

Noisy Subspace Clustering via Matching Pursuits

Michael Tschannen, *Student Member, IEEE*, and Helmut Bölcskei, *Fellow, IEEE*

Abstract—Sparsity-based subspace clustering algorithms have attracted significant attention thanks to their excellent performance in practical applications. A prominent example is the sparse subspace clustering (SSC) algorithm by Elhamifar and Vidal, which performs spectral clustering based on an adjacency matrix obtained by sparsely representing each data point in terms of all the other data points via the Lasso. When the number of data points is large or the dimension of the ambient space is high, the computational complexity of SSC quickly becomes prohibitive. Dyer et al. observed that SSC-OMP obtained by replacing the Lasso by the greedy orthogonal matching pursuit (OMP) algorithm results in significantly lower computational complexity, while often yielding comparable performance. The central goal of this paper is an analytical performance characterization of SSC-OMP for noisy data. Moreover, we introduce and analyze the SSC-MP algorithm, which employs matching pursuit (MP) in lieu of OMP. Both SSC-OMP and SSC-MP are proven to succeed even when the subspaces intersect and when the data points are contaminated by severe noise. The clustering conditions we obtain for SSC-OMP and SSC-MP are similar to those for SSC and for the thresholding-based subspace clustering (TSC) algorithm due to Heckel and Bölcskei. Analytical results in combination with numerical results indicate that both SSC-OMP and SSC-MP with a data-dependent stopping criterion automatically detect the dimensions of the subspaces underlying the data. Experiments on synthetic and on real data show that SSC-MP often matches or exceeds the performance of the computationally more expensive SSC-OMP algorithm. Moreover, SSC-MP compares very favorably to SSC, TSC, and the nearest subspace neighbor algorithm, both in terms of clustering performance and running time. In addition, we find that, in contrast to SSC-OMP, the performance of SSC-MP is very robust with respect to the choice of parameters in the stopping criteria.

Index Terms—Subspace clustering, matching pursuit algorithms, sparse signal representations, unions of subspaces, spectral clustering, noisy data.

I. INTRODUCTION

Extracting structural information from large high-dimensional data sets in a computationally efficient manner is a major challenge in many modern machine learning tasks. A structure widely encountered in practical applications is that of unions of (low-dimensional) subspaces. The problem of extracting the assignments of the data points in a given data set to the subspaces without prior knowledge of the number of subspaces, their orientations and dimensions is referred to as subspace clustering and has found applications in, e.g., image representation and segmentation [1], face clustering [2], motion segmentation [3], system identification [4], and genomic inference [5]. More formally, given a set $\mathcal{Y} = \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_L$ of N data points in \mathbb{R}^m , where the

points in \mathcal{Y}_ℓ lie in or near the d_ℓ -dimensional linear subspace $\mathcal{S}_\ell \subset \mathbb{R}^m$, we want to find the association of the points in \mathcal{Y} to the \mathcal{Y}_ℓ , without prior knowledge on the \mathcal{S}_ℓ .

The subspace clustering problem has been studied for more than two decades with a correspondingly sizeable body of literature. The algorithms available to date can roughly be categorized as algebraic, statistical, and spectral clustering-based; we refer to [6] for a review of the most prominent representatives of each class. While many subspace clustering algorithms exhibit good performance in practice, corresponding analytical results under non-restrictive conditions on the relative orientations of the subspaces are available only for a small set of algorithms. Specifically, during the past few years a number of new algorithms, which rely on sparse representations (of each data point in terms of all the other data points) followed by spectral clustering [7], were proposed and mathematically analyzed [8], [9], [10], [11], [12], [13], [14], [15]. These algorithms exhibit good empirical performance and succeed provably under quite generous conditions on the relative orientations of the subspaces. Almost all analytical performance results available to date apply, however, to the noiseless case, where the data points lie exactly in the union of the \mathcal{S}_ℓ . A notable exception is the sparse subspace clustering (SSC) algorithm by Elhamifar and Vidal [8], which was shown by Soltanolkotabi et al. [10] and Wang and Xu [11] to succeed for noisy data even when the subspaces intersect. SSC employs the Lasso¹ (or ℓ_1 -minimization in the noiseless case) to find a sparse representation (or, more precisely, approximation) of each data point in terms of all the other data points, then constructs an affinity graph based on the so-obtained sparse representations, and finally determines subspace assignments through spectral clustering of the affinity graph. To understand the intuition behind this approach, first note that in the noiseless case every data point \mathbf{y}_j in \mathcal{S}_ℓ can be represented by (at most d_ℓ) other data points in \mathcal{S}_ℓ provided that the points in \mathcal{Y}_ℓ are non-degenerate. In the noisy case, the hope is now that the sparse representation of $\mathbf{y}_j \in \mathcal{Y}_\ell$ in terms of $\mathcal{Y} \setminus \{\mathbf{y}_j\}$ delivered by SSC involves mostly points belonging to \mathcal{Y}_ℓ thanks to the sparsity-promoting nature of the Lasso. Of course, this will happen only if the subspaces \mathcal{S}_ℓ underlying the \mathcal{Y}_ℓ are sufficiently far apart. The analytical performance results in [9], [10], [11] quantify the impact of subspace dimensions and relative orientations, noise variance, and the number of data points on the performance of SSC.

When the data is high-dimensional or the number of data points is large, solving the N Lasso problems (each in $N - 1$

The authors are with the Department of Information Technology and Electrical Engineering, ETH Zurich, Switzerland (e-mail: michaelt@nari.ee.ethz.ch; boelcskei@nari.ee.ethz.ch).

Copyright (c) 2018 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

¹We note that the SSC formulation in [8] adds a term to the Lasso objective function to account for sparse corruptions of the data points. The performance guarantees in [10], [11] apply, however, to the “pure” Lasso version of SSC. Throughout this paper, unless explicitly stated otherwise, SSC will refer to the “pure” Lasso version.

variables) in SSC can be computationally challenging. Greedy algorithms for computing sparse representations of the data points (in terms of all the other data points) are therefore an interesting alternative. Three such alternatives were proposed in the literature, namely the SSC-orthogonal matching pursuit (SSC-OMP) algorithm by Dyer et al. [13], the thresholding-based subspace clustering (TSC) algorithm by Heckel and Bölcskei [12], and the nearest subspace neighbor (NSN) algorithm by Park et al. [14]. SSC-OMP employs OMP instead of the Lasso to compute sparse representations of the data points. TSC relies on the nearest neighbors—in spherical distance—of each data point to construct the affinity graph, and NSN greedily assigns to each data point a subset of the other data points by iteratively selecting the data point closest (in Euclidean distance) to the subspace spanned by the previously selected data points.

To the best of our knowledge, besides SSC, TSC is the only subspace clustering algorithm that was proven to succeed under noise. The performance guarantees available for SSC-OMP [13], [15], [16] all apply to the noiseless case.

Contributions: The main contributions of this paper are an analytical performance characterization of SSC-OMP in the noisy case, and of a new algorithm, termed SSC-matching pursuit (SSC-MP), which is obtained by replacing OMP in SSC-OMP by the MP algorithm [17], [18]. Matching pursuit algorithms per se have been studied extensively in the sparse signal representation literature [19] and the approximation theory literature [20]. Replacing OMP by MP is attractive as the per-iteration complexity of MP is smaller than that of OMP thanks to the absence of the orthogonalization step. On the other hand, the representation error (in ℓ_2 -norm) of MP may decay slower—as a function of the number of iterations—than that of OMP [20]. We shall see, however, that in the context of subspace clustering, in practice, the lower per-iteration cost of MP usually translates into lower overall running time, while delivering essentially the same clustering performance as OMP.

Our main results are sufficient conditions for SSC-OMP and SSC-MP to succeed in terms of the no false connections property (see Definition 1), a widely used [9], [10], [13], [12], [11], [16], [13], [15], [14] subspace clustering performance measure. Specifically, we find that both algorithms succeed even when the subspaces intersect and when the signal to noise ratio is as low as 0dB. Furthermore, the sufficient conditions we obtain point at an intuitively appealing tradeoff between the affinity of the subspaces (a similarity measure for pairs of subspaces defined later), the noise variance, and the number of points in the data set corresponding to each subspace. This “clustering condition” is structurally similar to those for SSC in [10, Thm. 3.1], [11, Thm. 10] and for TSC in [12, Thm. 3]. Moreover, numerical results indicate that our clustering condition is order-wise optimal. The main technical challenge in proving our results stems from the need to handle statistical dependencies between quantities computed in different iterations of the OMP and MP algorithms.

OMP and MP are commonly stopped either after a prescribed maximum number of iterations, which we henceforth call data-independent (DI)-stopping, or when the representa-

tion error falls below a threshold value, referred to as data-dependent (DD)-stopping. For a given data point to be represented, OMP is guaranteed to select a new data point in every iteration and the sparsity level of the resulting representation therefore equals the number of OMP iterations performed. MP, on the other hand, may select individual data points to participate repeatedly in the sparse representation of a given data point. The sparsity level of the representation computed by MP may therefore be smaller than the number of iterations performed. As it is important for subspace clustering purposes to be able to control the sparsity level, we propose a new hybrid stopping criterion for MP terminating the algorithm either when a given maximum number of iterations was performed or when a given target sparsity level is attained. We consider SSC-OMP and SSC-MP both with DI- and DD-stopping. For DI-stopping, we present numerical results which indicate that performing (order-wise) more than d_ℓ OMP iterations can severely compromise the performance of SSC-OMP. SSC-OMP with DI-stopping therefore requires fairly accurate knowledge of the subspace dimensions. SSC-MP, on the other hand, exhibits a much more robust behavior in this regard. For DD-stopping, we prove that taking the threshold value on the representation error to be linear in the noise standard deviation ensures that both OMP and MP select order-wise at least d_ℓ points from $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$ to represent $\mathbf{y}_j \in \mathcal{Y}_\ell$, provided that the noise variance is sufficiently small. Numerical results further indicate that both algorithms, indeed, select order-wise no more than d_ℓ points from $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$ and essentially no points from $\mathcal{Y} \setminus \mathcal{Y}_\ell$. This means that SSC-OMP and SSC-MP with DD-stopping implicitly estimate (again order-wise) the subspace dimensions d_ℓ . This can—in principle—also be accomplished by SSC with a selection procedure for the Lasso parameter that is based on solving an auxiliary (constrained Lasso) optimization problem for each data point [10]. This procedure imposes, however, significant computational burden; in contrast DD-stopping as performed here comes at essentially zero computational cost.

Finally, we present extensive numerical results comparing the performance of SSC-OMP, SSC-MP, SSC, TSC, and NSN for synthetic and real data. In particular, we find that SSC-MP outperforms SSC in the reference problem of face clustering on the Extended Yale B data set [21], [22] and does so at drastically lower running time.

Notation: We use lowercase boldface letters to denote (column) vectors and uppercase boldface letters to designate matrices. The superscript T stands for transposition. For the vector \mathbf{v} , $[\mathbf{v}]_i$ denotes its i th element, $\|\mathbf{v}\|_0$ is the number of non-zero entries, and $\|\mathbf{v}\|_\infty := \max_i |[\mathbf{v}]_i|$. For the matrix \mathbf{A} , \mathbf{A}_{-i} stands for the matrix obtained by removing the i th column from \mathbf{A} , $\mathbf{A}_\mathcal{T}$ is the submatrix of \mathbf{A} consisting of the columns with index in the set \mathcal{T} , $\mathcal{R}(\mathbf{A})$ is its range space, $\|\mathbf{A}\|_{2 \rightarrow 2} := \max_{\|\mathbf{v}\|_2=1} \|\mathbf{A}\mathbf{v}\|_2$ its spectral norm, $\|\mathbf{A}\|_F := (\sum_{i,j} |\mathbf{A}_{ij}|^2)^{1/2}$ its Frobenius norm, and $\sigma_{\min}(\mathbf{A})$ and $\sigma_{\max}(\mathbf{A})$ refer to its minimum and maximum singular value, respectively. For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, $m \geq n$, of full column rank, we denote its pseudoinverse by $\mathbf{A}^\dagger := (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T$. The identity matrix is \mathbf{I} . $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ stands for

the distribution of a Gaussian random vector with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. The expectation of the random variable X is written as $\mathbb{E}[X]$. For random variables X and Y , we indicate their equivalence in distribution by $X \sim Y$. The set $\{1, \dots, N\}$ is denoted by $[N]$. The cardinality of the set \mathcal{T} is $|\mathcal{T}|$ and its complement is $\bar{\mathcal{T}}$. The unit sphere in \mathbb{R}^m is $\mathbb{S}^{m-1} := \{\mathbf{x} \in \mathbb{R}^m : \|\mathbf{x}\|_2 = 1\}$. $\log(\cdot)$ refers to the natural logarithm.

We say that a subgraph H of a graph G is connected if every pair of nodes in H can be joined by a path with nodes exclusively in H . A connected subgraph H of G is called a connected component of G if there are no edges between H and the remaining nodes in G .

II. SUBSPACE CLUSTERING VIA MATCHING PURSUITS

OMP and MP per se were introduced in [23] and [17], respectively, and have been studied extensively in the sparse signal representation literature, see, e.g., [19], [24, Chap. 3]. In the context of subspace clustering the premise is that the (sparse) representations of each data point in terms of all the other data points delivered by OMP and MP contain predominantly data points that lie in the same subspace as the data point under consideration. We refer to [8], [12, Sec. 2.B] for a detailed discussion on the relation between sparse signal representation theory and subspace clustering.

A. The algorithms

We first briefly review the SSC-OMP algorithm, introduced in [13], and then present the novel SSC-MP algorithm. The ensuing formulations of SSC-OMP and SSC-MP assume that the data points are of comparable ℓ_2 -norm. This assumption is relevant in Step 1 in both algorithms, but is not restrictive as the data points can always be normalized prior to processing. Further, an estimate \hat{L} of the number of subspaces L is assumed to be available. The estimation of L from the data set under consideration is discussed below.

The SSC-OMP algorithm [13]. Given a set of N data points \mathcal{Y} in \mathbb{R}^m , an estimate of the number of subspaces \hat{L} , and

- a maximum number of iterations $s_{\max} \leq \min\{m, N-1\}$ for DI-stopping,
- a threshold τ on the representation error for DD-stopping,

perform the following steps:

Step 1: For every $\mathbf{y}_j \in \mathcal{Y}$, find a representation of \mathbf{y}_j in terms of $\mathcal{Y} \setminus \{\mathbf{y}_j\}$ using OMP as follows: Initialize the iteration counter $s = 0$, the residual $\mathbf{r}_0 = \mathbf{y}_j$, and the set of selected indices $\Lambda_0 = \emptyset$. Denote the data matrix containing the points in \mathcal{Y} by $\mathbf{Y} \in \mathbb{R}^{m \times N}$. For $s = 1, 2, \dots$, perform the updates

$$\lambda_s = \arg \max_{i \in [N] \setminus (\Lambda_{s-1} \cup \{j\})} |\langle \mathbf{y}_i, \mathbf{r}_{s-1} \rangle| \quad (1)$$

$$\Lambda_s = \Lambda_{s-1} \cup \lambda_s$$

$$\begin{aligned} \mathbf{r}_s &= \left(\mathbf{I} - \mathbf{Y}_{\Lambda_s} (\mathbf{Y}_{\Lambda_s})^\dagger \right) \mathbf{y}_j \\ &= \left(\mathbf{I} - \frac{\tilde{\mathbf{y}}_{\lambda_s} (\tilde{\mathbf{y}}_{\lambda_s})^T}{\|\tilde{\mathbf{y}}_{\lambda_s}\|_2^2} \right) \mathbf{r}_{s-1}, \end{aligned} \quad (2)$$

where $\tilde{\mathbf{y}}_{\lambda_s} = (\mathbf{I} - \mathbf{Y}_{\Lambda_{s-1}} (\mathbf{Y}_{\Lambda_{s-1}})^\dagger) \mathbf{y}_{\lambda_s}$, until at least one of the following criteria is met

- for DI-stopping:
 $s = s_{\max}, \max_{i \in [N] \setminus (\Lambda_s \cup \{j\})} |\langle \mathbf{y}_i, \mathbf{r}_s \rangle| = 0.2$
- for DD-stopping:
 $\|\mathbf{r}_s\|_2 \leq \tau, \max_{i \in [N] \setminus (\Lambda_s \cup \{j\})} |\langle \mathbf{y}_i, \mathbf{r}_s \rangle| = 0.$

Ties in the maximization (1) are broken arbitrarily.

Step 2: With the number of OMP iterations actually performed denoted by s_a , compute the representation coefficient vectors $\mathbf{b}_j \in \mathbb{R}^N$, $j \in [N]$, according to $(\mathbf{b}_j)_{\Lambda_{s_a}} = (\mathbf{Y}_{\Lambda_{s_a}})^\dagger \mathbf{y}_j$, $(\mathbf{b}_j)_{\bar{\Lambda}_{s_a}} = \mathbf{0}$, and construct the adjacency matrix $\mathbf{A} = \mathbf{B} + \mathbf{B}^T$, where $\mathbf{B} = \text{abs}([\mathbf{b}_1 \dots \mathbf{b}_N])$ with $\text{abs}(\cdot)$ denoting absolute values taken element-wise.

Step 3: Apply normalized spectral clustering [25], [7] to (\mathbf{A}, \hat{L}) .

The SSC-MP algorithm. Given a set of N data points \mathcal{Y} in \mathbb{R}^m , an estimate of the number of subspaces \hat{L} , and

- a maximum number of iterations s_{\max} and a target sparsity level p_{\max} for DI-stopping,
- a threshold τ on the representation error for DD-stopping,

perform the following steps:

Step 1: For every $\mathbf{y}_j \in \mathcal{Y}$, find a representation of \mathbf{y}_j in terms of $\mathcal{Y} \setminus \{\mathbf{y}_j\}$ using MP as follows: Initialize the iteration counter $s = 0$, the residual $\mathbf{q}_0 = \mathbf{y}_j$, and the coefficient vector $\mathbf{b}_j \in \mathbb{R}^N$ as $\mathbf{b}_j = \mathbf{0}$. For $s = 1, 2, \dots$, perform the updates

$$\omega_s = \arg \max_{i \in [N] \setminus \{j\}} |\langle \mathbf{y}_i, \mathbf{q}_{s-1} \rangle| \quad (3)$$

$$[\mathbf{b}_j]_{\omega_s} \leftarrow [\mathbf{b}_j]_{\omega_s} + \frac{\langle \mathbf{y}_{\omega_s}, \mathbf{q}_{s-1} \rangle}{\|\mathbf{y}_{\omega_s}\|_2^2} \quad (4)$$

$$\mathbf{q}_s = \left(\mathbf{I} - \frac{\mathbf{y}_{\omega_s} (\mathbf{y}_{\omega_s})^T}{\|\mathbf{y}_{\omega_s}\|_2^2} \right) \mathbf{q}_{s-1} \quad (5)$$

until at least one of the following criteria is met

- for DI-stopping:
 $s = s_{\max}, \|\mathbf{b}_j\|_0 = p_{\max}, \max_{i \in [N] \setminus \{j\}} |\langle \mathbf{y}_i, \mathbf{q}_s \rangle| = 0.$
- for DD-stopping:
 $\|\mathbf{q}_s\|_2 \leq \tau, \max_{i \in [N] \setminus \{j\}} |\langle \mathbf{y}_i, \mathbf{q}_s \rangle| = 0.$

Ties in the maximization (3) are broken arbitrarily.

Step 2: Construct the adjacency matrix $\mathbf{A} = \mathbf{B} + \mathbf{B}^T$, where $\mathbf{B} = \text{abs}([\mathbf{b}_1 \dots \mathbf{b}_N])$.

Step 3: Apply normalized spectral clustering [25], [7] to (\mathbf{A}, \hat{L}) .

Stopping criteria: We emphasize that OMP, thanks to the orthogonalization (2) of the residual \mathbf{r}_{s-1} w.r.t. all data points selected previously, is guaranteed to select a new data point in every iteration and hence the sparsity level of \mathbf{b}_j equals the number of OMP iterations performed. In contrast, MP orthogonalizes (see (5)) the residual \mathbf{q}_{s-1} w.r.t. the data point \mathbf{y}_{ω_s} selected in the current iteration s only and may therefore select the same data point to participate repeatedly in the representation of \mathbf{y}_j . The sparsity level of \mathbf{b}_j may hence be smaller than the number of MP iterations performed,

²Throughout the paper, we use the convention of maximization over the empty set evaluating to 0.

which is why the DI-stopping criterion for MP incorporates termination when a given target sparsity level, namely p_{\max} , is attained. Choosing s_{\max} large enough, stopping will, indeed, be activated by $\|\mathbf{b}_j\|_0 = p_{\max}$. Having control over the sparsity level of the coefficient vectors \mathbf{b}_j can be important to achieve good clustering performance as discussed below. Setting $p_{\max} = N$, on the other hand, guarantees that stopping is activated through $s = s_{\max}$ and thereby allows to control the maximum number of MP iterations through choice of s_{\max} . This hybrid stopping criterion does not seem to have been considered before in the literature.

For DD-stopping, OMP is guaranteed to stop as soon as a basis for the subspace \mathcal{Y}_j lies in has been found or, in case \mathbf{y}_j does not lie in the span of $\mathcal{Y} \setminus \{\mathbf{y}_j\}$, the best representation—in the least-squares sense—of \mathbf{y}_j in terms of the points in $\mathcal{Y} \setminus \{\mathbf{y}_j\}$. On the other hand, MP is, in general, not guaranteed to terminate after a finite number of iterations as it may fail to activate either of the conditions $\|\mathbf{q}_s\|_2 \leq \tau$ and $\max_{i \in [N] \setminus \{j\}} |\langle \mathbf{y}_i, \mathbf{q}_s \rangle| = 0$ if τ is chosen too small [18]. For most data sets encountered in practice this is not an issue. It can, however, become a problem when the data set contains outliers that cannot be represented sparsely by the other data points. In such cases it is advisable to employ the DI-stopping criterion which guarantees that at most s_{\max} iterations are performed.

Implementation aspects: As MP requires the computation of inner products only, whereas OMP contains a (least-squares) orthogonalization step (typically carried out by QR decomposition or Cholesky factorization [19]), the per-iteration computational cost of SSC-MP is lower than that of SSC-OMP. The numerical results in Section IV-B indicate that this typically also translates into a lower overall running time for SSC-MP at fixed performance.

Weak selection rules: For very large data sets one can often speed up OMP and MP by relaxing the selection rules (1) and (3) to so-called weak selection rules [19] as follows. Instead of (1), one determines λ_s such that $|\langle \mathbf{y}_{\lambda_s}, \mathbf{r}_{s-1} \rangle| \geq \alpha \max_{i \in [N] \setminus (\Lambda_{s-1} \cup \{j\})} |\langle \mathbf{y}_i, \mathbf{r}_{s-1} \rangle|$ and instead of (3), one finds ω_s according to $|\langle \mathbf{y}_{\omega_s}, \mathbf{q}_{s-1} \rangle| \geq \alpha \max_{i \in [N] \setminus \{j\}} |\langle \mathbf{y}_i, \mathbf{q}_{s-1} \rangle|$, in both cases for a fixed relaxation parameter $\alpha \in (0, 1]$. These weakened selection rules can be implemented efficiently using, e.g., locality-sensitive hashing [26], [27]. For conciseness, we shall not analyze weak selection rules here, but only note that our main results presented in Section III extend to weak selection rules with minor modifications.

B. Parameter selection

In both algorithms, \hat{L} may be estimated in Step 2 based on the adjacency matrix \mathbf{A} using the *eigengap heuristic* [7] (note that L is needed only in Step 3), which relies on the fact that the number of zero eigenvalues of the normalized Laplacian of the graph G with adjacency matrix \mathbf{A} corresponds to the number of connected components of G .

The spectral clustering step (Step 3 in both algorithms) recovers the oracle segmentation $\{\mathcal{Y}_1, \dots, \mathcal{Y}_L\}$ of \mathcal{Y} if $\hat{L} = L$ and if each connected component in G corresponds to exactly

one of the \mathcal{Y}_ℓ [7, Prop. 4; Sec. 7]. Choosing the parameters s_{\max} and p_{\max} in the case of DI-stopping, and τ for DD-stopping, appropriately is therefore crucial. Indeed, as s_{\max} , p_{\max} , and τ determine the sparsity level of the representation coefficient vectors \mathbf{b}_j , they control the number of edges in G and hence the connectivity properties of G . Specifically, taking s_{\max} , p_{\max} too small or τ too large results in weak connectivity and may hence cause the subgraphs of G corresponding to individual \mathcal{Y}_ℓ to split up into multiple connected components. Spectral clustering would then assign these components to different clusters, which means that a given set \mathcal{Y}_ℓ is divided up into multiple (disjoint) sets. On the other hand, choosing s_{\max} , p_{\max} too big or τ too small will result in strong connectivity and hence potentially in “false connections”, i.e., in edges in G between points that correspond to different \mathcal{Y}_ℓ . Spectral clustering exhibits, however, a certain amount of robustness vis-à-vis false connections with small corresponding weights in G . From what was just said it follows that ideally s_{\max} , p_{\max} , and τ should be chosen such that the sparsity levels of the \mathbf{b}_j are on the order of the subspace dimensions. This can be seen as follows. Assume that the points in \mathcal{Y}_ℓ are well spread out on the subspace \mathcal{S}_ℓ , $\ell \in [L]$, and perturbed by additive isotropic Gaussian noise. If the noise variance is not too large the noisy data points $\mathbf{y}_j \in \mathcal{Y}_\ell$ will remain close to \mathcal{S}_ℓ . In this case, roughly d_ℓ points from $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$ will suffice to represent $\mathbf{y}_j \in \mathcal{Y}_\ell$ with small representation error. Hence, if the subspaces \mathcal{S}_ℓ are sufficiently far apart, imposing a sparsity level of $\approx d_\ell$ for $\mathbf{y}_j \in \mathcal{Y}_\ell$ through suitable choice of s_{\max} , p_{\max} , τ will force OMP and MP to select points predominantly from $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$. If, however, the subspaces are too close to each other, OMP and MP are likely to select points from $\mathcal{Y} \setminus \mathcal{Y}_\ell$ (i.e., false connections) as well.

We next discuss the selection of the parameters s_{\max} , p_{\max} for DI- and τ for DD-stopping in the light of what was just said. For simplicity of exposition, in the noisy case we assume that the points in \mathcal{Y} are in general position, i.e., every subset of m or fewer points in \mathcal{Y} is linearly independent. Real-world data sets usually have this property and the statistical data model our analysis is based on (see Section III) conforms with this assumption as well. Note that if the points in \mathcal{Y} are in general position, OMP under DI-stopping is guaranteed to perform exactly s_{\max} iterations, whereas MP will perform at least $\min\{s_{\max}, p_{\max}, m, N-1\}$ iterations under DI-stopping.

DI-stopping: We first consider SSC-OMP. In the noiseless case, if the subspaces are sufficiently far apart, OMP under DI-stopping automatically stops after at most d_ℓ iterations for all $\mathbf{y}_j \in \mathcal{Y}_\ell$, $\ell \in [L]$ [13], [15]. Therefore, choosing $s_{\max} \geq \max_{\ell \in [L]} d_\ell$ guarantees that OMP automatically detects the dimensions of the subspaces the points \mathbf{y}_j reside in. In contrast, in the noisy case, performing more than $\approx d_\ell$ iterations “forces” OMP to select points from $\mathcal{Y} \setminus \mathcal{Y}_\ell$ (corresponding to false connections) for the representation of $\mathbf{y}_j \in \mathcal{Y}_\ell$ as after $\approx d_\ell$ iterations \mathbf{r}_s will roughly be orthogonal to \mathcal{S}_ℓ (see Figure 1). The choice of s_{\max} , in practice, therefore requires knowledge of the subspace dimensions. In certain applications such as, e.g., face clustering [28], information on the subspace dimensions may, indeed, be available a priori. When the subspace dimensions d_ℓ vary widely (across ℓ),

there may be no s_{\max} that ensures both connectivity of all the subgraphs of G corresponding to the \mathcal{Y}_ℓ , and at the same time guarantees a small number of false connections. In principle this problem could be mitigated by setting s_{\max} for each \mathbf{y}_j individually according to the dimension of the subspace it belongs to. This would, however, require knowledge of the assignments of the data points \mathbf{y}_j to the subspaces \mathcal{S}_ℓ , thereby performing the actual subspace clustering task.

We now turn to SSC-MP. MP with $p_{\max} = N$, i.e., the hybrid stopping criterion is activated once s_{\max} iterations have been performed, exhibits a stopping behavior that is fundamentally different from that of OMP. As already mentioned this is a consequence of MP orthogonalizing the residual only w.r.t. the data point selected in the current iteration, thereby allowing for repeated selection of individual data points. In the noiseless case, unlike SSC-OMP, SSC-MP can select more than d_ℓ data points from $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$ to represent $\mathbf{y}_j \in \mathcal{Y}_\ell$, thereby potentially producing a graph G with better connectivity than SSC-OMP. On the other hand, SSC-MP tends to assign smaller weights to the points in $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$ than SSC-OMP, which can lead to slightly inferior performance (see the experiments reported in Appendix A). For the noisy case, the experiments on synthetic and real-world data, reported in Section IV-D, show that in the first d_ℓ or so iterations MP adds predominantly new points from the correct cluster $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$ and thereafter tends to revisit previously selected data points or add new points stemming mostly from $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$. For fixed s_{\max} , this leads to a smaller number of false connections than what would be obtained by OMP. The example in Figure 1 illustrates this behavior. The data set \mathcal{Y} is drawn from the union of a two-dimensional subspace \mathcal{S}_1 and a one-dimensional subspace \mathcal{S}_2 , where \mathcal{S}_1 and \mathcal{S}_2 span a (principal) angle of approximately 50 degrees. We consider the data point $\mathbf{y}_j \in \mathcal{Y}_1$. In the first iteration, both OMP and MP select the same data point $\mathbf{y}_{\lambda_1} = \mathbf{y}_{\omega_1}$ and we have $\mathbf{r}_1 = \mathbf{q}_1$. In the second iteration, the two algorithms again select the same data point, namely $\mathbf{y}_{\lambda_2} = \mathbf{y}_{\omega_2}$, but \mathbf{r}_2 is now almost orthogonal to \mathcal{S}_1 whereas \mathbf{q}_2 remains close to \mathcal{S}_1 as MP orthogonalizes only w.r.t. \mathbf{y}_{ω_2} . Specifically, the (principal) angle between \mathbf{q}_2 and \mathcal{S}_1 is approximately 27 degrees, while \mathbf{q}_2 and \mathcal{S}_2 span an angle of approximately 50 degrees. In the third iteration MP therefore selects a point from the set $\mathcal{Y}_1 \setminus \{\mathbf{y}_j\}$, whereas OMP chooses a point from the wrong subspace \mathcal{Y}_2 .

The experiments reported in Section IV-D reveal the following remarkable property. Even when we discount the advantage of MP—owing to its ability to reselect data points—by forcing the sparsity levels of MP and OMP to be equal to, say s_{high} , through appropriate choice of s_{\max} and p_{\max} , with $s_{\text{high}} \gg d_\ell$ for at least one $\ell \in [L]$, MP still tends to select more points from \mathcal{Y}_ℓ to represent $\mathbf{y}_j \in \mathcal{Y}_\ell$ than OMP does. Moreover, these numerical results indicate that MP also tends to assign smaller (in absolute value) coefficients to false connections, i.e., to points in $\mathcal{Y} \setminus \mathcal{Y}_\ell$; this can have a favorable effect on performance thanks to the robustness of spectral clustering to false connections with small associated weights.

Our analytical results in Section III guarantee that SSC-OMP and SSC-MP succeed for s_{\max} and p_{\max} linear—up to log-terms—in the smallest subspace dimension, provided that




Fig. 1: Evolution of the (normalized) residuals \mathbf{r}_s and \mathbf{q}_s corresponding to OMP and MP, respectively, for the data point $\mathbf{y}_j \in \mathcal{Y}_1$. The data points belonging to \mathcal{S}_2 are marked by circles. Note that all data points were normalized to unit ℓ_2 -norm prior to clustering.

the subspaces are sufficiently far apart, the noise variance is sufficiently small, and the data set contains sufficiently many points from each subspace.

DD-stopping: We assume throughout that τ is sufficiently large for SSC-MP to terminate. For OMP in the context of sparse noisy signal recovery, taking τ linear in the noise standard deviation σ is known to lead to correct recovery of the sparse signal under certain technical conditions [29]. In the context of subspace clustering, where the problem is actually of a different nature, the analytical results in Section III indicate that such a choice for τ guarantees that both SSC-OMP and SSC-MP select order-wise at least d_ℓ points from \mathcal{Y}_ℓ to represent $\mathbf{y}_j \in \mathcal{Y}_\ell$, provided that the subspaces \mathcal{S}_ℓ are sufficiently far apart and the points in \mathcal{Y}_ℓ are well spread out on \mathcal{S}_ℓ , $\ell \in [L]$, and perturbed by additive isotropic Gaussian noise of sufficiently small variance. The numerical results in Section IV-C show that the graphs G generated by SSC-OMP and SSC-MP will also have a small number of false connections if τ , in addition, is not too small. Appropriate choice of τ therefore makes both OMP and MP automatically adjust the sparsity level for each data point according to the dimension of the subspace the data point lies in. We say that the algorithms detect the dimensions of the subspaces the data points reside in. Recall that under DI-stopping this has to be accomplished through suitable choice of s_{\max} , p_{\max} .

When the data set \mathcal{Y} contains outliers that cannot be represented sparsely in terms of the other points in \mathcal{Y} , DD-stopping usually leads to a high number of false connections as the representation error for outliers will decay slowly resulting in late activation of the DD-stopping criterion. In this case, we need to rely on DI-stopping.

Summary: We recommend DI-stopping with s_{\max} on the order of the subspace dimensions if the subspace dimensions are (approximately) known, and DD-stopping with τ^2 on the order of the noise variance if the noise variance is known and not too large. If no prior knowledge on the subspace dimensions or noise variance is available, or if the noise

variance is large, we recommend relying on DI-stopping with s_{\max} in the range $\{5, \dots, 10\}$. These values for s_{\max} turned out to work well in the relevant experiments conducted in this paper, and were also used in related works [13], [15].

III. MAIN RESULTS

Our analytical performance results are for a statistical data model, also employed in [10], [12]. Specifically, we take the subspaces \mathcal{S}_ℓ to be fixed and the points in the corresponding subsets \mathcal{Y}_ℓ of the data set $\mathcal{Y} = \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_L$ to be randomly distributed on $\mathcal{S}_\ell \cap \mathbb{S}^{m-1}$ and perturbed by additive random noise. Concretely, the points in \mathcal{Y}_ℓ , $\ell \in [L]$, are given by $\mathbf{y}_i^{(\ell)} = \mathbf{x}_i^{(\ell)} + \mathbf{z}_i^{(\ell)} = \mathbf{U}^{(\ell)} \mathbf{a}_i^{(\ell)} + \mathbf{z}_i^{(\ell)}$, $i \in [n_\ell]$, where the columns of $\mathbf{U}^{(\ell)} \in \mathbb{R}^{m \times d_\ell}$ constitute an orthonormal basis for \mathcal{S}_ℓ , the $\mathbf{a}_i^{(\ell)}$ are independently (across $i \in [n_\ell]$, $\ell \in [L]$) and uniformly distributed on $\mathbb{S}^{d_\ell-1}$, and the $\mathbf{z}_i^{(\ell)}$ are i.i.d. $\mathcal{N}(\mathbf{0}, (\sigma^2/m)\mathbf{I}_m)$. The factor $1/m$ in the noise covariance matrix ensures that $\|\mathbf{z}_i^{(\ell)}\|_2^2$ concentrates around $\mathbb{E}[\|\mathbf{z}_i^{(\ell)}\|_2^2] = \sigma^2$ for large m . Note further that the ℓ_2 -norm of the data points concentrates around $\sqrt{1 + \sigma^2}$ for large m and that they are hence of comparable ℓ_2 -norm, as required by the formulations of SSC-OMP and SSC-MP in Section II. Moreover, the data points are in general position w.p. 1 for $\sigma > 0$.

Prima facie assuming the noiseless data points $\mathbf{x}_i^{(\ell)}$ to be uniformly distributed on the subspaces \mathcal{S}_ℓ may appear overly stylized. However, for any algorithm to have a chance of producing correct assignments, we need the noiseless data points to be well spread out to a certain extent (albeit not necessarily around the origin as in our data model) on the subspaces. To see this, suppose for example, that the points in \mathcal{Y}_ℓ are concentrated on two distinct subspaces of \mathcal{S}_ℓ , say \mathcal{S}'_ℓ and \mathcal{S}''_ℓ . Then, one can assign the points in \mathcal{Y}_ℓ either to two clusters, one containing the points concentrated on \mathcal{S}'_ℓ and the other one those concentrated on \mathcal{S}''_ℓ , or one can assign all the points in \mathcal{Y}_ℓ to a single cluster.

Our results will depend on the affinity between pairs of subspaces which measures how far apart two subspaces are. The affinity between the subspaces \mathcal{S}_k and \mathcal{S}_ℓ is defined as [9, Def. 2.6], [10, Def. 1.2]

$$\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) := \frac{1}{\sqrt{\min\{d_k, d_\ell\}}} \left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F \quad (6)$$

and can equivalently be expressed in terms of the principal angles $\theta_1 \leq \dots \leq \theta_{\min\{d_k, d_\ell\}}$ between \mathcal{S}_k and \mathcal{S}_ℓ [30, Sec. 6.3.4] according to

$$\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) = \sqrt{\frac{\cos^2(\theta_1) + \dots + \cos^2(\theta_{\min\{d_k, d_\ell\}})}{\min\{d_k, d_\ell\}}}. \quad (7)$$

We have $0 \leq \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \leq 1$ and for subspaces intersecting in t dimensions, we get $\cos(\theta_1) = \dots = \cos(\theta_t) = 1$ and hence $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \geq \sqrt{t/\min\{d_k, d_\ell\}}$.

Recall that spectral clustering recovers the oracle segmentation $\{\mathcal{Y}_1, \dots, \mathcal{Y}_L\}$ if $\hat{L} = L$ and each connected component in G corresponds to one of the \mathcal{Y}_ℓ . Establishing conditions that guarantee zero clustering error is inherently difficult. To the best of our knowledge the only instances of such a result for spectral clustering-based subspace clustering algorithms are

[12, Thm. 2] for TSC in the noiseless case and a condition in [31] guaranteeing that a post-processing procedure for SSC yields correct clustering in the noisy case. We will rely on the following intermediate, albeit sensible, performance measure, which has become standard in the subspace clustering literature and was also employed in [9], [10], [13], [12], [11], [16], [13], [15], [14].

Definition 1 (No false connections (NFC) property). *The graph G satisfies the no false connections (NFC) property if, for all $\ell \in [L]$, the nodes corresponding to \mathcal{Y}_ℓ are connected to other nodes corresponding to \mathcal{Y}_ℓ only.*

In what follows, we often say ‘‘SSC-OMP/SSC-MP satisfies the NFC property’’ instead of ‘‘the graph G generated by SSC-OMP/SSC-MP satisfies the NFC property’’. To guarantee perfect clustering, we would need to ensure—in addition to the NFC property—that the subgraphs of G corresponding to the \mathcal{Y}_ℓ are connected. This would preclude split-ups of the subgraphs of G corresponding to the individual \mathcal{Y}_ℓ . Sufficient conditions guaranteeing this property for SSC were established in [32] for $m = 3$ in the noiseless case.

Note that the NFC property does not involve the parameter \hat{L} . The sufficient conditions for SSC-OMP and SSC-MP to satisfy the NFC property reported next, therefore, do not require $\hat{L} = L$.

Our main result for SSC-OMP with DI-stopping is the following.

Theorem 1 (SSC-OMP with DI-stopping). *Define the sampling density $\rho_\ell := (n_\ell - 1)/d_\ell$, and let $d_{\max} := \max_{\ell \in [L]} d_\ell$ and $\rho_{\min} = \min_{\ell \in [L]} \rho_\ell$. Assume that $m \geq 2d_{\max}$, $\rho_{\min} \geq c_\rho$, $\sigma \leq 1/2$, and $s_{\max} \leq \min_{\ell \in [L]} \{c_s d_\ell / \log((n_\ell - 1)e/s_{\max})\}$, where c_ρ and c_s are numerical constants satisfying $c_\rho > 1$, $0 < c_s \leq 1/10$. Then, the clustering condition*

$$\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \frac{10\sigma}{\sqrt{\log(N^3 s_{\max})}} \left(\frac{\sqrt{d_{\max}}}{\sqrt{m}} c(\sigma) + \frac{\sqrt{2}}{\sqrt{\rho_{\min}}} \left(1 + \frac{3}{2}\sigma \right) \right) \leq \frac{1}{8 \log(N^3 s_{\max})} \quad (8)$$

with $c(\sigma) = 10 + 13\sigma$ guarantees that the graph G generated by SSC-OMP under DI-stopping satisfies the NFC property w.p. at least

$$P^* := 1 - 6/N - 5Ne^{-c_m m} - 6 \sum_{\ell \in [L]} n_\ell e^{-c_d d_\ell} \quad (9)$$

for numerical constants c_d and c_m obeying $0 < c_d \leq 1/18$, $0 < c_m \leq 1/8$.

The main result for SSC-MP with DI-stopping is as follows.

Theorem 2 (SSC-MP with DI-stopping). *Define the sampling density $\rho_\ell := (n_\ell - 1)/d_\ell$, and let $d_{\max} := \max_{\ell \in [L]} d_\ell$ and $\rho_{\min} = \min_{\ell \in [L]} \rho_\ell$. Assume that $m \geq 2d_{\max}$, $\rho_{\min} \geq c_\rho$, $\sigma \leq 1/2$, $s_{\max} > 0$, and $p_{\max} \leq \min_{\ell \in [L]} \{c_s d_\ell / \log((n_\ell - 1)e/p_{\max})\}$, where c_ρ and c_s are numerical constants satisfying $c_\rho > 1$, $0 < c_s \leq 1/10$. Then, the clustering condition (8) with $c(\sigma) = 22 + 29\sigma$ guarantees that the graph G generated*

by SSC-MP under DI-stopping satisfies the NFC property w.p. at least P^* as defined in (9).

The proofs of Theorems 1 and 2 can be found in Appendices B and C, respectively.

Theorems 1 and 2 essentially state that SSC-OMP and SSC-MP satisfy the NFC property for s_{\max} and p_{\max} linear—up to log-terms—in $d_{\min} := \min_{\ell \in [L]} d_\ell$, provided that the subspaces are not too close (in terms of their pairwise affinities), the noise variance σ^2 is sufficiently small, and the data set \mathcal{Y} contains sufficiently many points from each subspace \mathcal{S}_ℓ . Specifically, the clustering condition (8) tells us that the subspaces \mathcal{S}_ℓ are allowed to be quite close to each other and can even intersect in a substantial fraction of their dimensions, all provided that σ^2 is not too large. Moreover, inspection of the second term on the left-hand side (LHS) of (8) shows that a higher noise variance σ^2 is tolerated when m becomes large relative to the largest subspace dimension d_{\max} and/or the data set \mathcal{Y} contains an increasing number of points in each of the subspaces, resulting in an increase in the minimum sampling density ρ_{\min} . The clustering condition (8) can hence be satisfied under the condition $\sigma \leq 1/2$ imposed by Theorems 1 and 2 if m is sufficiently large relative to d_{\max} and if ρ_{\min} is sufficiently large (but not too large, in order to prevent the right-hand side (RHS) of (8) from becoming too small; for example, ρ_{\min} should not scale exponentially in one of the d_ℓ). This shows that SSC-OMP and SSC-MP, indeed, satisfy the NFC property even when the noise variance σ^2 is on the order of the signal energy, i.e., when the signal to noise ratio $\text{SNR} := \mathbb{E}[\|\mathbf{x}_j\|_2^2] / \mathbb{E}[\|\mathbf{z}_j\|_2^2] = 1/\sigma^2$ satisfies $\text{SNR} \approx 0\text{dB}$ (recall that $\mathbf{y}_i^{(\ell)} = \mathbf{x}_i^{(\ell)} + \mathbf{z}_i^{(\ell)}$ with $\mathbb{E}[\|\mathbf{x}_i^{(\ell)}\|_2^2] = 1$).

The RHS of (8) going to zero as $N \rightarrow \infty$ may appear counter-intuitive as one would expect clustering to become easier when the number of data points increases. Note, however, that (8) allows the subspaces to intersect, and Theorems 1 and 2 guarantee the NFC property for *all data points*. Now, when N increases, owing to the statistical data model our analysis is based on, the number of data points that are close to the intersection of two subspaces also increases, which in turn leads to an increase in the probability of the NFC property being violated for at least one data point. This then results in the clustering condition becoming more restrictive. The clustering conditions for SSC in [10, Eq. (3.1)], [11, Thm. 10] and for TSC in [12, Eq. (8)] exhibit the same $O(1/\log(N))$ scaling and hence the same seemingly counter-intuitive behavior.

We hasten to add that the condition $\sigma \leq 1/2$ in Theorems 1 and 2 was imposed only to get clustering conditions that are of simple form. Removing the restriction $\sigma \leq 1/2$ (which is used to get the bounds (27) and (28) in Appendix B) would lead to clustering conditions allowing, in principle, for arbitrarily large σ (i.e., even for $\text{SNR} < 0\text{dB}$), provided that the d_ℓ are sufficiently small compared to m , and ρ_{\min} is sufficiently large. One might further expect that the upper bounds on s_{\max} and p_{\max} in Theorems 1 and 2, respectively, should depend on σ because the number of iterations for which SSC-OMP and SSC-MP are guaranteed to select points from $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$ for $\mathbf{y}_j \in \mathcal{Y}_\ell$ should decrease as σ increases. However, this

is not the case as the clustering condition (8) limits the noise variance (more precisely, the variance of the noise components on the subspaces) depending on $\max_{k,\ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell)$, N , ρ_{\min} , and s_{\max} . We furthermore note that the conditions in Theorems 1 and 2 (with different constants in (8)) continue to guarantee the NFC property for bounded noise or sub-gaussian noise, in both cases of isotropic distribution. It is interesting to see that SSC-MP satisfies the NFC property under virtually the same conditions as SSC-OMP, although in practice SSC-MP typically exhibits a lower running time at fixed performance.

Comparing the clustering condition (8) to those for SSC in [10, Thm. 3.1] and for TSC in [12, Thm. 3], both of which guarantee the NFC property and apply to the same data model as used here, we find that (8) exhibits the same structure (up to log-factors and constants) apart from the term proportional to $\sqrt{1/\rho_{\min}}$ on the LHS of (8). This term dominates the term proportional to $\sqrt{d_{\max}/m}$ only if $\sqrt{d_{\max}/m} \ll \sqrt{1/\rho_{\min}} = \sqrt{\max_{\ell \in [L]} (d_\ell / (n_\ell - 1))}$, i.e., if $\max_{\ell \in [L]} n_\ell \ll m$ (owing to $\max_{\ell \in [L]} (d_\ell / (n_\ell - 1)) \geq (\max_{\ell \in [L]} d_\ell) / (\max_{\ell \in [L]} n_\ell - 1) > d_{\max} / (\max_{\ell \in [L]} n_\ell)$). Similarly, the clustering condition in [11, Thm. 10] does not have a term proportional to $\sqrt{1/\rho_{\min}}$ as (8) does, but imposes a slightly more restrictive condition on σ , requiring $\sigma(c + \sigma)$ to be at most on the order of $\sqrt{m - d}/d$ instead of \sqrt{m}/\sqrt{d} (assuming $d_\ell = d$ for all $\ell \in [L]$ and neglecting log-terms for simplicity of exposition), where c is a constant. Numerical results in Section IV-A indicate that the term proportional to $\sqrt{1/\rho_{\min}}$ in (8) is not an artifact of our proof techniques, but rather fundamental. We further note that setting $\sigma = 0$, the second term on the LHS of (8) vanishes and we recover (up to log-factors and constants) the clustering condition

$$\max_{k,\ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \leq \frac{\sqrt{\log(\rho_{\min})}}{64 \log(N)}$$

found in [16, Cor. 1] for SSC-OMP in the noiseless case.

In summary, SSC-OMP, SSC-MP, TSC, and SSC all satisfy the NFC property under similar (sufficient) conditions, while differing considerably w.r.t. computational complexity. Specifically, SSC-OMP and SSC-MP, albeit greedy, are computationally more expensive than TSC, but significantly less expensive than SSC. On the other hand, SSC-MP can outperform TSC quite significantly in certain applications (see Section IV-B). A detailed comparison of SSC, SSC-OMP, SSC-MP, and TSC in terms of performance and running times is provided in Section IV-B. The performance of all four algorithms varies across data sets, and none of the algorithms consistently outperforms the other ones.

Recall that under DI-stopping the choice of the parameters s_{\max} , p_{\max} is critical for the success of SSC-OMP and SSC-MP. Taking s_{\max} , p_{\max} too small or too large may lead to cluster split-ups or to many false connections, respectively. The maximum range for s_{\max} , p_{\max} for our results to guarantee the NFC property is determined (up to log-factors) by the smallest subspace dimension d_{\min} , which is usually unknown. Furthermore, if d_{\min} is small, the range of admissible values for s_{\max} , p_{\max} will also be small. The clustering condition (8) is, however, only sufficient (for the NFC property to hold) and good clustering performance may be obtained in

practice for larger values of s_{\max} , p_{\max} than those identified by Theorems 1 and 2.

We proceed to our main result on DD-stopping, which indicates that the problems with choosing s_{\max} , p_{\max} for DI-stopping due to unknown d_ℓ can be mitigated—to a certain extent—through DD-stopping. Specifically, we show that SSC-OMP and SSC-MP under DD-stopping automatically select at least on the order of d_ℓ points from \mathcal{Y}_ℓ to represent $\mathbf{y}_j \in \mathcal{Y}_\ell$. We hasten to add, however, that Theorem 3 does not guarantee that no additional data points corresponding to false connections are selected and Theorem 3 hence does not guarantee the NFC property.

Theorem 3 (SSC-OMP and SSC-MP with DD-stopping). *Define the sampling density $\rho_\ell := (n_\ell - 1)/d_\ell$, and let $d_{\max} := \max_{\ell \in [L]} d_\ell$ and $\rho_{\min} = \min_{\ell \in [L]} \rho_\ell$. Suppose that $m \geq 2d_{\max}$, $\rho_{\min} \geq c_\rho$, and $\sigma \leq 1/2$, where c_ρ is a numerical constant satisfying $c_\rho > 1$. Pick $\tau \in [0, 2/3 - (\sqrt{d_{\max}}/\sqrt{m})\sigma]$. Then, the clustering condition (8) with s_{\max} on both sides replaced by $\max_{\ell \in [L]} \lfloor c_s d_\ell / \log((n_\ell - 1)e) \rfloor$ guarantees w.p. at least P^* as defined in (9), for all $\mathbf{y}_j \in \mathcal{Y}_\ell$, $j \in [n_\ell]$, $\ell \in [L]$, that the corresponding coefficient vectors \mathbf{b}_j computed by OMP and MP (if it terminates) have at least*

$$\left\lfloor \frac{d_\ell}{\log((n_\ell - 1)e)} \min \left\{ \frac{1}{3} \left(\frac{2}{3} - \frac{\tau}{1 - \frac{3}{2} \frac{\sqrt{d_\ell}}{\sqrt{m}} \sigma} \right)^2, c_s \right\} \right\rfloor \quad (10)$$

non-zero entries corresponding to points in $\mathcal{Y}_\ell \setminus \{\mathbf{y}_j\}$.

Note that MP is not guaranteed to terminate under the conditions of Theorem 3 as τ in the admissible range indicated by Theorem 3 could be too small for termination (see the corresponding discussion in Section II). The ensuing statements on SSC-MP all apply only if MP, indeed, terminates for all points $\mathbf{y}_j \in \mathcal{Y}$. Theorem 3 identifies a range for the threshold parameter τ guaranteeing that both SSC-OMP and SSC-MP deliver a graph G which has each $\mathbf{y}_j \in \mathcal{Y}_\ell$, $\ell \in [L]$, connected to at least $O(d_\ell / \log(n_\ell - 1))$ other points in \mathcal{Y}_ℓ . If σ increases, the probability of OMP and MP selecting false connections increases and more iterations need to be performed for a given number of true connections to be selected. As a consequence, the interval for τ specified in Theorem 3 decreases as σ increases. As already pointed out, Theorem 3 does not guarantee the NFC property, and choosing τ too small will result in entries in the coefficient vectors \mathbf{b}_j that correspond to false connections. Intuitively, we expect that choosing τ sufficiently large, OMP and MP should stop early enough so as to avoid false connections. More specifically, one would expect that τ needs to be chosen larger as σ increases so as to avoid OMP and MP selecting points from $\mathcal{Y} \setminus \mathcal{Y}_\ell$ to represent $\mathbf{y}_j \in \mathcal{Y}_\ell$. Unfortunately, it seems rather difficult, at least for the statistical data model at hand, to analytically characterize a range for τ that guarantees the NFC property and simultaneously on the order of d_ℓ connections between $\mathbf{y}_j \in \mathcal{Y}_\ell$ and other points in \mathcal{Y}_ℓ , for all $j \in [n_\ell]$, $\ell \in [L]$. Nonetheless, it turns out, that in practice G often exhibits both of these properties if τ is chosen appropriately. Numerical results in Section IV-C corroborate this claim. In

summary, SSC-OMP and SSC-MP under DD-stopping with appropriately chosen τ detect the dimensions of the subspaces \mathcal{S}_ℓ correctly and adapt the sparsity level of the individual representations according to the dimension of the subspace the data point at hand lies in.

The procedure in [10, Alg. 2] for the selection of a per-data-point Lasso parameter in SSC has a similar subspace dimension-detecting property but comes with stronger theoretical guarantees. Specifically, [10, Thm. 3.1] and [10, Thm. 3.2] taken together guarantee the NFC property and, concurrently, that each $\mathbf{y}_j \in \mathcal{Y}_\ell$ is connected to on the order of d_ℓ other points in \mathcal{Y}_ℓ , all this provided that the noise variance—assumed known—is small enough. The procedure in [10, Alg. 2] could also be employed to select s_{\max} in SSC-OMP and SSC-MP under DI-stopping for each data point individually. This emulates DD-stopping by employing DI-stopping together with a data-point-wise parameter selection procedure. More specifically, as shown in [10, Lem. A.2], the optimal cost of the auxiliary Lasso problem in [10, Eq. (2.4)] for $\mathbf{y}_j \in \mathcal{Y}_\ell$ is proportional to $\sqrt{d_\ell}$. Squaring the optimal cost therefore yields an estimate of d_ℓ , which can, in turn, be used to select the parameter s_{\max} in SSC-OMP such that it is on the order of d_ℓ for $\mathbf{y}_j \in \mathcal{Y}_\ell$ and thereby satisfies the condition in Theorem 1 (we can lower-bound the factor $1/\log((n_\ell - 1)e/s_{\max})$ in that condition by $1/\log(Ne)$). This would then guarantee, in addition to the NFC property, on the order of $d_\ell / \log(n_\ell)$ true connections for $\mathbf{y}_j \in \mathcal{Y}_\ell$, $j \in [n_\ell]$, $\ell \in [L]$, for both SSC-OMP and SSC-MP under the conditions of Theorems 1 and 2, and would hence realize the “many true discoveries” and the NFC property at the same time as guaranteed by [10, Thm. 3.1] and [10, Thm. 3.2] for SSC. We point out, however, that the selection procedure in [10, Alg. 2] results in considerable computational burden in addition to solving the (already computationally demanding) N Lasso problems required by SSC or running the OMP and MP routines in SSC-OMP and SSC-MP, respectively, to perform the actual clustering.

Finally, we note that determining the range of admissible threshold parameters $\tau \in [0, 2/3 - (\sqrt{d_{\max}}/\sqrt{m})\sigma]$ in Theorem 3 requires knowledge of the noise variance σ^2 . In principle, knowledge of d_{\max} is required as well. We can, however, upper-bound $\sqrt{d_{\max}}/\sqrt{m}$ by 1 thereby obviating the need for knowing d_{\max} at the cost of a reduced range for τ .

IV. NUMERICAL RESULTS

We compare the performance of SSC-OMP, SSC-MP, SSC, TSC, and NSN. SSC-OMP and SSC-MP were implemented in Matlab exactly following their descriptions in Section II. For SSC, TSC, and NSN, we used the implementations provided in the corresponding references [8], [12], and [14], respectively. The code for reproducing all experiments in this paper can be found at <http://www.nari.ee.ethz.ch/commth/research/>. This code also contains information on the number of Monte Carlo runs used in the individual experiments.

We provide the algorithms with the true number of subspaces L in all experiments. All running times were measured on a PC with 32 GB RAM and 4-core Intel Core i7-3770K

CPU clocked at 3.50 GHz. It is quite common in the computer vision literature to perform post-processing on the adjacency matrix \mathbf{A} generated by the individual clustering algorithms. This can improve the clustering performance, but will not be pursued here in order to simplify our comparisons.

A. Comparison of SSC-OMP and SSC-MP

As the clustering condition (8) is only a sufficient condition and guarantees the NFC property only, it is unclear to what extent the behavior predicted by (8) is reflected in the clustering error (CE), i.e., the fraction of misclustered data points. The following experiment is devoted to answering this question while comparing SSC-OMP and SSC-MP. We generate data sets \mathcal{Y} according to the statistical data model described in Section III. Specifically, we set $d_\ell = d$, $\ell \in [L]$, and we choose the bases $\mathbf{U}^{(\ell)} \in \mathbb{R}^{m \times d}$ to all intersect in a shared t -dimensional space and to be mutually orthogonal on the orthogonal complement of this intersection. More specifically, the bases $\mathbf{U}^{(\ell)}$ are obtained by choosing a matrix $\mathbf{U} \in \mathbb{R}^{m \times (L(d-t)+t)}$ uniformly at random from the set of all orthonormal matrices of dimension $m \times (L(d-t)+t)$ and setting $\mathbf{U}^{(\ell)} := [\mathbf{U}_{[t]} \mathbf{U}_{\mathcal{T}_\ell}]$, where $\mathcal{T}_\ell := \{t + (\ell - 1)(d - t) + 1, \dots, t + \ell(d - t)\}$. This results in $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) = \sqrt{t/d}$, $k, \ell \in [L]$, $k \neq \ell$. Varying the parameter t therefore allows us to vary the pairwise affinities. We furthermore set $n_\ell = n$, $\ell \in [L]$, and generate instances of \mathcal{Y} by sampling n data points uniformly at random from each subspace and adding $\mathcal{N}(\mathbf{0}, (\sigma^2/m)\mathbf{I})$ noise to each data point (the noise vectors are generated independently across data points). We let $L = 3$, $d = 20$, $m = 200$, and vary t , $\rho_{\min} = \rho = n/d$, and σ^2 . Furthermore, we employ DI-stopping and set $s_{\max} = d/2 = 10$ for both SSC-OMP and SSC-MP, and $p_{\max} = N$. Figure 2 shows the CE as a function of $\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) = \sqrt{t/d}$, ρ , and σ^2 . The results nicely reflect the qualitative behavior indicated by the clustering condition (8). Specifically, both SSC-OMP and SSC-MP tolerate higher noise variance as the affinities between the subspaces decrease and the number of points in \mathcal{Y} drawn from each subspace, n , increases. It is furthermore interesting to observe that the performance of SSC-OMP and SSC-MP is virtually identical.

Recall that the clustering condition (8), apart from the term proportional to $\sqrt{1/\rho_{\min}}$, exhibits the same scaling behavior as those guaranteeing the NFC property for SSC in [10, Thm. 3.1] and for TSC in [12, Thm. 3]. To find out whether this additional term is an artifact of our proof technique, we first note that the clustering condition (8) takes the form

$$\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \frac{c_1}{\sqrt{\rho}} \leq c_2, \quad (11)$$

for fixed σ , and

$$\sigma(c_3 + \sigma c_4 + \frac{1}{\sqrt{\rho}}(c_5 + \sigma c_6)) \leq c_7, \quad (12)$$

for fixed maximum affinity, where c_1 - $c_7 > 0$ are constants, d_{\max} and m were assumed constant in both cases, and factors logarithmic in any of the parameters (variable or fixed) were neglected. Rewriting (11) and (12) assuming equality, we get

$$\rho = \left(\frac{c_1}{c_2 - \max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell)} \right)^2 \quad (13)$$

and

$$\rho = \left(\frac{\sigma(c_5 + c_6\sigma)}{c_7 - \sigma(c_3 + \sigma c_4)} \right)^2, \quad (14)$$

respectively. In the top and bottom rows of Figure 2, we now fit (13) and (14), respectively (by manually adjusting the constants c_1 - c_7), to the boundaries between the regions of success and the regions of failure. The shape of the fitted curves follows the boundaries indicated by the numerical results closely. This can be taken as an indication—at least to a certain extent—of the term in (8) proportional to $\sqrt{1/\rho_{\min}}$ being fundamental.




Fig. 2: CE as a function of $\text{aff} := \max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell)$, $\rho = n/d$, and σ^2 . The (fitted) black curves in the top and bottom rows correspond to the curves (13) and (14), respectively, and delineate the boundary between success and failure.

B. Face clustering

Next, we consider the problem of clustering face images of different individuals taken under varying illumination conditions. The rationale for employing subspace clustering to solve this problem stems from the observation that vectorized images of a given individual taken under different lighting conditions lie (approximately) in a 9-dimensional linear subspace [28]. We consider the Extended Yale B data set [21], [22], which contains 192×168 pixel frontal face images of 38 individuals, each taken under 64 different illumination conditions. As SSC-MP does not seem to have been considered before in the literature, we compare its performance in terms of CE to that of SSC-OMP, SSC, TSC, and NSN. Similar experiments comparing the performance of SSC-OMP to that of other subspace clustering algorithms for face clustering were presented in [13], [15], [16]. To accommodate the memory requirements

incurred by the optimization problems in SSC, we apply SSC (and, to ensure a fair comparison, also the other algorithms) to the downsampled 48×42 pixel versions of the images in the Extended Yale B data set provided in [8]. Note, however, that SSC-OMP, SSC-MP, TSC, and NSN all could handle the original 192×168 pixel images. Instances of the data set \mathcal{Y} are obtained by first choosing a subset of L individuals uniformly at random from the set of all individuals and then collecting the 64 (vectorized) images corresponding to each of the chosen individuals. As DD-stopping leads to a large number of false connections—arguably due to (sparse) corruptions induced by shadows and specular reflections in the face images [8]—we rely on DI-stopping (i.e., $\tau = 0$) with $s_{\max} = 5$ (which corresponds to the choice made in [13] for SSC-OMP) for both SSC-OMP and SSC-MP, and we set $p_{\max} = N$. Note that both OMP and MP perform exactly s_{\max} iterations in this experiment as the points in \mathcal{Y} are in general position. For SSC, TSC, and NSN we use the parameter values employed for the face clustering experiments in [8], [12], and [14], respectively. We emphasize that the experiments in [8], [12], and [14] all also provide the true number of subspaces L to the individual algorithms. Note furthermore that we employ SSC with the objective function as formulated in [8] accounting for sparse corruptions of the data points.

Table I shows the CE for different choices of L , averaged over 100 instances of \mathcal{Y} for each L . In Table II, we report the corresponding average running times. SSC-MP outperforms SSC-OMP (and TSC) for all values of L , but $L = 2$, and does so at consistently lower running times (recall that SSC-OMP and SSC-MP both perform exactly s_{\max} iterations in this experiment, but SSC-MP has a lower per-iteration cost than SSC-OMP). Furthermore, SSC-OMP uniformly outperforms SSC. The lowest CE is obtained with NSN. We note, however, that the performance of SSC-MP, except for $L = 2$, almost matches that of NSN, and SSC-MP consistently has roughly half the running time of NSN. TSC has the lowest running time but yields the largest CE (which is particularly high for this very data set [12]). The running time of SSC is one to two orders of magnitude higher than that of all other algorithms. Also note that the difference between the running times of SSC-OMP and SSC-MP is small because s_{\max} is small. Finally, the difference between the results for SSC-OMP, SSC, TSC, and NSN reported here compared to the results reported in published works [15], [8], [12], [14] can be attributed to the fact that we do not perform post-processing on the adjacency matrix \mathbf{A} produced by the individual algorithms.

Finally, we emphasize that for other applications such as, e.g., handwritten digit clustering [12], SSC and TSC yield lower CE than the other algorithms considered here, and none of the algorithms outperforms the others uniformly across applications as seen from experiments in [8], [12], [14].

L	2	3	5	8	10
SSC-OMP	2.83	4.04	6.81	12.98	14.14
SSC-MP	4.16	3.72	5.24	9.09	11.36
SSC	2.88	4.59	8.65	16.95	21.28
TSC	10.71	16.81	29.73	38.34	41.22
NSN	1.81	2.89	5.37	8.15	10.05

TABLE I: Average CE (in percent) for face clustering.

L	2	3	5	8	10
SSC-OMP	0.21	0.36	0.80	2.01	3.13
SSC-MP	0.15	0.28	0.65	1.74	2.80
SSC	12.63	17.50	28.24	45.78	60.33
TSC	0.08	0.14	0.29	0.62	0.95
NSN	0.34	0.55	1.10	3.05	5.64

TABLE II: Average running times (in seconds) for face clustering.

C. True positives/false positives tradeoff in DD-stopping

Recall that for DD-stopping Theorem 3 guarantees that each data point $\mathbf{y}_j \in \mathcal{Y}_\ell$, $\ell \in [L]$, is connected (in G) to at least $O(d_\ell/\log(n_\ell - 1))$ other points in \mathcal{Y}_ℓ ; we here refer to such connections as true positives (TP). As already mentioned, Theorem 3, does, however, not guarantee the absence of false connections and provides a lower bound on the number of TP only. It is therefore not clear, a priori, to what extent SSC-OMP and SSC-MP under DD-stopping, indeed, do detect the subspace dimensions. First, recall that by “detecting the subspace dimensions” we mean that the coefficient vectors \mathbf{b}_j for all $\mathbf{y}_j \in \mathcal{Y}_\ell$, $\ell \in [L]$, have on the order of d_ℓ non-zero entries corresponding to TP and essentially no non-zero entries corresponding to false connections. We designate the number of TP as $\#\text{TP}_\ell$, and the number of false positives (FP) as $\#\text{FP}_\ell$, both averaged over the data points corresponding to the subspace \mathcal{S}_ℓ . We will also need the notions of true positive rate (TPR) and false positive rate (FPR) defined, in this experiment, as $\text{TPR}_\ell = \#\text{TP}_\ell/d_\ell$ and $\text{FPR}_\ell = \#\text{FP}_\ell/(m - d_\ell)$, respectively.

We set $m = 300$, $L = 4$, $\rho_\ell = 4$, $\ell \in [L]$, and choose the basis matrices $\mathbf{U}^{(\ell)}$ for the subspaces of dimensions 20, 40, 60, and 80, according to $\mathbf{U}^{(\ell)} := [\bar{\mathbf{U}} \ \tilde{\mathbf{U}}^{(\ell)}]$, where $\bar{\mathbf{U}}$ and the $\tilde{\mathbf{U}}^{(\ell)}$ are drawn uniformly at random from the set of all orthonormal matrices of dimensions 300×4 and $300 \times (d_\ell - 4)$, respectively. This guarantees that $\max_{k,\ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \geq \sqrt{1/5}$ as the subspaces all intersect in a shared 4-dimensional space and possibly overlap in the orthogonal complement of this shared space. The data points are then drawn according to the statistical data model described in Section III.

Figure 3 shows $\#\text{TP}_\ell$ along with TPR_ℓ and FPR_ℓ as a function of τ and for different values of σ^2 . The middle and bottom rows in Figure 3 show that the TPR curves corresponding to different ℓ are almost on top of each other, i.e., $\text{TPR}_1(\tau) \approx \dots \approx \text{TPR}_4(\tau) \approx c_{\text{TPR}}(\tau)$, which means that the number of TP for each subspace is, indeed, roughly proportional to the subspace dimension as $\#\text{TP}_\ell(\tau) = \text{TPR}_\ell(\tau)d_\ell \approx c_{\text{TPR}}(\tau)d_\ell$, $\ell \in [L]$. This indicates that the result in Theorem 3 is order-wise optimal in d_ℓ . As, in addition, for large enough τ , $\text{FPR}_\ell \approx 0$, $\ell \in [L]$, we conclude that SSC-OMP and SSC-MP, indeed, exhibit excellent subspace dimension detection properties provided that τ is chosen appropriately.

We finally note that a similar experiment investigating the TPR/FPR tradeoff as a function of the Lasso parameter in SSC was conducted in [10, Sec. 2.4.3] with the main conclusion that a Lasso parameter on the order of $1/\sqrt{d_\ell}$ extracts the subspace dimensions correctly order-wise with essentially no FP.

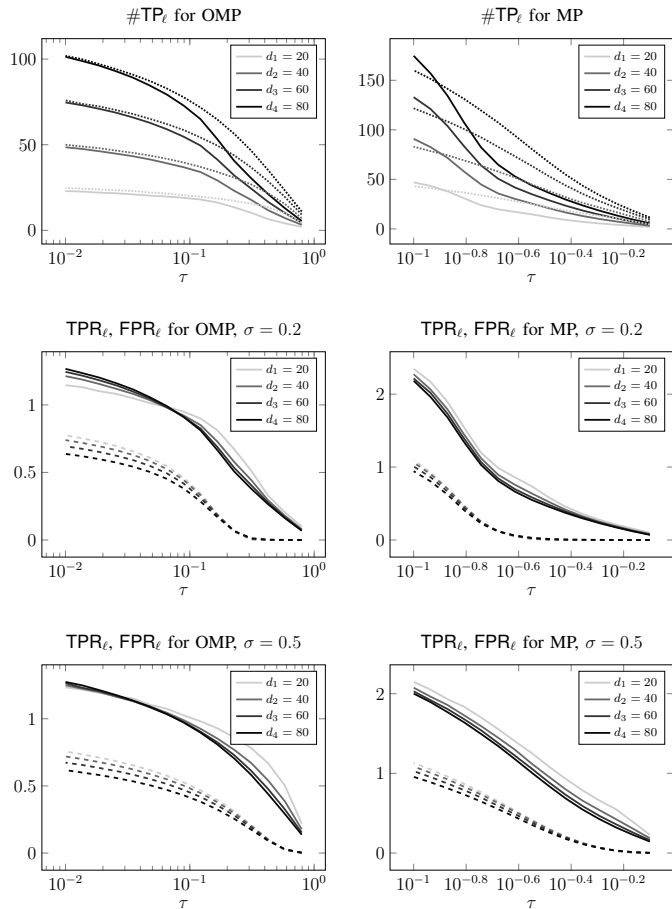


Fig. 3: TP/FP tradeoff in DD-stopping as a function of τ . Top row: Solid lines: $\sigma = 0.2$, dotted lines: $\sigma = 0.5$. Middle and bottom rows: Solid lines: TPR, dashed lines: FPR.

D. Influence of s_{\max} and p_{\max} in DI-stopping

Recall that Theorems 1 and 2 provide a range of admissible values for s_{\max} and p_{\max} in DI-stopping. For small d_{\min} , these ranges are, however, small. As already pointed out, taking s_{\max} , p_{\max} too small leads to cluster split-ups, whereas s_{\max} , p_{\max} too large results in a large number of false connections in G . It is therefore important to determine the sensitivity of SSC-OMP and SSC-MP performance w.r.t. the choice of s_{\max} and p_{\max} . The next experiment is devoted to this matter. We consider the CE as well as the average TPR and FPR which, in this experiment, are defined as $\text{TPR}'_{\ell} = \#\text{TP}_{\ell}/n_{\ell}$ and $\text{FPR}'_{\ell} = \#\text{FP}_{\ell}/(N - n_{\ell})$, respectively. These alternative normalizations are used as we are not interested in investigating the dependence of TP and FP on the d_{ℓ} , as was done in Section IV-C; rather, we want to ensure that $\text{TPR}'_{\ell}, \text{FPR}'_{\ell} \in [0, 1]$. Furthermore, we define the TP- ℓ_1 -norm and the FP- ℓ_1 -norm as the ℓ_1 -norm of the entries of the coefficient vectors \mathbf{b}_j corresponding to TP and FP, respectively, both averaged over all data points. The TP- and FP- ℓ_1 -norms hence correspond to half of the average weight associated with TP and FP, respectively; the factor $1/2$ stems from the adjacency matrix of G being given by $\mathbf{A} = \mathbf{B} + \mathbf{B}^T$. The motivation for considering TP-/FP- ℓ_1 -norms comes from the fact that the performance of spectral clustering is determined not only by the number of

TP and FP, but also by the weights associated with the TP and the FP. Specifically, even when the \mathbf{b}_j contain a considerable number of FP, good performance can still be obtained provided that the corresponding FP- ℓ_1 -norm is sufficiently small.

We cluster synthetic data generated according to the statistical data model described in Section III as well as images taken from the Extended Yale B data set (we use downsampled 48×42 pixel versions of the images, see Section IV-B for a detailed description). More specifically, in the case of synthetic data, we set $L = 3$, $m = 80$, $\sigma = 0.5$, $d_{\ell} = 15$, $\ell \in [L]$, and $\rho_{\ell} = 4$, $\ell \in [L]$, and we generate the bases $\mathbf{U}^{(\ell)}$, $\ell \in [L]$, to intersect in a shared 3-dimensional space by following the construction employed in Section IV-C such that $\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_{\ell}) = \sqrt{1/5}$. In the case of face clustering, we follow the procedure described in Section IV-B to obtain instances of \mathcal{Y} containing the face images of $L = 3$ randomly selected individuals.

Figure 4 shows the CE, TPR/FPR, and TP-/FP- ℓ_1 -norms as a function of s_{\max} and p_{\max} . For both clustering problems the CE of SSC-OMP is seen to increase rapidly as a function of s_{\max} , while for SSC-MP with $p_{\max} = N$ (i.e., stopping is activated by $s = s_{\max}$), the CE increases very slowly for the face clustering problem and does not increase at all in the case of synthetic data. This indicates that SSC-MP exhibits significantly smaller sensitivity to the choice of s_{\max} than SSC-OMP. As already pointed out in Section II, this is due to the ability of SSC-MP to select points $\mathbf{y}_i \in \mathcal{Y}_{\ell} \setminus \{\mathbf{y}_j\}$ repeatedly to participate in the representation of $\mathbf{y}_j \in \mathcal{Y}_{\ell}$; for given s_{\max} this results in SSC-MP producing fewer FP than SSC-OMP. Moreover, for SSC-MP, the TP- ℓ_1 -norm exceeds the FP- ℓ_1 -norm for all values of s_{\max} , while for SSC-OMP the FP- ℓ_1 -norm exceeds the TP- ℓ_1 -norm for large s_{\max} , and does so significantly. The coefficient vectors produced by SSC-MP hence lead to more favorable conditions for the spectral clustering step than those produced by SSC-OMP.

We further observe that the TPR and FPR of SSC-MP, with $s_{\max} = \infty$ (i.e., stopping is activated by $\|\mathbf{b}_j\|_0 = p_{\max}$), as a function of p_{\max} , increase at the same rate as the TPR and FPR of SSC-OMP as a function of s_{\max} . However, the ratio of TP- and FP- ℓ_1 -norms for SSC-MP significantly exceeds that for SSC-OMP for all values of u (u is the variable on the x -axis in Figure 4 and corresponds to s_{\max} for SSC-OMP and p_{\max} for SSC-MP). In other words, even when we force the representations computed by SSC-MP and SSC-OMP to have the same sparsity level, thereby discounting the advantage MP has through its ability to reselect data points, SSC-MP still produces weight assignments in G that are more favorable in terms of spectral clustering. Indeed, in both the face clustering and the synthetic data clustering problem the CE incurred by SSC-OMP significantly exceeds that of SSC-MP for most values of u . In summary, this indicates that when we enforce the same target sparsity level for SSC-MP and SSC-OMP, SSC-MP is less sensitive to the choice of the sparsity level than SSC-OMP.

We finally note that while in the noisy case SSC-MP is much more robust than SSC-OMP w.r.t. the choice of the parameters for DI-stopping, in the noiseless case SSC-OMP yields slightly lower CE than SSC-MP for DI-stopping if

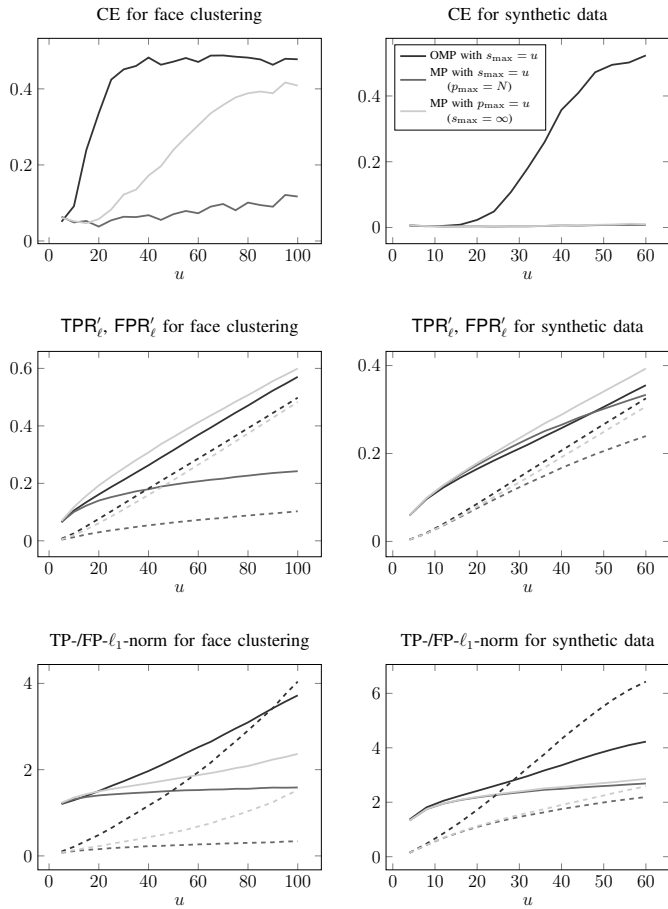


Fig. 4: Clustering performance of SSC-OMP and SSC-MP for DI-stopping, as a function of s_{\max} and p_{\max} , respectively. Middle row: solid lines: TPR, dashed lines: FPR. Bottom row: solid lines: TP- ℓ_1 -norm, dashed lines: FP- ℓ_1 -norm.

the subspace affinities are large. This matter is investigated numerically in Appendix A.

ACKNOWLEDGMENTS

The authors would like to thank Reinhard Heckel and Martin Jaggi for insightful discussions.

APPENDIX A

INFLUENCE OF s_{\max} AND p_{\max} IN DI-STOPPING FOR NOISELESS DATA

We compare the influence of s_{\max} and p_{\max} on the performance of SSC-OMP and SSC-MP with DI-stopping and for noiseless data. To this end, we generate data lying in a union of three subspaces as described in Section IV-D (with $\sigma = 0$), considering the pairs $(5, 3)$, $(10, 3)$, and $(10, 6)$ for (t, ρ) , where t denotes the number of dimensions which the three subspaces intersect in. Figure 5 shows the CE along with the quantities TPR/FPR and the TP-/FP- ℓ_1 -norm as a function of s_{\max} (or p_{\max} for SSC-MP if the maximum sparsity level is used as stopping criterion) in the range $\{1, \dots, 2d\}$. We observe that SSC-OMP yields a slightly lower CE than SSC-MP for $t = 10$. While SSC-MP yields a higher TPR and a lower FPR than SSC-OMP for most of the values of t, ρ ,

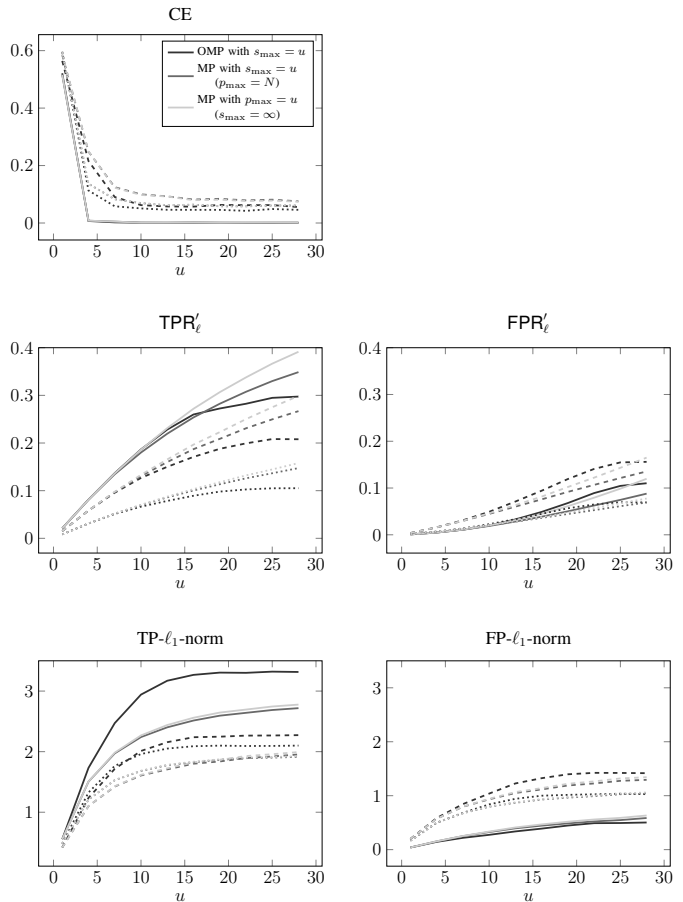


Fig. 5: Clustering performance of SSC-OMP and SSC-MP for DI-stopping, as a function of s_{\max} and p_{\max} for noiseless data. Solid line: $t = 5, \rho = 3$; dashed line: $t = 10, \rho = 3$; dotted line: $t = 10, \rho = 6$.

and s_{\max} or p_{\max} , SSC-MP assigns smaller values, than SSC-OMP, to entries in the adjacency matrix \mathbf{A} corresponding to the true connections (see the plot in the last row, left, in Figure 5). This arguably leads to the slightly higher CE of SSC-MP compared to SSC-OMP for $t = 10$.

APPENDIX B

PROOF OF THEOREM 1

Throughout the proof, we shall use $\mathbf{Y}^{(\ell)} := \mathbf{X}^{(\ell)} + \mathbf{Z}^{(\ell)} = \mathbf{U}^{(\ell)}\mathbf{A}^{(\ell)} + \mathbf{Z}^{(\ell)}$, $\mathbf{X}^{(\ell)} \in \mathbb{R}^{m \times n_\ell}$, $\mathbf{A}^{(\ell)} \in \mathbb{R}^{d_\ell \times n_\ell}$, and $\mathbf{Z}^{(\ell)} \in \mathbb{R}^{m \times n_\ell}$ to denote the matrices whose columns are the $\mathbf{y}_i^{(\ell)}$, $\mathbf{x}_i^{(\ell)}$, $\mathbf{a}_i^{(\ell)}$, and $\mathbf{z}_i^{(\ell)}$, $i \in [n_\ell]$, respectively. Furthermore, $\mathbf{P}_\parallel := \mathbf{U}^{(\ell)}\mathbf{U}^{(\ell)T}$ and $\mathbf{P}_\perp := \mathbf{I} - \mathbf{U}^{(\ell)}\mathbf{U}^{(\ell)T}$ stand for the orthogonal projection onto \mathcal{S}_ℓ and its orthogonal complement (in \mathbb{R}^m) \mathcal{S}_ℓ^\perp , respectively. We do not indicate the dependence of \mathbf{P}_\parallel and \mathbf{P}_\perp on ℓ as this is always clear from the context. For $\mathbf{v} \in \mathbb{R}^m$ and $\mathbf{A} \in \mathbb{R}^{m \times n}$, we use the shorthands $\mathbf{v}_\parallel := \mathbf{P}_\parallel \mathbf{v}$, $\mathbf{A}_\parallel := \mathbf{P}_\parallel \mathbf{A}$, $\mathbf{v}_\perp := \mathbf{P}_\perp \mathbf{v}$, and $\mathbf{A}_\perp := \mathbf{P}_\perp \mathbf{A}$. Further, $\tilde{\mathbf{a}}_i^{(\ell)} \in \mathbb{R}^{d_\ell}$ denotes the coefficients of $\mathbf{y}_{i\parallel}^{(\ell)}$ in the basis $\mathbf{U}^{(\ell)}$, i.e., $\mathbf{y}_{i\parallel}^{(\ell)} = \mathbf{U}^{(\ell)}\tilde{\mathbf{a}}_i^{(\ell)}$, and similarly $\mathbf{Y}_{\parallel}^{(\ell)} = \mathbf{P}_\parallel \mathbf{Y}^{(\ell)} = \mathbf{U}^{(\ell)}\tilde{\mathbf{A}}^{(\ell)}$. Note that the distribution of $\tilde{\mathbf{a}}_i^{(\ell)} = \mathbf{a}_i^{(\ell)} + \mathbf{U}^{(\ell)T} \mathbf{z}_i^{(\ell)}$ is rotationally invariant as $\mathbf{a}_i^{(\ell)}$ and $\mathbf{U}^{(\ell)T} \mathbf{z}_i^{(\ell)}$ are statistically independent and both have rotationally invariant distributions.

Finally, $\mathbf{r}_s(\mathbf{x}, \mathbf{D})$ and $\Lambda_s(\mathbf{x}, \mathbf{D})$ denote the residual and the index set, respectively, after iteration s , obtained by OMP applied to \mathbf{x} with the columns of \mathbf{D} as dictionary elements.

If $\min_{\ell \in [L]} \{c_s d_\ell / \log((n_\ell - 1)e/s_{\max})\} < 1$, then the condition in Theorem 1 admits zero OMP iterations, i.e., the graph G delivered by SSC-OMP has an empty edge set and thereby trivially no false connections. We therefore consider the case $1 \leq s_{\max} \leq \min_{\ell \in [L]} \{c_s d_\ell / \log((n_\ell - 1)e/s_{\max})\}$ in the remainder of the proof.

The graph G obtained by SSC-OMP has no false connections if for each $\mathbf{y}_i^{(\ell)} \in \mathcal{Y}_\ell$, for all $\ell \in [L]$, the OMP algorithm selects points from \mathcal{Y}_ℓ in all s_{\max} iterations.³ Now, the OMP selection rule (1) for $\mathbf{y}_i^{(\ell)}$ implies that OMP selects a point from \mathcal{Y}_ℓ in the $(s+1)$ -st iteration if

$$\max_{k \neq \ell, j} \left| \left\langle \mathbf{y}_j^{(k)}, \mathbf{r}_s \right\rangle \right| < \max_{j \in [n_\ell] \setminus (\Lambda_s \cup \{i\})} \left| \left\langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_s \right\rangle \right|, \quad (15)$$

where $\max_{k \neq \ell, j}$ denotes maximization over subspaces $k \in [L]$, $k \neq \ell$, and over the indices j of the points $\mathbf{y}_j^{(k)} \in \mathcal{Y}_k$ in these subspaces. Hence, the graph G obtained by SSC-OMP satisfies the NFC property if (15) holds for all s_{\max} iterations, for each $\mathbf{y}_i^{(\ell)} \in \mathcal{Y}_\ell$, for all $\ell \in [L]$. We now establish that this holds for our statistical data model w.p. at least P^* as defined in (9).

Our analysis will be based on an auxiliary algorithm termed “reduced OMP”, which has access to the reduced dictionary $\mathcal{Y}_\ell \setminus \{\mathbf{y}_i^{(\ell)}\}$ only—instead of the full dictionary $\mathcal{Y} \setminus \{\mathbf{y}_i^{(\ell)}\}$ —to represent $\mathbf{y}_i^{(\ell)}$. We henceforth use the shorthands $\mathbf{r}_s^{(\ell)}$ for the residuals $\mathbf{r}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)})$ corresponding to reduced OMP. The dependence of the $\mathbf{r}_s^{(\ell)}$ on the index i of the data point $\mathbf{y}_i^{(\ell)}$ to be represented is not made explicit for notational ease. If the $\mathbf{r}_s^{(\ell)}$ satisfy (15) for *all iterations* $s \in [s_{\max}]$, the reduced OMP algorithm and the original OMP algorithm select exactly the same data points and do so in exactly the same order. In this case, we also have $\mathbf{r}_s^{(\ell)} = \mathbf{r}_s$, for all $s \in [s_{\max}]$. As $\mathbf{r}_s^{(\ell)}$ satisfying (15) for all $s \in [s_{\max}]$ is necessary and sufficient for \mathbf{r}_s to satisfy (15) for all $s \in [s_{\max}]$, a lower bound P^* on the probability of $\mathbf{r}_s^{(\ell)}$ satisfying (15) for all $s \in [s_{\max}]$ also constitutes a lower bound on the probability of \mathbf{r}_s satisfying (15) for all $s \in [s_{\max}]$. Working with the reduced OMP algorithm is beneficial as $\mathbf{r}_s^{(\ell)}$ is a function of the points in \mathcal{Y}_ℓ only and is therefore statistically independent of the points in $\mathcal{Y} \setminus \mathcal{Y}_\ell$. This is significant as it will allow us to apply standard concentration of measure inequalities for independent random variables. In the remainder of the proof, we work with reduced OMP exclusively.

We start by providing intuition on the proof idea. To this end, we expand the inner products in (15) according to

$$\begin{aligned} \left\langle \mathbf{y}_j^{(k)}, \mathbf{r}_s^{(\ell)} \right\rangle &= \left\langle \mathbf{x}_j^{(k)} + \mathbf{z}_j^{(k)}, \mathbf{r}_{s\parallel}^{(\ell)} + \mathbf{r}_{s\perp}^{(\ell)} \right\rangle \\ &= \left\langle \mathbf{x}_j^{(k)}, \mathbf{r}_{s\parallel}^{(\ell)} \right\rangle + \left\langle \mathbf{x}_j^{(k)}, \mathbf{r}_{s\perp}^{(\ell)} \right\rangle \end{aligned}$$

³ For DI-stopping the OMP algorithm terminates w.p. 1 after exactly s_{\max} iterations as in our data model the points in \mathcal{Y} are in general position w.p. 1 and $s_{\max} < \min_{\ell \in [L]} c_s d_\ell < \min\{m, N-1\}$ by the condition on s_{\max} in Theorem 1.

$$+ \left\langle \mathbf{z}_j^{(k)}, \mathbf{r}_{s\parallel}^{(\ell)} \right\rangle + \left\langle \mathbf{z}_j^{(k)}, \mathbf{r}_{s\perp}^{(\ell)} \right\rangle. \quad (16)$$

The first term in (16) quantifies the similarity of the portions of $\mathbf{y}_j^{(k)}$ and $\mathbf{r}_s^{(\ell)}$ that lie in \mathcal{S}_k and \mathcal{S}_ℓ , respectively, i.e., the “signal components” of $\mathbf{y}_j^{(k)}$ and $\mathbf{r}_s^{(\ell)}$, while the other terms all account for interactions with or between “undesired components” residing in $\mathcal{S}_k^\perp, \mathcal{S}_\ell^\perp$. If the similarities—in absolute value—of the “signal components” for $k = \ell, j \in [n_\ell] \setminus (\Lambda_s \cup \{i\})$, are sufficiently large relative to those for $k \neq \ell, j \in [n_k]$, and if the interactions of all “undesired components” are sufficiently small, for $k, \ell \in [L]$, then (15) holds. Following [16, Proofs of Thm. 3, Cor. 1] this intuition will be made quantitative and rigorous by introducing events that, when conditioned on, yield bounds on the absolute values of the individual terms in (16) that are of analytically amenable form. These bounds will then be employed to derive an upper bound on the LHS and a lower bound on the RHS of (15) that both hold conditionally on the intersection of the underlying events. Based on these bounds, we will then show that the clustering condition (8) implies (15) w.p. at least P^* . The particular choice of the events that we condition on is delicate, but when done properly, allows us to make the statistical dependencies between $\mathbf{r}_s^{(\ell)}$ and $\mathbf{y}_j^{(\ell)}, j \in [n_\ell] \setminus \{i\}$, analytically tractable. We finally note that although the general idea of conditioning on suitably defined events is taken from previous work by the authors [16, Proofs of Thm. 3, Cor. 1], the choice of the specific events as well as other technical aspects of the present proof differ significantly from [16, Proofs of Thm. 3, Cor. 1].

We commence the formal proof by upper-bounding the LHS of (15) according to

$$\begin{aligned} \max_{k \neq \ell, j} \left| \left\langle \mathbf{x}_j^{(k)} + \mathbf{z}_j^{(k)}, \mathbf{r}_{s\parallel}^{(\ell)} + \mathbf{r}_{s\perp}^{(\ell)} \right\rangle \right| &\leq \max_{k \neq \ell, j} \left| \left\langle \mathbf{x}_j^{(k)}, \mathbf{r}_{s\parallel}^{(\ell)} \right\rangle \right| + \max_{k \neq \ell, j} \left| \left\langle \mathbf{x}_j^{(k)}, \mathbf{r}_{s\perp}^{(\ell)} \right\rangle \right| \\ &\quad + \max_{k \neq \ell, j} \left| \left\langle \mathbf{z}_j^{(k)}, \mathbf{r}_s^{(\ell)} \right\rangle \right| \\ &\leq 4 \log(N^3 s_{\max}) \frac{\left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F}{\sqrt{d_k} \sqrt{d_\ell}} \left\| \mathbf{r}_{s\parallel}^{(\ell)} \right\|_2 \\ &\quad + \frac{\sqrt{2 \log(N^3 s_{\max})}}{\sqrt{m-d_\ell}} \left\| \mathbf{r}_{s\perp}^{(\ell)} \right\|_2 \\ &\quad + \frac{\sqrt{2 \log(N^3 s_{\max})}}{\sqrt{m}} \frac{3}{2} \sigma \left(1 + \frac{3}{2} \sigma \right), \quad (17) \end{aligned}$$

where the second inequality holds on the event $\mathcal{E}_1^{(\ell, i, s)} \cap \mathcal{E}_2^{(\ell, i, s)} \cap \mathcal{E}_3^{(\ell, i, s)} \cap \mathcal{E}_4$ with

$$\begin{aligned} \mathcal{E}_1^{(\ell, i, s)} &:= \left\{ \max_{k \neq \ell, j} \left| \left\langle \mathbf{x}_j^{(k)}, \mathbf{r}_{s\parallel}^{(\ell)} \right\rangle \right| \right. \\ &\quad \left. \leq 4 \log(N^3 s_{\max}) \frac{\left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F}{\sqrt{d_k} \sqrt{d_\ell}} \left\| \mathbf{r}_{s\parallel}^{(\ell)} \right\|_2 \right\}, \quad (18) \end{aligned}$$

$$\mathcal{E}_2^{(\ell, i, s)} := \left\{ \max_{k \neq \ell, j} \left| \left\langle \mathbf{x}_j^{(k)}, \mathbf{r}_{s\perp}^{(\ell)} \right\rangle \right| \leq \frac{\sqrt{2 \log(N^3 s_{\max})}}{\sqrt{m-d_\ell}} \left\| \mathbf{r}_{s\perp}^{(\ell)} \right\|_2 \right\}, \quad (19)$$

$$\begin{aligned} \mathcal{E}_3^{\ell,i,s} &:= \left\{ \max_{k \neq \ell, j} \left| \langle \mathbf{z}_j^{(k)}, \mathbf{r}_s^{(\ell)} \rangle \right| \right. \\ &\leq \left. \frac{\sqrt{2 \log(N^3 s_{\max})}}{\sqrt{m}} \left(1 + \|\mathbf{z}_i^{(\ell)}\|_2 \right) \max_{k \neq \ell, j} \|\mathbf{z}_j^{(k)}\|_2 \right\}, \end{aligned} \quad (20)$$

$$\begin{aligned} \mathcal{E}_4 &:= \left\{ \left\{ \|\mathbf{z}_j^{(\ell)}\|_2 \leq \frac{3}{2} \frac{\sqrt{d_\ell}}{\sqrt{m}} \sigma \right\} \cap \left\{ \|\mathbf{z}_j^{(\ell)}\|_2 \leq \frac{3}{2} \sigma \right\}, \right. \\ &\quad \left. \forall \ell \in [L], j \in [n_\ell] \right\}. \end{aligned} \quad (21)$$

Note that the dependence of $\mathcal{E}_1^{\ell,i,s}$ and $\mathcal{E}_2^{\ell,i,s}$ on i is due to $\mathbf{r}_{s\parallel}^{(\ell)}$ and $\mathbf{r}_{s\perp}^{(\ell)}$, both of which are functions of $\mathbf{y}_i^{(\ell)}$. Here, $\mathcal{E}_1^{\ell,i,s}$ pertains to the similarities of the ‘‘signal components’’ for $k \neq \ell$, $\mathcal{E}_2^{\ell,i,s}$ and $\mathcal{E}_3^{\ell,i,s}$ quantify the similarity of ‘‘undesired components’’, and \mathcal{E}_4 controls the magnitude of the ‘‘undesired components’’ of the $\mathbf{y}_i^{(\ell)}$.

We proceed by lower-bounding the RHS of (15). Using $\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2 = \|(\mathbf{I} - \mathbf{Y}_{\Lambda_s}^{(\ell)}(\mathbf{Y}_{\Lambda_s}^{(\ell)})^\dagger)\mathbf{y}_i^{(\ell)}\|_2 \leq \|\mathbf{y}_i^{(\ell)}\|_2 \leq 1 + \|\mathbf{z}_i^{(\ell)}\|_2$ (where the inequality is thanks to $\|\mathbf{x}_i^{(\ell)}\|_2 = 1$ and $\mathbf{I} - \mathbf{Y}_{\Lambda_s}^{(\ell)}(\mathbf{Y}_{\Lambda_s}^{(\ell)})^\dagger$ being an orthogonal projection matrix), we find that on the event $\mathcal{E}_4 \cap \mathcal{E}_5^{\ell,i}$ with

$$\begin{aligned} \mathcal{E}_5^{\ell,i} &:= \left\{ \max_{j \in [n_\ell] \setminus \{\Lambda_s \cup \{i\}\}} \left| \langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_s^{(\ell)} \rangle \right| \right. \\ &\geq \left(1 - \frac{c_4 + 1}{\sqrt{\rho_\ell}} \right) \frac{\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2}{\sqrt{d_\ell}} \\ &\quad \left. - \sigma \left(\frac{1}{\sqrt{m}} + \frac{2}{\sqrt{n_\ell - 1}} \right) \|\mathbf{r}_s^{(\ell)}\|_2 \right\}, \end{aligned} \quad (22)$$

where $c_4 > 0$ is the numerical constant in Lemma 9, the RHS of (15) obeys

$$\begin{aligned} \max_{j \in [n_\ell] \setminus \{\Lambda_s \cup \{i\}\}} \left| \langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_s^{(\ell)} \rangle \right| &\geq \left(1 - \frac{c_4 + 1}{\sqrt{\rho_\ell}} \right) \frac{\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2}{\sqrt{d_\ell}} \\ &\quad - \sigma \left(\frac{1}{\sqrt{m}} + \frac{2}{\sqrt{n_\ell - 1}} \right) \left(1 + \frac{3}{2} \sigma \right). \end{aligned} \quad (23)$$

On $\mathcal{E}_1^{\ell,i,s} \cap \mathcal{E}_2^{\ell,i,s} \cap \mathcal{E}_3^{\ell,i,s} \cap \mathcal{E}_4 \cap \mathcal{E}_5^{\ell,i}$, (15) is now implied by [RHS of (17)] < [RHS of (23)]; multiplying this inequality by $\sqrt{d_\ell}/(4 \log(N^3 s_{\max}))\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2$, we get

$$\begin{aligned} &\underbrace{\frac{\|\mathbf{U}^{(k)T} \mathbf{U}^{(\ell)}\|_F}{\sqrt{d_k}}}_{\leq \max_{k: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell)} + \frac{1}{\sqrt{8 \log(N^3 s_{\max})}} \left(\frac{\sqrt{d_\ell}}{\sqrt{m - d_\ell}} \frac{\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2}{\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2} \right) \\ &\quad + \frac{\sigma}{\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2} \frac{\sqrt{d_\ell}}{\sqrt{m}} \frac{3}{2} \left(1 + \frac{3}{2} \sigma \right) \\ &< \frac{1}{4 \log(N^3 s_{\max})} \underbrace{\left(\left(1 - \frac{c_4 + 1}{\sqrt{\rho_\ell}} \right) \right)}_{\geq 1/2} \end{aligned}$$

$$- \frac{\sigma}{\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2} \left(\frac{\sqrt{d_\ell}}{\sqrt{m}} + \frac{2\sqrt{d_\ell}}{\sqrt{n_\ell - 1}} \right) \left(1 + \frac{3}{2} \sigma \right), \quad (24)$$

where $1 - (c_4 + 1)/\sqrt{\rho_\ell} \geq 1/2$ follows from $\rho_\ell \geq \rho_{\min} \geq c_\rho := 4(c_4 + 1)^2$, for all $\ell \in [L]$, and $\sqrt{d_\ell}/\sqrt{m - d_\ell} \leq \sqrt{2d_\ell}/\sqrt{m}$ is by $m \geq 2d_{\max} \geq 2d_\ell$, for all $\ell \in [L]$. Rearranging terms in (24) and using $\sqrt{8 \log(N^3 s_{\max})} < 4 \log(N^3 s_{\max})$, for $N \geq 2$, we can see that (24) is implied by

$$\begin{aligned} &\max_{k: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \frac{1}{\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2 \sqrt{8 \log(N^3 s_{\max})}} \left(\frac{\sqrt{2d_\ell}}{\sqrt{m}} \|\mathbf{r}_{s\perp}^{(\ell)}\|_2 \right. \\ &\quad \left. + \sigma \left(\frac{5}{2} \frac{\sqrt{d_\ell}}{\sqrt{m}} + \frac{2}{\sqrt{\rho_\ell}} \right) \left(1 + \frac{3}{2} \sigma \right) \right) \leq \frac{1}{8 \log(N^3 s_{\max})}. \end{aligned} \quad (25)$$

To further simplify (25), we upper-bound $\|\mathbf{r}_{s\perp}^{(\ell)}\|_2$ and lower-bound $\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2$. To this end, we introduce the events

$$\begin{aligned} \mathcal{E}_6^{\ell,i} &:= \left\{ \|\mathbf{r}_{s\perp}^{(\ell)}\|_2 \leq \|\mathbf{z}_{i\perp}^{(\ell)}\|_2 + \frac{3\sigma}{\tilde{a}} \|\mathbf{y}_i^{(\ell)}\|_2, \forall s \leq s_{\max} \right\}, \\ \mathcal{E}_7^{\ell,i} &:= \left\{ \|\mathbf{r}_{s\parallel}^{(\ell)}\|_2 \right. \\ &\quad \left. > \|\mathbf{y}_{i\parallel}^{(\ell)}\|_2 \left(\frac{2}{3} - \sqrt{\frac{3\bar{s} \log((n_\ell - 1)e/\bar{s})}{d_\ell}} \right), \forall s \leq \bar{s} \right\}, \end{aligned} \quad (26)$$

where $\tilde{a} := \min_{j \in [n_\ell] \setminus \{i\}} \|\mathbf{y}_j^{(\ell)}\|_2 \geq \min_{j \in [n_\ell] \setminus \{i\}} (\|\mathbf{x}_j^{(\ell)}\|_2 - \|\mathbf{z}_j^{(\ell)}\|_2) \geq 1 - \max_{j \in [n_\ell] \setminus \{i\}} \|\mathbf{z}_j^{(\ell)}\|_2$. Setting $\bar{s} = s_{\max}$ in (26), on $\mathcal{E}_4 \cap \mathcal{E}_6^{\ell,i} \cap \mathcal{E}_7^{\ell,i}$, we have

$$\|\mathbf{r}_{s\perp}^{(\ell)}\|_2 \leq \frac{3}{2} \sigma + 3\sigma \frac{1 + \frac{3}{2} \sigma}{1 - \frac{3}{2} \frac{\sqrt{d_\ell}}{\sqrt{m}} \sigma} \leq \sigma(8 + 10\sigma), \quad (27)$$

$$\begin{aligned} \|\mathbf{r}_{s\parallel}^{(\ell)}\|_2 &> \left(1 - \frac{3}{2} \frac{\sqrt{d_\ell}}{\sqrt{m}} \sigma \right) \left(\frac{2}{3} - \sqrt{3c_s} \right) \\ &> \left(1 - \frac{3}{2} \frac{\sqrt{d_\ell}}{\sqrt{m}} \sigma \right) \frac{1}{9} > \frac{1}{20}, \end{aligned} \quad (28)$$

where we employed the assumptions $m \geq 2d_{\max} \geq 2d_\ell$, for all $\ell \in [L]$, and $\sigma \leq 1/2$ to get (27), and used $s_{\max} \leq c_s d_\ell / \log((n_\ell - 1)e/s_{\max})$, for all $\ell \in [L]$, and $c_s := \min\{1/10, c_1\}$ (with c_1 the constant in Lemma 7 below), to arrive at (28). With (27) and (28), it follows that (25) is implied by

$$\begin{aligned} &\max_{k: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \frac{10\sigma}{\sqrt{\log(N^3 s_{\max})}} \left(\frac{\sqrt{d_\ell}}{\sqrt{m}} (10 + 13\sigma) \right. \\ &\quad \left. + \frac{\sqrt{2}}{\sqrt{\rho_\ell}} \left(1 + \frac{3}{2} \sigma \right) \right) \leq \frac{1}{8 \log(N^3 s_{\max})}. \end{aligned}$$

This inequality holds for all $\ell \in [L]$ by the clustering condition (8) with $c(\sigma) = 10 + 13\sigma$. Hence, on the event

$$\mathcal{E}^* := \bigcap_{\ell, i, s} (\mathcal{E}_1^{\ell,i,s} \cap \mathcal{E}_2^{\ell,i,s} \cap \mathcal{E}_3^{\ell,i,s})$$

$$\cap \mathcal{E}_4 \cap \mathcal{E}_5^{(\ell,i)} \cap \mathcal{E}_6^{(\ell,i)} \cap \mathcal{E}_7^{(\ell,i)}, \quad (29)$$

(8) implies (15) for every $\mathbf{y}_i^{(\ell)} \in \mathcal{Y}_\ell$, for all $\ell \in [L]$, and the graph G obtained by SSC-OMP has no false connections. It remains to lower-bound $P[\mathcal{E}^*]$. By the union bound, we have

$$\begin{aligned} P[\mathcal{E}^*] &= 1 - P[\bar{\mathcal{E}}^*] \\ &\geq 1 - P[\bar{\mathcal{E}}_4] \\ &\quad - \sum_{\ell \in [L], i \in [n_\ell]} \left(P[\bar{\mathcal{E}}_5^{(\ell,i)}] + P[\bar{\mathcal{E}}_6^{(\ell,i)}] + P[\bar{\mathcal{E}}_7^{(\ell,i)}] \right) \\ &\quad - \sum_{\substack{\ell \in [L], i \in [n_\ell], \\ s \in [s_{\max}]}} \left(P[\bar{\mathcal{E}}_1^{(\ell,i,s)}] + P[\bar{\mathcal{E}}_2^{(\ell,i,s)}] + P[\bar{\mathcal{E}}_3^{(\ell,i,s)}] \right) \\ &\geq 1 - \sum_{\ell \in [L]} n_\ell (e^{-d_\ell/8} + e^{-m/8}) \\ &\quad - \sum_{\ell \in [L], i \in [n_\ell]} \left(2e^{-c_5 d_\ell} + 2e^{-2m} \right. \\ &\quad \quad \left. + 2e^{-c_2 m} + 2e^{-c_3 d_\ell} + e^{-d_\ell/18} \right) \\ &\quad - \sum_{\substack{\ell \in [L], i \in [n_\ell], \\ s \in [s_{\max}]}} \left(\frac{2}{N^2 s_{\max}} + \frac{2}{N^2 s_{\max}} + \frac{2}{N^2 s_{\max}} \right) \quad (30) \\ &\geq 1 - \frac{6}{N} - \sum_{\ell \in [L]} n_\ell (6e^{-c_d d_\ell} + 5e^{-c_m m}), \end{aligned}$$

where $c_d := \min\{1/18, c_3, c_5\}$, $c_m := \min\{1/8, c_2\}$, and (30) follows from Lemmata 2, 5, 7, 9, and 10.

The proofs of Lemmata 2 and 5 rely on the rotational invariance of the distributions of $\mathbf{r}_{s\parallel}^{(\ell)}$ and $\mathbf{r}_{s\perp}^{(\ell)}$ on \mathcal{S}_ℓ and \mathcal{S}_ℓ^\perp , respectively, which is inherited from the rotational invariance of the distributions of the $\mathbf{x}_j^{(\ell)}$ and $\mathbf{z}_j^{(\ell)}$ characterized next.

Lemma 1. *The distributions of $\mathbf{r}_{s\parallel}^{(\ell)}$ and $\mathbf{r}_{s\perp}^{(\ell)}$ are rotationally invariant on \mathcal{S}_ℓ and \mathcal{S}_ℓ^\perp , respectively, for all $i \in [n_\ell]$, $\ell \in [L]$, i.e., we have $\mathbf{V}^\parallel \mathbf{r}_{s\parallel}^{(\ell)} \sim \mathbf{r}_{s\parallel}^{(\ell)}$ and $\mathbf{V}^\perp \mathbf{r}_{s\perp}^{(\ell)} \sim \mathbf{r}_{s\perp}^{(\ell)}$ for all unitary matrices \mathbf{V}^\parallel and \mathbf{V}^\perp of the form $\mathbf{V}^\parallel = \mathbf{U}^{(\ell)} \mathbf{W}^\parallel \mathbf{U}^{(\ell)T} + \mathbf{P}_\perp$ and $\mathbf{V}^\perp = \mathbf{P}_\parallel + \mathbf{U}_o^{(\ell)} \mathbf{W}^\perp \mathbf{U}_o^{(\ell)T}$, respectively, where $\mathbf{W}^\parallel \in \mathbb{R}^{d_\ell \times d_\ell}$, $\mathbf{W}^\perp \in \mathbb{R}^{(m-d_\ell) \times (m-d_\ell)}$ are unitary and the columns of $\mathbf{U}_o^{(\ell)} \in \mathbb{R}^{m \times (m-d_\ell)}$ form an orthonormal basis for \mathcal{S}_ℓ^\perp .*

Note that the unitary transformations \mathbf{V}^\parallel and \mathbf{V}^\perp act only on \mathcal{S}_ℓ and \mathcal{S}_ℓ^\perp , respectively, and leave components in \mathcal{S}_ℓ^\perp and \mathcal{S}_ℓ , respectively, unchanged.

Proof. We first show that the reduced OMP residual $\mathbf{r}_s^{(\ell)}$ is covariant w.r.t. transformations of the points in \mathcal{Y}_ℓ by a unitary matrix $\mathbf{V} \in \mathbb{R}^{m \times m}$, i.e., we establish that $\mathbf{r}_s(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) = \mathbf{V}\mathbf{r}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)})$, for all $s \in [s_{\max}]$, $i \in [n_\ell]$, $\ell \in [L]$. Combining this covariance property of $\mathbf{r}_s^{(\ell)}$ with the rotational invariance on \mathcal{S}_ℓ and \mathcal{S}_ℓ^\perp of the distributions of $\mathbf{y}_{j\parallel}^{(\ell)} = \mathbf{U}^{(\ell)} \tilde{\mathbf{a}}_j^{(\ell)}$ and $\mathbf{y}_{j\perp}^{(\ell)} = \mathbf{P}_\perp \mathbf{z}_j^{(\ell)}$, respectively, then establishes the desired result.

We prove $\mathbf{r}_s(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) = \mathbf{V}\mathbf{r}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)})$ by induction and start with the inductive step. Assume that after some iteration $s' < s_{\max}$, the index set $\Lambda_{s'}$ corresponding to the transformed data $(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)})$ is identical to

the index set $\Lambda_{s'}$ associated with the original data, i.e., $\Lambda_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) = \Lambda_{s'}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)})$. Using the shorthands $\Lambda_{s'}(\mathbf{V})$ for $\Lambda_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)})$ and $\Lambda_{s'}$ for $\Lambda_{s'}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)})$, it then follows that

$$\begin{aligned} \mathbf{r}_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) &= \left(\mathbf{I} - \mathbf{V}\mathbf{Y}_{\Lambda_{s'}(\mathbf{V})}^{(\ell)} \left(\mathbf{V}\mathbf{Y}_{\Lambda_{s'}(\mathbf{V})}^{(\ell)} \right)^\dagger \right) \mathbf{V}\mathbf{y}_i^{(\ell)} \\ &= \left(\mathbf{I} - \mathbf{V}\mathbf{Y}_{\Lambda_{s'}}^{(\ell)} \left(\mathbf{V}\mathbf{Y}_{\Lambda_{s'}}^{(\ell)} \right)^\dagger \right) \mathbf{V}\mathbf{y}_i^{(\ell)} \\ &= \left(\mathbf{I} - \mathbf{V}\mathbf{Y}_{\Lambda_{s'}}^{(\ell)} \left(\mathbf{Y}_{\Lambda_{s'}}^{(\ell)T} \mathbf{V}^T \mathbf{V}\mathbf{Y}_{\Lambda_{s'}}^{(\ell)} \right)^{-1} \mathbf{Y}_{\Lambda_{s'}}^{(\ell)T} \mathbf{V}^T \right) \mathbf{V}\mathbf{y}_i^{(\ell)} \\ &= \mathbf{V} \left(\mathbf{I} - \mathbf{Y}_{\Lambda_{s'}}^{(\ell)} \left(\mathbf{Y}_{\Lambda_{s'}}^{(\ell)T} \mathbf{Y}_{\Lambda_{s'}}^{(\ell)} \right)^{-1} \mathbf{Y}_{\Lambda_{s'}}^{(\ell)T} \right) \mathbf{y}_i^{(\ell)} \\ &= \mathbf{V}\mathbf{r}_{s'}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}). \quad (31) \end{aligned}$$

For the index $\lambda_{s'+1}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)})$ selected for the \mathbf{V} -transformed data set in iteration $s' + 1$, (31) implies

$$\begin{aligned} \lambda_{s'+1}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) &= \arg \max_{j \in [n_\ell] \setminus (\Lambda_{s'} \cup \{i\})} \left| \left\langle \mathbf{V}\mathbf{y}_j^{(\ell)}, \mathbf{r}_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) \right\rangle \right| \\ &= \arg \max_{j \in [n_\ell] \setminus (\Lambda_{s'} \cup \{i\})} \left| \left\langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_{s'}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) \right\rangle \right| \\ &= \lambda_{s'+1}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}), \quad (32) \end{aligned}$$

i.e., the index selected in iteration $s' + 1$ by operating on the \mathbf{V} -transformed data set is identical to that obtained for the original data set. It remains to establish the base case. This is done by noting that thanks to $\mathbf{r}_0(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) = \mathbf{V}\mathbf{y}_i^{(\ell)}$, we have

$$\begin{aligned} \lambda_1(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) &= \arg \max_{j \in [n_\ell] \setminus (\Lambda_{s'} \cup \{i\})} |\langle \mathbf{V}\mathbf{y}_j^{(\ell)}, \mathbf{V}\mathbf{y}_i^{(\ell)} \rangle| \\ &= \arg \max_{j \in [n_\ell] \setminus (\Lambda_{s'} \cup \{i\})} |\langle \mathbf{y}_j^{(\ell)}, \mathbf{y}_i^{(\ell)} \rangle| \\ &= \lambda_1(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}). \end{aligned}$$

We next establish the rotational invariance of $\mathbf{r}_{s\parallel}^{(\ell)}$ and $\mathbf{r}_{s\perp}^{(\ell)}$. Note that $\mathbf{V}^\parallel \mathbf{y}_j^{(\ell)} = \mathbf{U}^{(\ell)} \mathbf{W}^\parallel \mathbf{U}^{(\ell)T} \mathbf{y}_{j\parallel}^{(\ell)} + \mathbf{y}_{j\perp}^{(\ell)} = \mathbf{U}^{(\ell)} \mathbf{W}^\parallel \tilde{\mathbf{a}}_j^{(\ell)} + \mathbf{y}_{j\perp}^{(\ell)} \sim \mathbf{U}^{(\ell)} \tilde{\mathbf{a}}_j^{(\ell)} + \mathbf{y}_{j\perp}^{(\ell)} = \mathbf{y}_j^{(\ell)}$, $j \in [n_\ell]$, and $\mathbf{V}^\perp \mathbf{y}_j^{(\ell)} = \mathbf{y}_{j\parallel}^{(\ell)} + \mathbf{U}_o^{(\ell)} \mathbf{W}^\perp \mathbf{U}_o^{(\ell)T} \mathbf{y}_{j\perp}^{(\ell)} = \mathbf{y}_{j\parallel}^{(\ell)} + \mathbf{U}_o^{(\ell)} \mathbf{W}^\perp \mathbf{U}_o^{(\ell)T} \mathbf{z}_j^{(\ell)} \sim \mathbf{y}_{j\parallel}^{(\ell)} + \mathbf{z}_j^{(\ell)} = \mathbf{y}_j^{(\ell)}$, $j \in [n_\ell]$, by unitarity of \mathbf{V}^\parallel and \mathbf{V}^\perp . Together with $\mathbf{r}_s(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) = \mathbf{V}\mathbf{r}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)})$, this yields

$$\begin{aligned} \mathbf{r}_{s\parallel}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) &= \mathbf{P}_\parallel \mathbf{r}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) \quad (33) \\ &\sim \mathbf{P}_\parallel \mathbf{r}_s(\mathbf{V}^\parallel \mathbf{y}_i^{(\ell)}, \mathbf{V}^\parallel \mathbf{Y}_{-i}^{(\ell)}) \\ &= \mathbf{P}_\parallel \mathbf{V}^\parallel \mathbf{r}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) \\ &= \underbrace{(\mathbf{P}_\parallel \mathbf{U}^{(\ell)} \mathbf{W}^\parallel \mathbf{U}^{(\ell)T})}_{=\mathbf{U}^{(\ell)} \mathbf{W}^\parallel \mathbf{U}^{(\ell)T} \mathbf{P}_\parallel} \mathbf{r}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) + \underbrace{\mathbf{P}_\parallel \mathbf{P}_\perp}_{=0} \mathbf{r}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) \\ &= \mathbf{V}^\parallel \mathbf{P}_\parallel \mathbf{r}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) \\ &= \mathbf{V}^\parallel \mathbf{r}_{s\parallel}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}), \quad (34) \end{aligned}$$

and establishes $\mathbf{V}^\parallel \mathbf{r}_{s^\parallel} \sim \mathbf{r}_{s^\parallel}$. Repeating the steps leading from (33) to (34) with \mathbf{P}_\perp and \mathbf{V}^\perp in place of \mathbf{P}^\parallel and \mathbf{V}^\parallel , respectively, we analogously obtain $\mathbf{r}_{s^\perp} \sim \mathbf{V}^\perp \mathbf{r}_{s^\perp}$, thereby finishing the proof. \square

We next derive lower bounds on $\mathbb{P}[\mathcal{E}_1^{(\ell,i,s)}]$, $\mathbb{P}[\mathcal{E}_2^{(\ell,i,s)}]$, and $\mathbb{P}[\mathcal{E}_3^{(\ell,i,s)}]$.

Lemma 2. *We have*

$$\begin{aligned} \mathbb{P}[\mathcal{E}_1^{(\ell,i,s)}] &\geq 1 - \frac{2}{N^2 s_{\max}}, & \mathbb{P}[\mathcal{E}_2^{(\ell,i,s)}] &\geq 1 - \frac{2}{N^2 s_{\max}}, \\ \mathbb{P}[\mathcal{E}_3^{(\ell,i,s)}] &\geq 1 - \frac{2}{N^2 s_{\max}}. \end{aligned} \quad (35)$$

Proof. First note that $\mathbf{r}_{s^\parallel}^{(\ell)} / \|\mathbf{r}_{s^\parallel}^{(\ell)}\|_2$ and $\mathbf{r}_{s^\perp}^{(\ell)} / \|\mathbf{r}_{s^\perp}^{(\ell)}\|_2$ are distributed uniformly at random on $\mathbb{S}^{m-1} \cap \mathcal{S}_\ell$ and $\mathbb{S}^{m-1} \cap \mathcal{S}_\ell^\perp$, respectively, as a consequence of rotational invariance (Lemma 1) and normalization [33, Thm. 1.5.6]. This allows us to apply Lemma 3 below with $\mathbf{L} = \mathbf{A}^{(k)}$, $\mathbf{C} = \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)}$, $\mathbf{a} = \mathbf{r}_{s^\parallel}^{(\ell)} / \|\mathbf{r}_{s^\parallel}^{(\ell)}\|_2$, and $\alpha = 4 \log(N^3 s_{\max})$ (note that the condition $\alpha > 12$ is satisfied as $N \geq 3$ and $s_{\max} \geq 1$ by the assumptions of Theorem 1) to get a lower bound on $\mathbb{P}[\mathcal{E}_1^{(\ell,i,s)}]$ according to

$$\begin{aligned} &\mathbb{P} \left[\max_{j \in [n_k]} \left| \left\langle \mathbf{x}_j^{(k)}, \mathbf{r}_{s^\parallel}^{(\ell)} \right\rangle \right| \right. \\ &\quad \left. > 4 \log(N^3 s_{\max}) \frac{\|\mathbf{U}^{(k)T} \mathbf{U}^{(\ell)}\|_F}{\sqrt{d_k} \sqrt{d_\ell}} \|\mathbf{r}_{s^\parallel}^{(\ell)}\|_2 \right] \leq \frac{n_k + 1}{N^3 s_{\max}} \\ &\leq \frac{2n_k}{N^3 s_{\max}}, \end{aligned}$$

for $k \neq \ell$. The desired bound on $\mathbb{P}[\mathcal{E}_1^{(\ell,i,s)}]$ then follows by a union bound over $k \in [L] \setminus \{\ell\}$.

The lower bound on $\mathbb{P}[\mathcal{E}_2^{(\ell,i,s)}]$ is obtained by invoking Lemma 4 below with $\mathbf{a} = \mathbf{r}_{s^\perp}^{(\ell)} / \|\mathbf{r}_{s^\perp}^{(\ell)}\|_2 \in \mathcal{S}_\ell^\perp$ (hence replacing \mathbb{S}^{m-1} by $\mathcal{S}_\ell^\perp \cap \mathbb{S}^{m-1}$), $\mathbf{b} = \mathbf{x}_j^{(k)}$, and $\beta = \sqrt{2 \log(N^3 s_{\max})}$, which yields

$$\mathbb{P} \left[\left| \left\langle \mathbf{r}_{s^\perp}^{(\ell)}, \mathbf{x}_j^{(k)} \right\rangle \right| > \frac{\sqrt{2 \log(N^3 s_{\max})}}{\sqrt{m - d_\ell}} \|\mathbf{r}_{s^\perp}^{(\ell)}\|_2 \right] \leq \frac{2}{N^3 s_{\max}}, \quad (36)$$

for $k \neq \ell$. Again, a union bound over $k \in [L] \setminus \{\ell\}$, $j \in [n_k]$, gives the desired bound on $\mathbb{P}[\mathcal{E}_2^{(\ell,i,s)}]$.

Finally, for $\mathbb{P}[\mathcal{E}_3^{(\ell,i,s)}]$, we set $\mathbf{a} = \mathbf{z}_j^{(k)} / \|\mathbf{z}_j^{(k)}\|_2$, $\mathbf{b} = \mathbf{r}_s^{(\ell)}$, and $\beta = \sqrt{2 \log(N^3 s_{\max})}$ in Lemma 4, and use $\|\mathbf{r}_s^{(\ell)}\|_2 \leq 1 + \|\mathbf{z}_i^{(\ell)}\|_2$, to obtain

$$\begin{aligned} \mathbb{P} \left[\left| \left\langle \mathbf{z}_j^{(k)}, \mathbf{r}_s^{(\ell)} \right\rangle \right| > \frac{\sqrt{2 \log(N^3 s_{\max})}}{\sqrt{m}} \left(1 + \|\mathbf{z}_i^{(\ell)}\|_2 \right) \|\mathbf{z}_j^{(k)}\|_2 \right] \\ \leq \frac{2}{N^3 s_{\max}}, \end{aligned}$$

for all $k \neq \ell$. Again, the lower bound on $\mathbb{P}[\mathcal{E}_3^{(\ell,i,s)}]$ follows from a union bound over $k \in [L] \setminus \{\ell\}$, $j \in [n_k]$.

Lemma 3 (Extracted from the proof of [9, Lem. 7.5]). *Let $\mathbf{a} \in \mathbb{R}^{d_2}$ be distributed uniformly at random on \mathbb{S}^{d_2-1} and let*

the columns of $\mathbf{L} \in \mathbb{R}^{d_1 \times n_1}$ be independent and distributed uniformly on \mathbb{S}^{d_1-1} . Let $\mathbf{C} \in \mathbb{R}^{d_1 \times d_2}$. Then, for $\alpha \geq 12$, we have

$$\mathbb{P} \left[\|\mathbf{L} \mathbf{C} \mathbf{a}\|_\infty \geq \frac{\alpha}{\sqrt{d_1} \sqrt{d_2}} \|\mathbf{C}\|_F \right] \leq (n_1 + 1) e^{-\alpha/4}. \quad (37)$$

Lemma 4 (E.g., [34, Ex. 5.25]). *Let \mathbf{a} be uniformly distributed on \mathbb{S}^{m-1} and fix $\mathbf{b} \in \mathbb{R}^m$. Then, for $\beta \geq 0$, we have*

$$\mathbb{P} \left[|\langle \mathbf{a}, \mathbf{b} \rangle| > \frac{\beta}{\sqrt{m}} \|\mathbf{b}\|_2 \right] \leq 2e^{-\beta^2/2}. \quad \square$$

Next, we lower-bound $\mathbb{P}[\mathcal{E}_7^{(\ell,i)}]$.

Lemma 5. *Let $n_\ell \geq d_\ell + 1$ and $\bar{s} \leq d_\ell$. We have*

$$\begin{aligned} \mathbb{P}[\mathcal{E}_7^{(\ell,i)}] &= \mathbb{P} \left[\left\| \mathbf{r}_{s^\parallel}^{(\ell)} \right\|_2 \right. \\ &\quad \left. > \left\| \mathbf{y}_{i^\parallel}^{(\ell)} \right\|_2 \left(\frac{2}{3} - \sqrt{\frac{3\bar{s} \log((n_\ell - 1)e/\bar{s})}{d_\ell}} \right), \forall s \leq \bar{s} \right] \\ &\geq 1 - e^{-d_\ell/18}. \end{aligned} \quad (38)$$

Proof. The bound obviously holds for $s = 0$ as $\|\mathbf{r}_{0^\parallel}^{(\ell)}\| = \|\mathbf{y}_{i^\parallel}^{(\ell)}\|_2$. For $1 \leq s \leq \bar{s}$ the outline of the proof is as follows. As $\|\mathbf{r}_{s^\parallel}^{(\ell)}\|_2 = \|\mathbf{P}_\parallel (\mathbf{I} - \mathbf{Y}_{\Lambda_s}^{(\ell)} \mathbf{Y}_{\Lambda_s}^{(\ell)\dagger}) \mathbf{y}_i^{(\ell)}\|_2$ is hard to analyze directly owing to statistical dependencies between $\mathbf{y}_i^{(\ell)}$ and the columns of $\mathbf{Y}_{\Lambda_s}^{(\ell)}$ induced by the dependence of Λ_s on $\mathbf{y}_i^{(\ell)}$, we rely on an auxiliary quantity, namely $\|\mathbf{P}_\parallel (\mathbf{I} - \mathbf{Y}_\Gamma^{(\ell)} \mathbf{Y}_\Gamma^{(\ell)\dagger}) \mathbf{y}_i^{(\ell)}\|_2$ for a fixed index set $\Gamma \subset [n_\ell] \setminus \{i\}$ with cardinality satisfying $1 \leq |\Gamma| \leq \bar{s}$. We start the proof by deriving a lower bound $\varphi_{\Gamma'}$ on $\|\mathbf{P}_\parallel (\mathbf{I} - \mathbf{Y}_\Gamma^{(\ell)} \mathbf{Y}_\Gamma^{(\ell)\dagger}) \mathbf{y}_i^{(\ell)}\|_2$ and then show that $\mathcal{E}_7^{(\ell,i)} \supseteq \mathcal{F}_7^{(\ell,i)}$, where

$$\mathcal{F}_7^{(\ell,i)} := \left\{ \varphi_{\Gamma'} > \left\| \mathbf{y}_{i^\parallel}^{(\ell)} \right\|_2 \left(\frac{2}{3} - \sqrt{\frac{3\bar{s} \log((n_\ell - 1)e/\bar{s})}{d_\ell}} \right), \forall \Gamma' \in \mathcal{I} \right\}$$

with

$$\mathcal{I} := \{\Gamma' \subset [n_\ell] \setminus \{i\} : |\Gamma'| = \bar{s}\}, \quad (39)$$

which implies $\mathbb{P}[\mathcal{E}_7^{(\ell,i)}] \geq \mathbb{P}[\mathcal{F}_7^{(\ell,i)}]$. The proof is then completed by establishing a lower bound on $\mathbb{P}[\mathcal{F}_7^{(\ell,i)}]$ using a version of Borell's inequality.

We proceed by lower-bounding $\|\mathbf{P}_\parallel (\mathbf{I} - \mathbf{Y}_\Gamma^{(\ell)} \mathbf{Y}_\Gamma^{(\ell)\dagger}) \mathbf{y}_i^{(\ell)}\|_2$ for $1 \leq |\Gamma| \leq \bar{s}$. Define the orthogonal projection matrices $\mathbf{P}_\Gamma := \mathbf{Y}_\Gamma^{(\ell)} \mathbf{Y}_\Gamma^{(\ell)\dagger}$ and $\mathbf{P}_\Gamma^\parallel := \mathbf{Y}_{\Gamma^\parallel}^{(\ell)} \mathbf{Y}_{\Gamma^\parallel}^{(\ell)\dagger}$ and note that $\mathbf{P}_\parallel \mathbf{P}_\Gamma^\parallel = \mathbf{P}_\Gamma^\parallel = \mathbf{P}_\Gamma^\parallel \mathbf{P}_\parallel$. We now get

$$\begin{aligned} &\left\| \mathbf{P}_\parallel (\mathbf{I} - \mathbf{P}_\Gamma) \mathbf{y}_i^{(\ell)} \right\|_2^2 \\ &= \left\| \mathbf{P}_\parallel (\mathbf{I} - \mathbf{P}_\Gamma^\parallel + \mathbf{P}_\Gamma^\parallel - \mathbf{P}_\Gamma) \mathbf{y}_i^{(\ell)} \right\|_2^2 \end{aligned} \quad (40)$$

$$= \left\| \mathbf{P}_\parallel (\mathbf{I} - \mathbf{P}_\Gamma^\parallel) \mathbf{y}_i^{(\ell)} + (\mathbf{P}_\Gamma^\parallel - \mathbf{P}_\parallel \mathbf{P}_\Gamma) \mathbf{y}_i^{(\ell)} \right\|_2^2 \quad (41)$$

$$\begin{aligned}
&= \left\| \mathbf{P}_{\parallel} (\mathbf{I} - \mathbf{P}_{\Gamma}^{\parallel}) \mathbf{y}_i^{(\ell)} \right\|_2^2 + \left\| (\mathbf{P}_{\Gamma}^{\parallel} - \mathbf{P}_{\parallel} \mathbf{P}_{\Gamma}) \mathbf{y}_i^{(\ell)} \right\|_2^2 \quad (42) \\
&\geq \left\| \mathbf{P}_{\parallel} (\mathbf{I} - \mathbf{P}_{\Gamma}^{\parallel}) \mathbf{y}_i^{(\ell)} \right\|_2^2 \\
&= \left\| \mathbf{y}_{i\parallel}^{(\ell)} - \mathbf{P}_{\Gamma}^{\parallel} \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2^2, \quad (43)
\end{aligned}$$

where the last equality is thanks to $\mathbf{P}_{\parallel} \mathbf{P}_{\Gamma}^{\parallel} = \mathbf{P}_{\Gamma}^{\parallel} \mathbf{P}_{\parallel}$ and the step leading from (41) to (42) follows from

$$\begin{aligned}
&(\mathbf{P}_{\parallel} (\mathbf{I} - \mathbf{P}_{\Gamma}^{\parallel}))^T (\mathbf{P}_{\Gamma}^{\parallel} - \mathbf{P}_{\parallel} \mathbf{P}_{\Gamma}) \\
&= (\mathbf{I} - \mathbf{P}_{\Gamma}^{\parallel}) \mathbf{P}_{\parallel} (\mathbf{P}_{\Gamma}^{\parallel} - \mathbf{P}_{\parallel} \mathbf{P}_{\Gamma}) \\
&= (\mathbf{I} - \mathbf{P}_{\Gamma}^{\parallel}) \mathbf{P}_{\parallel} (\mathbf{P}_{\Gamma}^{\parallel} \mathbf{P}_{\Gamma}^{\parallel} - \mathbf{P}_{\parallel} \mathbf{P}_{\Gamma} \mathbf{P}_{\Gamma}) \\
&= (\mathbf{I} - \mathbf{P}_{\Gamma}^{\parallel}) \mathbf{P}_{\parallel} \mathbf{P}_{\Gamma}^{\parallel} (\mathbf{P}_{\Gamma}^{\parallel} - \mathbf{P}_{\parallel} \mathbf{P}_{\Gamma}) \\
&= \underbrace{(\mathbf{I} - \mathbf{P}_{\Gamma}^{\parallel}) \mathbf{P}_{\Gamma}^{\parallel}}_{=0} (\mathbf{P}_{\Gamma}^{\parallel} - \mathbf{P}_{\parallel} \mathbf{P}_{\Gamma}) = \mathbf{0}.
\end{aligned}$$

Using (40)–(43), we further have

$$\begin{aligned}
\left\| \mathbf{P}_{\parallel} (\mathbf{I} - \mathbf{P}_{\Gamma}) \mathbf{y}_i^{(\ell)} \right\|_2 &\geq \left\| \mathbf{y}_{i\parallel}^{(\ell)} - \mathbf{P}_{\Gamma}^{\parallel} \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2 \quad (44) \\
&\geq \left\| \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2 - \left\| \mathbf{P}_{\Gamma}^{\parallel} \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2 \\
&\geq \underbrace{\left\| \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2 - \left\| \mathbf{P}_{\Gamma}^{\parallel} \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2}_{=: \varphi_{\Gamma'}}, \quad (45)
\end{aligned}$$

for all $\Gamma \subseteq \Gamma' \in \mathcal{I}$, where the second inequality is by the reverse triangle inequality and the third is a consequence of $\mathcal{R}(\mathbf{P}_{\Gamma}^{\parallel}) \subseteq \mathcal{R}(\mathbf{P}_{\Gamma'}^{\parallel}) \subset \mathcal{S}_{\ell}$. It follows from (2) and (44)–(45) that $\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2 = \|\mathbf{P}_{\parallel} (\mathbf{I} - \mathbf{P}_{\Lambda_s}) \mathbf{y}_i^{(\ell)}\|_2 \geq \varphi_{\Gamma'}$ for $\Lambda_s \subset \mathcal{I}$, which implies $\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2 \geq \min_{\Gamma' \in \mathcal{I}} \varphi_{\Gamma'}$, and thus, indeed, $\mathcal{E}_7^{(\ell, i)} \supseteq \mathcal{F}_7^{(\ell, i)}$. It remains to lower-bound $\mathbb{P}[\mathcal{F}_7^{(\ell, i)}]$.

To this end, we first work on the second term in (45) and note that

$$\begin{aligned}
\frac{\left\| \mathbf{P}_{\Gamma'}^{\parallel} \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2}{\left\| \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2} &= \frac{\left\| \mathbf{Y}_{\Gamma'}^{(\ell)} \mathbf{Y}_{\Gamma'}^{(\ell)\dagger} \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2}{\left\| \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2} \\
&= \frac{\left\| \mathbf{U}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} (\mathbf{U}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)})^{\dagger} \mathbf{U}^{(\ell)} \tilde{\mathbf{a}}_i^{(\ell)} \right\|_2}{\left\| \mathbf{U}^{(\ell)} \tilde{\mathbf{a}}_i^{(\ell)} \right\|_2} \\
&= \frac{\left\| \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger} \tilde{\mathbf{a}}_i^{(\ell)} \right\|_2}{\left\| \tilde{\mathbf{a}}_i^{(\ell)} \right\|_2}. \quad (46)
\end{aligned}$$

Since the columns of $\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)}$, i.e., the vectors $\tilde{\mathbf{a}}_j^{(\ell)}$, $j \in \Gamma'$, are i.i.d. and of rotationally invariant distribution, $\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger}$ is the projector onto a subspace of $\mathbb{R}^{d_{\ell}}$ that is \bar{s} -dimensional w.p. 1. In particular, this subspace is distributed uniformly at random on the set of all \bar{s} -dimensional subspaces of $\mathbb{R}^{d_{\ell}}$. Indeed, note that we have, w.p. 1,

$$\begin{aligned}
\mathcal{R}(\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)}) &= \mathcal{R}(\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \text{diag}(1/\|\tilde{\mathbf{a}}_{\gamma_1'}^{(\ell)}\|_2, \dots, 1/\|\tilde{\mathbf{a}}_{\gamma_{\bar{s}}'}^{(\ell)}\|_2)) \\
&\sim \mathcal{R}(\mathbf{G} \text{diag}(1/\|\mathbf{g}_1\|_2, \dots, 1/\|\mathbf{g}_{\bar{s}}\|_2)) \quad (47) \\
&= \mathcal{R}(\mathbf{G}) \quad (48)
\end{aligned}$$

$$= \mathcal{R}(\mathbf{G}(\mathbf{G}^T \mathbf{G})^{-1/2}), \quad (49)$$

where γ_j' , $j \in [\bar{s}]$, denotes the elements of Γ' and $\mathbf{G} = [\mathbf{g}_1 \dots \mathbf{g}_{\bar{s}}] \in \mathbb{R}^{d_{\ell} \times \bar{s}}$ has i.i.d. standard normal random variables as entries. Here, to obtain (47), we used that the $\tilde{\mathbf{a}}_{\gamma_j'}^{(\ell)}/\|\tilde{\mathbf{a}}_{\gamma_j'}^{(\ell)}\|_2$, $j \in [\bar{s}]$, and the $\mathbf{g}_j/\|\mathbf{g}_j\|_2$, $j \in [\bar{s}]$, are all i.i.d. uniform on $\mathbb{S}^{d_{\ell}-1}$, and for (48) and (49) we exploit that $\text{diag}(1/\|\mathbf{g}_1\|_2, \dots, 1/\|\mathbf{g}_{\bar{s}}\|_2)$ and \mathbf{G} , respectively, have full rank w.p. 1. The claim now follows by noting that $\mathbf{G}(\mathbf{G}^T \mathbf{G})^{-1/2}$ is distributed uniformly at random on the set of all orthonormal matrices in $\mathbb{R}^{d_{\ell} \times \bar{s}}$ [35, Thm. 2.2.1 iii]).

Next, we note that conditioning on $\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger}$ does not change the distribution of $\|\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger} \tilde{\mathbf{a}}_i^{(\ell)}\|_2^2 / \|\tilde{\mathbf{a}}_i^{(\ell)}\|_2^2$. To see this, consider $\mathbf{a} \in \mathbb{R}^m$ distributed uniformly at random on $\mathbb{S}^{d_{\ell}-1}$ and choose \mathbf{V} uniformly at random from the set of all orthonormal matrices in $\mathbb{R}^{d_{\ell} \times d_{\ell}}$. Further, let $\mathbf{P}_{\bar{s}}$ be a projector onto an arbitrary, but fixed \bar{s} -dimensional subspace of $\mathbb{R}^{d_{\ell}}$. Then, we have $\|\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger} \tilde{\mathbf{a}}_i^{(\ell)}\|_2^2 / \|\tilde{\mathbf{a}}_i^{(\ell)}\|_2^2 \sim \|\mathbf{P}_{\bar{s}} \mathbf{V}^T \mathbf{a}\|_2^2 \sim \|\mathbf{P}_{\bar{s}} \mathbf{a}\|_2^2$, where the first distributional equivalence follows from $\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger} \sim \mathbf{V} \mathbf{P}_{\bar{s}} \mathbf{V}^T$ (by [35, Thm. 2.2.1 ii]) and the second from $\mathbf{V}^T \mathbf{a} \sim \mathbf{a}$ (by rotational invariance of the distributions of \mathbf{V} and \mathbf{a}).

Now, using $\|\mathbf{P}_{\Gamma'}^{\parallel} \mathbf{y}_{i\parallel}^{(\ell)}\|_2 = \|\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger} \tilde{\mathbf{a}}_i^{(\ell)}\|_2$ and conditioning $\|\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger} \tilde{\mathbf{a}}_i^{(\ell)}\|_2^2$ on $\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger}$ allows us to apply the following version of Borell's inequality to get an upper bound on the second term on the RHS of (45).

Lemma 6 (Extracted from the proof of [9, Lem. 7.5]). *Let $\Sigma \in \mathbb{R}^{d_1 \times d_2}$ be a deterministic matrix and take $\lambda \in \mathbb{R}^{d_2}$ to be distributed uniformly at random on \mathbb{S}^{d_2-1} . Then, we have*

$$\mathbb{P}\left[\|\Sigma \lambda\|_2 - \frac{\|\Sigma\|_F}{\sqrt{d_2}} \geq \varepsilon\right] < e^{-d_2 \varepsilon^2 / (2\sigma_{\max}(\Sigma)^2)}.$$

Setting $\Sigma = \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger}$ and $\lambda = \tilde{\mathbf{a}}_i^{(\ell)} / \|\tilde{\mathbf{a}}_i^{(\ell)}\|_2$ in Lemma 6 and noting that $\|\Sigma\|_F = \|\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger}\|_F = \sqrt{\bar{s}}$ and $\sigma_{\max}(\Sigma) = 1$ yields

$$\mathbb{P}\left[\left\|\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger} \tilde{\mathbf{a}}_i^{(\ell)}\right\|_2 \geq \|\tilde{\mathbf{a}}_i^{(\ell)}\|_2 \left(\sqrt{\frac{\bar{s}}{d_{\ell}}} + \varepsilon\right)\right] < e^{-d_{\ell} \varepsilon^2 / 2}. \quad (50)$$

We now have

$$\begin{aligned}
\mathbb{P}\left[\varphi_{\Gamma'} \geq \left\|\mathbf{y}_{i\parallel}^{(\ell)}\right\|_2 \left(1 - \sqrt{\frac{\bar{s}}{d_{\ell}}} - \varepsilon\right), \forall \Gamma' \in \mathcal{I}\right] \\
= \mathbb{P}\left[\left\|\mathbf{P}_{\parallel \Gamma'} \mathbf{y}_{i\parallel}^{(\ell)}\right\|_2 < \left\|\mathbf{y}_{i\parallel}^{(\ell)}\right\|_2 \left(\sqrt{\frac{\bar{s}}{d_{\ell}}} + \varepsilon\right), \forall \Gamma' \in \mathcal{I}\right] \quad (51)
\end{aligned}$$

$$\geq 1 - \sum_{\Gamma' \in \mathcal{I}} \mathbb{P}\left[\left\|\mathbf{P}_{\parallel \Gamma'} \mathbf{y}_{i\parallel}^{(\ell)}\right\|_2 \geq \left\|\mathbf{y}_{i\parallel}^{(\ell)}\right\|_2 \left(\sqrt{\frac{\bar{s}}{d_{\ell}}} + \varepsilon\right)\right] \quad (52)$$

$$\geq 1 - \binom{n_{\ell} - 1}{\bar{s}} e^{-d_{\ell} \varepsilon^2 / 2} \quad (53)$$

$$\geq 1 - e^{\bar{s} \log((n_{\ell} - 1)e/\bar{s})} e^{-d_{\ell} \varepsilon^2 / 2}, \quad (54)$$

where we used the inequality (44)–(45) to get (51), a union bound for the step leading from (51) to (52), (50) and $|\mathcal{I}| = \binom{n_{\ell} - 1}{\bar{s}}$ to obtain (53) (recall that $\|\mathbf{P}_{\parallel \Gamma'} \mathbf{y}_{i\parallel}^{(\ell)}\|_2 =$

$\|\tilde{\mathbf{A}}_{\Gamma'}^{(\ell)} \tilde{\mathbf{A}}_{\Gamma'}^{(\ell)\dagger} \tilde{\mathbf{a}}_i^{(\ell)}\|_2$ and $\|\mathbf{y}_i^{(\ell)}\|_2 = \|\tilde{\mathbf{a}}_i^{(\ell)}\|_2$, and $\binom{n_\ell-1}{\bar{s}} \leq ((n_\ell-1)e/\bar{s})^{\bar{s}}$ to get (54). Finally, setting

$$\varepsilon = \sqrt{\frac{1}{9} + \frac{2\bar{s} \log((n_\ell-1)e/\bar{s})}{d_\ell}} < \frac{1}{3} + \sqrt{\frac{2\bar{s} \log((n_\ell-1)e/\bar{s})}{d_\ell}}$$

in (54) and noting that

$$\begin{aligned} 1 - \sqrt{\frac{\bar{s}}{d_\ell}} - \varepsilon &> \frac{2}{3} - \sqrt{\frac{\bar{s}}{d_\ell}} - \sqrt{\frac{2\bar{s} \log((n_\ell-1)e/\bar{s})}{d_\ell}} \\ &> \frac{2}{3} - \sqrt{\frac{3\bar{s} \log((n_\ell-1)e/\bar{s})}{d_\ell}}, \end{aligned}$$

where we used $\log((n_\ell-1)e/\bar{s}) \geq 1$ (as $n_\ell-1 \geq d_\ell$ and $\bar{s} \leq d_\ell$) for the last inequality, we have

$$\begin{aligned} \mathbb{P}[\mathcal{F}_7^{(\ell,i)}] &= \mathbb{P}\left[\varphi_{\Gamma'} > \|\mathbf{y}_{i\parallel}^{(\ell)}\|_2 \left(\frac{2}{3} - \sqrt{\frac{3\bar{s} \log((n_\ell-1)e/\bar{s})}{d_\ell}}\right), \right. \\ &\quad \left. \forall \Gamma' \in \mathcal{I}\right] \geq 1 - e^{-d_\ell/18}. \end{aligned}$$

This completes the proof of Lemma 5. \square

We continue by deriving a lower bound on $\mathbb{P}[\mathcal{E}_6^{(\ell,i)}]$.

Lemma 7. *Set $\tilde{a} := \min_{j \in [n_\ell] \setminus \{i\}} \|\mathbf{y}_{j\parallel}^{(\ell)}\|_2$ and suppose that $s_{\max} \leq c_1 d_\ell / \log(e(n_\ell-1)/s_{\max})$ for a numerical constant $c_1 > 0$. Then, we have*

$$\begin{aligned} \mathbb{P}[\mathcal{E}_6^{(\ell,i)}] &= \mathbb{P}\left[\|\mathbf{r}_{s\perp}^{(\ell)}\|_2 \leq \|\mathbf{z}_{i\perp}^{(\ell)}\|_2 + \frac{3\sigma}{\tilde{a}} \|\mathbf{y}_i^{(\ell)}\|_2, \forall s \leq s_{\max}\right] \\ &\geq 1 - 2e^{-c_2 m} - 2e^{-c_3 d_\ell}, \end{aligned} \quad (55)$$

where $c_2, c_3 > 0$ are numerical constants.

Proof. First note that for $s = 0$, $\mathbf{r}_{0\perp}^{(\ell)} = \mathbf{y}_{i\perp}^{(\ell)} = \mathbf{z}_{i\perp}^{(\ell)}$ and the inequality $\|\mathbf{r}_{s\perp}^{(\ell)}\|_2 \leq \|\mathbf{z}_{i\perp}^{(\ell)}\|_2 + (3\sigma/\tilde{a})\|\mathbf{y}_i^{(\ell)}\|_2$ holds trivially. For $1 \leq s \leq s_{\max}$, as in the proof of Lemma 5, we consider fixed index sets $\Gamma \in \mathcal{J}$, with

$$\mathcal{J} := \{\Gamma' \subset [n_\ell] \setminus \{i\} : |\Gamma'| = s \in [s_{\max}]\}, \quad (56)$$

to resolve the issue of statistical dependencies (between the columns of $\mathbf{Y}_{\Lambda_s}^{(\ell)}$ and $\mathbf{y}_i^{(\ell)}$) in $\|\mathbf{r}_{s\perp}^{(\ell)}\|_2 = \|\mathbf{P}_\perp(\mathbf{I} - \mathbf{Y}_{\Lambda_s}^{(\ell)} \mathbf{Y}_{\Lambda_s}^{(\ell)\dagger}) \mathbf{y}_i^{(\ell)}\|_2$. Specifically, this is accomplished by upper-bounding $\|\mathbf{P}_\perp(\mathbf{I} - \mathbf{Y}_\Gamma^{(\ell)} \mathbf{Y}_\Gamma^{(\ell)\dagger}) \mathbf{y}_i^{(\ell)}\|_2$ according to (57) and establishing that the submatrix $\mathbf{Y}_\Gamma^{(\ell)}$ of $\mathbf{Y}^{(\ell)}$ is well-conditioned with high probability for all $\Gamma \in \mathcal{J}$, in particular for the sets $\Lambda_s \in \mathcal{J}$, $s \in [s_{\max}]$, which determine $\mathbf{r}_{s\perp}^{(\ell)}$. This will be accomplished by employing the restricted isometry property (RIP) [34, Sec. 5.6] and standard concentration of measure results from random matrix theory.

By the triangle inequality and the submultiplicativity of the operator norm, we have, for every $\Gamma \in \mathcal{J}$, that

$$\begin{aligned} &\left\| \mathbf{P}_\perp(\mathbf{I} - \mathbf{Y}_\Gamma^{(\ell)} \mathbf{Y}_\Gamma^{(\ell)\dagger}) \mathbf{y}_i^{(\ell)} \right\|_2 \\ &\leq \left\| \mathbf{y}_{i\perp}^{(\ell)} \right\|_2 + \left\| \mathbf{P}_\perp \mathbf{Y}_\Gamma^{(\ell)} \right\|_{2 \rightarrow 2} \left\| \mathbf{Y}_\Gamma^{(\ell)\dagger} \right\|_{2 \rightarrow 2} \left\| \mathbf{y}_i^{(\ell)} \right\|_2 \end{aligned} \quad (57)$$

$$= \left\| \mathbf{z}_{i\perp}^{(\ell)} \right\|_2 + \frac{1}{\sigma_{\min}(\mathbf{Y}_\Gamma^{(\ell)})} \left\| \mathbf{P}_\perp \mathbf{Z}_\Gamma^{(\ell)} \right\|_{2 \rightarrow 2} \left\| \mathbf{y}_i^{(\ell)} \right\|_2 \quad (58)$$

$$\leq \left\| \mathbf{z}_{i\perp}^{(\ell)} \right\|_2 + \frac{1}{\sigma_{\min}(\tilde{\mathbf{A}}_\Gamma^{(\ell)})} \left\| \mathbf{Z}_\Gamma^{(\ell)} \right\|_{2 \rightarrow 2} \left\| \mathbf{y}_i^{(\ell)} \right\|_2, \quad (59)$$

where we used $\sigma_{\min}(\mathbf{Y}_\Gamma^{(\ell)}) = 1/\|\mathbf{Y}_\Gamma^{(\ell)\dagger}\|_{2 \rightarrow 2}$ [34, Sec. 5.2.1] to get (58) and $\sigma_{\min}(\mathbf{Y}_\Gamma^{(\ell)}) \geq \sigma_{\min}(\tilde{\mathbf{A}}_\Gamma^{(\ell)}) > 0$ w.p. 1 (where the first inequality is a consequence of $\|\mathbf{Y}_\Gamma^{(\ell)} \mathbf{v}\|_2 \geq \|\mathbf{P}_\perp \mathbf{Y}_\Gamma^{(\ell)} \mathbf{v}\|_2 = \|\tilde{\mathbf{A}}_\Gamma^{(\ell)} \mathbf{v}\|_2$, for all $\mathbf{v} \in \mathbb{R}^s$, and the second stems from the fact that $\tilde{\mathbf{A}}_\Gamma^{(\ell)}$ has full column rank w.p. 1) to get (59). Denote the elements of Γ by γ_j , $j \in [s]$. We continue by decomposing $\tilde{\mathbf{A}}_\Gamma^{(\ell)}$ according to $\tilde{\mathbf{A}}_\Gamma^{(\ell)} = \mathbf{E}_\Gamma \mathbf{D}_\Gamma$, where $\mathbf{E} := [\tilde{\mathbf{a}}_1^{(\ell)}/\|\tilde{\mathbf{a}}_1^{(\ell)}\|_2 \quad \tilde{\mathbf{a}}_2^{(\ell)}/\|\tilde{\mathbf{a}}_2^{(\ell)}\|_2 \quad \dots \quad \tilde{\mathbf{a}}_s^{(\ell)}/\|\tilde{\mathbf{a}}_s^{(\ell)}\|_2]$, and $\mathbf{D}_\Gamma := \text{diag}(\|\tilde{\mathbf{a}}_{\gamma_1}^{(\ell)}\|_2, \|\tilde{\mathbf{a}}_{\gamma_2}^{(\ell)}\|_2, \dots, \|\tilde{\mathbf{a}}_{\gamma_s}^{(\ell)}\|_2)$. Note that the columns of \mathbf{E} are distributed i.i.d. uniformly at random on $\mathbb{S}^{d_\ell-1}$ and $\sigma_{\min}(\mathbf{D}_\Gamma) \geq \tilde{a}$. We next establish that $\tilde{\mathbf{A}}_\Gamma^{(\ell)}$ and $\mathbf{Z}^{(\ell)}$ satisfy the RIP with high probability, which will then allow us to bound $\sigma_{\min}(\tilde{\mathbf{A}}_\Gamma^{(\ell)})$ and $\sigma_{\max}(\mathbf{Z}_\Gamma^{(\ell)})$, respectively, in (59), for all $\Gamma \in \mathcal{J}$.

We start by recalling the definition of the RIP.

Definition 2. *A matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ satisfies the RIP of order $p \geq 1$ if there exists $\delta_p > 0$ such that*

$$(1 - \delta_p) \|\mathbf{v}\|_2^2 \leq \|\mathbf{A}\mathbf{v}\|_2^2 \leq (1 + \delta_p) \|\mathbf{v}\|_2^2 \quad (60)$$

holds for all $\mathbf{v} \in \mathbb{R}^n$ with $\|\mathbf{v}\|_0 \leq p$. The smallest number $\delta_p = \delta_p(\mathbf{A})$ satisfying (60) is called the restricted isometry constant of \mathbf{A} .

If $\mathbf{A} \in \mathbb{R}^{m \times n}$ satisfies the RIP of order p , it follows from [34, Lem. 5.36] that for $\delta \in [\delta_p, 1]$,

$$\begin{aligned} 1 - \delta &\leq \sigma_{\min}(\mathbf{A}_\mathcal{T}) \leq \sigma_{\max}(\mathbf{A}_\mathcal{T}) \leq 1 + \delta, \\ &\text{for all } \mathcal{T} \subseteq [n] \text{ with } |\mathcal{T}| \leq p. \end{aligned} \quad (61)$$

By [34, Ex. 5.25] the rows of $(\sqrt{m}/\sigma) \mathbf{Z}_{-i}^{(\ell)} \in \mathbb{R}^{m \times (n_\ell-1)}$ are independent sub-gaussian isotropic random vectors [34, Def. 5.19, Def. 5.22], and by [34, Ex. 5.25] the columns of $\sqrt{d_\ell} \mathbf{E}_{-i} \in \mathbb{R}^{d_\ell \times (n_\ell-1)}$ are independent sub-gaussian isotropic random vectors with ℓ_2 -norm $\sqrt{d_\ell}$ a.s. We can therefore apply the next lemma to show that $(\sqrt{m}/\sigma) \mathbf{Z}_{-i}^{(\ell)}$ and $\sqrt{d_\ell} \mathbf{E}_{-i}$ satisfy the RIP for suitable p and δ with high probability. This will then allow us to bound $\sigma_{\min}(\tilde{\mathbf{A}}_\Gamma^{(\ell)})$ and $\sigma_{\max}(\mathbf{Z}_\Gamma^{(\ell)})$ for all $\Gamma \in \mathcal{J}$.

Lemma 8 [34, Thm. 5.65]. *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ be a random matrix with independent sub-gaussian isotropic random vectors as rows or independent sub-gaussian isotropic random vectors with ℓ_2 -norm \sqrt{m} a.s. as columns. Select p with $1 \leq p \leq n$ and let $\delta \in (0, 1)$. If*

$$m \geq C\delta^{-2} p \log(en/p),$$

then, w.p. at least $1 - 2e^{-c\delta^2 m}$, the normalized matrix $\bar{\mathbf{A}} := (1/\sqrt{m})\mathbf{A}$ satisfies the RIP of order p with $\delta_p(\bar{\mathbf{A}}) \leq \delta$. Here, the constants $c, C > 0$ depend only on the sub-gaussian norm⁴ of the rows or columns of \mathbf{A} .

⁴ The sub-gaussian norm of a random variable X is defined as $\|X\|_{\Psi_2} := \sup_{p \geq 1} p^{-1/2} (\mathbb{E}[|X|^p])^{1/p}$ [34, Def. 5.7].

By Lemma 8 with $\mathbf{A} = (\sqrt{m}/\sigma)\mathbf{Z}_{-i}^{(\ell)}$ and $\delta = 1/2$, and (61), if $s_{\max} \leq C_2 m / \log(e(n_\ell - 1)/s_{\max})$, there exist constants $c_2, C_2 > 0$ such that

$$\begin{aligned} \mathbb{P}\left[\frac{1}{2}\sigma \leq \sigma_{\min}(\mathbf{Z}_\Gamma^{(\ell)}) \leq \sigma_{\max}(\mathbf{Z}_\Gamma^{(\ell)}) \leq \frac{3}{2}\sigma, \forall \Gamma \in \mathcal{J}\right] \\ \geq 1 - 2e^{-c_2 m}. \end{aligned} \quad (62)$$

Setting $\mathbf{A} = \sqrt{d_\ell}\mathbf{E}_{-i}$ and $\delta = 1/2$ in Lemma 8, we have similarly

$$\mathbb{P}\left[\frac{1}{2} \leq \sigma_{\min}(\mathbf{E}_\Gamma) \leq \sigma_{\max}(\mathbf{E}_\Gamma) \leq \frac{3}{2}, \forall \Gamma \in \mathcal{J}\right] \geq 1 - 2e^{-c_3 d_\ell}, \quad (63)$$

if $s_{\max} \leq C_3 d_\ell / \log(e(n_\ell - 1)/s_{\max})$, for constants $c_3, C_3 > 0$. Putting (62) and (63) together, we get (55) as follows

$$\begin{aligned} \mathbb{P}[\mathcal{E}_6^{(\ell, i)}] &= \mathbb{P}\left[\left\|\mathbf{P}_\perp(\mathbf{I} - \mathbf{Y}_{\Lambda_s}^{(\ell)}\mathbf{Y}_{\Lambda_s}^{(\ell)\dagger})\mathbf{y}_i^{(\ell)}\right\|_2 \right. \\ &\quad \left. \leq \left\|\mathbf{z}_{i\perp}^{(\ell)}\right\|_2 + \frac{3\sigma}{\tilde{a}}\left\|\mathbf{y}_i^{(\ell)}\right\|_2\right] \\ &\geq \mathbb{P}\left[\frac{\sigma_{\max}(\mathbf{Z}_{\Lambda_s}^{(\ell)})}{\sigma_{\min}(\tilde{\mathbf{A}}_{\Lambda_s}^{(\ell)})} \leq \frac{3\sigma}{\tilde{a}}\right] \end{aligned} \quad (64)$$

$$\geq \mathbb{P}\left[\frac{\sigma_{\max}(\mathbf{Z}_{\Lambda_s}^{(\ell)})}{\sigma_{\min}(\mathbf{E}_{\Lambda_s})} \leq 3\sigma\right] \quad (65)$$

$$\begin{aligned} &\geq \mathbb{P}\left[\left\{\sigma_{\min}(\mathbf{E}_{\Lambda_s}) \geq \frac{1}{2}\right\} \cap \left\{\sigma_{\max}(\mathbf{Z}_{\Lambda_s}^{(\ell)}) \leq \frac{3}{2}\sigma\right\}\right] \\ &\geq 1 - \mathbb{P}\left[\sigma_{\min}(\mathbf{E}_{\Lambda_s}) < \frac{1}{2}\right] - \mathbb{P}\left[\sigma_{\max}(\mathbf{Z}_{\Lambda_s}^{(\ell)}) > \frac{3}{2}\sigma\right] \\ &\geq 1 - 2e^{-c_2 m} - 2e^{-c_3 d_\ell}, \end{aligned} \quad (66)$$

$$\geq 1 - 2e^{-c_2 m} - 2e^{-c_3 d_\ell}, \quad (67)$$

for all $s \leq s_{\max}$. Here (64) follows from (59) with $\Gamma = \Lambda_s$, (65) is by

$$\frac{1}{\sigma_{\min}(\tilde{\mathbf{A}}_{\Lambda_s}^{(\ell)})} \leq \frac{1}{\tilde{a}\sigma_{\min}(\mathbf{E}_{\Lambda_s})},$$

where we used

$$\begin{aligned} \sigma_{\min}(\tilde{\mathbf{A}}_{\Lambda_s}^{(\ell)}) &= \min_{\mathbf{v} \in \mathbb{R}^s} \|\mathbf{E}_{\Lambda_s} \mathbf{D}_{\Lambda_s} \mathbf{v}\|_2 \\ &= \min_{\mathbf{v} \in \mathbb{R}^s} \left\| \mathbf{E}_{\Lambda_s} \frac{\mathbf{D}_{\Lambda_s} \mathbf{v}}{\|\mathbf{D}_{\Lambda_s} \mathbf{v}\|_2} \right\|_2 \|\mathbf{D}_{\Lambda_s} \mathbf{v}\|_2 \\ &\geq \min_{\mathbf{v} \in \mathbb{R}^s} \left\| \mathbf{E}_{\Lambda_s} \frac{\mathbf{D}_{\Lambda_s} \mathbf{v}}{\|\mathbf{D}_{\Lambda_s} \mathbf{v}\|_2} \right\|_2 \tilde{a} \\ &= \tilde{a}\sigma_{\min}(\mathbf{E}_{\Lambda_s}), \end{aligned}$$

(66) is by a union bound, and (67) follows from (62) and (63) with $\Lambda_s \in \mathcal{J}$ for all $s \leq s_{\max}$. Finally, letting $c_1 := \min\{2C_2, C_3\} \leq \min\{(m/d_\ell)C_2, C_3\}$ (using the assumption $m \geq 2d_{\max}$) ensures that s_{\max} , and thereby $|\Gamma|$, is small enough for both (62) and (63) to hold. This concludes the proof of Lemma 7. \square

We proceed by establishing a lower bound on $\mathbb{P}[\mathcal{E}_5^{(\ell, i)}]$.

Lemma 9. For numerical constants $c_4, c_5 > 0$ it holds that

$$\mathbb{P}[\mathcal{E}_5^{(\ell, i)}] = \mathbb{P}\left[\max_{j \in [n_\ell] \setminus (\Lambda_s \cup \{i\})} \left| \langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_s^{(\ell)} \rangle \right| \right]$$

$$\begin{aligned} &\geq \left(1 - \frac{c_4 + 1}{\sqrt{\rho_\ell}}\right) \frac{\|\mathbf{r}_s^{(\ell)}\|_2}{\sqrt{d_\ell}} \\ &\quad - \sigma \left(\frac{1}{\sqrt{m}} + \frac{2}{\sqrt{n_\ell - 1}} \right) \|\mathbf{r}_s^{(\ell)}\|_2 \\ &\geq 1 - 2e^{-c_5 d_\ell} - 2e^{-m/2}. \end{aligned} \quad (68)$$

Proof. We first lower-bound the maximum in (68) by a term proportional to $\|\mathbf{Y}_{-i}^{(\ell)T} \mathbf{r}_s^{(\ell)}\|_2$ and then establish (68) by leveraging standard bounds on the singular values of random matrices. We have

$$\begin{aligned} \max_{j \in [n_\ell] \setminus (\Lambda_s \cup \{i\})} \left| \langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_s^{(\ell)} \rangle \right| &= \max_{j \in [n_\ell] \setminus \{i\}} \left| \langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_s^{(\ell)} \rangle \right| \\ &= \left\| \mathbf{Y}_{-i}^{(\ell)T} \mathbf{r}_s^{(\ell)} \right\|_\infty \\ &\geq \frac{\left\| \mathbf{Y}_{-i}^{(\ell)T} \mathbf{r}_s^{(\ell)} \right\|_2}{\sqrt{n_\ell - 1}} \\ &= \frac{\left\| \mathbf{A}_{-i}^{(\ell)T} \mathbf{U}^{(\ell)T} \mathbf{r}_s^{(\ell)} + \mathbf{Z}_{-i}^{(\ell)T} \mathbf{r}_s^{(\ell)} \right\|_2}{\sqrt{n_\ell - 1}} \\ &\geq \frac{\left\| \mathbf{A}_{-i}^{(\ell)T} \mathbf{U}^{(\ell)T} \mathbf{r}_s^{(\ell)} \right\|_2}{\sqrt{n_\ell - 1}} - \frac{\left\| \mathbf{Z}_{-i}^{(\ell)T} \mathbf{r}_s^{(\ell)} \right\|_2}{\sqrt{n_\ell - 1}} \\ &\geq \frac{\sigma_{\min}(\mathbf{A}_{-i}^{(\ell)T})}{\sqrt{n_\ell - 1}} \|\mathbf{r}_s^{(\ell)}\|_2 - \frac{\sigma_{\max}(\mathbf{Z}_{-i}^{(\ell)T})}{\sqrt{n_\ell - 1}} \|\mathbf{r}_s^{(\ell)}\|_2, \end{aligned} \quad (69)$$

where the first equality is thanks to orthogonality of $\mathbf{r}_s^{(\ell)}$ and $\mathbf{y}_j^{(\ell)}$, for all $j \in \Lambda_s$, the first inequality follows from $\|\mathbf{v}\|_2 \leq \sqrt{n_\ell - 1} \|\mathbf{v}\|_\infty$, for all $\mathbf{v} \in \mathbb{R}^{n_\ell - 1}$, and the second inequality is by the reverse triangle inequality.

Noting that $\sqrt{d_\ell} \mathbf{A}_{-i}^{(\ell)T}$ is a $(n_\ell - 1) \times d_\ell$ matrix whose rows are independent isotropic subgaussian random vectors (as defined in [34, Def. 5.19, Def. 5.22]), it follows from [34, Thm. 5.39] (see Theorem 5 in Appendix E) that

$$\mathbb{P}\left[\sqrt{d_\ell} \sigma_{\min}(\mathbf{A}_{-i}^{(\ell)T}) < \sqrt{n_\ell - 1} - c_4 \sqrt{d_\ell} - t\right] < 2e^{-c_5 t^2}, \quad (70)$$

where the constants $c_4, c_5 > 0$ depend only on the subgaussian norm of the rows of $\mathbf{A}_{-i}^{(\ell)T}$. Setting $t = \sqrt{d_\ell}$ in (70), we get

$$\begin{aligned} \mathbb{P}\left[\frac{\sigma_{\min}(\mathbf{A}_{-i}^{(\ell)T})}{\sqrt{n_\ell - 1}} \|\mathbf{r}_s^{(\ell)}\|_2 < \left(1 - \frac{c_4 + 1}{\sqrt{\rho_\ell}}\right) \frac{\|\mathbf{r}_s^{(\ell)}\|_2}{\sqrt{d_\ell}}\right] \\ < 2e^{-c_5 d_\ell}. \end{aligned} \quad (71)$$

Since $(\sqrt{m}/\sigma)\mathbf{Z}_{-i}^{(\ell)T}$ is a $(n_\ell - 1) \times m$ matrix with i.i.d. standard normal entries, it follows from [34, Cor. 5.35] (see Corollary 6 in Appendix E) that

$$\mathbb{P}\left[\frac{\sqrt{m}}{\sigma} \sigma_{\max}(\mathbf{Z}_{-i}^{(\ell)T}) > \sqrt{n_\ell - 1} + \sqrt{m} + t\right] < 2e^{-t^2/2}. \quad (72)$$

Setting $t = \sqrt{m}$ in (72), we obtain

$$\mathbb{P} \left[\frac{\sigma_{\max}(\mathbf{Z}_{-i}^{(\ell)T})}{\sqrt{n_\ell - 1}} \|\mathbf{r}_s^{(\ell)}\|_2 > \sigma \left(\frac{1}{\sqrt{m}} + \frac{2}{\sqrt{n_\ell - 1}} \right) \|\mathbf{r}_s^{(\ell)}\|_2 \right] < 2e^{-m/2}. \quad (73)$$

The claim in Lemma 9 now follows by lower-bounding the first term in (69) using (71), by upper-bounding the second term in (69) using (73), and by application of a union bound. \square

Finally, we derive a lower bound on $\mathbb{P}[\mathcal{E}_4]$.

Lemma 10. *We have*

$$\mathbb{P}[\mathcal{E}_4] \geq 1 - \sum_{\ell \in [L]} n_\ell (e^{-d_\ell/8} + e^{-m/8}). \quad (74)$$

Proof. The proof is effected by applying the following well-known concentration result.

Theorem 4 ([36]). *Let $f: \mathbb{R}^m \rightarrow \mathbb{R}$ be a Lipschitz function with Lipschitz constant K , i.e., $|f(\mathbf{a}) - f(\mathbf{b})| \leq K\|\mathbf{a} - \mathbf{b}\|_2$, for all $\mathbf{a}, \mathbf{b} \in \mathbb{R}^m$. Let $\mathbf{z} \in \mathbb{R}^m$ be a $\mathcal{N}(\mathbf{0}, \mathbf{I}_m)$ vector. Then, for $t \geq 0$, we have*

$$\mathbb{P}[f(\mathbf{z}) - \mathbb{E}[f(\mathbf{z})] > t] \leq e^{-t^2/(2K^2)}. \quad (75)$$

The functions $f(\mathbf{z}) = \|\mathbf{z}\|_2$ and $f_{\parallel}(\mathbf{z}) = \|\mathbf{P}_{\parallel}\mathbf{z}\|_2$ both have Lipschitz constant $K = 1$ ($|f_{\parallel}(\mathbf{a}) - f_{\parallel}(\mathbf{b})| = \|\mathbf{P}_{\parallel}\mathbf{a}\|_2 - \|\mathbf{P}_{\parallel}\mathbf{b}\|_2 \leq \|\mathbf{P}_{\parallel}(\mathbf{a} - \mathbf{b})\|_2 \leq \|\mathbf{a} - \mathbf{b}\|_2$, for all $\mathbf{a}, \mathbf{b} \in \mathbb{R}^m$, by the reverse triangle inequality and $\|\mathbf{P}_{\parallel}\|_{2 \rightarrow 2} = 1$). For $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$, we get by Jensen's inequality $\mathbb{E}[\|\mathbf{z}\|_2] \leq \sqrt{\mathbb{E}[\|\mathbf{z}\|_2^2]} = \sqrt{m}$ and $\mathbb{E}[\|\mathbf{P}_{\parallel}\mathbf{z}\|_2] \leq \sqrt{\mathbb{E}[\|\mathbf{P}_{\parallel}\mathbf{z}\|_2^2]} = \sqrt{d_\ell}$ (where the equality follows from $\mathbf{P}_{\parallel}\mathbf{z} = \mathbf{U}^{(\ell)}\mathbf{U}^{(\ell)T}\mathbf{z}$ and the fact that $\mathbf{U}^{(\ell)T}\mathbf{z}$ is $\mathcal{N}(0, \mathbf{I}_{d_\ell})$ -distributed). Noting that $\mathbf{z}_i^{(\ell)} \sim (\sigma/\sqrt{m})\mathbf{z}$, application of (75) to \mathbf{z} and $\mathbf{P}_{\parallel}\mathbf{z}$ yields

$$\mathbb{P} \left[\left\| \mathbf{z}_i^{(\ell)} \right\|_2 > \frac{3}{2}\sigma \right] \leq e^{-m/8} \quad \text{and} \\ \mathbb{P} \left[\left\| \mathbf{P}_{\parallel}\mathbf{z}_i^{(\ell)} \right\|_2 > \frac{3\sqrt{d_\ell}}{2\sqrt{m}}\sigma \right] \leq e^{-d_\ell/8},$$

where we set $t = \sqrt{m}/2$ and $t = \sqrt{d_\ell}/2$, respectively. A union bound over $\ell \in [L]$, $i \in [n_\ell]$, yields the desired lower bound (74). \square

APPENDIX C PROOF OF THEOREM 2

Most steps of the proof of Theorem 2 are almost identical to those in the proof of Theorem 1. We therefore elaborate only on the arguments the proofs differ in significantly.

Analogously to the proof for SSC-OMP, we will henceforth work with the ‘‘reduced MP’’ algorithm which, for the representation of $\mathbf{y}_i^{(\ell)}$ selects elements from the reduced dictionary $\mathcal{Y}_\ell \setminus \{\mathbf{y}_i^{(\ell)}\}$ only, instead of the full dictionary $\mathcal{Y} \setminus \{\mathbf{y}_i^{(\ell)}\}$. The justification for relying on reduced MP to establish the desired result is identical to that for OMP. Throughout the proof the

residual of the reduced MP algorithm will be denoted by $\mathbf{q}_s^{(\ell)}$ and the number of iterations actually performed when a stopping condition has been met by s_a . As in the OMP case, for expositional convenience, the quantities $\mathbf{q}_s^{(\ell)}$ and s_a do not reflect the dependence on the index i of the data point $\mathbf{y}_i^{(\ell)}$.

Note that the selection rules for OMP and MP are equivalent in the following sense. As the OMP residual $\mathbf{r}_s^{(\ell)}$ is orthogonal to $\mathbf{y}_j^{(\ell)}$, for all $j \in \Lambda_s$, we can replace $\max_{j \in [N] \setminus (\Lambda_s \cup \{i\})} |\langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_s \rangle|$ on the RHS of (15) by $\max_{j \in [N] \setminus \{i\}} |\langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_s \rangle|$, i.e., we can take the maximization over $j \in [N] \setminus \{i\}$ as in MP. We therefore need to show that (15) with \mathbf{r}_s replaced by \mathbf{q}_s and $\max_{j \in [N] \setminus (\Lambda_s \cup \{i\})}$ replaced by $\max_{j \in [N] \setminus \{i\}}$ holds for all MP iterations, and for every $\mathbf{y}_i^{(\ell)} \in \mathcal{Y}_\ell$, $\ell \in [L]$, w.p. at least P^* .

Next, we systematically revisit the events $\mathcal{E}_1^{(\ell,i,s)} - \mathcal{E}_7^{(\ell,i)}$ and adapt the corresponding bounds for MP where needed. Recall that the bounds on $\mathbb{P}[\mathcal{E}_1^{(\ell,i,s)}]$ and $\mathbb{P}[\mathcal{E}_2^{(\ell,i,s)}]$ in Lemma 2 rely on the rotational invariance—as expressed in Lemma 1—of the distributions of $\mathbf{r}_{s\parallel}^{(\ell)}$ and $\mathbf{r}_{s\perp}^{(\ell)}$, respectively. As the residual update rule (5) for MP differs from that for OMP in (2), we need to establish rotational invariance for $\mathbf{q}_{s\parallel}^{(\ell)}$ and $\mathbf{q}_{s\perp}^{(\ell)}$, which will be done in Lemma 11 below. The bounds on $\mathbb{P}[\mathcal{E}_3^{(\ell,i,s)}]$, $\mathbb{P}[\mathcal{E}_4]$, and $\mathbb{P}[\mathcal{E}_5^{(\ell,i)}]$ in Lemmata 2, 10, and 9, respectively, do not depend on a particular property of $\mathbf{r}_s^{(\ell)}$ apart from $\|\mathbf{r}_s^{(\ell)}\|_2 \leq \|\mathbf{y}_i^{(\ell)}\|_2$ in the case of $\mathbb{P}[\mathcal{E}_3^{(\ell,i,s)}]$ (we have $\|\mathbf{q}_s^{(\ell)}\|_2 \leq \|\mathbf{y}_i^{(\ell)}\|_2$ as a consequence of [18, Eq. 13]). Thus, the bounds on $\mathbb{P}[\mathcal{E}_1^{(\ell,i,s)}] - \mathbb{P}[\mathcal{E}_5^{(\ell,i)}]$ continue to hold for $\mathbf{r}_s^{(\ell)}$, $\mathbf{r}_{s\parallel}^{(\ell)}$, and $\mathbf{r}_{s\perp}^{(\ell)}$ in $\mathcal{E}_1^{(\ell,i,s)} - \mathcal{E}_5^{(\ell,i)}$ replaced by $\mathbf{q}_s^{(\ell)}$, $\mathbf{q}_{s\parallel}^{(\ell)}$, and $\mathbf{q}_{s\perp}^{(\ell)}$, respectively, and we readily get the upper bound (17) on the LHS of (15) and the lower bound (23) on the RHS of (15). The bounds on $\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2$ and $\|\mathbf{r}_{s\perp}^{(\ell)}\|_2$ in $\mathcal{E}_6^{(\ell,i)}$ and $\mathcal{E}_7^{(\ell,i)}$, respectively, in the proof of Theorem 1 require more work. In particular, as the stopping behavior of MP is different from that of OMP, we need the corresponding bounds on $\|\mathbf{q}_{s\perp}^{(\ell)}\|_2$ and $\|\mathbf{q}_{s\parallel}^{(\ell)}\|_2$ to hold for all MP iterations $s \in [s_a]$ and for a maximum sparsity level of p_{\max} . This will be accomplished by deriving an upper bound on $\|\mathbf{q}_{s\perp}^{(\ell)}\|_2$ and a lower bound on $\|\mathbf{q}_{s\parallel}^{(\ell)}\|_2$ leading to suitably modified events $\tilde{\mathcal{E}}_6^{(\ell,i)}$ and $\tilde{\mathcal{E}}_7^{(\ell,i)}$, respectively, as defined below. Specifically, the resulting upper bound on $\|\mathbf{q}_{s\perp}^{(\ell)}\|_2$ is slightly weaker than that on $\|\mathbf{r}_{s\perp}^{(\ell)}\|_2$ in $\mathcal{E}_6^{(\ell,i)}$, but exhibits the same scaling behavior in σ , whereas the resulting lower bound on $\|\mathbf{q}_{s\parallel}^{(\ell)}\|_2$ is identical to the one on $\|\mathbf{r}_{s\parallel}^{(\ell)}\|_2$ in $\mathcal{E}_7^{(\ell,i)}$.

We proceed by introducing the modified events

$$\tilde{\mathcal{E}}_6^{(\ell,i)} := \left\{ \left\| \mathbf{q}_{s\perp}^{(\ell)} \right\|_2 \leq \left\| \mathbf{z}_{i\perp}^{(\ell)} \right\|_2 + \frac{6\sigma}{\tilde{a}} \left\| \mathbf{y}_i^{(\ell)} \right\|_2, \quad \forall s \leq s_a \right\}, \quad \text{and} \\ \tilde{\mathcal{E}}_7^{(\ell,i)} := \left\{ \left\| \mathbf{q}_{s\parallel}^{(\ell)} \right\|_2 > \left\| \mathbf{y}_{i\parallel}^{(\ell)} \right\|_2 \left(\frac{2}{3} - \sqrt{\frac{3p_{\max} \log((n_\ell - 1)e/p_{\max})}{d_\ell}} \right), \quad \forall s \leq s_a \right\}, \quad (76)$$

where $\tilde{a} := \min_{j \in [n_\ell] \setminus \{i\}} \|\mathbf{y}_{j\parallel}^{(\ell)}\|_2$, and deriving lower bounds

on $P[\tilde{\mathcal{E}}_6^{(\ell,i)}]$ and $P[\tilde{\mathcal{E}}_7^{(\ell,i)}]$ in Lemmata 12 and 13, respectively. These lower bounds will turn out to be identical to those for SSC-OMP. Although the corresponding proofs rely on arguments similar in spirit to those used for OMP, the technical details are dissimilar enough to warrant detailed presentation.

We continue by establishing the bounds on $\|\mathbf{q}_{s\parallel}^{(\ell)}\|_2$ and $\|\mathbf{q}_{s\perp}^{(\ell)}\|_2$, as announced. Using the assumptions $m \geq 2d_{\max}$, $\sigma \leq 1/2$, and $p_{\max} \leq \min_{\ell \in [L]} \{c_s d_\ell / \log((n_\ell - 1)e/p_{\max})\}$, we have on $\mathcal{E}_4 \cap \tilde{\mathcal{E}}_6^{(\ell,i)}$ that

$$\|\mathbf{q}_{s\perp}^{(\ell)}\|_2 \leq \frac{3}{2}\sigma + 6\sigma \frac{1 + \frac{3}{2}\sigma}{1 - \frac{3}{2}\frac{\sqrt{d_\ell}}{\sqrt{m}}\sigma} \leq \sigma(15 + 20\sigma) \quad (77)$$

and, by repeating the steps in (28), we get $\|\mathbf{q}_{s\parallel}^{(\ell)}\|_2 > 1/20$ on $\tilde{\mathcal{E}}_7^{(\ell,i)}$. It therefore follows that (15) for MP is implied by (8) with $c(\sigma) = 17 + 23\sigma$ on $\tilde{\mathcal{E}}^* := \bigcap_{\ell,i,s} (\mathcal{E}_1^{(\ell,i,s)} \cap \mathcal{E}_2^{(\ell,i,s)} \cap \mathcal{E}_3^{(\ell,i,s)} \cap \mathcal{E}_4 \cap \mathcal{E}_5^{(\ell,i)} \cap \tilde{\mathcal{E}}_6^{(\ell,i)} \cap \tilde{\mathcal{E}}_7^{(\ell,i)})$. The proof is completed by lower-bounding $P[\tilde{\mathcal{E}}^*]$ via a union bound.

We proceed by establishing the rotational invariance properties of $\mathbf{q}_{s\parallel}^{(\ell)}$ and $\mathbf{q}_{s\perp}^{(\ell)}$ needed to establish the lower bounds on $P[\tilde{\mathcal{E}}_6^{(\ell,i)}]$ and $P[\tilde{\mathcal{E}}_7^{(\ell,i)}]$ in Lemmata 12 and 13, respectively.

Lemma 11. *The distributions of $\mathbf{q}_{s\parallel}^{(\ell)}$ and $\mathbf{q}_{s\perp}^{(\ell)}$ are rotationally invariant on \mathcal{S}_ℓ and \mathcal{S}_ℓ^\perp , respectively, i.e., for unitary transformations $\mathbf{V}^\parallel, \mathbf{V}^\perp \in \mathbb{R}^{m \times m}$ of the form specified in Lemma 1, we have $\mathbf{V}^\parallel \mathbf{q}_{s\parallel}^{(\ell)} \sim \mathbf{q}_{s\parallel}^{(\ell)}$ and $\mathbf{V}^\perp \mathbf{q}_{s\perp}^{(\ell)} \sim \mathbf{q}_{s\perp}^{(\ell)}$.*

Proof. The arguments employed in this proof are similar to those in the proof of the corresponding result for OMP, Lemma 1, but the structure of the proof differs as the MP residual $\mathbf{q}_s^{(\ell)}$ can only be expressed recursively, i.e., as a function of previous residuals. In contrast, the reduced OMP residual $\mathbf{r}_s^{(\ell)}$ can be written as the projection of $\mathbf{y}_i^{(\ell)}$ onto the orthogonal complement of the span of the data points indexed by Λ_s . Throughout the proof, $\omega_s(\mathbf{x}, \mathbf{D})$ denotes the index obtained by the MP algorithm in iteration s when applied to \mathbf{x} with the columns of \mathbf{D} as dictionary elements, and $\mathbf{q}_s(\mathbf{x}, \mathbf{D})$ is the corresponding residual.

We first establish results analogous to (31) and (32). Again, the proof is effected through induction. We start with the inductive step and assume that

$$\mathbf{q}_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) = \mathbf{V}\mathbf{q}_{s'}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) \quad (78)$$

for fixed $s' < s_a$, for all unitary matrices $\mathbf{V} \in \mathbb{R}^{m \times m}$. For iteration $s' + 1$ we then have

$$\begin{aligned} \omega_{s'+1}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) &= \arg \max_{j \in [n_\ell] \setminus \{i\}} \left| \left\langle \mathbf{V}\mathbf{y}_j^{(\ell)}, \mathbf{q}_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) \right\rangle \right| \\ &= \arg \max_{j \in [n_\ell] \setminus \{i\}} \left| \left\langle \mathbf{y}_j^{(\ell)}, \mathbf{q}_{s'}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) \right\rangle \right| \\ &= \omega_{s'+1}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}). \end{aligned} \quad (79)$$

Now, using the shorthands $\omega_{s'+1}(\mathbf{V})$ and $\omega_{s'+1}$ for $\omega_{s'+1}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)})$ and $\omega_{s'+1}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)})$, respectively, we get

$$\mathbf{q}_{s'+1}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) = \mathbf{q}_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)})$$

$$\begin{aligned} & - \left\langle \mathbf{q}_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}), \frac{\mathbf{V}\mathbf{y}_{\omega_{s'+1}(\mathbf{V})}^{(\ell)}}{\|\mathbf{V}\mathbf{y}_{\omega_{s'+1}(\mathbf{V})}^{(\ell)}\|_2} \right\rangle \frac{\mathbf{V}\mathbf{y}_{\omega_{s'+1}(\mathbf{V})}^{(\ell)}}{\|\mathbf{V}\mathbf{y}_{\omega_{s'+1}(\mathbf{V})}^{(\ell)}\|_2} \\ &= \mathbf{q}_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) \\ & \quad - \left\langle \mathbf{q}_{s'}(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}), \frac{\mathbf{V}\mathbf{y}_{\omega_{s'+1}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'+1}}^{(\ell)}\|_2} \right\rangle \frac{\mathbf{V}\mathbf{y}_{\omega_{s'+1}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'+1}}^{(\ell)}\|_2} \\ &= \mathbf{V} \left(\mathbf{q}_{s'}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}) \right. \\ & \quad \left. - \left\langle \mathbf{q}_{s'}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}), \frac{\mathbf{y}_{\omega_{s'+1}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'+1}}^{(\ell)}\|_2} \right\rangle \frac{\mathbf{y}_{\omega_{s'+1}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'+1}}^{(\ell)}\|_2} \right) \\ &= \mathbf{V}\mathbf{q}_{s'+1}(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}), \end{aligned} \quad (80)$$

where the second and third equality follow from (79) and (78), respectively. Now, the base case is $\mathbf{q}_0(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) = \mathbf{V}\mathbf{y}_i^{(\ell)} = \mathbf{V}\mathbf{q}_0(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)})$, and we therefore established that

$$\mathbf{q}_s(\mathbf{V}\mathbf{y}_i^{(\ell)}, \mathbf{V}\mathbf{Y}_{-i}^{(\ell)}) = \mathbf{V}\mathbf{q}_s(\mathbf{y}_i^{(\ell)}, \mathbf{Y}_{-i}^{(\ell)}), \quad (81)$$

for all $s \leq s_a$ and all unitary $\mathbf{V} \in \mathbb{R}^{m \times m}$.

Finally, repeating the steps leading from (33) to (34) for \mathbf{q}_s and $\mathbf{q}_{s\parallel}$ instead of \mathbf{r}_s and $\mathbf{r}_{s\parallel}$, respectively, yields the desired result. \square

We continue with the lower bound on $P[\tilde{\mathcal{E}}_7^{(\ell,i)}]$.

Lemma 12. *Let $n_\ell > 1$. We have*

$$\begin{aligned} P[\tilde{\mathcal{E}}_7^{(\ell,i)}] &= P \left[\|\mathbf{q}_{s\parallel}^{(\ell)}\|_2 \right. \\ & \left. > \|\mathbf{y}_{i\parallel}^{(\ell)}\|_2 \left(\frac{2}{3} - \sqrt{\frac{3p_{\max} \log((n_\ell - 1)e/p_{\max})}{d_\ell}} \right), \forall s \leq s_a \right] \\ &\geq 1 - e^{-d_\ell/18}. \end{aligned} \quad (82)$$

Proof. We start by recalling that the reduced MP algorithm decomposes $\mathbf{y}_i^{(\ell)}$ according to (see, e.g., [18])

$$\mathbf{y}_i^{(\ell)} = \sum_{s'=1}^s \left\langle \mathbf{q}_{s'-1}^{(\ell)}, \frac{\mathbf{y}_{\omega_{s'}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'}}^{(\ell)}\|_2} \right\rangle \frac{\mathbf{y}_{\omega_{s'}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'}}^{(\ell)}\|_2} + \mathbf{q}_s^{(\ell)}. \quad (83)$$

Denote by Ω_s the set containing the indices of the points in $\mathcal{Y}_\ell \setminus \{\mathbf{y}_i^{(\ell)}\}$ selected during the first s iterations, i.e.,

$$\Omega_s := \{\omega_{s'} : s' \in [s]\}. \quad (84)$$

Note that Ω_s may contain fewer than s indices as reduced MP may select one or more points (from $\mathcal{Y}_\ell \setminus \{\mathbf{y}_i^{(\ell)}\}$) repeatedly.

With $\mathbf{P}_{\Omega_s}^\parallel := \mathbf{Y}_{\Omega_s}^{(\ell)} \mathbf{Y}_{\Omega_s}^{(\ell)\dagger}$, using (83), we get

$$\begin{aligned} \|\mathbf{P}_{\parallel} \mathbf{q}_s^{(\ell)}\|_2 &\geq \|(\mathbf{I} - \mathbf{P}_{\Omega_s}^\parallel) \mathbf{P}_{\parallel} \mathbf{q}_s^{(\ell)}\|_2 \\ &= \left\| (\mathbf{I} - \mathbf{P}_{\Omega_s}^\parallel) \mathbf{P}_{\parallel} \left(\mathbf{y}_i^{(\ell)} \right) \right\| \end{aligned}$$

$$\begin{aligned}
& - \sum_{s'=1}^s \left\langle \mathbf{q}_{s'-1}^{(\ell)}, \frac{\mathbf{y}_{\omega_{s'}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'}}^{(\ell)}\|_2} \right\rangle \frac{\mathbf{y}_{\omega_{s'}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'}}^{(\ell)}\|_2} \Bigg\|_2 \\
&= \left\| (\mathbf{I} - \mathbf{P}_{\Omega_s}^{\parallel}) \mathbf{y}_i^{(\ell)} \right. \\
&\quad \left. - (\mathbf{I} - \mathbf{P}_{\Omega_s}^{\parallel}) \underbrace{\left(\sum_{s'=1}^s \left\langle \mathbf{q}_{s'-1}^{(\ell)}, \frac{\mathbf{y}_{\omega_{s'}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'}}^{(\ell)}\|_2} \right\rangle \frac{\mathbf{y}_{\omega_{s'}}^{(\ell)}}{\|\mathbf{y}_{\omega_{s'}}^{(\ell)}\|_2} \right)}_{\in \mathcal{R}(\mathbf{P}_{\Omega_s}^{\parallel})} \right\|_2 \\
&= \left\| (\mathbf{I} - \mathbf{P}_{\Omega_s}^{\parallel}) \mathbf{y}_i^{(\ell)} \right\|_2 \\
&\geq \left\| (\mathbf{I} - \mathbf{P}_{\Omega_{s_a}}^{\parallel}) \mathbf{y}_i^{(\ell)} \right\|_2, \tag{85}
\end{aligned}$$

where the last inequality is by $\mathcal{R}(\mathbf{I} - \mathbf{P}_{\Omega_{s_a}}^{\parallel}) \subseteq \mathcal{R}(\mathbf{I} - \mathbf{P}_{\Omega_s}^{\parallel})$, for $s \leq s_a$. The proof is now completed by replacing Ω_{s_a} in (85) by a fixed $\Gamma \in \mathcal{I}$ (recall the definition of \mathcal{I} from (39), with \bar{s} in (39) replaced by p_{\max}) and by lower-bounding $\|(\mathbf{I} - \mathbf{P}_{\Gamma}^{\parallel}) \mathbf{y}_i^{(\ell)}\|_2$ for all $\Gamma \in \mathcal{I}$ as in the proof of Lemma 5 (following the steps starting from (44)). \square

Next, we lower-bound $\mathbb{P}[\tilde{\mathcal{E}}_6^{(\ell,i)}]$.

Lemma 13. *Set $\tilde{a} := \min_{j \in [n_\ell] \setminus \{i\}} \|\mathbf{y}_j^{(\ell)}\|_2$ and assume that $p_{\max} \leq c_1 d_\ell / \log(e(n_\ell - 1)/p_{\max})$ for a numerical constant $c_1 > 0$. Then, we have*

$$\begin{aligned}
\mathbb{P}[\tilde{\mathcal{E}}_6^{(\ell,i)}] &= \mathbb{P}\left[\left\|\mathbf{q}_{s_\perp}^{(\ell)}\right\|_2 \leq \left\|\mathbf{z}_{i_\perp}^{(\ell)}\right\|_2 + \frac{6\sigma}{\tilde{a}} \left\|\mathbf{y}_i^{(\ell)}\right\|_2, \forall s \leq s_a\right] \\
&\geq 1 - 2e^{-c_2 m} - 2e^{-c_3 d_\ell}, \tag{86}
\end{aligned}$$

where $c_2, c_3 > 0$ are numerical constants.

Proof. We start by rewriting (83) as

$$\mathbf{y}_i^{(\ell)} = \mathbf{Y}_{\Omega_s}^{(\ell)} \mathbf{b}_{\Omega_s} + \mathbf{q}_s^{(\ell)}, \tag{87}$$

where Ω_s was defined in (84) and \mathbf{b}_{Ω_s} contains the coefficients of the representation of $\mathbf{y}_i^{(\ell)}$ computed by MP according to (4), i.e., $[\mathbf{b}]_\omega = \sum_{s': \omega_{s'} = \omega} \langle \mathbf{y}_{\omega_{s'}}^{(\ell)}, \mathbf{q}_{s'-1}^{(\ell)} \rangle / \|\mathbf{y}_{\omega_{s'}}^{(\ell)}\|_2^2$, $\omega \in \Omega_s$ (this is a direct consequence of (4); we do not reflect dependence of \mathbf{b} on i, ℓ for expositional ease). Next, note that

$$\begin{aligned}
\sigma_{\min}(\tilde{\mathbf{A}}_{\Omega_s}^{(\ell)}) \|\mathbf{b}_{\Omega_s}\|_2 &\leq \sigma_{\min}(\mathbf{Y}_{\Omega_s}^{(\ell)}) \|\mathbf{b}_{\Omega_s}\|_2 \\
&\leq \left\| \mathbf{Y}_{\Omega_s}^{(\ell)} \mathbf{b}_{\Omega_s} \right\|_2 = \left\| \mathbf{y}_i^{(\ell)} - \mathbf{q}_s^{(\ell)} \right\|_2 \\
&\leq \left\| \mathbf{y}_i^{(\ell)} \right\|_2 + \left\| \mathbf{q}_s^{(\ell)} \right\|_2 \leq 2 \left\| \mathbf{y}_i^{(\ell)} \right\|_2, \tag{88}
\end{aligned}$$

where the first inequality is a consequence of $\|\mathbf{Y}_{\Omega_s}^{(\ell)} \mathbf{v}\|_2 \geq \|\mathbf{P}_{\parallel} \mathbf{Y}_{\Omega_s}^{(\ell)} \mathbf{v}\|_2 = \|\mathbf{P}_{\parallel} \mathbf{U}^{(\ell)} \tilde{\mathbf{A}}_{\Omega_s}^{(\ell)} \mathbf{v}\|_2 = \|\tilde{\mathbf{A}}_{\Omega_s}^{(\ell)} \mathbf{v}\|_2$, for all $\mathbf{v} \in \mathbb{R}^{|\Omega_s|}$, and the last inequality follows from $\|\mathbf{q}_s^{(\ell)}\|_2 \leq \|\mathbf{y}_i^{(\ell)}\|_2$ [18, Eq. 13]. For $\sigma_{\min}(\tilde{\mathbf{A}}_{\Omega_s}^{(\ell)}) > 0$ (we will justify below that $\sigma_{\min}(\tilde{\mathbf{A}}_{\Omega_s}^{(\ell)})$ is, indeed, bounded away from 0 with high probability) it hence follows that $\|\mathbf{b}_{\Omega_s}\|_2 \leq (2/\sigma_{\min}(\tilde{\mathbf{A}}_{\Omega_s}^{(\ell)})) \|\mathbf{y}_i^{(\ell)}\|_2$ and therefore, together with (87), we get

$$\begin{aligned}
\left\| \mathbf{q}_{s_\perp}^{(\ell)} \right\|_2 &= \left\| \mathbf{y}_{i_\perp}^{(\ell)} - \mathbf{Y}_{\Omega_s \perp}^{(\ell)} \mathbf{b}_{\Omega_s} \right\|_2 \\
&\leq \left\| \mathbf{y}_{i_\perp}^{(\ell)} \right\|_2 + \left\| \mathbf{Y}_{\Omega_s \perp}^{(\ell)} \mathbf{b}_{\Omega_s} \right\|_2
\end{aligned}$$

$$\begin{aligned}
&\leq \left\| \mathbf{y}_{i_\perp}^{(\ell)} \right\|_2 + \left\| \mathbf{Y}_{\Omega_s \perp}^{(\ell)} \right\|_{2 \rightarrow 2} \|\mathbf{b}_{\Omega_s}\|_2 \\
&= \left\| \mathbf{z}_{i_\perp}^{(\ell)} \right\|_2 + \left\| \mathbf{Z}_{\Omega_s \perp}^{(\ell)} \right\|_{2 \rightarrow 2} \|\mathbf{b}_{\Omega_s}\|_2 \\
&\leq \left\| \mathbf{z}_{i_\perp}^{(\ell)} \right\|_2 + \frac{2}{\sigma_{\min}(\tilde{\mathbf{A}}_{\Omega_s}^{(\ell)})} \left\| \mathbf{Z}_{\Omega_s \perp}^{(\ell)} \right\|_{2 \rightarrow 2} \left\| \mathbf{y}_i^{(\ell)} \right\|_2. \tag{89}
\end{aligned}$$

Replacing Ω_s in (89) by a fixed $\Gamma \in \mathcal{J}$ (recall the definition of \mathcal{J} from (56), with s_{\max} in (56) replaced by p_{\max}), (89) and (59) are equal up to the factor 2 in the second term of (89). This factor-of-two difference stems from the update rule for $\mathbf{q}_s^{(\ell)}$ differing from that for $\mathbf{r}_s^{(\ell)}$ and leads to the difference in $c(\sigma)$ between Theorems 2 and 1. We proceed as in the proof of Lemma 13 (starting from (57)–(59)) by establishing bounds on the tail probabilities of $\|\mathbf{Z}_{\Gamma \perp}^{(\ell)}\|_{2 \rightarrow 2}$ and $\sigma_{\min}(\tilde{\mathbf{A}}_{\Gamma}^{(\ell)})$, for all $\Gamma \in \mathcal{J}$, via (62) and (63), respectively, to obtain the result in Lemma 13. \square

APPENDIX D PROOF OF THEOREM 3

We prove the result for OMP. The proof for MP follows simply by replacing the lower bound on $\|\mathbf{r}_{s_\perp}^{(\ell)}\|_2$ in $\mathcal{E}_7^{(\ell,i)}$ by the lower bound on $\|\mathbf{q}_{s_\perp}^{(\ell)}\|_2$ in $\tilde{\mathcal{E}}_7^{(\ell,i)}$ (defined in (76)), and by noting that $\mathbb{P}[\mathcal{E}_7^{(\ell,i)}] = \mathbb{P}[\tilde{\mathcal{E}}_7^{(\ell,i)}]$.

We only need to address the case

$$\frac{d_\ell}{\log((n_\ell - 1)e)} \min \left\{ \frac{1}{3} \left(\frac{2}{3} - \frac{\tau}{1 - \frac{3}{2} \frac{\sqrt{d_\ell}}{\sqrt{m}} \sigma} \right)^2, c_s \right\} \geq 1, \tag{90}$$

as otherwise there is nothing to prove. We start by noting that under the conditions of Theorem 3 (which are identical to the conditions of Theorem 1 minus the condition $s_{\max} \leq \min_{\ell \in [L]} \{c_s d_\ell / \log((n_\ell - 1)e/s_{\max})\}$), it follows from the proof of Theorem 1 that, conditionally on \mathcal{E}^* as defined in (29), reduced OMP and OMP are guaranteed to select the same points to represent $\mathbf{y}_i^{(\ell)}$ during the first $\lfloor c_s d_\ell / \log((n_\ell - 1)e) \rfloor \geq (10)$ iterations, for all $i \in [n_\ell]$, $\ell \in [L]$. Therefore, conditionally on \mathcal{E}^* , for τ small enough reduced OMP will perform a number of iterations lower-bounded by (10), for all $\mathbf{y}_i^{(\ell)}$, $i \in [n_\ell]$, $\ell \in [L]$, which implies that the number of points from $\mathcal{Y}_\ell \setminus \{\mathbf{y}_i^{(\ell)}\}$ selected by OMP is lower-bounded by (10) as well, for all $\mathbf{y}_i^{(\ell)}$, $i \in [n_\ell]$, $\ell \in [L]$. Specifically, we will show that for $\tau \in [0, 2/3 - (\sqrt{d_{\max}}/\sqrt{m})\sigma]$ the number of reduced OMP iterations is lower-bounded by (10). This will be accomplished by establishing that for $\tau \in [0, 2/3 - (\sqrt{d_{\max}}/\sqrt{m})\sigma]$, on \mathcal{E}^* , we have $\|\mathbf{r}_{s_\tau}^{(\ell)}\|_2 > \tau$ for all $\mathbf{y}_i^{(\ell)}$, $i \in [n_\ell]$, $\ell \in [L]$, where

$$s_\tau := \left\lceil \frac{d_\ell}{\log((n_\ell - 1)e)} \frac{1}{3} \left(\frac{2}{3} - \frac{\tau}{1 - \frac{3}{2} \frac{\sqrt{d_\ell}}{\sqrt{m}} \sigma} \right)^2 \right\rceil. \tag{91}$$

Indeed, as the stopping criterion $\max_{\mathbf{y}_j^{(\ell)} \in \mathcal{Y}_\ell \setminus (\Lambda_s \cup \{\mathbf{y}_i^{(\ell)}\})} |\langle \mathbf{y}_j^{(\ell)}, \mathbf{r}_s^{(\ell)} \rangle| = 0$ is activated only after $\min\{m, n_\ell - 1\} > d_\ell > s_\tau$ iterations (as the points in \mathcal{Y}_ℓ are in general position w.p. 1), $\|\mathbf{r}_{s_\tau}^{(\ell)}\|_2 > \tau$ implies that reduced OMP performs at least s_τ iterations. On the event

$\mathcal{E}_4 \cap \mathcal{E}_7^{(\ell,i)} \supset \mathcal{E}^*$ (\mathcal{E}_4 and $\mathcal{E}_7^{(\ell,i)}$ are defined in (21) and (26), respectively), setting $\bar{s} = s_r$ in $\mathcal{E}_7^{(\ell,i)}$, we have

$$\begin{aligned}
\left\| \mathbf{r}_{\bar{s}}^{(\ell)} \right\|_2 &\geq \left\| \mathbf{r}_{\bar{s}}^{(\ell)} \right\|_2 \\
&> \left\| \mathbf{y}_i^{(\ell)} \right\|_2 \left(\frac{2}{3} - \sqrt{\frac{3\bar{s} \log((n_\ell - 1)e/\bar{s})}{d_\ell}} \right) \\
&\geq \left(1 - \frac{3\sqrt{d_\ell}}{2\sqrt{m}}\sigma \right) \left(\frac{2}{3} - \sqrt{\frac{3\bar{s} \log((n_\ell - 1)e/\bar{s})}{d_\ell}} \right) \\
&\geq \left(1 - \frac{3\sqrt{d_\ell}}{2\sqrt{m}}\sigma \right) \left(\frac{2}{3} \right. \\
&\quad \left. - \left(\left[\frac{d_\ell}{\log((n_\ell - 1)e)} \frac{1}{3} \left(\frac{2}{3} - \frac{\tau}{1 - \frac{3}{2}\frac{\sqrt{d_\ell}}{\sqrt{m}}\sigma} \right)^2 \right] \right. \right. \\
&\quad \left. \left. \cdot \frac{3 \log((n_\ell - 1)e)}{d_\ell} \right)^{\frac{1}{2}} \right) \\
&\geq \left(1 - \frac{3\sqrt{d_\ell}}{2\sqrt{m}}\sigma \right) \left(\frac{2}{3} - \left(\frac{2}{3} - \frac{\tau}{1 - \frac{3}{2}\frac{\sqrt{d_\ell}}{\sqrt{m}}\sigma} \right) \right) \\
&= \tau, \tag{92}
\end{aligned}$$

where for the second inequality we used that $\left\| \mathbf{y}_i^{(\ell)} \right\|_2 \geq \left\| \mathbf{x}_i^{(\ell)} \right\|_2 - \left\| \mathbf{z}_i^{(\ell)} \right\|_2 \geq 1 - 3\sqrt{d_\ell}/(2\sqrt{m})$ on \mathcal{E}_4 , $\log((n_\ell - 1)e/\bar{s}) \leq \log((n_\ell - 1)e)$, for $\bar{s} \geq 1$, for the third inequality, and $\tau \leq 2/3 - (\sqrt{d_{\max}}/\sqrt{m})\sigma$ for the last inequality. This completes the proof.

APPENDIX E SUPPLEMENTARY NOTES

Lemma 14 ([34, Lem. 5.24]). *Let X_1, \dots, X_n be independent centered sub-gaussian random variables (see footnote 4). Then, $X = (X_1, \dots, X_n)$ is a centered sub-gaussian random vector in \mathbb{R}^n , and*

$$\left\| X \right\|_{\Psi_2} \leq C \max_{i \in [n]} \left\| X_i \right\|_{\Psi_2}, \tag{93}$$

where C is a numerical constant.

Theorem 5 ([34, Lem. 5.39]). *Let \mathbf{A} be an $m \times n$ matrix whose rows are independent sub-gaussian isotropic random vectors [34, Def. 5.19, Def. 5.22] in \mathbb{R}^n . Then, for $t \geq 0$, w.p. at least $1 - 2e^{-ct^2}$, we have*

$$\sqrt{m} - C\sqrt{n} - t \leq \sigma_{\min}(\mathbf{A}) \leq \sigma_{\max}(\mathbf{A}) \leq \sqrt{m} + C\sqrt{n} + t. \tag{94}$$

Here $C = C_K, c = c_K > 0$ depend only on the sub-gaussian norm $K = \max_i \left\| \mathbf{A}_{\cdot, i} \right\|_{\Psi_2}$ of the rows of \mathbf{A} .

Theorem 6 ([34, Cor. 5.35]). *Let \mathbf{A} be an $m \times n$ matrix with i.i.d. standard normal entries. Then, for $t \geq 0$, w.p. at least $1 - 2e^{-t^2/2}$, we have*

$$\sqrt{m} - \sqrt{n} - t \leq \sigma_{\min}(\mathbf{A}) \leq \sigma_{\max}(\mathbf{A}) \leq \sqrt{m} + \sqrt{n} + t. \tag{95}$$

REFERENCES

- [1] W. Hong, J. Wright, K. Huang, and Y. Ma, "Multiscale hybrid linear models for lossy image representation," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3655–3671, 2006.
- [2] J. Ho, M. Yang, J. Lim, K. Lee, and D. Kriegman, "Clustering appearances of objects under varying illumination conditions," in *Proc. IEEE Conf. on Comput. Vision and Pattern Recogn.*, vol. 1, 2003, pp. 11–18.
- [3] J. P. Costeira and T. Kanade, "A multibody factorization method for independently moving objects," *International Journal of Computer Vision*, vol. 29, no. 3, pp. 159–179, 1998.
- [4] R. Vidal, S. Soatto, Y. Ma, and S. Sastry, "An algebraic geometric approach to the identification of a class of linear hybrid systems," in *Proc. IEEE Conf. on Decision and Control*, vol. 1, 2003, pp. 167–172.
- [5] D. Jiang, C. Tang, and A. Zhang, "Cluster analysis for gene expression data: A survey," *IEEE Trans. Knowl. Data Eng.*, vol. 16, no. 11, pp. 1370–1386, 2004.
- [6] R. Vidal, "Subspace clustering," *IEEE Signal Process. Mag.*, vol. 28, no. 2, pp. 52–68, 2011.
- [7] U. von Luxburg, "A tutorial on spectral clustering," *Stat. Comput.*, vol. 17, no. 4, pp. 395–416, 2007.
- [8] E. Elhamifar and R. Vidal, "Sparse subspace clustering: Algorithm, theory, and applications," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 35, no. 11, pp. 2765–2781, 2013.
- [9] M. Soltanolkotabi and E. J. Candès, "A geometric analysis of subspace clustering with outliers," *The Annals of Statistics*, vol. 40, no. 4, pp. 2195–2238, 2012.
- [10] M. Soltanolkotabi, E. Elhamifar, and E. J. Candès, "Robust subspace clustering," *The Annals of Statistics*, vol. 42, no. 2, pp. 669–699, 2014.
- [11] Y.-X. Wang and H. Xu, "Noisy sparse subspace clustering," *Journal of Mach. Learn. Research*, vol. 17, no. 12, pp. 1–41, 2016.
- [12] R. Heckel and H. Bölcskei, "Robust subspace clustering via thresholding," *IEEE Trans. Inf. Theory*, vol. 61, no. 11, pp. 6320–6342, 2015.
- [13] E. L. Dyer, A. C. Sankaranarayanan, and R. G. Baraniuk, "Greedy feature selection for subspace clustering," *Journal of Mach. Learn. Research*, vol. 14, pp. 2487–2517, 2013.
- [14] D. Park, C. Caramanis, and S. Sanghavi, "Greedy subspace clustering," in *Advances in Neural Information Processing Systems*, 2014, pp. 2753–2761.
- [15] C. You, D. Robinson, and R. Vidal, "Scalable sparse subspace clustering by orthogonal matching pursuit," in *IEEE Conf. on Comp. Vision and Pattern Recogn.*, 2016, pp. 3918–3927.
- [16] R. Heckel, M. Tschannen, and H. Bölcskei, "Dimensionality-reduced subspace clustering," *Information and Inference: A Journal of the IMA*, vol. 6, no. 3, pp. 246–283, 2017.
- [17] J. H. Friedman and W. Stuetzle, "Projection pursuit regression," *Journal of the American Statistical Association*, vol. 76, no. 376, pp. 817–823, 1981.
- [18] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [19] T. Blumensath, M. E. Davies, and G. Rilling, "Greedy algorithms for compressed sensing," in *Compressed Sensing: Theory and Applications*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge University Press, 2012, pp. 348–393.
- [20] V. N. Temlyakov, "Nonlinear methods of approximation," *Foundations of Computational Mathematics*, vol. 3, no. 1, pp. 33–107, 2003.
- [21] A. S. Georghiadis, P. N. Belhumeur, and D. J. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE Trans. Pattern Anal. Mach. Intelligence*, vol. 23, no. 6, pp. 643–660, 2001.
- [22] K. C. Lee, J. Ho, and D. J. Kriegman, "Acquiring linear subspaces for face recognition under variable lighting," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 5, pp. 684–698, 2005.
- [23] S. Chen, S. A. Billings, and W. Luo, "Orthogonal least squares methods and their application to non-linear system identification," *International Journal of Control*, vol. 50, no. 5, pp. 1873–1896, 1989.
- [24] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Springer, Berlin, Heidelberg, 2013.
- [25] A. Ng, I. M. Jordan, and W. Yair, "On spectral clustering: Analysis and an algorithm," in *Advances in Neural Information Processing Systems*, 2001, pp. 849–856.
- [26] P. Jain, A. Tewari, and I. S. Dhillon, "Orthogonal matching pursuit with replacement," in *Advances in Neural Information Processing Systems*, 2011, pp. 1215–1223.

- [27] S. N. Vitaladevuni, P. Natarajan, R. Prasad, and P. Natarajan, "Efficient orthogonal matching pursuit using sparse random projections for scene and video classification," in *Proc. of IEEE Int. Conf. on Computer Vision*, 2011, pp. 2312–2319.
- [28] R. Basri and D. Jacobs, "Lambertian reflectance and linear subspaces," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 2, pp. 218–233, 2003.
- [29] T. T. Cai and L. Wang, "Orthogonal matching pursuit for sparse signal recovery with noise," *IEEE Trans. Inf. Theory*, vol. 57, no. 7, pp. 4680–4688, 2011.
- [30] G. H. Golub and C. F. Van Loan, *Matrix Computations*. JHU Press, 1996.
- [31] Y. Wang, Y.-X. Wang, and A. Singh, "Graph connectivity in noisy sparse subspace clustering," in *Proc. Int. Conf. on Artificial Intelligence and Statistics*, 2016, pp. 538–546.
- [32] B. Nasihatkon and R. Hartley, "Graph connectivity in sparse subspace clustering," in *Proc. of IEEE Conf. Comput. Vision Pattern Recogn.*, 2011, pp. 2137–2144.
- [33] R. J. Muirhead, *Aspects of multivariate statistical theory*, ser. Wiley Series in Probability and Statistics. John Wiley & Sons, 2009, vol. 197.
- [34] R. Vershynin, "Non-asymptotic random matrix theory," in *Compressed Sensing: Theory and Applications*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge University Press, 2012.
- [35] Y. Chikuse, *Statistics on special manifolds*, ser. Lecture Notes in Statistics. Springer Science & Business Media, 2003, vol. 174.
- [36] M. Ledoux, *The concentration of measure phenomenon*, ser. Mathematical Surveys and Monographs. American Mathematical Soc., 2005, no. 89.