity and using Proposition 3, we obtain that $\lim_{t\to\infty} g(S^t, y^t) \leq \|w\|$, which implies the desired result. This completes the proof. ∎

Theorem 2 guarantees that the sequence $H^t = y^{t\mathsf{T}} \otimes S^t$ generated by our proposed alternating minimization algorithm converges to a feasible point in (17) (with $\beta = 1$) as in Remark 1 for which the recovery performance is guaranteed by Theorem 1. However, the implicit assumption is that a *good* initialization with a $\nu_{S^1,\check{S}} \geq \nu_0 > 0$ is possible. In this paper, we use a random initialization for S^1. We discuss further the effect of initialization in Section VI and VII.

VI. SIMULATION RESULTS

We run numerical simulation to assess the performance of our proposed recovery algorithm. We consider a signal $\check{y}$ of dimension k. For each simulation, we generate a Gaussian $n \times k$ measurement matrix B, where we select m out of n measurements $B\check{y}$ completely randomly via a random sampling matrix $\check{S} \in \mathcal{S}_{m,n}$. The resulting noisy measurement is given by $\check{x} = \check{S}B\check{y} + w$, where w is the additive white Gaussian measurement noise with $\|w\|^2 \approx \frac{\|\check{S}B\check{y}\|^2}{\text{snr}}$, where snr denotes the measurement Signal-to-Noise Ratio (SNR). For the simulations, we assume an SNR of $20\,\text{dB}$.

A. Probability of Success of the Algorithm

We run our proposed algorithm with the noisy input $\check{x}$, where we initialize the algorithm with a random $S^1 \in \mathcal{S}_{m,n}$. To see the effect of the initialization, we repeat the simulation with a genie-aided initialization S^1 in which $\nu_{S^1,\check{S}} = \nu_0 = 0.2$, where 20% of the rows of S^1 are set equal to the corresponding rows of $\check{S}$ while the remaining rows are selected completely randomly among the remaining possible rows. In both cases, we define the output of the alternating projection algorithm by $S^1 \mapsto y^1 \mapsto S^2 \mapsto y^2 \cdots$ and denote the final output produced by the algorithm by y^∞. We call the recovery successful if the relative error satisfies $\frac{\|\check{y}-y^\infty\|^2}{\|\check{y}\|^2} \approx O(\frac{1}{\text{snr}})$. Note that we cannot guarantee successful recovery by

merely checking the weaker condition $\frac{\|H-\widehat{H}\|_F^2}{\|H\|_F^2} = O(\frac{1}{\mathsf{snr}})$ as suggested by Theorem 1 since for some range of parameters m, n, k the RIP over $\mathcal{A}-\mathcal{A}$ might not hold. For simulations, we set $n = 1000$ and define parameters $\delta = \frac{k}{n}$ as the measurement ratio and $\rho = \frac{m}{n}$ as the sampling ratio as before. For each δ and ρ, we run simulations for 1000 independent realizations of the measurement matrix B and $\check{S}$ to obtain an estimate of the success probability of the algorithm. Fig. 4 illustrates the probability of success as a function of $\delta, \rho \in (0, 1)$ for the fully random initialization in (4a) and for the genie-aided initialization in (4b).

It is seen that, in both cases, when $\delta > 0.5$, thus, $n \leq 2k$, the algorithm recovers the signal only when $m = n$ (full sampling) and totally fails when $m = (1 - \theta)n$ even for a small fractional loss $\theta \in (0, 1)$. For $\delta \ll 0.5$, however, the algorithm is more resilient to the fractional sampling loss θ and tolerates larger θ for smaller δ. Also, comparing the random initialization in Fig. 4a with the genie-aided one in Fig. 4b reveals the importance of a good initialization in the recovery performance of the algorithm. Moreover, it is seen that the genie-aided case undergoes a much sharper phase-transition in $\delta - \rho$ plane since it is less affected by the *adverse* random initialization.

B. Application to System Identification

In this section, as an example, we study a practical signal processing problem that can be well modeled by Unlabeled Ordered Sampling. In this example, we use our proposed alternating minimization algorithm to estimate the impulse response of a linear time-invariant system. We assume that the output of the system is observed through a deletion channel and is contaminated with noise. This is illustrated in Fig. 5, where a known pre-designed training sequence $\{b_l\}_{l=0}^{\tau-1}$ of length τ is applied to the input of a linear system with an impulse response of length at most k given by $\{y_l\}_{l=0}^{k-1}$. We assume that an estimate of the delay spread of the channel k is a priori known. The output of the linear system is given by $z = b*y$, where $*$ denotes the convolution operation, where the output z is given by $z_l = \sum_{t=0}^{k-1} y_t b_{l-t}$ for $l = 0, 1, \ldots, n-1$, where $n = k + \tau - 1$ denotes the length of the output z and where $b_r = 0$ for $r < 0$. We consider a scenario in which the output $\{z_i\}_{i=0}^{n-1}$ can be observed only through a deletion channel, which deletes some of the output samples $\{z_i\}_{i=0}^{n-1}$ but preserves their underlying order. Denoting by $y = (y_0, \ldots, y_{k-1})^\mathsf{T}$ and $z = (z_0, \ldots, z_{n-1})^\mathsf{T}$, we can write $z = By$ with a measurement matrix B given by

$$B = \begin{pmatrix} b_0 & 0 & \cdots & 0 \\ b_1 & b_0 & \ddots & 0 \\ b_2 & b_1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ b_{\tau-1} & b_{\tau-2} & \cdots & b_{\tau-k} \\ 0 & b_{\tau-1} & \ddots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & b_{\tau-1} \end{pmatrix}, \qquad (29)$$
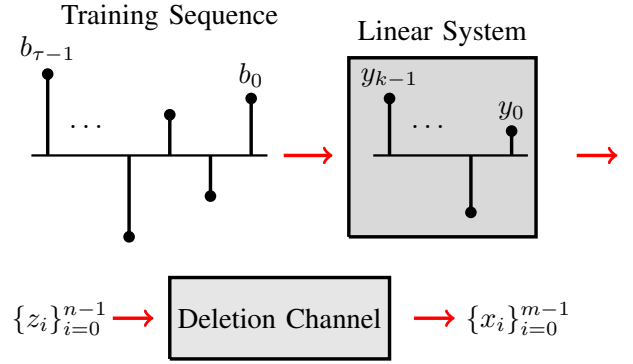
Fig. 5: Identifying a linear system with an impulse response $\{y_l\}_{l=0}^{k-1}$ of length k via a training sequence $\{b_l\}_{l=0}^{\tau-1}$ of length τ. We assume that some of the output samples $\{z_i\}_{i=0}^{n-1}$ (at unknown positions) are deleted, thus, only a limited number of samples are available for system identification.
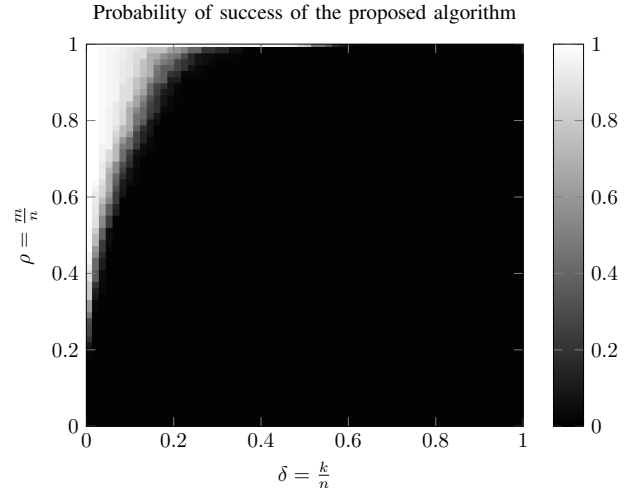
Fig. 6: Probability of success of the alternating minimization algorithm for the estimation of a dispersive channel as a function of $\delta = \frac{k}{n}$ and $\rho = \frac{m}{n}$ with an initialization S^1 with $\nu_{S^1, \check{S}} = \nu_0 = 0.2$.

where it is seen that B is an $n \times k$ matrix that depends on the training sequence $\{b_l\}_{l=0}^{\tau-1}$. We denote the final set of m samples, for some $m \leq n$, available for system identification by $x = (x_0, \ldots, x_{m-1})^\mathsf{T}$, where $x = \check{S}z + w = \check{S}By + w$, where $\check{S}$ is a selection matrix representing the location of those samples in z that are not deleted by the deletion channel, and where w is the additive measurement noise. It is seen that the system identification in the scenario illustrated in Fig. 5 boils down to the Unlabeled Ordered Sampling problem with a measurement matrix B given by (29), to which we can apply our proposed alternating minimization algorithm.

For simulation, we assume that the training sequence $\{b_l\}_{l=0}^{\tau-1}$ has i.i.d. $\mathsf{N}(0, 1)$ samples and is known to the system identification algorithm. Note that although each row of B in (29) still consists of Gaussian variables, due to the special structure of convolution operation, they are highly correlated. Nevertheless, we can still run our alternating minimization algorithm with B as the measurement matrix. Fig. 6 illustrates the simulation results for $n = 1000$ and for an SNR of

20 dB. We also assume that the selection matrix S^1 for the alternating minimization algorithm is properly initialized such that $\nu_{S^1,\check{S}} = \nu_0 = 0.2$. It is seen that, as expected, for a given delay spread k, the performance always improves by increasing the length of the training sequence τ (equivalently n) and the number of available samples (those samples not deleted by the deletion channel) m. We also observe that, in comparison with Fig. 4b in which B has fully i.i.d. components across different rows, the correlation among the rows of B degrades the performance of the alternating minimization algorithm only slightly.

VII. DISCUSSION AND FURTHER REMARKS

A. Improving the Recovery Performance

In Fig. 4, we illustrated the success probability of a single round of alternating minimization. The difficulty arises when the probability of success is quite small, where in that case the alternating minimization algorithm hits a local minimum and fails to find a solution $\widehat{H} = \widehat{y}^\mathsf{T} \otimes \widehat{S}$ in the feasible region $\{H' : \|H'\mathbb{b} - \check{x}\| \le \beta\|w\|\}$, for a $\beta = O(1)$, as explained in Remark 1. We can improve the recovery performance by running alternating minimization several times each time with a different random initialization S^1. This requires a procedure to certify whether the algorithm succeeds to find a solution, where in that case we output the solution and quit running the algorithm again. Since the only information about the signal $\check{y}$ comes from the unlabeled measurements $\check{S}B\check{y}$, developing such a procedure in the presence of measurement noise seems to require the RIP over $\mathcal{A} - \mathcal{A}$.

Another direction to improve the performance of the alternating minimization is to find a *good* initialization for S^1. The hope is that such an S^1 lie in the basin of attraction of the desired signal $(\check{S}, \check{y})$ such that the solution path generated by alternating minimization converges to $\check{y}$ without hitting any local minima. Notice that in our proposed algorithm, we use a random initialization for S^1. Assuming that B satisfies the RIP over $\mathcal{A} - \mathcal{A}$, as far as $\nu_{S^1,\check{S}} \ge \nu_0 > 0$ holds for the initialization S^1, Theorem 2 together with Theorem 1 guarantee a suitable recovery of the target signal $(\check{S}, \check{y})$. However, the requirement $\nu_{S^1,\check{S}} \ge \nu_0 > 0$ is quite difficult to meet for some $\check{S}$. For example, if $\check{S}$ is the selection matrix that samples the first m measurements in $B\check{y}$, then under random uniform sampling of $S^1 \in \mathcal{S}$, we have that $\mathbb{P}[\nu_{S^1,\check{S}} = 0] = 1 - \frac{1}{\binom{n}{m}}$, thus the initialization fails with a very high probability for large n and m unless $m = n$. Of course, when the RIP holds and an estimate of the ℓ_2 norm of the noise $\|w\|$ is available, one can repeat the alternating minimization with several random initializations until one finds a good initialization for which the estimate $\widehat{H}$ obtained from the algorithm satisfies $\|\widehat{H}\mathbb{b} - \check{x}\| \le \beta\|w\|$ with some $\beta = O(1)$ for which the recovery guarantee follows from Theorem 1. However, for the example mentioned, this requires trying approximately $\binom{n}{m}$ random initializations, which is infeasible for large m and n unless $n - m = O(1)$. This implies that a *good* initialization method is necessary even when the RIP holds. Running a nonconvex optimization problem with

a good initialization has recently been of interest in other problems in Compressed Sensing such as phase retrieval [17], blind deconvolution [18], and blind calibration [19]. We leave developing a good initialization scheme for our algorithm as a future work.

B. Necessity of RIP on $\mathcal{A} - \mathcal{A}$

In this paper, our goal was to recover the signal y via unlabeled samples $x = SBy + w$. From a statistical point of view, the selection matrix S plays the role of an ancillary parameter, i.e., it affects the statistics of the observations in x, but it is not the desired parameter to be estimated. The RIP over $\mathcal{A} - \mathcal{A}$ defined in this paper, however, puts an equal emphasis on y and S. Note that from (12) and (11), it is seen that for a given signal $H = y^\mathsf{T} \otimes S$ and any other signal $H' = y'^\mathsf{T} \otimes S'$ in $\mathcal{A}$, the only way to control $\|y - y'\|$ statistically based on $(H - H')\mathbb{b}$ is to have $\nu_{S,S'} \approx 1$. This implies that the selection matrices S and S' should have almost identical rows, thus, the recovery of S seems to be necessary for a suitable recovery of y.

C. Extension to other Signal Models

The unlabeled sensing problem $x = SBy + w$ studied in this paper can be extended to cases in which the main signal y belongs to a structured class of signals $\mathcal{Y} \subset \mathbb{R}^k$. Moreover, depending on the applications, the matrix S can, in general, belong to a class of selection matrices $\mathcal{S}$ other than $\mathcal{S}_{m,n}$, where each $S \in \mathcal{S}$ still keeps m out of n measurements in By but might not necessarily preserve their underlying order. We denote again by $\mathcal{A} = \{y^\mathsf{T} \otimes S : y \in \mathcal{Y}, S \in \mathcal{S}\}$ the set of all possible signals.

As in the UOS, in terms of signal recovery, we need to check two main requirements. The first is to develop an RIP under a suitable metric (e.g., ℓ_2 distance as in this paper) for the difference set $\mathcal{A} - \mathcal{A}$. Following Proposition 2, when B has i.i.d. Gaussian components, such an RIP with a constant $\epsilon \in (0, 1)$ can be generally derived for a collection of selection matrices $\mathcal{S}$ with a probability larger than $1 - 2|\mathcal{S}|^2 N^2 e^{-cm\epsilon^2}$, where $|\mathcal{S}|$ denotes the cardinality of $\mathcal{S}$, and where N denotes the size of a minimal $\frac{\epsilon}{2}$-net $\mathcal{N}_{\mathcal{Y},\epsilon} = \{g^1, \ldots, g^N\}$ over the signal set $\mathcal{Y}$. This provides a theoretical lower bound on the the number of measurements m for a given precision ϵ.

The second requirement is to develop an algorithm to recover the signal $\check{H} = \check{y}^\mathsf{T} \otimes \check{S} \in \mathcal{A}$ from the noisy unlabeled measurements $\check{x} = \check{H}\mathbb{b} + w = \check{S}B\check{y} + w$. In particular, under the RIP, such an algorithm needs to recover an estimate $\widehat{H} \in \mathcal{A}$ satisfying $\|\widehat{H}\mathbb{b} - \check{x}\| \le \beta\|w\|$ for some $\beta = O(1)$, where $\|.\|$ here denotes the metric with respect to which the RIP is derived. For such an estimate, one should be able to derive a performance guarantee similar to that obtained for UOS in Theorem 1. For the UOS studied in this paper, we used the alternating minimization over y and S, where the latter minimization was done with a feasible complexity by using the ordered structure of the matrices in $\mathcal{S}_{m,n}$ and applying the dynamic programming. Deriving such an algorithm for a general signal set $\mathcal{A}$ (for a general signal structure in $\mathcal{Y}$ and unlabeled sampling structure in $\mathcal{S}$) requires exploiting the algebraic as well as the geometric structure of $\mathcal{A}$.

VIII. Conclusion

In this paper, we studied the Unlabeled Ordered Sampling problem, where the goal was to recover a signal from a set of unlabeled linear measurements taken via a known measurement matrix. We defined a restricted isometry property (RIP) for the measurement matrix over the signal set and proved that under the proposed RIP every arbitrary signal can be suitably recovered from a set of unlabeled and possibly noisy samples. We proposed a low-complexity recovery algorithm and studied its performance via numerical simulations. We also provided guidelines for extending the results in the paper to other signal models with a different signal structure and different class of selection matrices governing the structure of missing labels in the measurements.

Appendix A
Proof of Proposition 2

Without loss of generality, we prove the RIP over the normalized signal set $\mathcal{A}^\circ$. Since $d_{H,H'}^2 = \mathbb{E}[\mathscr{D}_{H,H'}] \leq 4$ for any $H, H' \in \mathcal{A}^\circ$, to derive the RIP over $\mathcal{A}^\circ - \mathcal{A}^\circ$, we need to prove that $\mathbb{P}[\mathscr{R}_{\mathcal{A}^\circ - \mathcal{A}^\circ} \leq 4\epsilon]$ is larger than $1 - 4(1 + \frac{2}{\epsilon})^{2k} \binom{n}{m}^2 e^{-cm\epsilon^2}$ as in the statement of the Proposition, where

$$\mathscr{R}_{\mathcal{A}^\circ - \mathcal{A}^\circ} := \sup_{H, H' \in \mathcal{A}^\circ} \left| \mathscr{D}_{H,H'} - \mathbb{E}[\mathscr{D}_{H,H'}] \right| \quad (30)$$

is as in (30). The proof follows by an argument similar to that in the proof of Proposition 1, with the only difference that we need to apply the union bound over a joint net for both $y \in \mathbb{S}^{k-1}$ and $y' \in \mathbb{S}^{k-1}$ and joint selection matrices S and S' in $\mathcal{S}_{m,n}$.

Consider two arbitrary signals $H = y^\mathsf{T} \otimes S$ and $y'^\mathsf{T} \otimes S'$ belonging to $\mathcal{A}^\circ$, where $S, S' \in \mathcal{S}_{m,n}$ and $y, y' \in \mathcal{N}_\epsilon$ with $\|y\| = \|y'\| = 1$, where $\mathcal{N}_\epsilon$ is the minimal $\frac{\epsilon}{2}$-net defined in the proof of Proposition 1 and has a size of at most $N = (1 + \frac{2}{\epsilon})^k$. We have $(H - H')\flat = SBy - S'By'$. Note that although the elements of the vector $(H - H')\flat$ are still Gaussian, they are correlated unless $S = S'$. Hence, the i.i.d. Gaussian concentration in Proposition 1 does not immediately apply, and we need to extend the concentration result to the correlated case that we have here.

Let us denote by $\mathscr{R} := r_1 < \cdots < r_m$ and $\mathscr{R}' := r_1' < \cdots < r_m'$ the ordered sequences consisting of indices of those rows of B selected by S and S', where $r_i, r_i' \in [n]$, and let $\mathscr{S} := \{i \in [m] : r_i = r_i'\}$ be the index set of similar elements (i.e., similar rows) in the sequence $\mathscr{R}$ and $\mathscr{R}'$. We denote by $s = |\mathscr{S}|$ the number of the similar rows in S and S'. We decompose $\|(H - H')\flat\|^2$ into two terms as follows

$$\|(H - H')\flat\|^2 = \sum_{i \in \mathscr{S}} |B_{r_i,.}(y - y')|^2 \quad (31)$$

$$+ \sum_{i \in \mathscr{S}^c} |B_{r_i,.}y - B_{r_i',.}y'|^2, \quad (32)$$

where $\mathscr{S}^c = [m] \backslash \mathscr{S}$ denotes the complement of $\mathscr{S}$ and where in (31) we used the fact that $r_i = r_i'$ and $B_{r_i,.} = B_{r_i',.}$ for $i \in \mathscr{S}$. Note that (31) consists of the summation of the ℓ_2

norm of a sequence of size $s = |\mathscr{S}|$ of i.i.d. Gaussian variables $B_{r_i,.}(y - y')$, each having a zero mean and a variance $\|y - y'\|^2$, to which the Gaussian concentration inequality can be immediately applied. Thus, we have

$$\mathbb{P}\left[\left| \sum_{i \in \mathscr{S}} |B_{r_i,.}(y - y')|^2 - s\|y - y'\|^2 \right| \geq m\epsilon\right] \leq 2e^{-cm\epsilon^2}, \quad (33)$$

where c is as in Proposition 1. Note that in (33), if $s = o(m)$ is a vanishing function of m for a large m, the concentration bound trivially holds.

Now consider the term in (32) consisting of those elements of $(H - H')\flat$ whose indices belong to $\mathscr{S}^c$. We represent those elements by the vector $h := (H_{\mathscr{S}^c} - H'_{\mathscr{S}^c})\flat$, where $H_{\mathscr{S}^c}$ and $H'_{\mathscr{S}^c}$ denote the submatrices of H and H' consisting of rows $\{r_i : i \in \mathscr{S}^c\}$ and $\{r_i' : i \in \mathscr{S}^c\}$ respectively. Note that h is a zero-mean Gaussian variable with a covariance matrix

$$\mathsf{C}_h := \mathbb{E}[hh^\mathsf{T}] = (H_{\mathscr{S}^c} - H'_{\mathscr{S}^c})(H_{\mathscr{S}^c} - H'_{\mathscr{S}^c})^\mathsf{T}$$
$$= (\|y\|^2 + \|y'\|^2)I_{m-s} - \langle y, y'\rangle\Gamma, \quad (34)$$

where I_{m-s} denotes the identity matrix of order $m - s$, where $\Gamma := (S_{\mathscr{S}^c}S'^\mathsf{T}_{\mathscr{S}^c} + S'_{\mathscr{S}^c}S^\mathsf{T}_{\mathscr{S}^c})$, and where we used the properties of the vec operator to show that $H_{\mathscr{S}^c}H^\mathsf{T}_{\mathscr{S}^c} = \|y\|^2 I_{m-s}$, $H'_{\mathscr{S}^c}H'^\mathsf{T}_{\mathscr{S}^c} = \|y'\|^2 I_{m-s}$, and $H_{\mathscr{S}^c}H'^\mathsf{T}_{\mathscr{S}^c} = \langle y, y'\rangle S_{\mathscr{S}^c}S'^\mathsf{T}_{\mathscr{S}^c}$. Note that Γ is a symmetric matrix of order $m - s$. Let us denote by $\tilde{r}_1 < \cdots < \tilde{r}_{m-s}$ and $\tilde{r}_1' < \cdots < \tilde{r}_{m-s}'$ the subsequence of $\mathscr{R}$ and $\mathscr{R}'$ with indices belonging to $\mathscr{S}^c$. We have

$$\Gamma_{i,j} = \begin{cases} 1 & \text{if } \tilde{r}_i = \tilde{r}_j' \text{ or } \tilde{r}_i' = \tilde{r}_j \\ 0 & \text{otherwise.} \end{cases} \quad (35)$$

Note that Γ has all 0 diagonal elements since $\tilde{r}_i \neq \tilde{r}_i'$ for $i \in \mathscr{S}^c$, thus, $\mathsf{Tr}(\Gamma) = 0$. Also, it has at most two 1s in each row and in each column. Since C_h is *positive semi-definite* (PSD) for all choices of y and y', setting $y = y'$ and $y = -y'$ in (34), it results that the eigen-values of Γ (which are all real-valued due to symmetry of Γ) should lie in $[-2, 2]$. This implies that for a fixed y, y', the maximum eigen-value of C_h is at most $\|y\|^2 + \|y'\|^2 + 2|\langle y, y'\rangle|$. Thus, using the inequality $2|\langle y, y'\rangle| \leq \|y\|^2 + \|y'\|^2$, all the eigen-values of C_h are upper bounded by $\tau := 2(\|y\|^2 + \|y'\|^2) = 4$, where we used $\|y\| = \|y'\| = 1$ for the signals belonging to $\mathcal{A}^\circ$. Also, from (34) and $\mathsf{Tr}(\Gamma) = 0$, we have that $\mathsf{Tr}(\mathsf{C}_h) = 2(m - s)$.

Let us denote the eigen-values of C_h by $0 \leq \lambda_1 \leq \cdots \leq \lambda_{m-s} \leq \tau$ and let $\mathsf{C}_h = Q\Lambda Q^\mathsf{T}$ be the spectral decomposition of C_h where $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_{m-s})$ denotes the diagonal matrix consisting of the eigen-values of C_h. Define the normalized Gaussian random vector $\tilde{h} = \Lambda^{-1/2}Q^\mathsf{T}h$. We can check that $\tilde{h}$ has i.i.d. $\mathsf{N}(0, 1)$ components. We can bound the second term in (32) as follows:

$$\mathbb{P}\left[\left| \sum_{i \in \mathscr{S}^c} |B_{r_i,.}y - B_{r_i',.}y'|^2 - (m - s)(\|y\|^2 + \|y'\|^2) \right| \geq m\epsilon\right]$$

$$= \mathbb{P}\left[\left| \|h\|^2 - \mathbb{E}[\|h\|^2] \right| \geq m\epsilon\right]$$

$$= \mathbb{P}\left[\left| \sum_{i=1}^{m-s} \lambda_i(\tilde{h}_i^2 - \mathbb{E}[\tilde{h}_i^2]) \right| \geq m\epsilon\right] \quad (36)$$

$$\overset{(a)}{\leq} 2\mathbb{P}\Big[\sum_{i=1}^{m-s} \lambda_i \tilde{h}_i^2 \geq \mathsf{Tr}(\mathsf{C}_h) + m\epsilon\Big] \tag{37}$$

$$\overset{(b)}{\leq} 2e^{-\frac{1}{2}\sum_{i=1}^{m-s}\log(1-2\mu\lambda_i)-\mu\mathsf{Tr}(\mathsf{C}_h)-\mu m\epsilon} \tag{38}$$

$$\overset{(c)}{\leq} 2e^{2\mu^2\sum_{i=1}^{m-s}\lambda_i^2-\mu m\epsilon}, \tag{39}$$

where in (a) and (b) we used the exponential Markov's inequality [20] with a parameter $\mu > 0$ and the identity $\mathbb{E}[e^{s\tilde{h}_i^2}] = \frac{1}{\sqrt{1-2s}}$ for a normalized Gaussian random variable $\tilde{h}_i$ and $s \in (0, \frac{1}{2})$, and in (c) we used the inequality $\log(1-\theta) \geq -\theta - \theta^2$ which holds for $\theta \in (0, \frac{1}{2})$. Thus, the inequality (39) holds for any $\mu \in (0, \frac{1}{4\lambda_{m-s}})$ and especially for $\mu \in (0, \frac{1}{4\tau}) = (0, \frac{1}{16})$. Note that from (34), we have that $\mathsf{C}_h = 2I_{m-s} - \langle y, y'\rangle\Gamma$, thus,

$$\sum_{i=1}^{m-s}\lambda_i^2 = \mathsf{Tr}(\mathsf{C}_h^2) = \mathsf{Tr}(4I_{m-s} - 4\langle y, y'\rangle\Gamma + |\langle y, y'\rangle|^2\Gamma^2)$$

$$\overset{(a)}{=} 4(m-s) + 0 + |\langle y, y'\rangle|^2\mathsf{Tr}(\Gamma^2)$$

$$\overset{(b)}{\leq} 4(m-s) + 2|\langle y, y'\rangle|^2(m-s)$$

$$\leq 6(m-s), \tag{40}$$

where in (a) we used $\mathsf{Tr}(\Gamma) = 0$, and in (b) we used the fact that Γ has at most two 1s in each row and in each column, thus, from the symmetry of Γ, we have $\mathsf{Tr}(\Gamma^2) = \mathsf{Tr}(\Gamma\Gamma^\mathsf{T}) = \|\Gamma\|_\mathsf{F}^2 \leq 2(m-s)$. Finally, setting $\mu = \frac{\epsilon}{24}$, which lies in the allowed region $(0, \frac{1}{16})$ for $\epsilon \in (0,1)$, and upper bounding the term $\sum_{i=1}^{m-s}\lambda_i^2$ by $6(m-s)$ from (40), we obtain the upper bound $2e^{-c'm\epsilon^2}$ for $c' = \frac{1}{48}$. From (33) and (39), we obtain

$$\mathbb{P}\Big[\Big|\mathscr{D}_{H,H'} - \mathbb{E}[\mathscr{D}_{H,H'}]\Big| \geq 2\epsilon\Big]$$

$$= \mathbb{P}\Big[\Big|\|(H-H')\mathbb{b}\|^2 - \mathbb{E}[\|(H-H')\mathbb{b}\|^2]\Big| \geq 2m\epsilon\Big]$$

$$= \mathbb{P}\Big[\Big|\sum_{i\in\mathscr{S}}|B_{r_i,.}(y-y')|^2 - s\|y-y'\|^2$$

$$+ \sum_{i\in\mathscr{S}^c}|B_{r_i,.}y - B_{r_i',.}y'|^2 - (m-s)(\|y\|^2 + \|y'\|^2)\Big| \geq 2m\epsilon\Big]$$

$$\leq \mathbb{P}\Big[\Big|\sum_{i\in\mathscr{S}}|B_{r_i,.}(y-y')|^2 - s\cdots\Big| \geq m\epsilon\Big]$$

$$+ \mathbb{P}\Big[\Big|\sum_{i\in\mathscr{S}^c}|B_{r_i,.}y - B_{r_i',.}y'|^2 - (m-s)\cdots\Big| \geq m\epsilon\Big]$$

$$\leq 4e^{-c''m\epsilon^2},$$

where $c'' = \min\{c, c'\}$, with $c' = \frac{1}{48}$ and c as in Proposition 1. This gives the desired concentration bound for a fixed $H = y^\mathsf{T} \otimes S$ and $H' = {y'}^\mathsf{T} \otimes S'$ in $\mathcal{A}^\circ$. Finally, following a similar argument as that in the proof of Proposition 1, namely, taking the union bound over the joint net $(y, y') \in \mathcal{N}_\epsilon \times \mathcal{N}_\epsilon$ and all selection matrices $S, S' \in \mathcal{S}_{m,n}$, we obtain the desired result. This completes the proof.

REFERENCES

[1] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[2] E. J. Candes and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.

[3] ——, "Decoding by linear programming," *IEEE transactions on information theory*, vol. 51, no. 12, pp. 4203–4215, 2005.

[4] M. A. Herman and T. Strohmer, "General deviants: An analysis of perturbations in compressed sensing," *IEEE Journal of Selected topics in signal processing*, vol. 4, no. 2, pp. 342–349, 2010.

[5] M. A. Herman and D. Needell, "Mixed operators in compressed sensing," in *Information Sciences and Systems (CISS), 2010 44th Annual Conference on*. IEEE, 2010, pp. 1–6.

[6] J. Unnikrishnan, S. Haghighatshoar, and M. Vetterli, "Unlabeled sensing with random linear measurements," *arXiv preprint arXiv:1512.00115*, 2015.

[7] ——, "Unlabeled sensing: Solving a linear system with unordered measurements," in *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2015, pp. 786–793.

[8] A. Pananjady, M. J. Wainwright, and T. A. Courtade, "Linear regression with an unknown permutation: Statistical and computational limits," *arXiv preprint arXiv:1608.02902*, 2016.

[9] A. v. Balakrishnan, "On the problem of time jitter in sampling," *IRE Transactions on Information Theory*, vol. 8, no. 3, pp. 226–236, 1962.

[10] C. Rose, I. S. Mian, and R. Song, "Timing channels with multiple identical quanta," *arXiv preprint arXiv:1208.1070*, 2012.

[11] R. Fung, A. Hanna, O. Vendrell, S. Ramakrishna, T. Seideman, R. Santra, and A. Ourmazd, "Dynamics from noisy data with extreme timing uncertainty," *Nature*, vol. 532, no. 7600, pp. 471–475, 2016.

[12] M. I. Skolnik, *Introduction to Radar Systems*, 2nd ed. New York: McGraw Hill Book Co., 1980.

[13] X. Huang and A. Madan, "Cap3: A dna sequence assembly program," *Genome research*, vol. 9, no. 9, pp. 868–877, 1999.

[14] G. Bresler, M. Bresler, and D. Tse, "Optimal assembly for high throughput shotgun sequencing," *BMC bioinformatics*, vol. 14, no. 5, p. S18, 2013.

[15] M. Mitzenmacher *et al.*, "A survey of results for deletion channels and related synchronization channels," *Probability Surveys*, vol. 6, no. 1-33, p. 1, 2009.

[16] R. Vershynin, "Estimation in high dimensions: a geometric perspective," in *Sampling Theory, a Renaissance*. Springer, 2015, pp. 3–66.

[17] E. J. Candes, X. Li, and M. Soltanolkotabi, "Phase retrieval via wirtinger flow: Theory and algorithms," *IEEE Transactions on Information Theory*, vol. 61, no. 4, pp. 1985–2007, 2015.

[18] X. Li, S. Ling, T. Strohmer, and K. Wei, "Rapid, robust, and reliable blind deconvolution via nonconvex optimization," *arXiv preprint arXiv:1606.04933*, 2016.

[19] V. Cambareri and L. Jacques, "Through the haze: A non-convex approach to blind calibration for linear random sensing models," *arXiv preprint arXiv:1610.09028*, 2016.

[20] G. Grimmett and D. Stirzaker, *Probability and random processes*. Oxford university press, 2001.