

Deterministic Conditions for Subspace Identifiability from Incomplete Sampling

Daniel L. Pimentel-Alarcón, Nigel Boston, Robert D. Nowak
University of Wisconsin-Madison

Abstract—Consider an r -dimensional subspace of \mathbb{R}^d , $r < d$, and suppose that we are only given projections of this subspace onto small subsets of the canonical coordinates. The paper establishes necessary and sufficient deterministic conditions on the subsets for subspace identifiability. The results also shed new light on low-rank matrix completion.

I. INTRODUCTION

Subspace identification arises in a wide variety of signal and information processing applications. In many cases, especially high-dimensional situations, it is common to encounter missing data. Hence the growing literature concerning the estimation of low-dimensional subspaces and matrices from incomplete data in theory [1–7] and applications [8, 9].

This paper considers the problem of identifying an r -dimensional subspace of \mathbb{R}^d from projections of the subspace onto small subsets of the canonical coordinates. The main contribution of this paper is to establish deterministic necessary and sufficient conditions on such subsets that guarantee that there is *only* one r -dimensional subspace consistent with all the projections. These conditions also have implications for low-rank matrix completion and related problems.

Organization of the paper

In Section II we formally state the problem and our main results. We present the proof of our main theorem in Section III. Section IV illustrates the implications of our results for low-rank matrix completion. Section V presents the graphical interpretation of the problem and another necessary condition based on this viewpoint.

II. MODEL AND MAIN RESULTS

Let S^* denote an r -dimensional subspace of \mathbb{R}^d . Define Ω as a $d \times N$ binary matrix and let ω_i denote the i^{th} column of Ω . The nonzero entries of ω_i indicate the canonical coordinates involved in the i^{th} projection.

Since S^* is r -dimensional, the restriction of S^* onto $\ell \leq r$ coordinates will be \mathbb{R}^ℓ (in general), and hence such a projection will provide no information specific to S^* . Therefore, without loss of generality (see the appendix for immediate generalizations) we will assume that:

A1 Ω has exactly $r + 1$ nonzero entries per column.

Given an r -dimensional subspace S , let $S_{\omega_i} \subset \mathbb{R}^{r+1}$ denote the restriction of S to the nonzero coordinates in ω_i . The question addressed in this paper is whether the restrictions $\{S_{\omega_i}\}_{i=1}^N$ uniquely determine S^* . This depends on the sampling pattern in Ω .




Fig. 1. When can S^* be identified from its canonical projections $\{S_{\omega_i}^*\}_{i=1}^N$?

We will see that identifiability of this sort can only be possible if $N \geq d - r$, since $\ker S^*$ is $(d - r)$ -dimensional. Thus, unless otherwise stated, we will also assume that:

A2 Ω has exactly $N = d - r$ columns.

Let $\text{Gr}(r, \mathbb{R}^d)$ denote the Grassmannian manifold of r -dimensional subspaces in \mathbb{R}^d . Define $\mathcal{S}(S^*, \Omega) \subset \text{Gr}(r, \mathbb{R}^d)$ such that every $S \in \mathcal{S}(S^*, \Omega)$ satisfies $S_{\omega_i} = S_{\omega_i}^* \forall i$. In words, $\mathcal{S}(S^*, \Omega)$ is the set of all r -dimensional subspaces matching S^* on Ω .

Example 1. Let $d = 5$, $r = 1$,

$$S^* = \text{span} \begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 4 \end{bmatrix} \quad \text{and} \quad \Omega = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Then, for example,

$$S_{\omega_3}^* = \text{span} \begin{bmatrix} 1 \\ 3 \end{bmatrix}.$$

It is easy to see that there are infinitely many 1-dimensional subspaces that match S^* on Ω . In fact,

$$\mathcal{S}(S^*, \Omega) = \left\{ \text{span}[1 \ 2 \ 3 \ \alpha \ \alpha]^T : \alpha \in \mathbb{R} \setminus \{0\} \right\}.$$

However, if we instead had $\omega_3 = [0 \ 0 \ 1 \ 1 \ 0]^T$, then S^* would be the only subspace in $\mathcal{S}(S^*, \Omega)$.

The main result of this paper is the following theorem, which gives necessary and sufficient conditions on Ω to guarantee that $\mathcal{S}(S^*, \Omega)$ contains no subspace other than S^* . Our results hold for (a.e.) S^* , with respect to the uniform measure over $\text{Gr}(r, \mathbb{R}^d)$.

Given a matrix, let $n(\cdot)$ denote its number of columns, and $m(\cdot)$ the number of its *nonzero* rows.

Theorem 1. *Let A1 and A2 hold. For almost every S^* , S^* is the only subspace in $\mathcal{S}(S^*, \Omega)$ if and only if every matrix Ω' formed with a subset of the columns in Ω satisfies*

$$m(\Omega') \geq n(\Omega') + r. \quad (1)$$

The proof of Theorem 1 is given in Section III. In words, Theorem 1 is stating that S^* is the only subspace that matches S^* in Ω if and only if every subset of n columns of Ω has at least $n + r$ nonzero rows.

Example 2. *The following matrix, where $\mathbf{1}$ denotes a block of all 1's and \mathbf{I} denotes the identity matrix, satisfies the conditions of Theorem 1:*

$$\Omega = \left[\begin{array}{c} \mathbf{1} \\ \mathbf{I} \end{array} \right] \left\{ \begin{array}{l} r \\ d-r \end{array} \right\}.$$

When the conditions of Theorem 1 are satisfied, identifying S^* becomes a trivial task: $S^* = \ker \mathbf{A}^T$, with \mathbf{A} as defined in Section III.

In general, verifying the conditions on Ω in Theorem 1 may be computationally prohibitive, especially for large d . However, as the next theorem states, uniform random sampling patterns will satisfy the conditions in Theorem 1 with high probability (w.h.p.).

Theorem 2. *Assume A2 and let $0 < \epsilon \leq 1$ be given. Suppose $r \leq \frac{d}{6}$ and that each column of Ω contains at least ℓ nonzero entries, selected uniformly at random and independently across columns, with*

$$\ell \geq \max \left\{ 9 \log \left(\frac{d}{\epsilon} \right) + 12, 2r \right\}. \quad (2)$$

Then Ω will satisfy the conditions of Theorem 1 with probability at least $1 - \epsilon$.

Theorem 2 is proved in the appendix. Notice that $\mathcal{O}(r \log d)$ nonzero entries per column is a typical requirement of LRMC methods, while $\mathcal{O}(\max\{r, \log d\})$ is sufficient for subspace identifiability.

III. PROOF OF THEOREM 1

For any subspace, matrix or vector that is compatible with a binary vector \mathbf{v} , we will use the subscript \mathbf{v} to denote its restriction to the nonzero coordinates/rows in \mathbf{v} . For a.e. S^* , $S_{\omega_i}^*$ is an r -dimensional subspace of \mathbb{R}^{r+1} , and the kernel of $S_{\omega_i}^*$ is a 1-dimensional subspace of \mathbb{R}^{r+1} .

Lemma 1. *Let $\mathbf{a}_{\omega_i} \in \mathbb{R}^{r+1}$ be a nonzero element of $\ker S_{\omega_i}^*$. All entries of \mathbf{a}_{ω_i} are nonzero for a.e. S^* .*

Proof. Suppose \mathbf{a}_{ω_i} has at least one zero entry. Use \mathbf{v} to denote the binary vector of the nonzero entries of \mathbf{a}_{ω_i} . Since \mathbf{a}_{ω_i} is orthogonal to $S_{\omega_i}^*$, for every $\mathbf{u}_{\omega_i} \in S_{\omega_i}^*$ we have that $\mathbf{a}_{\omega_i}^T \mathbf{u}_{\omega_i} = \mathbf{a}_{\mathbf{v}}^T \mathbf{u}_{\mathbf{v}} = 0$. Then $S_{\mathbf{v}}^*$ satisfies

$$\dim S_{\mathbf{v}}^* \leq \dim \ker \mathbf{a}_{\mathbf{v}}^T = \|\mathbf{v}\|_1 - 1 < \|\mathbf{v}\|_1. \quad (3)$$

Observe that for every binary vector \mathbf{v} with $\|\mathbf{v}\|_1 \leq r$, a.e. r -dimensional subspace S satisfies $\dim S_{\mathbf{v}} = \|\mathbf{v}\|_1$. Thus (3) holds only in a set of measure zero. \square

Define \mathbf{a}_i as the vector in \mathbb{R}^d with the entries of \mathbf{a}_{ω_i} in the nonzero positions of ω_i and zeros elsewhere. Then $S \subset \ker \mathbf{a}_i^T$ for every $S \in \mathcal{S}(S^*, \Omega)$ and every i . Letting \mathbf{A} be the $d \times (d-r)$ matrix formed with $\{\mathbf{a}_i\}_{i=1}^{d-r}$ as columns, we have that $S \subset \ker \mathbf{A}^T$ for every $S \in \mathcal{S}(S^*, \Omega)$. Note that if $\dim \ker \mathbf{A}^T = r$, then $\mathcal{S}(S^*, \Omega)$ contains just one element, S^* , which is the identifiability condition of interest. Thus, we will establish conditions on Ω guaranteeing that the $d-r$ columns of \mathbf{A} are linearly independent.

Recall that for any matrix \mathbf{A}' formed with a subset of the columns in \mathbf{A} , $n(\mathbf{A}')$ denotes the number of columns in \mathbf{A}' , and $m(\mathbf{A}')$ denotes the number of *nonzero* rows in \mathbf{A}' .

Lemma 2. *For a.e. S^* , the columns of \mathbf{A} are linearly dependent if and only if $n(\mathbf{A}') > m(\mathbf{A}') - r$ for some matrix \mathbf{A}' formed with a subset of the columns in \mathbf{A} .*

We will show Lemma 2 using Lemmas 3 and 4 below. Let $\aleph(\mathbf{A}')$ be the largest number of linearly independent columns in \mathbf{A}' , i.e., the column rank of \mathbf{A}' .

Lemma 3. *For a.e. S^* , $\aleph(\mathbf{A}') \leq m(\mathbf{A}') - r$.*

Proof. Let \mathbf{v} be the binary vector of nonzero rows of \mathbf{A}' , and $\mathbf{A}'_{\mathbf{v}}$ be the $m(\mathbf{A}') \times n(\mathbf{A}')$ matrix formed with these rows.

For a.e. S^* , $\dim S_{\mathbf{v}}^* = r$. Since $S_{\mathbf{v}}^* \subset \ker \mathbf{A}'_{\mathbf{v}}^T$, $r = \dim S_{\mathbf{v}}^* \leq \dim \ker \mathbf{A}'_{\mathbf{v}}^T = m(\mathbf{A}') - \aleph(\mathbf{A}')$. \square

We say \mathbf{A}' is *minimally linearly dependent* if the columns in \mathbf{A}' are linearly dependent, but every proper subset of the columns in \mathbf{A}' is linearly independent.

Lemma 4. *Let \mathbf{A}' be minimally linearly dependent. Then for a.e. S^* , $n(\mathbf{A}') = m(\mathbf{A}') - r + 1$.*

Proof. Let $\mathbf{A}' = [\mathbf{A}'' \mid \mathbf{a}_i]$ be minimally linearly dependent. Let $m = m(\mathbf{A}'')$, $n = n(\mathbf{A}'')$, and $\aleph = \aleph(\mathbf{A}'')$. Define $\beta \in \mathbb{R}^n$ such that

$$\mathbf{A}'' \beta = \mathbf{a}_i. \quad (4)$$

Note that because \mathbf{A}' is minimally linearly dependent, all entries in β are nonzero. Since the columns of \mathbf{A}'' are linearly independent, $n = \aleph$. Thus, by Lemma 3, $n \leq m - r$. We want to show that $n = m - r$, so suppose for contradiction that $n < m - r$.

We can assume without loss of generality that \mathbf{A}'' has all its zero rows (if any) in the first positions. In that case, since \mathbf{A}' is minimally linearly dependent, it follows that the nonzero entries of \mathbf{a}_i cannot be in the corresponding rows.

Thus, without loss of generality, assume that \mathbf{a}_i has its first r nonzero entries in the first r nonzero rows of \mathbf{A}'' , and that the last nonzero entry of \mathbf{a}_i is 1 (i.e., rescale \mathbf{a}_i if needed), and is located in the last row. Let $\hat{\mathbf{a}}_i \in \mathbb{R}^r$ denote the vector with the first nonzero entries of \mathbf{a}_i , such that we can write:

$$[\mathbf{A}'' \mid \mathbf{a}_i] = \left[\begin{array}{c|c} \mathbf{0} & \mathbf{0} \\ \hline \mathbf{C} & \hat{\mathbf{a}}_i \\ \hline \mathbf{B} & \mathbf{0} \\ \hline & 1 \end{array} \right] \begin{array}{l} \left. \vphantom{\begin{array}{c} \mathbf{0} \\ \mathbf{C} \\ \mathbf{B} \end{array}} \right\} d-m \\ \left. \vphantom{\begin{array}{c} \mathbf{C} \\ \mathbf{B} \end{array}} \right\} r \\ \left. \vphantom{\begin{array}{c} \mathbf{B} \end{array}} \right\} m-r-1 \\ \left. \vphantom{\begin{array}{c} \mathbf{B} \\ 1 \end{array}} \right\} 1 \end{array} \quad (5)$$

where \mathbf{C} and \mathbf{B} are submatrices used to denote the blocks of \mathbf{A}'' corresponding to the partition of \mathbf{a}_i .

The columns of \mathbf{B} are linearly independent. To see this, suppose for contradiction that they are not. This means that there exists some nonzero $\gamma \in \mathbb{R}^n$, such that $\mathbf{B}\gamma = \mathbf{0}$. Let $\mathbf{c} = \mathbf{A}''\gamma$ and note that only the r rows in \mathbf{c} corresponding to the block \mathbf{C} may be nonzero. Let \mathbf{v} denote the binary vector of these nonzero entries. Since S^* is orthogonal to every column of \mathbf{A}'' and \mathbf{c} is a linear combination of the columns in \mathbf{A}'' , it follows that $S_v^* \subset \ker \mathbf{c}_v^T$. This implies that $\dim S_v^* \leq \dim \ker \mathbf{c}_v^T = \|\mathbf{v}\|_1 - 1$. As in the proof of Lemma 1, this implies that the columns of \mathbf{B} are linearly dependent only in a set of measure zero.

Going back to (5), since the n columns of \mathbf{B} are linearly independent and because we are assuming that $n < m - r$, it follows that \mathbf{B} has n linearly independent rows. Let \mathbf{B}_1 denote the $n \times n$ block of \mathbf{B} that contains n linearly independent rows, and \mathbf{B}_2 the $(m - n - r) \times n$ remaining block of \mathbf{B} .

Notice that the row of \mathbf{B} corresponding to the 1 in \mathbf{a}_i must belong to \mathbf{B}_1 , since otherwise, we have that $\mathbf{B}_1\beta = \mathbf{0}$, with β as in (4), which implies that \mathbf{B}_1 is rank deficient, in contradiction to its construction.

We can further assume without loss of generality that the first nonzero entry of every column of \mathbf{B} is 1 (otherwise we may just rescale each column), and that these nonzero entries are in the first columns (otherwise we may just permute the columns accordingly). We will also let $\tilde{\mathbf{B}}_2$ denote all but the first row of \mathbf{B}_2 . Thus, our matrix is organized as

$$[\mathbf{A}'' \mid \mathbf{a}_i] = \mathbf{B}_2 \left\{ \begin{array}{c|c} \mathbf{0} & \mathbf{0} \\ \hline \mathbf{C} & \hat{\mathbf{a}}_i \\ \hline \mathbf{1} & \mathbf{0} & \mathbf{0} \\ \hline \tilde{\mathbf{B}}_2 & \mathbf{0} \\ \hline \mathbf{B}_1 & \mathbf{0} \\ \hline & 1 \end{array} \right\} \begin{array}{l} \left. \vphantom{\begin{array}{c} \mathbf{0} \\ \mathbf{C} \end{array}} \right\} d-m \\ \left. \vphantom{\begin{array}{c} \mathbf{C} \end{array}} \right\} r \\ \left. \vphantom{\begin{array}{c} \mathbf{1} \\ \tilde{\mathbf{B}}_2 \end{array}} \right\} 1 \\ \left. \vphantom{\begin{array}{c} \tilde{\mathbf{B}}_2 \end{array}} \right\} m-n-r-1 \\ \left. \vphantom{\begin{array}{c} \mathbf{B}_1 \end{array}} \right\} n-1 \\ \left. \vphantom{\begin{array}{c} \mathbf{B}_1 \\ 1 \end{array}} \right\} 1 \end{array} \quad (6)$$

Now (4) implies $\mathbf{B}_1\beta = [\mathbf{0} \mid 1]^T$, and since \mathbf{B}_1 is full rank, we may write

$$\beta = \mathbf{B}_1^{-1} \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix},$$

i.e., β is the the last column of the inverse of \mathbf{B}_1 , which is a rational function in the elements of \mathbf{B}_1 .

Next, let us look back at (4). If $n < m - r$, then using the additional row $[\mathbf{1} \mid \mathbf{0}]$ of (6) (which does not appear if $m = n + r$) we obtain $[\mathbf{1} \mid \mathbf{0}]\beta = 0$. Recall that all the entries of β are nonzero. Thus, the last equation defines the following nonzero rational function in the elements of \mathbf{B}_1 :

$$[\mathbf{1} \mid \mathbf{0}]\mathbf{B}_1^{-1} \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix} = 0. \quad (7)$$

Equivalently, (7) is a polynomial equation in the elements of \mathbf{B}_1 , which we will denote as $f(\mathbf{B}_1) = 0$.

Next note that for a.e. S^* , we can write $S^* = \ker \mathbf{A}^{*T}$ for a unique $\mathbf{A}^* \in \mathbb{R}^{d \times (d-r)}$ in column echelon form¹:

$$\mathbf{A}^* = \left[\begin{array}{c|c} \mathbf{I} & \\ \hline & \mathbf{D}^* \end{array} \right] \begin{array}{l} \left. \vphantom{\begin{array}{c} \mathbf{I} \end{array}} \right\} d-r \\ \left. \vphantom{\begin{array}{c} \mathbf{D}^* \end{array}} \right\} r \end{array}. \quad (8)$$

On the other hand every $\mathbf{D}^* \in \mathbb{R}^{r \times (d-r)}$ defines a unique r -dimensional subspace of \mathbb{R}^d , via (8). Thus, we have a bijection between $\mathbb{R}^{r \times (d-r)}$ and a dense open subset of $\text{Gr}(r, \mathbb{R}^d)$.

Since the columns of \mathbf{A}'' must be linear combinations of the columns of \mathbf{A}^* , the elements of \mathbf{B}_1 are linear functions in the entries of \mathbf{D}^* . Therefore, we can express $f(\mathbf{B}_1)$ as a nonzero polynomial function g in the entries of \mathbf{D}^* and rewrite (7) as $g(\mathbf{D}^*) = 0$. But we know that $g(\mathbf{D}^*) \neq 0$ for almost every $\mathbf{D}^* \in \mathbb{R}^{r \times (d-r)}$, and hence for almost every $S^* \in \text{Gr}(r, \mathbb{R}^d)$. We conclude that almost every subspace in $\text{Gr}(r, \mathbb{R}^d)$ will not satisfy (7), and thus $n = m - r$. \square

We are now ready to present the proofs of Lemma 2 and Theorem 1.

Proof. (Lemma 2)

- (\Rightarrow) Suppose \mathbf{A}' is minimally linearly dependent. By Lemma 4, $n(\mathbf{A}') = m(\mathbf{A}') - r + 1 > m(\mathbf{A}') - r$, and we have the first implication.
- (\Leftarrow) Suppose there exists an \mathbf{A}' with $n(\mathbf{A}') > m(\mathbf{A}') - r$. By Lemma 3, $n(\mathbf{A}') > \kappa(\mathbf{A}')$, which implies the columns in \mathbf{A}' , and hence \mathbf{A} , are linearly dependent. \square

Proof. (Theorem 1) Lemma 1 shows that for a.e. S^* , the $(j, i)^{th}$ entry of \mathbf{A} is nonzero if and only if the $(j, i)^{th}$ entry of Ω is nonzero.

- (\Rightarrow) Suppose there exists an Ω' such that $m(\Omega') < n(\Omega') + r$. Then $m(\mathbf{A}') < n(\mathbf{A}') + r$ for some \mathbf{A}' . Lemma 2

¹Certain S^* may not admit this representation, e.g., if S^* is orthogonal to certain canonical coordinates, which, as discussed in Lemma 1, is not the case for almost every S^* in $\text{Gr}(r, \mathbb{R}^d)$.

implies that the columns of \mathbf{A}' , and hence \mathbf{A} , are linearly dependent. This implies $\dim \ker \mathbf{A}^\top > r$.

(\Leftarrow) Suppose every Ω' satisfies $m(\Omega') \geq n(\Omega') + r$. Then $m(\mathbf{A}') \geq n(\mathbf{A}') + r$ for every \mathbf{A}' , including \mathbf{A} . Therefore, by Lemma 2, the $d - r$ columns in \mathbf{A} are linearly independent, hence $\dim \ker \mathbf{A}^\top = r$. \square

IV. IMPLICATIONS FOR LOW-RANK MATRIX COMPLETION

Subspace identifiability is closely related to the low-rank matrix completion (LRMC) problem [4]: given a subset of entries in a rank- r matrix, exactly recover *all* of the missing entries. This requires, implicitly, identification of the subspace spanned by the complete columns of the matrix. We use this section to present the implications of our results for LRMC.

Let \mathbf{X} be a $d \times N$, rank- r matrix and assume that

A3 The columns of \mathbf{X} are drawn independently according to ν , an absolutely continuous distribution with respect to the Lebesgue measure on S^* .

Let \mathbf{X}_Ω be the incomplete version of \mathbf{X} , observed only in the nonzero positions of Ω .

Necessary and sufficient conditions for LRMC

To relate the LRMC problem to our main results, define \tilde{N} as the number of *distinct* columns (sampling patterns) in Ω , and let $\tilde{\Omega}$ denote a $d \times \tilde{N}$ matrix composed of these columns.

Corollary 1. *If $\tilde{\Omega}$ does not contain a $d \times (d - r)$ submatrix satisfying the conditions of Theorem 1, then \mathbf{X} cannot be uniquely recovered from \mathbf{X}_Ω .*

Since \mathbf{X} is rank- r , a column with fewer than r observed entries cannot be completed (in general). We will thus assume without loss of generality the following relaxation of **A1**:

A1' Ω has at least r nonzero entries per column.

Corollary 2. *Let **A1'** and **A3** hold. Suppose $\tilde{\Omega}$ contains a $d \times (d - r)$ submatrix satisfying the conditions of Theorem 1, and that for every column ω_i in this submatrix, at least r columns in \mathbf{X}_Ω are observed at the nonzero locations of ω_i . Then for a.e. S^* , and almost surely with respect to ν , \mathbf{X} can be uniquely recovered from \mathbf{X}_Ω .*

Proofs of these results are given in the appendix. The intuition behind Corollary 1 is simply that identifying a subspace from its projections onto sets of canonical coordinates is *easier* than LRMC, and so the necessary condition of Theorem 1 is also necessary for LRMC. Corollary 2 follows from the fact that S^* (or its projections) can be determined from r or more observations drawn from ν .

Validating LRMC

Under certain assumptions on the subset of observed entries (e.g., random sampling) and S^* (e.g., incoherence), existing methods, for example nuclear norm minimization [4], succeed with high probability in completing the matrix exactly and

thus identifying S^* . These assumptions are sufficient, but not necessary, and are sometimes unverifiable or unjustified in practice. Therefore, the result of an LRMC algorithm can be suspect. Simply finding a low-rank matrix that agrees with the observed data does not guarantee that it is the correct completion. It is possible that there exist other r -dimensional subspaces different from S^* that agree with the observed entries.

Example 3. *Suppose we run an LRMC algorithm on a matrix observed on the support of Ω , with Ω and S^* as in Example 1 in Section II. Suppose that the algorithm produces a completion with columns from $S = \text{span}[1 \ 2 \ 3 \ 5 \ 5]^\top$ instead of S^* . It is clear that the residual of the projection of any vector from $S_{\omega_i}^*$ onto S_{ω_i} will be zero, despite the fact that $S \neq S^*$.*

In other words, if the residuals are nonzero, we can discard an incorrect solution, but if the residuals are zero, we cannot validate whether our solution is correct or not.

Corollary 3, below, allows one to drop the sampling and incoherence assumptions, and validate the result of any LRMC algorithm deterministically.

Let \mathbf{x}_i denote the i^{th} column of \mathbf{X} , and \mathbf{x}_{ω_i} be the restriction of \mathbf{x}_i to the nonzero coordinates of ω_i . We say that a subspace S fits \mathbf{X}_Ω if $\mathbf{x}_{\omega_i} \in S_{\omega_i}$ for every i .

Corollary 3. *Let **A3** hold, and suppose \mathbf{X}_Ω contains two disjoint sets of columns, \mathbf{X}_{Ω_1} and \mathbf{X}_{Ω_2} , such that Ω_2 is a $d \times (d - r)$ matrix satisfying the conditions of Theorem 1. Let S be the subspace spanned by the columns of a completion of \mathbf{X}_{Ω_1} . Then for a.e. S^* , and almost surely with respect to ν , S fits \mathbf{X}_{Ω_2} if and only if $S = S^*$.*

The proof of Corollary 3 is given in the appendix. In words, Corollary 3 states that if one runs an LRMC algorithm on \mathbf{X}_{Ω_1} , then the uniqueness and correctness of the resulting low-rank completion can be verified by testing whether it agrees with the validation set \mathbf{X}_{Ω_2} .

Example 4. *Consider a 1000×2000 matrix \mathbf{X}_Ω with $r = 30$ and ideal incoherence. In this case, the best sufficient conditions for LRMC that we are aware of [5] require that all entries are observed. Simulations show that alternating minimization [7] can exactly complete such matrices when fewer than half of the entries are observed, and only using half of the columns. While previous theory for matrix completion gives no guarantees in scenarios like this, our new results do.*

To see this, split \mathbf{X}_Ω into two 1000×1000 submatrices \mathbf{X}_{Ω_1} and \mathbf{X}_{Ω_2} . Use nuclear norm, alternating minimization, or any LRMC method, to find a completion of \mathbf{X}_{Ω_1} . Theorem 2 can be used to show that the sampling of \mathbf{X}_{Ω_2} will satisfy the conditions of Theorem 1 w.h.p. even when only half the entries are observed randomly. We can then use Corollary 3 to show that if \mathbf{X}_{Ω_2} is consistent with the completion of \mathbf{X}_{Ω_1} , then the completion is unique and correct.

Remarks

Observe that the necessary and sufficient conditions in Corollaries 1 and 2 and the validation in Corollary 3 do not require the incoherence assumptions typically needed in LRMC results in order to guarantee correctness and uniqueness.

Another advantage of results above is that they work for matrices of any rank, while standard LRMC results only hold for ranks significantly smaller than the dimension d .


Finally, the results above hold with probability 1, as opposed to standard LRMC statements, that hold w.h.p. On the other hand, verifying whether Ω_2 meets the conditions of Theorem 1 may be difficult. Nevertheless, if the entries in our data matrix are sampled randomly with rates comparable to standard conditions in LRMC, we know by Theorem 2 that w.h.p. Ω_2 will satisfy such conditions.

V. GRAPHICAL INTERPRETATION OF THE PROBLEM

The problem of LRMC has also been studied from the graph theory perspective. For example, it has been shown that graph connectivity is a necessary condition for completion [6]. Being subspace identifiability so tightly related to LRMC, it comes as no surprise that there also exist graph conditions for subspace identifiability. In this section we draw some connections between subspace identifiability and graph theory that give insight on the conditions in Theorem 1. We use this interpretation to show that graph connectivity is a necessary yet insufficient condition for subspace identification.

Define $\mathcal{G}(\Omega)$ as the bipartite graph with disjoint sets of row and column vertices, where there is an edge between row vertex j and column vertex i if the $(j, i)^{th}$ entry of Ω is nonzero.

Example 5. With $d = 5$, $r = 1$ and




Recall that the *neighborhood* of a set of vertices is the collection of all their adjacent vertices.

The graph theoretic interpretation of the condition on Ω in Theorem 1 is that every set of n column vertices in $\mathcal{G}(\Omega)$ must have a neighborhood of at least $n + r$ row vertices.

Example 6. One may verify that every set of n column vertices in $\mathcal{G}(\Omega)$ from Example 5 has a neighborhood of at least $n + r$ row vertices. On the other hand, if we consider Ω as in

Example 1, the neighborhood of the column vertices $\{1, 2, 3\}$ in $\mathcal{G}(\Omega)$ contains fewer than $n + r$ row vertices:



With this interpretation of Theorem 1, we can extend terms and results from graph theory to our context. One example is the next corollary, which states that r -row-connectivity is a necessary but insufficient condition for subspace identifiability.

We say $\mathcal{G}(\Omega)$ is *r-row-connected* if $\mathcal{G}(\Omega)$ remains a connected graph after removing any set of $r - 1$ row vertices and all their adjacent edges.

Corollary 4. For a.e. S^* , $|\mathcal{S}(S^*, \Omega)| > 1$ if $\mathcal{G}(\Omega)$ is not r -row-connected. The converse is only true for $r = 1$.

Corollary 4 is proved in the appendix.

VI. CONCLUSIONS

In this paper we determined when and only when can one identify a subspace from its projections onto subsets of the canonical coordinates. We show that the conditions for identifiability hold w.h.p. under standard random sampling schemes, and that when these conditions are met, identifying the subspace becomes a trivial task.

This gives new necessary and sufficient conditions for LRMC, and allows one to verify whether the result of *any* LRMC algorithm is unique and correct without prior incoherence or sampling assumptions.

REFERENCES

- [1] L. Balzano, B. Recht and R. Nowak, *High-dimensional matched subspace detection when data are missing*, IEEE International Symposium on Information Theory, 2010.
- [2] Y. Chi, Y. Eldar and R. Calderbank, *PETRELS: Subspace estimation and tracking from partial observations*, IEEE International Conference on Acoustics, Speech and Signal Processing, 2012.
- [3] M. Mardani, G. Mateos and G. Giannakis, *Rank minimization for subspace tracking from incomplete data*, IEEE International Conference on Acoustics, Speech and Signal Processing, 2013.
- [4] E. Candès and B. Recht, *Exact matrix completion via convex optimization*, Foundations of Computational Mathematics, 2009.
- [5] B. Recht, *A simpler approach to matrix completion*, Journal of Machine Learning Research, 2011.
- [6] F. Király and R. Tomioka, *A combinatorial algebraic approach for the identifiability of low-rank matrix completion*, International Conference on Machine Learning, 2012.
- [7] P. Jain, P. Netrapalli and S. Sanghavi, *Low-rank matrix completion using alternating minimization*, ACM Symposium on Theory Of Computing, 2013.
- [8] B. Eriksson, P. Barford and R. Nowak, *Network discovery from passive measurements*, ACM SIGCOMM, 2008.
- [9] J. He, L. Balzano and A. Szlam, *Incremental gradient on the grassmannian for online foreground and background separation in subsampled video*, Conference on Computer Vision and Pattern Recognition, 2012.
- [10] B. Bollobás, *Extremal graph theory*, Dover Publications, 2004.

Generalization of Our Results

Since the restriction of S^* onto $\ell \leq r$ coordinates will be \mathbb{R}^ℓ (in general), such a projection will provide no information specific to S^* . We will thus assume without loss of generality that:

A1'' Ω' has at least $r+1$ nonzero entries per column.

Under **A1''**, a column with ℓ observed entries restricts $\mathcal{S}(S^*, \Omega)$ just as $\ell-r$ columns under **A1**. Thus in general, if there are columns in Ω with more than $r+1$ nonzero entries, we can *split* them to obtain an *expanded* matrix $\check{\Omega}$ (defined below), with exactly $r+1$ nonzero entries per column, and use Theorem 1 directly on this expanded matrix.

More precisely, let k_1, \dots, k_{ℓ_i} denote the indices of the ℓ_i nonzero entries in the i^{th} column of Ω . Define Ω_i as the $d \times (\ell_i - r)$ matrix, whose j^{th} column has the value 1 in rows k_1, \dots, k_r, k_{r+j} , and zeros elsewhere. For example, if $k_1 = 1, \dots, k_{\ell_i} = \ell_i$, then

$$\Omega_i = \left[\begin{array}{c} \mathbf{1} \\ \mathbf{I} \\ \mathbf{0} \end{array} \right] \begin{array}{l} \left. \vphantom{\begin{array}{c} \mathbf{1} \\ \mathbf{I} \\ \mathbf{0} \end{array}} \right\} r \\ \left. \vphantom{\begin{array}{c} \mathbf{1} \\ \mathbf{I} \\ \mathbf{0} \end{array}} \right\} \ell_i - r \\ \left. \vphantom{\begin{array}{c} \mathbf{1} \\ \mathbf{I} \\ \mathbf{0} \end{array}} \right\} d - \ell_i, \end{array}$$

$\ell_i - r$

where $\mathbf{1}$ denotes a block of all 1's and \mathbf{I} the identity matrix. Finally, define $\check{\Omega} := [\Omega_1 \dots \Omega_N]$.

The following is a generalization of Theorem 1 to an arbitrarily number of projections and an arbitrary number of canonical coordinates involved in each projection. It states that S^* will be the only subspace in $\mathcal{S}(S^*, \Omega)$ if and only if there is a matrix $\check{\Omega}$, formed with $d-r$ columns of $\check{\Omega}$, that satisfies the conditions of Theorem 1.

Theorem 3. *Let A1'' hold. For almost every S^* , S^* is the only subspace in $\mathcal{S}(S^*, \Omega)$ if and only if there is a matrix $\check{\Omega}$, formed with $d-r$ columns of $\check{\Omega}$, such that every matrix Ω' formed with a subset of the columns in $\check{\Omega}$ satisfies (1).*

Proof. It suffices to show that $\mathcal{S}(S^*, \Omega) = \mathcal{S}(S^*, \check{\Omega})$. Let ω_{ij} denote the j^{th} column of Ω_i .

- (c) Let $S \in \mathcal{S}(S^*, \Omega)$. By definition, $S_{\omega_i} = S^*_{\omega_i}$, which trivially implies $\{S_{\omega_{ij}} = S^*_{\omega_{ij}}\}_{j=1}^{\ell_i-r}$. Since this is true for every i , we conclude $S \in \mathcal{S}(S^*, \check{\Omega})$.
- (c) Let $S \in \mathcal{S}(S^*, \check{\Omega})$. By definition, $\{S_{\omega_{ij}} = S^*_{\omega_{ij}}\}_{j=1}^{\ell_i-r}$. Notice that Ω_i satisfies the conditions of Theorem 1 restricted to the nonzero rows in ω_i , which implies $S_{\omega_i} = S^*_{\omega_i}$. Since this is true for every i , we conclude $S \in \mathcal{S}(S^*, \Omega)$. \square

Proof of Theorem 2

Let \mathcal{E} be the event that Ω fails to satisfy the conditions of Theorem 1. It is easy to see that this may only occur if there is a matrix formed with n columns from Ω that has all its nonzero entries in the same $n+r-1$ rows. Let \mathcal{E}_n denote the event that the matrix formed with the first n columns from Ω has all its nonzero entries in the first $n+r-1$ rows. Then

$$\mathbb{P}(\mathcal{E}) \leq \sum_{n=1}^{d-r} \binom{d-r}{n} \binom{d}{n+r-1} \mathbb{P}(\mathcal{E}_n) \quad (9)$$

If each column of Ω contains at least ℓ nonzero entries, distributed uniformly and independently at random with ℓ as in (2), it is easy to see that $\mathbb{P}(\mathcal{E}_n) = 0$ for $n \leq \ell-r$, and for $\ell-r < n \leq d-r$,

$$\mathbb{P}(\mathcal{E}_n) \leq \left(\frac{\binom{n+r-1}{\ell}}{\binom{d}{\ell}} \right)^n < \left(\frac{n+r-1}{d} \right)^{\ell n}.$$

Since $\binom{d-r}{n} < \binom{d}{n+r-1}$, continuing with (9) we obtain:

$$\begin{aligned} \mathbb{P}(\mathcal{E}) &< \sum_{n=\ell-r+1}^{d-r} \binom{d}{n+r-1}^2 \left(\frac{n+r-1}{d} \right)^{\ell n} \\ &< \sum_{n=\ell}^{\frac{d}{2}} \binom{d}{n}^2 \left(\frac{n}{d} \right)^{\ell(n-r+1)} \end{aligned} \quad (10)$$

$$+ \sum_{n=1}^{\frac{d}{2}} \binom{d}{d-n}^2 \left(\frac{d-n}{d} \right)^{\ell(d-n-r+1)}. \quad (11)$$

For the terms in (10), write

$$\binom{d}{n}^2 \left(\frac{n}{d} \right)^{\ell(n-r+1)} \leq \left(\frac{de}{n} \right)^{2n} \left(\frac{n}{d} \right)^{\ell(n-r+1)}. \quad (12)$$

Since $n \geq \ell \geq 2r$,

$$(12) < \left(\frac{de}{n} \right)^{2n} \left(\frac{n}{d} \right)^{\frac{\ell n}{2}} = e^{2n} \left(\frac{n}{d} \right)^{(\frac{\ell}{2}-2)n}, \quad (13)$$

and since $n \leq \frac{d}{2}$,

$$(13) \leq e^{2n} \left(\frac{1}{2} \right)^{(\frac{\ell}{2}-2)n} = \left(e^2 \cdot 2^{-\frac{\ell}{2}+2} \right)^n < \frac{\epsilon}{d}, \quad (14)$$

where the last step follows because $\ell > 2 \log_2(\frac{de^2}{\epsilon}) + 4$.

For the terms in (11), write

$$\binom{d}{d-n}^2 \left(\frac{d-n}{d} \right)^{\ell(d-n-r+1)} \leq \left(\frac{de}{n} \right)^{2n} \left(\frac{d-n}{d} \right)^{\ell(d-n-r+1)}. \quad (15)$$

In this case, since $1 \leq n \leq \frac{d}{2}$ and $r \leq \frac{d}{6}$, we have

$$\begin{aligned} (15) &< (de)^{2n} \left(\frac{d-n}{d} \right)^{\frac{\ell d}{3}} = (de)^{2n} \left[\left(1 - \frac{n}{d} \right)^d \right]^{\frac{\ell}{3}} \\ &\leq (de)^{2n} [e^{-n}]^{\frac{\ell}{3}}, \end{aligned}$$

which we may rewrite as

$$(e^{2 \log d})^n (e^2)^n \left(e^{-\frac{\ell}{3}} \right)^n = \left(e^{2 \log d + 2 - \frac{\ell}{3}} \right)^n < \frac{\epsilon}{d}, \quad (16)$$

where the last step follows because $\ell > 3\log(\frac{d}{\epsilon}) + 6\log d + 6$. Substituting (14) and (16) in (10) and (11), we have that $P(\mathcal{E}) < \epsilon$, as desired. \square

Proof of Corollary 1

A subspace satisfying $S_{\omega_i} = S_{\omega_i}^*$ will fit all the columns of \mathbf{X}_Ω observed in the nonzero positions of ω_i . Therefore, any subspace that satisfies $S_{\omega_i} = S_{\omega_i}^*$ for every ω_i in $\tilde{\Omega}$ will fit all the columns in \mathbf{X}_Ω . If $\tilde{\Omega}$ does not satisfy the conditions of Theorem 1, there will exist multiple subspaces that fit \mathbf{X}_Ω , whence \mathbf{X} cannot be uniquely recovered from \mathbf{X}_Ω . \square

Proof of Corollary 2

Suppose there are at least r columns in \mathbf{X}_Ω observed in the nonzero positions of ω_i . Then almost surely with respect to ν , the restrictions of such columns form a basis for $S_{\omega_i}^*$. Therefore, any subspace S that fits such columns must satisfy $S_{\omega_i} = S_{\omega_i}^*$. If this is true for every ω_i in a $d \times (d-r)$ submatrix of $\tilde{\Omega}$, then any subspace that fits \mathbf{X}_Ω must satisfy $S_{\omega_i} = S_{\omega_i}^*$ for every ω_i in this submatrix.

There will be only one subspace that satisfies this condition if this submatrix satisfies the conditions in Theorem 1. Finally, observe that under $\mathbf{A}\mathbf{1}'$, the condition that \mathbf{X} can be uniquely recovered from \mathbf{X}_Ω is equivalent to saying that S^* is the only r -dimensional subspace that fits \mathbf{X}_Ω . \square

Proof of Corollary 3

- (\Leftarrow) $\mathbf{x}_{\omega_i} \in S_{\omega_i}^*$ by assumption, so if $S = S^*$, it is trivially true that $\mathbf{x}_{\omega_i} \in S_{\omega_i}$.
- (\Rightarrow) Use $i = 1, \dots, (d-r)$ to index the columns in \mathbf{X}_{Ω_2} . Since S fits \mathbf{X}_{Ω_2} , by definition $\mathbf{x}_{\omega_i} \in S_{\omega_i}$. On the other hand, $\mathbf{x}_{\omega_i} \in S_{\omega_i}^*$ by assumption, which implies that for every i , \mathbf{x}_{ω_i} lies in the intersection of S_{ω_i} and $S_{\omega_i}^*$. Recall that \mathbf{x}_i is sampled independently according to ν , an absolutely continuous distribution with respect to the Lebesgue measure on S^* . Since $\dim S_{\omega_i} \leq r$, and for a.e. S^* , $\dim S_{\omega_i}^* = r$, the event


$$\bigcap_{i=1}^{d-r} \{\mathbf{x}_{\omega_i} \in S_{\omega_i} \cap S_{\omega_i}^*\}$$

will (almost surely with respect to ν) only happen if $S_{\omega_i} = S_{\omega_i}^* \forall i$, that is, if $S \in \mathcal{S}(S^*, \Omega_2)$. Since Ω_2 satisfies the conditions of Theorem 1, S^* is the only subspace in $\mathcal{S}(S^*, \Omega_2)$. This implies $S = S^*$, which concludes the proof. \square

Proof of Corollary 4

- (\Leftarrow) Suppose $\mathcal{G}(\Omega)$ is not r -row-connected. This means there exists a set Υ of $r-1$ row vertices such that if removed with their respective edges, $\mathcal{G}(\Omega)$ becomes a disconnected graph. Let Υ' , Υ'' and Υ be a partition of the row vertices in $\mathcal{G}(\Omega)$ such that Υ' and Υ'' become disconnected when Υ is removed. Similarly, let Ω' and Ω'' be a partition of the columns in Ω such that the column vertices corresponding to Ω'

are disconnected from the row vertices in Υ'' , and the column vertices corresponding to Ω'' are disconnected from the row vertices in Υ' .



Let $m' = m(\Omega')$, $m'' = m(\Omega'')$, $n' = n(\Omega')$ and $n'' = n(\Omega'')$. It is easy to see that m' denotes the number of row vertices that Ω' is connected to. Then

$$|\Upsilon'| + r - 1 = |\Upsilon'| + |\Upsilon| \geq m'. \quad (17)$$

Now suppose for contradiction that $|\mathcal{S}(S^*, \Omega)| = 1$. By Theorem 1, $m' \geq n' + r$. Substituting this into (17) we obtain

$$|\Upsilon'| \geq n' + 1, \quad (18)$$

$$|\Upsilon''| \geq n'' + 1, \quad (19)$$

where (19) follows by symmetry.

Now observe that since Υ' , Υ'' and Υ form a partition of the row vertices, $d = |\Upsilon'| + |\Upsilon''| + |\Upsilon|$, so using (18) and (19) we obtain

$$d - r \geq n' + n'' + 1. \quad (20)$$

On the other hand, since Ω' and Ω'' form a partition of the $d-r$ columns in Ω ,

$$d - r = n' + n''.$$


Plugging this in (20), we obtain $0 \geq 1$, which is a contradiction. We thus conclude that $|\mathcal{S}(S^*, \Omega)| > 1$.

- (\Rightarrow) For $r = 1$, we prove the converse by contrapositive. Suppose $|\mathcal{S}(S^*, \Omega)| > 1$. By Theorem 1 there exists a matrix Ω' formed with a subset of the columns of Ω with $m' < n' + 1$. Let Ω'' be the matrix formed with the remaining columns of Ω .

If $m' + m'' < d$, there is at least one row in $\mathcal{G}(\Omega)$ that is disconnected, and the converse follows trivially, so suppose $m' + m'' = d$. Observe that $n' + n'' = d - 1$. Putting these two equations together, we obtain

$$m' + m'' - n' - 1 = n''. \quad (21)$$

Since $m' \leq n'$, we obtain $m'' > n''$. Let Υ' and Υ'' be the row vertices connected to the column vertices in Ω' and Ω'' respectively.



Now observe that since each column vertex only has two edges, the column vertices in Ω'' may connect at most $n'' + 1$ row vertices. Since $m'' > n''$, either (i) the edges of Ω'' connect only vertices in Υ'' , leaving Υ' and Υ'' disconnected, or (ii) the edges of Ω'' connect a vertex in Υ' with a vertex in Υ'' , leaving at least one vertex in Υ'' disconnected. Either case, $\mathcal{G}(\Omega)$ is disconnected, as claimed.

For $r = 2$, consider the following sampling:

$$\Omega = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix}.$$

One may verify that $\mathcal{G}(\Omega)$ is r -row-connected, yet it does not satisfy the conditions of Theorem 1. For instance the first 3 columns of Ω , fail to satisfy (1). This example can be easily generalized for $r > 2$. \square