

Dimensionality-reduced subspace clustering

REINHARD HECKEL^{†,‡}, MICHAEL TSCHANNEN AND HELMUT BÖLCSKEI

*Department of Information Technology and Electrical Engineering, ETH Zurich,
Sternwartstr. 7, CH-8092 Zürich, Switzerland*

[†]Corresponding author: heckel@berkeley.edu

[Received on 13 December 2015; revised on 1 August 2016; accepted on 20 September 2016]

Subspace clustering refers to the problem of clustering unlabeled high-dimensional data points into a union of low-dimensional linear subspaces, whose number, orientations and dimensions are all unknown. In practice, one may have access to dimensionality-reduced observations of the data only, resulting, e.g., from undersampling due to complexity and speed constraints on the acquisition device or mechanism. More pertinently, even if the high-dimensional data set is available, it is often desirable to first project the data points into a lower-dimensional space and to perform clustering there; this reduces storage requirements and computational cost. The purpose of this article is to quantify the impact of dimensionality reduction through random projection on the performance of three subspace clustering algorithms, all of which are based on principles from sparse signal recovery. Specifically, we analyze the thresholding based subspace clustering (TSC) algorithm, the sparse subspace clustering (SSC) algorithm and an orthogonal matching pursuit variant thereof (SSC-OMP). We find, for all three algorithms, that dimensionality reduction down to the order of the subspace dimensions is possible without incurring significant performance degradation. Moreover, these results are order-wise optimal in the sense that reducing the dimensionality further leads to a fundamentally ill-posed clustering problem. Our findings carry over to the noisy case as illustrated through analytical results for TSC and simulations for SSC and SSC-OMP. Extensive experiments on synthetic and real data complement our theoretical findings.

Keywords: subspace clustering; dimensionality reduction; random projection; sparse signal recovery.

1. Introduction

One of the major challenges in modern data analysis is to find low-dimensional structure in large high-dimensional data sets. A prevalent low-dimensional structure is that of data points lying in a union of (low-dimensional) subspaces. The problem of extracting such a structure from a given data set can be formalized as follows. Consider the (high-dimensional) set \mathcal{Y} of points in \mathbb{R}^m and assume that $\mathcal{Y} = \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_L$, where the points in \mathcal{Y}_ℓ lie in a linear subspace \mathcal{S}_ℓ of \mathbb{R}^m . The association of the data points to the sets \mathcal{Y}_ℓ , the orientations, dimensions and the number of subspaces \mathcal{S}_ℓ are all unknown. The problem of identifying the assignments of the points in \mathcal{Y} to the \mathcal{Y}_ℓ is referred to as subspace clustering [34] or hybrid linear modeling and has applications, *inter alia*, in unsupervised learning, image representation and segmentation, computer vision and disease detection.

In practice, one may have access to dimensionality-reduced observations of \mathcal{Y} only, resulting, e.g., from “undersampling” due to complexity and speed constraints on the acquisition device or mechanism. More pertinently, even if the data points in \mathcal{Y} are directly accessible, it is often desirable to work on

[‡]Corresponding author. Present address: Department of Electrical Engineering and Computer Sciences, UC Berkeley, Berkeley, CA, USA.

a dimensionality-reduced version of \mathcal{Y} as this reduces data storage cost and leads to computational complexity savings. The idea of reducing computational complexity through dimensionality reduction appears, e.g., in [32] in a general context, and for subspace clustering in the experiments reported in [9,38]. Dimensionality reduction also has a privacy-enhancing effect in the sense that no access to the original data is needed for processing [25].

Dimensionality reduction will, in general, come at the cost of clustering performance. The purpose of this article is to analytically characterize this performance degradation for three subspace clustering algorithms, namely thresholding-based subspace clustering (TSC) [16], sparse subspace clustering (SSC) [8,9] and SSC-orthogonal matching pursuit (SSC-OMP) [7]. The common theme underlying these three algorithms is that they apply spectral clustering to an adjacency matrix constructed from sparse representations of the data points, obtained through a nearest neighbor search in the case of TSC, through ℓ_1 -minimization for SSC, and through OMP in the case of SSC-OMP. While there are numerous further approaches to subspace clustering (see [34] for an overview), we chose to study TSC, SSC and SSC-OMP, as they belong to the small group of subspace clustering algorithms that are computationally tractable and succeed provably under non-restrictive conditions [7,9,16,28,29,37]. Specifically, the results in [16] for TSC, and in [28,29] for SSC show that TSC and SSC can succeed even when the subspaces \mathcal{S}_ℓ intersect. The corresponding proof techniques, together with analytical performance guarantees for SSC-OMP developed in this article, form the basis for our analytical characterization of the impact of dimensionality reduction on subspace clustering performance.

1.1 Formal problem statement and contributions

Consider a set of N data points $\mathcal{Y} \in \mathbb{R}^m$, and assume that $\mathcal{Y} = \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_L$, where the points $\mathbf{y}_i^{(\ell)} \in \mathcal{Y}_\ell, i \in \{1, \dots, n_\ell\}$, lie in a d_ℓ -dimensional linear subspace of \mathbb{R}^m , denoted by \mathcal{S}_ℓ . The assignments of the points in \mathcal{Y} to the sets \mathcal{Y}_ℓ , the subspaces \mathcal{S}_ℓ , and the number of subspaces L are all assumed unknown. Traditional subspace clustering operates on the data \mathcal{Y} with the goal of segmenting it into the sets \mathcal{Y}_ℓ . Here, we assume, however, that clustering is performed on a dimensionality-reduced version of the points in \mathcal{Y} . Specifically, we employ the random projection method [32] by first applying the (same) realization of a random projection matrix $\Phi \in \mathbb{R}^{p \times m}$ (typically $p \ll m$) to each point in \mathcal{Y} to obtain the set of dimensionality-reduced data points \mathcal{X} . Then, we declare the segmentation obtained by operating on \mathcal{X} to be the segmentation of the data points in \mathcal{Y} . The realization of Φ does not need to be known. There are two sources of error that determine the performance of this approach, first, the error that would be obtained even if clustering was performed on the high-dimensional data set \mathcal{Y} directly, second, and more pertinently, the error incurred by operating on dimensionality-reduced data. The former is quantified for TSC in [16], for SSC in [28,29], and for SSC-OMP this article develops corresponding new results. Characterizing the error incurred by dimensionality reduction analytically is the main contribution of this article.

While it is conceivable that TSC, which is based on thresholding inner products, exhibits graceful performance degradation as the data set's dimensionality is reduced through random projection, this is far from obvious for the ℓ_1 -minimization based SSC algorithm and the iterative SSC-OMP algorithm. We prove our main results by first deriving conditions for TSC, SSC and SSC-OMP to ensure correct clustering of dimensionality-reduced data. While these conditions are general, they only become amenable to insightful interpretation once particularized for a random data model, also used in [16,28], that takes the subspace structure of the data set into account. The resulting clustering conditions make the impact of dimensionality reduction explicit, and reveal a tradeoff between the affinity of the subspaces \mathcal{S}_ℓ and the amount of dimensionality reduction possible. Specifically, we find that all three algorithms succeed

provably under quite generous conditions on the relative orientations of the subspaces \mathcal{S}_ℓ , provided that the dimensionality is reduced no more than down to the largest subspace dimension $d_{\max} = \max_\ell d_\ell$. As the computational complexity associated with the construction of the adjacency matrix is essentially linear in the dimension of the ambient space, m , for all three algorithms, random projection reduces the complexity of this step by a factor of m/d_{\max} . These complexity savings translate into, possibly significant, run-time savings for the overall clustering algorithms (which include the spectral clustering step), in particular when m is sufficiently large relative to N .

We study the impact of noise—added to the high-dimensional data points—on clustering performance. For TSC, we derive a clustering condition which quantifies the tradeoff between the affinity of the subspaces \mathcal{S}_ℓ and the amount of dimensionality reduction possible, as a function of noise variance. Specifically, this condition allows us to conclude that TSC succeeds provably provided that—as in the noiseless case—the dimensionality is reduced to no more than down to the largest subspace dimension d_{\max} , and the noise variance is sufficiently small. An approach akin to that used for TSC can be applied to establish a similar clustering condition for SSC-OMP. The corresponding technical details are, however, significantly more involved and cumbersome. We therefore decided not to state the formal result. Regarding SSC, we remark that Wang *et al.* [36] reported deterministic clustering conditions for the Lasso-version of SSC [29] applied to dimensionality-reduced noisy data. However, the corresponding results [36, Lemma 16, Theorem 18] make the critical assumption of the signal part of the *projected* noisy data being normalized, whereas the noise component remains un-normalized. It is difficult to see how one would realize this in practice, unless the noise realization is known perfectly, in which case the noise component could be removed which would take us back to the noiseless case. The results in [36] for noisy data therefore appear to be of limited practicality. While the statements in [36] may be particularized to the noiseless case, we note that corresponding results appeared in the conference version [17] of this article before the publication of [36].

We note that our results, both for the noiseless and the noisy case, apply even when the vectors in the subspaces \mathcal{S}_ℓ , collectively across ℓ , span the ambient space \mathbb{R}^m . This follows from our clustering conditions depending on the *pairwise* affinities between subspaces only, and pairwise affinities changing only moderately if the dimensionality is reduced down to no more than the order of the individual subspace dimensions.

Another popular dimensionality reduction method is principal component analysis (PCA). However, when used in the context of subspace clustering, PCA, in general, affords dimensionality reduction down to the dimension of the union $\bigcup_\ell \mathcal{S}_\ell$ of the subspaces \mathcal{S}_ℓ only, as opposed to the order of the dimensions of the individual subspaces. To see this, consider the L subspaces of dimension one that correspond to the standard basis in \mathbb{R}^m , i.e., the ℓ th subspace is spanned by the vector \mathbf{e}_ℓ given by $[\mathbf{e}_\ell]_\ell = 1$ and $[\mathbf{e}_\ell]_i = 0$, for $i \neq \ell$. Assuming that each of the data points in the data set under consideration, denoted by $\mathbf{Y} \in \mathbb{R}^{m \times N}$, lies in one of these L subspaces, the corresponding sample covariance matrix $\mathbf{Y}\mathbf{Y}^T$ has non-zero entries only in its first L main diagonal entries. The first L principal components are therefore given by the vectors \mathbf{e}_ℓ . Reducing the dimensionality of the data set to below L will result in certain data points being mapped to zero (owing to the orthogonality of the \mathbf{e}_ℓ). Moreover, PCA has computational complexity $O(Nm^2 + m^3)$, while random projection through Gaussian matrices and fast random projection matrices [1] has complexity $O(pmN)$ and $O(\log(m)mN)$, respectively, and is therefore computationally much less demanding. This is an important aspect as computational complexity is a major motivation for dimensionality reduction.

1.2 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{x} , x_q denotes its q th entry and \mathbf{x}_S

is the subvector of \mathbf{x} with entries corresponding to the indices in the set S . For the matrix \mathbf{A} , \mathbf{A}_{ij} designates the entry in its i th row and j th column, \mathbf{A}_S the matrix containing the columns of \mathbf{A} with indices in the set S , $\|\mathbf{A}\|_{2 \rightarrow 2} := \max_{\|\mathbf{v}\|_2=1} \|\mathbf{A}\mathbf{v}\|_2$ its spectral norm, $\sigma_{\min}(\mathbf{A})$ its minimum singular value, and $\|\mathbf{A}\|_F := (\sum_{i,j} |\mathbf{A}_{ij}|^2)^{1/2}$ its Frobenius norm. If \mathbf{A} has full column rank $\mathbf{A}^\dagger := (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T$ stands for its (left) pseudoinverse, and for \mathbf{A} with full row rank, $\mathbf{A}^\dagger := \mathbf{A}^T (\mathbf{A} \mathbf{A}^T)^{-1}$ is the (right) pseudoinverse. The identity matrix is denoted by \mathbf{I} . $\log(\cdot)$ refers to the natural logarithm, $\arccos(\cdot)$ is the inverse function of $\cos(\cdot)$, and $x \wedge y$ denotes the minimum of x and y . The set $\{1, \dots, N\}$ is written as $[N]$. The cardinality of the set S is designated by $|S|$ and its complement is \bar{S} . $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ stands for the distribution of a real Gaussian random vector with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. We write $X \sim Y$ to indicate that the random variables X and Y are equally distributed. For notational convenience, we use the following shorthands: \max_ℓ for $\max_{\ell \in [L]}$, $\max_{k \neq \ell}$ for $\max_{k \in [L]: k \neq \ell}$, and $\max_{k, \ell: k \neq \ell}$ for $\max_{k, \ell \in [L]: k \neq \ell}$. The unit sphere in \mathbb{R}^m is $\mathbb{S}^{m-1} := \{\mathbf{x} \in \mathbb{R}^m: \|\mathbf{x}\|_2 = 1\}$. A subgraph H of a graph G is said to be connected if every pair of nodes in H can be joined by a path along edges with nodes exclusively in H . A subgraph H of G is called a connected component of G if H is connected and if there are no edges between nodes in H and the remaining nodes in G .

2. A brief review of TSC, SSC and SSC-OMP

We next briefly summarize the TSC [16], SSC [8,9] and SSC-OMP [7] algorithms. All three algorithms apply normalized spectral clustering [35] to an adjacency matrix \mathbf{A} built by finding a sparse representation of each data point in terms of the other data points. Specifically, TSC is based on least-squares representations in terms of nearest neighbors, while SSC and SSC-OMP construct \mathbf{A} by finding sparse representations via ℓ_1 -minimization and OMP, respectively. Note that the focus in [16] is on a version of TSC that uses a spherical distance measure between data points instead of least-squares regression coefficients to determine the entries of \mathbf{A} . The analytical results presented here apply to both versions of TSC. We decided, however, to explicitly work with the least-squares version as this formulation better elucidates the sparsity aspect and thereby the relationship to SSC and SSC-OMP.

In order to emphasize that we consider the three algorithms applied to dimensionality-reduced data, their descriptions will be in terms of the dimensionality-reduced data set $\mathcal{X} \subset \mathbb{R}^p$. We furthermore assume that an estimate \hat{L} of the number of subspaces L is available. The estimation of L from \mathcal{X} is discussed later. We also note that the formulations of the TSC and SSC-OMP algorithms below assume that the data points in \mathcal{X} are of comparable ℓ_2 -norm. This assumption is relevant for Step 1 in both cases and is not restrictive as the data points can be normalized prior to clustering.

The TSC algorithm: Given a set of N data points \mathcal{X} in \mathbb{R}^p , an estimate of the number of subspaces \hat{L} , and the parameter q (the choice of q is discussed below), perform the following steps:

Step 1: For every $\mathbf{x}_j \in \mathcal{X}$, find the set $S_j \subset [N] \setminus \{j\}$ of cardinality q defined by

$$|\langle \mathbf{x}_j, \mathbf{x}_i \rangle| \geq |\langle \mathbf{x}_j, \mathbf{x}_k \rangle|, \text{ for all } i \in S_j \text{ and all } k \notin S_j,$$

and let \mathbf{z}_j be the coefficient vector corresponding to the minimum least-squares representation of \mathbf{x}_j in terms of $\mathbf{x}_i, i \in S_j$. Specifically, set $(\mathbf{z}_j)_{S_j} = \arg \min_{\mathbf{z}} \|\mathbf{x}_j - \mathbf{X}_{S_j} \mathbf{z}\|_2$ (if multiple solutions exist, choose, e.g., the \mathbf{z} with minimum ℓ_2 -norm), and $(\mathbf{z}_j)_{\bar{S}_j} = \mathbf{0}$. Construct the adjacency matrix \mathbf{A} according to $\mathbf{A} = \mathbf{Z} + \mathbf{Z}^T$, where $\mathbf{Z} = \text{abs}([\mathbf{z}_1 \dots \mathbf{z}_N])$ and $\text{abs}(\cdot)$ takes absolute values element-wise.

Step 2: Apply normalized spectral clustering [26,35] to (\mathbf{A}, \hat{L}) .

The SSC algorithm: Given a set of N data points \mathcal{X} in \mathbb{R}^p and an estimate of the number of subspaces \hat{L} , perform the following steps:

Step 1: Let $\mathbf{X} \in \mathbb{R}^{p \times N}$ be the matrix whose columns are the points in \mathcal{X} . For every $\mathbf{x}_j \in \mathcal{X}$ determine \mathbf{z}_j as a solution of

$$\underset{\mathbf{z}}{\text{minimize}} \|\mathbf{z}\|_1 \text{ subject to } \mathbf{x}_j = \mathbf{X}\mathbf{z} \text{ and } z_j = 0. \quad (2.1)$$

Construct the adjacency matrix \mathbf{A} according to $\mathbf{A} = \mathbf{Z} + \mathbf{Z}^T$, where $\mathbf{Z} = \text{abs}([\mathbf{z}_1 \dots \mathbf{z}_N])$.

Step 2: Apply normalized spectral clustering [26,35] to (\mathbf{A}, \hat{L}) .

The SSC-OMP algorithm: Given a set of N data points \mathcal{X} in \mathbb{R}^p , an estimate of the number of subspaces \hat{L} , and a maximum number of OMP iterations s_{\max} , perform the following steps:

Step 1: For every $\mathbf{x}_j \in \mathcal{X}$, find a sparse representation of \mathbf{x}_j in terms of $\mathcal{X} \setminus \{\mathbf{x}_j\}$ using OMP as follows: Initialize the iteration counter $s = 0$, the residual $\mathbf{r}_0 = \mathbf{x}_j$, and the set of selected indices $\Lambda_0 = \emptyset$. For $s = 1, 2, \dots$ perform updates according to

$$\Lambda_s = \Lambda_{s-1} \cup \underset{i \in [N]: i \neq j}{\text{argmax}} |\langle \mathbf{x}_i, \mathbf{r}_{s-1} \rangle| \quad (2.2)$$

$$\mathbf{r}_s = (\mathbf{I} - \mathbf{X}_{\Lambda_s} \mathbf{X}_{\Lambda_s}^\dagger) \mathbf{x}_j \quad (2.3)$$

until $\mathbf{r}_s = \mathbf{0}$ or $s = s_{\max}$ (when the maximizer in (2.2) is not unique, select any of the solutions). With the number of OMP iterations actually performed denoted by s_{end} , set $(\mathbf{z}_j)_{\Lambda_{s_{\text{end}}}} = \mathbf{X}_{\Lambda_{s_{\text{end}}}}^\dagger \mathbf{x}_j$, $(\mathbf{z}_j)_{\bar{\Lambda}_{s_{\text{end}}}} = \mathbf{0}$, and construct the adjacency matrix \mathbf{A} according to $\mathbf{A} = \mathbf{Z} + \mathbf{Z}^T$, where $\mathbf{Z} = \text{abs}([\mathbf{z}_1 \dots \mathbf{z}_N])$.

Step 2: Apply normalized spectral clustering [26,35] to (\mathbf{A}, \hat{L}) .

All three algorithms have an estimate of the number of subspaces, \hat{L} , as input parameter. This estimate can be generated from the adjacency matrix computed in Step 1 of the algorithms (note that \hat{L} is needed only in Step 2) through the *eigengap heuristic* [35], or it can be made available through prior knowledge on the clustering problem at hand. For expositional simplicity, we decided to consistently treat \hat{L} as an input parameter. The eigengap heuristic is based on the insight that the number of zero eigenvalues of the normalized Laplacian of the graph G with adjacency matrix \mathbf{A} , henceforth simply referred to as “the graph G ”, is equal to the number of connected components of G [30]. Formally, denoting the eigenvalues of the normalized Laplacian of G by $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$, the eigengap heuristic computes $\hat{L} = \arg \max_{i \in [N-1]} (\lambda_{i+1} - \lambda_i)$.

Let the oracle segmentation of \mathcal{X} be given by $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_L$. If each connected component in the graph G corresponds exclusively to points from one of the sets \mathcal{X}_ℓ , spectral clustering will deliver the oracle segmentation [35, Proposition 4] and the clustering error (CE), i.e., the fraction of misclustered points, will be zero. Since conditions guaranteeing zero CE are inherently hard to obtain, we will work with an intermediate, albeit sensible, performance measure, also employed in [7,16,28,29]. Specifically, this measure, termed the no false connections property, declares success if the graph G has no false connections, i.e., if each $\mathbf{x}_j \in \mathcal{X}_\ell$ is connected to points in \mathcal{X}_ℓ only, for all ℓ . Guaranteeing the absence of false connections, does, however, not guarantee that the connected components of G correspond to the \mathcal{X}_ℓ , as the points in a given set \mathcal{X}_ℓ may form two (or more) distinct connected components in G .

To counter this problem, sufficiently many entries in each row/column of the adjacency matrix \mathbf{A} have to be non-zero. Specifically, for the subgraphs of G corresponding to the \mathcal{X}_ℓ to be connected, each row/column of \mathbf{A} corresponding to a point in \mathcal{X}_ℓ needs to have between $O(\log n_\ell)$ and $O(n_\ell)$

non-zero entries. As the solutions \mathbf{z} to $\arg \min_{\mathbf{z}} \|\mathbf{x}_j - \mathbf{X}_{\mathcal{S}_j} \mathbf{z}\|_2$ are typically dense, TSC is likely to select a representation of \mathbf{x}_j in terms of points in $\mathcal{X}_\ell \setminus \{\mathbf{x}_j\}$ with on the order of q non-zero coefficients. Choosing q large enough therefore ensures sufficient connectivity of the graph G generated by TSC. On the other hand, taking q to be large increases the probability of false connections. The performance guarantee we obtain for TSC therefore requires q to be sufficiently small relative to the n_ℓ . For SSC and SSC-OMP, the number of non-zero entries in each row/column of \mathbf{A} turns out to be tied to d_ℓ , rather than n_ℓ . To see this, suppose that both algorithms exclusively select data points from $\mathcal{X}_\ell \setminus \{\mathbf{x}_j\}$ to represent \mathbf{x}_j . Moreover, assume that the \mathcal{X}_ℓ are non-degenerate in the sense that, indeed, d_ℓ points are needed to represent $\mathbf{x}_j \in \mathcal{X}_\ell$ through points in $\mathcal{X}_\ell \setminus \{\mathbf{x}_j\}$; this precludes, e.g., that \mathcal{X}_ℓ contains multiple copies of the same data point. The OMP algorithm in SSC-OMP then terminates after $\min(d_\ell, s_{\max})$ (recall that $d_\ell = \dim(\mathcal{S}_\ell)$) iterations for $\mathbf{x}_j \in \mathcal{X}_\ell$, and hence results in exactly $\min(d_\ell, s_{\max})$ non-zero entries in the corresponding column of \mathbf{Z} (recall that $\mathbf{A} = \mathbf{Z} + \mathbf{Z}^T$). For SSC, we simply note that d_ℓ points are enough to represent $\mathbf{x}_j \in \mathcal{X}_\ell$ through other points in \mathcal{X}_ℓ , and we cannot guarantee more than d_ℓ non-zero entries in the corresponding column of \mathbf{Z} , in general. This will lead to insufficient connectivity for SSC and SSC-OMP when d_ℓ is not in the range $O(\log n_\ell) - O(n_\ell)$. The problem is exacerbated when the data set is degenerate. To counter insufficient connectivity in SSC, a modification which adds an ℓ_2 -penalty to the cost function in (2.1) was proposed in [9, Section 5]. Such a modification is not known for SSC-OMP, and this may be considered a limitation of SSC-OMP.

We remark that TSC and SSC-OMP can be made essentially parameterless, like SSC. Specifically, a procedure for choosing the TSC parameter q in a data-driven fashion is described in [15], and for SSC-OMP we can get rid of the parameter s_{\max} by stopping the OMP step once the ℓ_2 -norm of the residual \mathbf{r}_s falls below a threshold value.

3. Main results

We start by specifying the statistical data model used throughout the article. The subspaces \mathcal{S}_ℓ are taken to be deterministic and the points within the \mathcal{S}_ℓ are chosen randomly. Specifically, the elements of the set \mathcal{Y}_ℓ in $\mathcal{Y} = \mathcal{Y}_1 \cup \dots \cup \mathcal{Y}_L$ are obtained by choosing n_ℓ points at random according to $\mathbf{y}_j^{(\ell)} = \mathbf{U}^{(\ell)} \mathbf{a}_j^{(\ell)}, j \in [n_\ell]$, where the columns of $\mathbf{U}^{(\ell)} \in \mathbb{R}^{m \times d_\ell}$ form an orthonormal basis for the d_ℓ -dimensional subspace \mathcal{S}_ℓ , and the $\mathbf{a}_j^{(\ell)}$ are i.i.d. uniform on $\mathbb{S}^{d_\ell-1}$. As the $\mathbf{U}^{(\ell)}$ are orthonormal, the data points $\mathbf{y}_j^{(\ell)}$ are distributed uniformly on the set $\{\mathbf{y} \in \mathcal{S}_\ell: \|\mathbf{y}\|_2 = 1\} = \mathcal{S}_\ell \cap \mathbb{S}^{m-1}$, which avoids degenerate situations where the data points lie in preferred directions. To see why such degeneracies can lead to ambiguous results, consider a two-dimensional subspace and assume that the data points in this subspace are skewed towards two distinct directions. Then, there are two sensible segmentations. One is to assign the points corresponding to each direction to separate clusters, the other to assign all points to one cluster.

The dimensionality-reduced data set $\mathcal{X} \subset \mathbb{R}^p$ is obtained by applying the (same) realization of a random matrix $\Phi \in \mathbb{R}^{p \times m}$ ($p \geq \max_\ell d_\ell$) to each point in \mathcal{Y} . The elements of the sets \mathcal{X}_ℓ in $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_L$ are hence given by $\mathbf{x}_j^{(\ell)} = \Phi \mathbf{y}_j^{(\ell)}, j \in [n_\ell]$. We take Φ as a random matrix satisfying the following concentration inequality

$$\mathbb{P}[\|\Phi \mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2 \geq t \|\mathbf{x}\|_2^2] \leq 2e^{-\tilde{c} t^2 p}, \quad \forall t > 0, \forall \mathbf{x} \in \mathbb{R}^m, \quad (3.1)$$

where \tilde{c} is either a numerical constant or a parameter mildly dependent on m . Random matrices satisfying (3.1) realize, with high probability, linear embeddings in the sense of the Johnson–Lindenstrauss (JL)

Lemma, see, e.g., [32], [10, Section 9.5]. The JL Lemma says that every set of N points in Euclidean space can be embedded in an $O(\epsilon^{-2} \log N)$ -dimensional space without perturbing the pairwise Euclidean distances between the points by more than a factor of $1 \pm \epsilon$. A similar statement on random projections preserving affinities between subspaces—as defined in (3.2)—is used in our proofs. Specifically, we show that randomly projecting a set of d -dimensional subspaces into p -dimensional space does not increase their pairwise affinities by more than $\text{const.} \sqrt{d/p}$, with high probability (cf. (B.8)). The concentration inequality (3.1) holds, *inter alia*, for matrices with i.i.d. subgaussian¹ entries [10, Lemma 9.8]; this includes $\mathcal{N}(0, 1/p)$ entries and entries that are uniformly distributed on $\{-1/\sqrt{p}, 1/\sqrt{p}\}$. Such matrices may, however, be costly to generate, store and apply to high-dimensional data points. In order to reduce these costs, structured random matrices satisfying (3.1) (with \tilde{c} possibly mildly dependent on m) were proposed in [1, 21]. For example, the structured random matrix proposed in [1] (and described in detail in Section 5) satisfies (3.1) with $\tilde{c} = c_2 \log^{-4}(m)$, where c_2 is a numerical constant [21, Proposition 3.2], and can be applied in time $O(m \log m)$ as opposed to time $O(mp)$ for the realizations of general subgaussian random matrices.

The clustering performance guarantees we obtain below are all in terms of the affinity between the subspaces \mathcal{S}_k and \mathcal{S}_ℓ defined as [28, Definition 2.6], [29, Definition 1.2]

$$\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) := \frac{1}{\sqrt{d_k \wedge d_\ell}} \|\mathbf{U}^{(k)T} \mathbf{U}^{(\ell)}\|_F. \quad (3.2)$$

Note that $0 \leq \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \leq 1$, with $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) = 1$ if $\mathcal{S}_k \subseteq \mathcal{S}_\ell$ or $\mathcal{S}_\ell \subseteq \mathcal{S}_k$ and $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) = 0$ if \mathcal{S}_k and \mathcal{S}_ℓ are orthogonal to each other. Moreover, we have

$$\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) = \sqrt{\cos^2(\theta_1) + \dots + \cos^2(\theta_{d_k \wedge d_\ell})} / \sqrt{d_k \wedge d_\ell}, \quad (3.3)$$

where $\theta_1 \leq \dots \leq \theta_{d_k \wedge d_\ell}$ are the principal angles between \mathcal{S}_k and \mathcal{S}_ℓ [12, Section 6.3.4]. If \mathcal{S}_k and \mathcal{S}_ℓ intersect in t dimensions, i.e., if $\mathcal{S}_k \cap \mathcal{S}_\ell$ is t -dimensional, then $\cos(\theta_1) = \dots = \cos(\theta_t) = 1$, and hence $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \geq \sqrt{t/(d_k \wedge d_\ell)}$. The affinity between subspaces plays an important role in subspace classification [27] as well, see [19, Theorems 2 and 3].

We start with our main result for TSC.

THEOREM 3.1 Choose q such that $q \leq \min_\ell n_\ell/6$. If

$$\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \frac{\sqrt{11}}{\sqrt{3\tilde{c}}} \frac{\sqrt{d_{\max}}}{\sqrt{p}} \leq \frac{1}{15 \log N}, \quad (3.4)$$

where $d_{\max} := \max_\ell d_\ell$ and \tilde{c} is the constant in the concentration inequality (3.1), then the graph G obtained by applying TSC to \mathcal{X} has no false connections with probability at least $1 - 7N^{-1} - \sum_{\ell=1}^L n_\ell e^{-c(n_\ell-1)}$, where $c > 1/20$ is a numerical constant.

Our main result for SSC is the following.

¹ A random variable x is subgaussian [10, Section 7.4] if its tail probability satisfies $\mathbb{P}[|x| > t] \leq c_1 e^{-c_2 t^2}$ for constants $c_1, c_2 > 0$.

THEOREM 3.2 Let $\rho_\ell := (n_\ell - 1)/d_\ell$, $\ell \in [L]$, $\rho_{\min} := \min_\ell \rho_\ell \geq \rho_0$, where $\rho_0 > 1$ is a numerical constant, and pick any $\tau > 0$. Set $d_{\max} := \max_\ell d_\ell$ and suppose that

$$\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \sqrt{\frac{28d_{\max} + 8 \log L + 2\tau}{3\tilde{c}p}} \leq \frac{\sqrt{\log \rho_{\min}}}{65 \log N}, \quad (3.5)$$

where \tilde{c} is the constant in (3.1). Then, the graph G obtained by applying SSC to \mathcal{X} has no false connections with probability at least $1 - 4e^{-\tau/2} - N^{-1} - \sum_{\ell=1}^L n_\ell e^{-\sqrt{\rho_\ell} d_\ell}$.

Finally, for SSC-OMP we obtain the following statement.

THEOREM 3.3 Let $\rho_\ell := (n_\ell - 1)/d_\ell$, $\ell \in [L]$, $\rho_{\min} := \min_\ell \rho_\ell \geq \rho_0$, where $\rho_0 > 1$ is a numerical constant, and pick any $\tau > 0$. Set $d_{\min} := \min_\ell d_\ell$, $d_{\max} := \max_\ell d_\ell$, and suppose that Φ has (in addition to satisfying the concentration inequality (3.1)) a rotationally invariant distribution, i.e., $\Phi \mathbf{V} \sim \Phi$ for all unitary matrices $\mathbf{V} \in \mathbb{R}^{m \times m}$. If

$$\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \sqrt{\frac{28d_{\max} + 8 \log L + 2\tau}{12\tilde{c}p}} \sqrt{\frac{d_{\max}}{d_{\min}}} \leq \frac{3}{200} \frac{\sqrt{\log \rho_{\min}}}{\log N}, \quad (3.6)$$

where \tilde{c} is the constant in (3.1), then, irrespectively² of the choice of the maximum number of OMP-iterations s_{\max} , the graph G obtained by applying SSC-OMP to \mathcal{X} has no false connections with probability at least $1 - 4e^{-\tau/2} - 4N^{-1} - \sum_{\ell=1}^L n_\ell e^{-\sqrt{\rho_\ell} d_\ell}$.

The proofs of Theorems 3.1–3.3 are provided in Appendices A, B and C, respectively, and are established by first deriving deterministic clustering conditions that are then evaluated for our statistical data model.

Theorems 3.1 and 3.2 essentially say that even when p is on the order of d_{\max} , TSC and SSC succeed (in the sense of the no false connections property being satisfied) with high probability if the affinities between the subspaces \mathcal{S}_ℓ are sufficiently small and if \mathcal{X} contains sufficiently many points from each subspace. The same conclusion applies to SSC-OMP provided that the term $\sqrt{d_{\max}/d_{\min}}$ is not too large, which is the case if the dimensions d_ℓ , $\ell \in [L]$, of the subspaces are of the same order. This condition is satisfied in many practical applications, such as, e.g., for the face clustering and the handwritten digit clustering problems described in Section 5. We believe the occurrence of the factor $\sqrt{d_{\max}/d_{\min}}$ in (3.6) to be an artifact of our proof technique. Also note that Theorem 3.3 imposes more restrictive conditions on Φ than Theorems 3.1 and 3.2, namely the distribution of Φ has to be rotationally invariant. This is a technical condition and it is not implied by (3.1). Examples of rotationally invariant matrices satisfying (3.1) include matrices with i.i.d. $\mathcal{N}(0, 1/p)$ entries.

Theorems 3.1–3.3 apply even when the union of the subspaces \mathcal{S}_ℓ equals the ambient space \mathbb{R}^m . This follows by virtue of our clustering conditions depending only on the pairwise affinities between subspaces, and pairwise affinities changing only moderately if the dimensionality is reduced down to no more than the order of the individual subspace dimensions. In contrast, results on the performance

² While the statement holds irrespectively of s_{\max} , recall from Section 2 that choosing s_{\max} too small may result in too few non-zero entries in \mathbf{A} for successful clustering.

of subspace clustering algorithms available in the literature often assume the subspaces to be independent³ [34]. Now, it turns out that if the \mathcal{S}_ℓ collectively span \mathbb{R}^m , the corresponding subspaces after dimensionality reduction cannot be independent.

Theorems 3.1–3.3 show that for all three algorithms, p may be taken to be linear (up to log-factors) in d_{\max} . We can therefore conclude that the dimensionality of the data set \mathcal{Y} can be reduced down to the order of the largest subspace dimension without affecting clustering performance significantly. This has important practical ramifications as, for all three algorithms, the computational complexity associated with the construction of the adjacency matrix is essentially linear in the dimension of the ambient space the data points “live in”. To get an idea of the resulting overall complexity savings, let us consider the TSC algorithm and assume that the (high-dimensional) data set $\mathcal{Y} \subset \mathbb{R}^m$ is projected down to \mathbb{R}^p , with $p = O(d_{\max} \log^2(N))$, via a Gaussian random projection; this choice of p guarantees, by Theorem 3.1, that clustering performance is not affected significantly by dimensionality reduction. The complexity associated with the construction of the adjacency matrix for TSC is given by the cost of computing the inner products between all pairs of data points, and is therefore $O(mN^2)$ for the original data set $\mathcal{Y} \subset \mathbb{R}^m$ and $O(pN^2)$ for the projected data set $\mathcal{X} \subset \mathbb{R}^p$. Adding the cost for applying the Gaussian random projection results in an overall cost of $O(pN^2) + O(pNm) = O(d_{\max} \log^2(N)N(N + m))$ for building the adjacency matrix associated with \mathcal{X} . The resulting complexity savings for TSC are therefore given by $O(\min(m, N)/(d_{\max} \log^2(N)))$. The absolute run-time savings are even more pronounced for SSC-OMP and SSC, as the corresponding costs for building the adjacency matrix are larger than $O(mN^2)$. Further gains can be obtained by employing fast random projections (FRPs) [1].

Dimensionality reduction affects the computational cost associated with the construction of the adjacency matrix only. The spectral clustering step, which when naïvely implemented has complexity $O(N^3)$, may be the dominating factor in the overall computational cost, in particular when m is small relative to N^3 . Notwithstanding, dimensionality reduction can still lead to significant total run-time savings. Our numerical results in Section 5 demonstrate this for SSC. To see savings on the same order for SSC-OMP and TSC, we would have to consider problems with N small relative to m .

The probability lower bounds in Theorems 3.1–3.3 are independent of p and m , and require the total number of data points N to be large in absolute terms in order to ensure a success probability close to one.

Theorems 3.1–3.3 are order-optimal in the following sense. If dimensionality is reduced to below d_{\max} , then, in general, there are points from different subspaces that are projected into the same lower-dimensional subspace, which renders the resulting clustering problem fundamentally ill-posed. To see this, take $d_\ell = d$, for all ℓ , and assume that $p \leq d$. Next, note that the (randomly projected) points \mathcal{X}_ℓ lie in the column span of $\Phi \mathbf{U}^{(\ell)}$. As $\mathbf{U}^{(\ell)}$ is a basis for the d_ℓ -dimensional subspace $\mathcal{S}_\ell \subset \mathbb{R}^m$, the span of $\Phi \mathbf{U}^{(\ell)}$ is \mathbb{R}^p , for all ℓ , and therefore all points in the projected data set $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_L$ lie in the same p -dimensional subspace, which renders the clustering problem ill-posed.

While our results pertain to the particular data model used in this article, the qualitative conclusions we draw remain valid under much more general conditions. Specifically, any model that has the data points spread out sufficiently on the subspaces \mathcal{S}_ℓ would lead to similar conclusions. In particular, we could take the vectors $\mathbf{a}_j^{(\ell)}$ to be independent zero-mean Gaussian with (full-rank) covariance matrices $\Sigma^{(\ell)}$. The corresponding analytical results would then additionally depend on the spreads of the singular values of the $\Sigma^{(\ell)}$; larger spreads would correspond to a harder clustering problem, reflected in a more stringent clustering condition which would necessitate the affinities to be smaller.

³ The linear subspaces \mathcal{S}_ℓ are said to be independent if the dimension of their (set) sum is equal to the sum of their dimensions.

We next compare the clustering conditions (3.4)–(3.6) in Theorems 3.1–3.3 with their counterparts for clustering of the original, high-dimensional data set \mathcal{Y} . Specifically, such reference conditions can be found in [16, Theorem 2] for TSC and in [28, Theorem 2.8] for SSC, but do not seem to be available for SSC-OMP for the statistical data model considered in this article. However, setting $\Phi = \mathbf{I}$ in the proof of Theorem 3.3, we can easily get a reference condition for SSC-OMP. Rather than providing the details of this simple modification, we refer the reader to the proof in [31, Chapter 4].

COROLLARY 3.1 Let $\rho_\ell := (n_\ell - 1)/d_\ell$, $\ell \in [L]$, and suppose that $\rho_{\min} := \min_\ell \rho_\ell \geq \rho_0$, where $\rho_0 > 1$ is a numerical constant. If

$$\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \leq \frac{\sqrt{\log \rho_{\min}}}{64 \log N}, \quad (3.7)$$

then the graph G obtained by applying SSC-OMP to the original, high-dimensional data set \mathcal{Y} has no false connections with probability at least $1 - 2N^{-1} - \sum_{\ell=1}^L n_\ell e^{-\sqrt{\rho_\ell} d_\ell}$.

We conclude that for all three algorithms, the impact of dimensionality reduction is essentially quantified through a term proportional to $\sqrt{d_{\max}/p}$ that adds to the maximum affinity between the subspaces \mathcal{S}_ℓ in the clustering conditions (3.4)–(3.6). These clustering conditions nicely reflect the intuition that the smaller the affinities between the subspaces \mathcal{S}_ℓ , the more aggressively we can reduce the dimensionality of the data set without compromising clustering performance.

As the result in Corollary 3.1 is new, a few comments on its relation to existing results, specifically those in [7] and [37], are in order. Corollary 3.1 imposes less restrictive conditions on the relative orientations of the subspaces than [7, Theorem 3], [37, Theorem 2, Corollary 1], but makes stronger assumptions on the data model. The result in [37, Theorem 3] applies to subspaces with random orientations, and therefore does not allow for statements involving subspace affinities. We refer the reader to the thesis [31, Section 4.1] for a more detailed comparison of Corollary 3.1 above to [7, Theorem 3]. Finally, numerical results corroborating the fundamental nature of the clustering condition (3.7) can be found in [31, Section 5.1].

4. Impact of noise

In many practical applications, the data points to be clustered are corrupted by noise, typically modeled as additive Gaussian noise. Additionally, in practice, the subspace model will, in general, not be an exact fit for the data at hand, so that additive noise can be seen as modeling this deviation. In this section, we study the interplay between dimensionality reduction and additive noise for the TSC algorithm. Specifically, we let the high-dimensional data points be corrupted by Gaussian noise according to

$$\tilde{\mathbf{y}}_i^{(\ell)} = \mathbf{y}_i^{(\ell)} + \mathbf{e}_i^{(\ell)},$$

where $\mathbf{e}_i^{(\ell)} \sim \mathcal{N}(\mathbf{0}, (\sigma^2/m)\mathbf{I})$, and assume, as before, that $\mathbf{y}_i^{(\ell)}$ is drawn i.i.d. uniformly from the intersection of the d_ℓ -dimensional subspace \mathcal{S}_ℓ with the unit sphere. The dimensionality-reduced noisy data set $\tilde{\mathcal{X}} \subset \mathbb{R}^p$ is obtained by applying the same realization of the random projection matrix $\Phi \in \mathbb{R}^{p \times m}$ to all (noisy) data points $\tilde{\mathbf{y}}_i^{(\ell)}$. The elements of the sets $\tilde{\mathcal{X}}_\ell$ in $\tilde{\mathcal{X}} = \tilde{\mathcal{X}}_1 \cup \dots \cup \tilde{\mathcal{X}}_L$ are hence given by

$$\tilde{\mathbf{x}}_j^{(\ell)} = \Phi(\mathbf{y}_i^{(\ell)} + \mathbf{e}_i^{(\ell)}), \quad j \in [n_\ell]. \quad (4.1)$$

THEOREM 4.1 Choose q such that $q \leq \min_{\ell} n_{\ell}/6$, and let $m \geq 6 \log N$. If

$$\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_{\ell}) + \frac{\sqrt{11}}{\sqrt{3\tilde{c}}} \frac{\sqrt{d_{\max}}}{\sqrt{p}} + \frac{\sigma(1+\sigma)\sqrt{6}}{\sqrt{\tilde{c}} \log N} \frac{\sqrt{d_{\max}}}{\sqrt{p}} \leq \frac{1}{15 \log N}, \quad (4.2)$$

where $d_{\max} := \max_{\ell} d_{\ell}$ and $\tilde{c} := \min(6, \tilde{c})$ with \tilde{c} the constant in the concentration inequality (3.1), then the graph G obtained by applying TSC to \mathcal{X} has no false connections with probability at least $1 - 14N^{-1} - 2N e^{-m} - \sum_{\ell=1}^L n_{\ell} e^{-c(n_{\ell}-1)}$, where $c > 1/20$ is a numerical constant.

Theorem 4.1 states that in the noisy case—just as in the noiseless case—TSC succeeds for p order-wise as small as d_{\max} , provided that the affinities between the subspaces \mathcal{S}_{ℓ} are sufficiently small and \mathcal{X} contains sufficiently many points from each subspace. More specifically, comparing the noiseless clustering condition (3.4) to (4.2), we can see that the impact of noise is simply to add the offset $\frac{\sigma(1+\sigma)\sqrt{6}}{\sqrt{\tilde{c}} \log N} \frac{\sqrt{d_{\max}}}{\sqrt{p}}$ to the LHS of the clustering condition. For fixed σ , owing to the factor $\sqrt{d_{\max}/p}$, the impact of noise on the effective affinity as quantified by the LHS of (4.2) becomes more pronounced when the dimensionality is reduced more aggressively.

Theorem 4.1 continues to hold (with \tilde{c} in the term $\frac{\sigma(1+\sigma)\sqrt{6}}{\sqrt{\tilde{c}} \log N} \frac{\sqrt{d_{\max}}}{\sqrt{p}}$ replaced by a numerical constant, and e^{-m} in the success probability replaced by e^{-p}), if noise $\tilde{\mathbf{e}}_i^{(\ell)} \sim \mathcal{N}(\mathbf{0}, (\sigma^2/p)\mathbf{I})$ is added *after* random projection according to $\tilde{\mathbf{x}}_j^{(\ell)} = \Phi \mathbf{y}_i^{(\ell)} + \tilde{\mathbf{e}}_i^{(\ell)}$. This is not surprising, as the absolute amount of noise injected remains the same, i.e., $\mathbb{E} \left[\left\| \tilde{\mathbf{e}}_i^{(\ell)} \right\|_2^2 \right] = \mathbb{E} \left[\left\| \mathbf{e}_i^{(\ell)} \right\|_2^2 \right] = \sigma^2$.

We finally note that an approach similar to that used for TSC can be applied to extend our result for SSC-OMP to the noisy case resulting in clustering conditions analogous to those for TSC. The corresponding technical details are, however, significantly more involved and cumbersome. We therefore decided not to state the formal result. We expect that a similar result can be proven for (a robust version of) SSC as our simulation results in Section 5.1.2 indicate that the qualitative behavior of all three algorithms in the presence of noise is essentially identical, and, in addition, is qualitatively accurately predicted by Theorem 4.1.

5. Numerical results

We evaluate the impact of dimensionality reduction on the CE, i.e., the fraction of misclustered points, for TSC, SSC and SSC-OMP applied to synthetic data as well as to publicly available standard data sets widely used in the subspace clustering literature. Specifically, we consider the problems of clustering faces, handwritten digits and gene expression data. All three algorithms, TSC, SSC and SSC-OMP, were observed to tolerate massive dimensionality reduction in all experiments. While our sufficient conditions for the no false connections property to hold do not allow to infer performance differences between the three algorithms, the algorithms' performance ranking varies across datasets, and no single algorithm is found to uniformly outperform the others. We also compare the algorithms in terms of their running times on a PC with 32 GB RAM and 8-core Intel Core i7-3770K CPU clocked at 3.50 GHz.

TSC and SSC-OMP were implemented in Matlab following the specifications in Section 2. For SSC, we used the Matlab implementation provided in [9], which is based on Lasso (instead of ℓ_1 -minimization) and uses the alternating direction method of multipliers. Code to reproduce the experiments in this section is available at <http://www.nari.ee.ethz.ch/commth/research/>. Information on the number of Monte Carlo runs used in our experiments is contained in this code.

Unless stated otherwise, we select the Lasso parameter λ in SSC from the set $\{0.001, 0.002, 0.004, 0.008, 0.01, 0.02, 0.04, 0.08, 0.1, 0.2\}$ such that the lowest CE is obtained on the original high-dimensional data set \mathcal{Y} . The parameters q and s_{\max} for TSC and SSC-OMP, respectively, are chosen in an analogous manner from the set $\{2, 4, \dots, 18\}$.

As projection matrices we consider i.i.d. $\mathcal{N}(0, 1/p)$ Gaussian random matrices (referred to as GRP) and FRP matrices [1] given by the real part of $\mathbf{FD} \in \mathbb{C}^{p \times m}$, where $\mathbf{D} \in \mathbb{R}^{m \times m}$ is diagonal with main diagonal elements drawn i.i.d. uniformly from $\{-1, 1\}$, and $\mathbf{F} \in \mathbb{C}^{p \times m}$ is obtained by choosing a set of p rows uniformly at random from the rows of an $m \times m$ discrete Fourier transform matrix. In all experiments, the dimensionality-reduced data set \mathcal{X} is obtained by applying the (same) realization of either a GRP or an FRP matrix to all data points in \mathcal{Y} . The FRP can be implemented efficiently by premultiplying \mathcal{Y} by \mathbf{D} and then applying the fast fourier transform (FFT) to each data point. With regards to storage space, we note that the FRP only requires the storage of a binary m -dimensional vector (namely the diagonal entries of \mathbf{D}), in contrast to mp real numbers for GRPs.

5.1 Synthetic data

5.1.1 Comparison of TSC, SSC and SSC-OMP We use the data model described in Section 3 with $m = 2^{15} = 32,768$ and generate $L = 3$ subspaces \mathcal{S}_ℓ of \mathbb{R}^m of dimension $d = 20$ at random such that every pair of subspaces intersects in at least r dimensions; this implies $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \geq \sqrt{r/d}$, for all $k, \ell \in [L], k \neq \ell$. More specifically, we take the basis matrices to be given by $\mathbf{U}^{(\ell)} = [\mathbf{U} \ \tilde{\mathbf{U}}^{(\ell)}]$, where $\mathbf{U} \in \mathbb{R}^{m \times r}$ and the $\tilde{\mathbf{U}}^{(\ell)} \in \mathbb{R}^{m \times (d-r)}$, $\ell \in [L]$, are chosen uniformly at random among all orthonormal matrices of dimensions $m \times r$ and $m \times (d-r)$, respectively. We sample $n_\ell = 80$ data points, for each $\ell \in [L]$, resulting in a total of $N = 240$ data points.

In Fig. 1, we plot the CE as a function of p for TSC, SSC and SSC-OMP applied to the dimensionality-reduced data set \mathcal{X} with $r = 4$ and $r = 8$. Fig. 2 shows the running times corresponding to the application of the FRP and the GRP matrix to the (entire) data set \mathcal{Y} along with the running times of the clustering algorithms alone.

The results show, as predicted by Theorems 3.1–3.3, that TSC, SSC and SSC-OMP, indeed, succeed provided that $\sqrt{d/p}$ is sufficiently small. Specifically, we observe a transition to $\text{CE} \approx 0$ for p between 20 and 100. As the subspaces \mathcal{S}_ℓ are of dimension 20, this corroborates the fact that the dimensionality of the data can be reduced down to the dimension of the subspaces without compromising clustering performance significantly. Equivalently, we accomplish a dimensionality reduction by a factor of about 1600–320.

For all three algorithms, the numerical results further confirm the tradeoff between the affinities of the \mathcal{S}_ℓ and the amount of dimensionality reduction possible as quantified by the clustering conditions (3.4)–(3.6). Specifically, the CE increases as r and hence $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell)$ increases. In this example, SSC consistently outperforms TSC and SSC-OMP, albeit at the cost of significantly longer running time (see Fig. 2). While the running time of SSC exhibits very pronounced increasing behavior in p , that of SSC-OMP shows much less pronounced increases, and that of TSC does not increase notably in p . It is furthermore interesting to see that the clustering performance is essentially identical for FRP and GRP. This is remarkable as the application of FRP requires only $O(m \log m)$ operations (per data point) and therefore its running time does not depend on p . Application of the GRP, in contrast, requires $O(mp)$ operations (per data point), which results in a running time that is linear in p .

5.1.2 Impact of noise In the next experiment, we study the interplay between noise and dimensionality reduction. We use the data model described in Section 3 with $m = 100$ and generate $L = 2$ orthogonal




FIG. 1. CE for synthetic data as a function of p using GRP (left) and FRP (right). Recall that for $r = 4$ and $r = 8$, we have $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \geq \sqrt{1/5}$ and $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \geq \sqrt{2/5}$, respectively, for all $k, \ell \in [L]$, $k \neq \ell$.




FIG. 2. Running times (in seconds) for clustering synthetic data.

subspaces \mathcal{S}_ℓ of \mathbb{R}^m of dimension $d = 10$. This ensures that the affinity between the subspaces equals 0 (fixing the affinity to some other constant would not change the qualitative conclusions). We generate the noisy data set \mathcal{Y} by sampling $n_\ell = 30$ points from each of the two subspaces and adding $\mathcal{N}(\mathbf{0}, (\sigma^2/m)\mathbf{I})$ noise according to $\tilde{\mathbf{y}}_i^{(\ell)} = \mathbf{y}_i^{(\ell)} + \mathbf{e}_i^{(\ell)}$. Fig. 3 shows the CE as a function of $\sqrt{d/p}$ and σ for dimensionality reduction via GRP.

The clustering condition in Theorem 4.1 guarantees that TSC succeeds as long as $\sqrt{d/p}(c_1 + \sigma(c_2 + \sigma)) \leq c_3$, where c_1, c_2, c_3 are independent of d, p, m and σ^2 . In order to find out whether this sufficient condition for the no false connections property to hold predicts the fundamental clustering behavior qualitatively correctly, we test whether a phase transition, separating the region where clustering succeeds from that where it fails, indeed, occurs at

$$\sqrt{d/p}(c_1 + \sigma(c_2 + \sigma)) = c_3. \quad (5.1)$$




FIG. 3. CE (color coded) as a function of $\sqrt{d/p}$ and σ for $L = 2$ subspaces of \mathbb{R}^{100} with $\text{aff}(\mathcal{S}_1, \mathcal{S}_2) = 0$ (top row) and $\text{aff}(\mathcal{S}_1, \mathcal{S}_2) = \sqrt{1/5}$ (bottom row). The black lines correspond to the curve $\sqrt{d/p}(c_1 + \sigma(c_2 + \sigma)) \leq c_3$ for suitably chosen constants c_1, c_2, c_3 , and roughly separate the success regions from the failure regions.

To this end, we fit (5.1) manually—by choosing c_1, c_2, c_3 —to the plots in Fig. 3 and observe that the answer is in the affirmative. Moreover, our numerical results show that the phase transition behavior of SSC and SSC-OMP is essentially identical to that of TSC, which provides evidence for SSC and SSC-OMP behaving similarly to TSC in the noisy case.

5.1.3 Dimensionality reduction when the union of the subspaces equals the ambient space As noted in Section 3, Theorems 3.1–3.3 and Theorem 4.1 indicate that dimensionality reduction down to the order of the subspace dimensions is possible even when the union of the subspaces \mathcal{S}_ℓ equals the ambient space. To verify this observation empirically, we perform the following experiment. We draw a random Gaussian matrix $\mathbf{V} \in \mathbb{R}^{200 \times 200}$. With probability one, the columns of \mathbf{V} span \mathbb{R}^{200} . We then extract the 200×20 matrices $\mathbf{V}^{(\ell)}$ from \mathbf{V} according to $[\mathbf{V}^{(1)} \dots \mathbf{V}^{(10)}] = \mathbf{V}$, and let the subspace \mathcal{S}_ℓ be given by the span of $\mathbf{V}^{(\ell)}$, $\ell = 1, \dots, 10$. This guarantees that the union of the \mathcal{S}_ℓ is \mathbb{R}^{200} . Note, however, that the affinities between pairs of the resulting subspaces will be small with high probability. We again use the data model described in Section 3 and sample $n_\ell = 60$ points on $\mathcal{S}_\ell \cap \mathbb{S}^{d_\ell-1}$, for all $\ell \in [L]$, to obtain a data set \mathcal{Y} with a total of $N = 600$ points. We select the values for q , λ and s_{\max} that yield the lowest CE for the majority of values for p .

Fig. 4 shows the CE as a function of p for TSC, SSC and SSC-OMP. The CE starts to be non-zero for $p < 60$ for TSC and SSC-OMP, and for $p < 40$ for SSC. We therefore conclude that the dimensionality can, as indicated by Theorems 3.1 and 4.1, indeed, be reduced quite significantly, even when the union of the subspaces equals the ambient space.




FIG. 4. CE as a function of p for a union of $L = 10$ subspaces; this union equals the ambient space \mathbb{R}^{200} .

5.2 Clustering faces

We next evaluate the impact of dimensionality reduction in the problem of clustering face images taken from the Extended Yale B data set [11,24], which contains 192×168 pixel ($m = 32,256$) frontal face images of 38 individuals, with 64 images per individual, each acquired under different illumination conditions. The motivation for applying subspace clustering algorithms to this problem stems from the insight that the vectorized images of a given face taken under varying illumination conditions lie approximately in a nine-dimensional linear subspace [3]. Each nine-dimensional subspace \mathcal{S}_i would then contain the images corresponding to a given person.

We generate \mathcal{V} by first selecting a subset of $L = 2$ individuals uniformly at random from the set of all $\binom{38}{2}$ pairs and then collecting all images corresponding to the two selected individuals. In Fig. 5, we plot the corresponding CE and the running times as a function of p . Again, for each p , the CE and the running times are obtained by averaging over 500 problem instances generated by randomly choosing 100 instances of \mathcal{V} and 5 realizations of the projection matrix per chosen data set \mathcal{V} . In contrast to the preceding experiment, here, SSC-OMP consistently outperforms TSC and SSC. For all three algorithms, the dimensionality of the data can be reduced by a factor of about 100 without notably increasing the CE. Note, however, that in this experiment the dimensionality cannot be reduced as aggressively as in the preceding synthetic data experiment. Specifically, here the data points lie in nine-dimensional subspaces and dimensionality reduction by a factor of 100 corresponds to $p \approx 322$. One possible explanation for this observation is that the principal angles between the subspaces spanned by the face images of different subjects are typically small (see [9, Section 7]), which means that the subspace affinities in this data set are large. Another possible explanation is that the data points only approximately lie in nine-dimensional subspaces. The conclusions regarding running times and choice of the random projection matrix are analogous to those reported for synthetic data above.

5.3 Clustering handwritten digits

In this experiment, we investigate the impact of dimensionality reduction in the context of clustering images of handwritten digits. We use the MNIST data set [23] containing 10,000 images of (horizontally and vertically) aligned handwritten digits of size 28×28 pixels ($m = 784$). The motivation for employing subspace clustering in this context stems from the observation that vectorized images of different handwritten versions of the same digit tend to lie near a low-dimensional subspace [14].

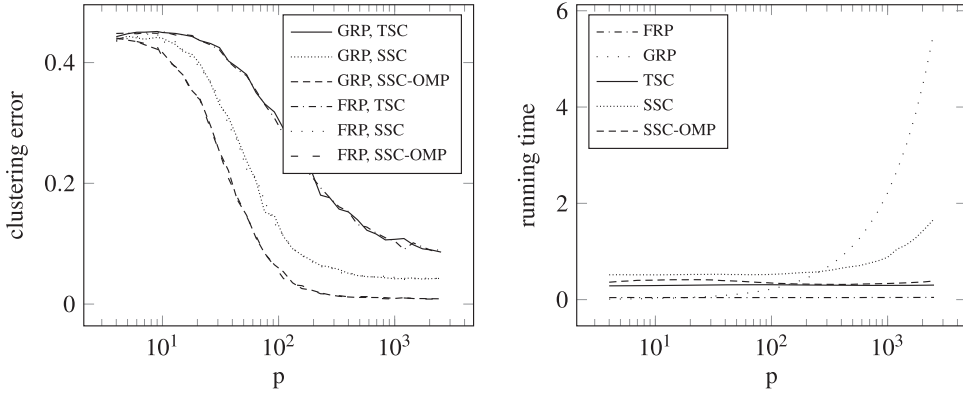


FIG. 5. CE and running times (in seconds) for clustering $L = 2$ faces from the Extended Yale B data set.




FIG. 6. CE for handwritten digits 2, 4 and 8 from the MNIST data set.

We generate the data sets \mathcal{V} by selecting 250 images (out of 1000) uniformly at random from each of the sets corresponding to the digits 2, 4 and 8. There is no specific reason for our choice of the digits 2, 4 and 8; other combinations of three digits yield similar results. However, some combinations of digits are more difficult to cluster than others, e.g., 1 and 7 are “closer” (in terms of the affinities between the subspaces the corresponding images approximately lie in) than 1 and 8; clustering 1 and 7 therefore typically results in a larger error than clustering 1 and 8. The results depicted in Fig. 6 show that the dimensionality of the data set can be reduced from $m = 784$ to $p = 200$, i.e., by a factor of 3.9, without notably increasing the CE incurred by TSC and SSC. For sufficiently large p , TSC yields a slightly lower CE than SSC. SSC-OMP is outperformed considerably by the other two algorithms.

5.4 Clustering gene expression data

Finally, we consider clustering of gene expression level data—originating from different types of cancer cells—according to cancer type. This problem is of significant practical relevance as it helps, *inter alia*, to identify genes that are involved in the same cellular process [20]. The use of subspace clustering in this context was suggested in [20]. We use the publicly available Novartis multi-tissue data set from the Broad Institute Cancer Program database [5]. This data set contains the 1000-dimensional gene expression level data of $n = 103$ tissue samples taken from $L = 4$ different cancer types. In order to




FIG. 7. CE for gene expression level data of $L = 4$ cancer types (left). Singular values of data matrices corresponding to a single cancer type (right).

illustrate that the gene expression level vectors of a single cancer type, indeed, lie near a low-dimensional subspace, we plot, in Fig. 7, the singular values of the data matrices corresponding to a single cancer type. We observe that the singular values decay rapidly and for every cancer type, more than 94% of the energy of the corresponding data vectors is concentrated in a six-dimensional subspace of the 1000-dimensional ambient space.

We cluster all $n = 103$ available samples. The CE obtained by averaging, for each p , over 200 realizations of the random projection matrix is shown in Fig. 7. For $p \approx 100$, which corresponds to dimensionality reduction by a factor of 10, the CEs of TSC and SSC are comparable to those obtained when operating on the original high-dimensional data set. SSC is seen to consistently (across p) perform best, followed by TSC and SSC-OMP. As in previous experiments the CEs observed for GRP and FRP, for each of the three algorithms, are virtually identical.

Acknowledgments

The authors would like to thank R. Calderbank and the anonymous reviewers for very helpful comments on earlier versions of this manuscript.

REFERENCES

1. AILON, N. & LIBERTY, E. (2013) An almost optimal unrestricted fast Johnson-Lindenstrauss transform. *ACM Trans. Algorithms*, **9**, 1–12.
2. ALONSO-GUTIÉRREZ, D. (2008) On the isotropy constant of random convex sets. *Proc. Am. Math. Soc.*, **136**, 3293–3300.
3. BASRI, R. & JACOBS, D. W. (2003) Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, **25**, 218–233.
4. BOYD, S. & VANDENBERGHE, L. (2004) *Convex Optimization*. Cambridge, UK: Cambridge University Press.
5. BROAD-INSTITUTE (2013) Cancer program data sets. <http://www.broadinstitute.org/cgi-bin/cancer/datasets.cgi>, last accessed 3 January 2017.
6. CANDÈS, E. J., ROMBERG, J. & TAO, T. (2006) Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theory*, **52**, 489–509.
7. DYER, E. L., SANKARANARAYANAN, A. C. & BARANIUK, R. G. (2013) Greedy feature selection for subspace clustering. *J. Mach. Learn. Res.*, **14**, 2487–2517.

8. ELHAMIFAR, E. & VIDAL, R. (2009) Sparse subspace clustering. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Miami Beach, Florida, pp. 2790–2797.
9. ELHAMIFAR, E. & VIDAL, R. (2013) Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Trans. Pattern Anal. Mach. Intell.*, **35**, 2765–2781.
10. FOUCART, S. & RAUHUT, H. (2013) *A Mathematical Introduction to Compressive Sensing*. Berlin: Springer.
11. GEORGHIADES, A. S., BELHUMEUR, P. N. & KRIEGMAN, D. J. (2001) From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Mach. Intell.*, **23**, 643–660.
12. GOLUB, G. H. & VAN LOAN, C. F. (1996) *Matrix Computations*. Baltimore: JHU Press.
13. GRITZMANN, P. & KLEE, V. (1992) Inner and outer j -radii of convex bodies in finite-dimensional normed spaces. *Discrete Comput. Geom.*, **7**, 255–280.
14. HASTIE, T. & SIMARD, P. Y. (1998) Metrics and models for handwritten character recognition. *Stat. Sci.*, **13**, 54–65.
15. HECKEL, R., AGUSTSSON, E. & BÖLCSKEI, H. (2014) Neighborhood selection for thresholding based subspace clustering. *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Florence, Italy, pp. 6761–6765.
16. HECKEL, R. & BÖLCSKEI, H. (2015) Robust subspace clustering via thresholding. *IEEE Trans. Inform. Theory*, **61**, 6320–6342.
17. HECKEL, R., TSCHANNEN, M. & BÖLCSKEI, H. (2014) Subspace clustering of dimensionality-reduced data. *Proceedings of IEEE International Symposium on Information Theory*, Honolulu, HI, pp. 2997–3001.
18. HORN, R. A. & JOHNSON, C. R. (2012) *Matrix Analysis*. Cambridge, UK: Cambridge University Press.
19. HUANG, J., QIU, Q. & CALDERBANK, R. (2016) The role of principal angles in subspace classification. *IEEE Trans. Signal Process.*, **64**, 1933–1945.
20. JIANG, D., TANG, C. & ZHANG, A. (2004) Cluster analysis for gene expression data: A survey. *IEEE Trans. Knowl. Data Eng.*, **16**, 1370–1386.
21. KRAHMER, F. & WARD, R. (2011) New and improved Johnson-Lindenstrauss embeddings via the restricted isometry property. *SIAM J. Math. Anal.*, **43**, 1269–1281.
22. LAPIDOTH, A. (2009) *A Foundation in Digital Communication*. Cambridge, UK: Cambridge University Press.
23. LECUN, Y. & CORTES, C. (2013) The MNIST database. <http://yann.lecun.com/exdb/mnist/>, last accessed 3 January 2017.
24. LEE, K. C., HO, J. & KRIEGMAN, D. J. (2005) Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. Pattern Anal. Mach. Intell.*, **27**, 684–698.
25. LIU, K., KARGUPTA, H. & RYAN, J. (2006) Random projection-based multiplicative data perturbation for privacy preserving distributed data mining. *IEEE Trans. Knowl. Data Eng.*, **18**, 92–106.
26. NG, A., JORDAN, I. M. & YAIR, W. (2001) On spectral clustering: Analysis and an algorithm. *Advances in Neural Information Processing Systems*, Vancouver, Canada, pp. 849–856.
27. NOKLEBY, M., RODRIGUES, M. & CALDERBANK, R. (2015) Discrimination on the Grassmann manifold: Fundamental limits of subspace classifiers. *IEEE Trans. Inform. Theory*, **61**, 2133–2147.
28. SOLTANOLKOTABI, M. & CANDÈS, E. J. (2012) A geometric analysis of subspace clustering with outliers. *Ann. Stat.*, **40**, 2195–2238.
29. SOLTANOLKOTABI, M., ELHAMIFAR, E. & CANDÈS, E. J. (2014) Robust subspace clustering. *Ann. Stat.*, **42**, 669–699.
30. SPIELMAN, D. (2012) Spectral graph theory. Lecture notes. <http://www.cs.yale.edu/homes/spielman/561/>, last accessed 3 January 2015.
31. TSCHANNEN, M. (2014) Dimensionality reduction for sparse subspace clustering. *MS thesis*, Zurich, Switzerland: ETH Zurich. <http://dx.doi.org/10.3929/ethz-a-010481065>.
32. VEMPALA, S. S. (2005) *The Random Projection Method*. American Mathematical Society.
33. VERSHYNIN, R. (2012) Introduction to the non-asymptotic analysis of random matrices. *Compressed Sensing: Theory and Applications* (Y. C. Eldar & G. Kutyniok eds). Cambridge, UK: Cambridge University Press, pp. 210–268.

34. VIDAL, R. (2011) Subspace clustering. *IEEE Signal Process. Mag.*, **28**, 52–68.
35. VON LUXBURG, U. (2007) A tutorial on spectral clustering. *Stat. Comput.*, **17**, 395–416.
36. WANG, Y., WANG, Y. & SINGH, A. (2015) A deterministic analysis of noisy sparse subspace clustering for dimensionality-reduced data. *Proceedings of International Conference on Machine Learning (ICML)*, Lille, France, pp. 1422–1431.
37. YOU, C., ROBINSON, D. P. & VIDAL, R. (2016) Scalable sparse subspace clustering by orthogonal matching pursuit. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV, pp. 3918–3927.
38. ZHANG, T., SZLAM, A., WANG, Y. & LERMAN, G. (2012) Hybrid linear modeling via local best-fit flats. *Int. J. Comput. Vision*, **100**, 217–240.

Appendix A. Proof of Theorems 3.1 and 4.1

The proof idea for Theorem 3.1 is to turn the effect of the random projection into an additive perturbation and to show that this perturbation is small for all values of p down to the order of d_{\max} . In the noisy case, addressed by Theorem 4.1, we have an additional perturbation due to noise. We detail the proof of the more general Theorem 4.1 below, and explain in Appendix A.3 the simple changes that yield Theorem 3.1. The proof of Theorem 4.1 follows closely that of [16, Theorem 3], which quantifies the performance of TSC under additive Gaussian noise alone. We therefore elaborate only on the steps that are new relative to [16] and encourage the interested reader to consult [16] for the arguments not repeated here.

The graph G obtained by applying TSC to the dimensionality-reduced noisy data set $\tilde{\mathcal{X}}$ has no false connections, i.e., each $\tilde{\mathbf{x}}_i^{(\ell)}$ is connected to points in $\tilde{\mathcal{X}}_\ell$ only, if for each $\tilde{\mathbf{x}}_i^{(\ell)} \in \tilde{\mathcal{X}}_\ell$ the associated set S_i corresponds to points in $\tilde{\mathcal{X}}_\ell$ only, for all ℓ . This is the case if

$$z_{(n_\ell - q)}^{(\ell)} > \max_{k \neq \ell, j} z_j^{(k)}, \quad (\text{A.1})$$

where $z_j^{(k)} := |\langle \tilde{\mathbf{x}}_j^{(k)}, \tilde{\mathbf{x}}_i^{(\ell)} \rangle|$ and $z_{(1)}^{(\ell)} \leq z_{(2)}^{(\ell)} \leq \dots \leq z_{(n_\ell - 1)}^{(\ell)}$ are the order statistics of $\{z_j^{(\ell)}\}_{j \in [n_\ell] \setminus \{i\}}$ and $\max_{k \neq \ell, j}$ denotes maximization over $k \in [L]$, $k \neq \ell$, and over the indices j of the corresponding points $\tilde{\mathbf{x}}_j^{(k)} \in \tilde{\mathcal{X}}_k$. Note that, for simplicity of exposition, the notation $z_j^{(k)}$ does not reflect dependence on $\tilde{\mathbf{x}}_i^{(\ell)}$. The proof is established by upper-bounding the probability of (A.1) being violated for a given data point $\tilde{\mathbf{x}}_i^{(\ell)}$. A union bound over all N points $\tilde{\mathbf{x}}_i^{(\ell)}$, $i \in [n_\ell]$, $\ell \in [L]$, then yields the final result. We start by setting $\tilde{z}_j^{(k)} := |\langle \mathbf{y}_j^{(k)}, \mathbf{y}_i^{(\ell)} \rangle|$, where $\mathbf{y}_j^{(k)} = \mathbf{U}^{(k)} \mathbf{a}_j^{(k)}$ are the original data points in the (high-dimensional) space \mathbb{R}^m , and noting that $z_j^{(k)} = |\langle \tilde{\mathbf{x}}_j^{(k)}, \tilde{\mathbf{x}}_i^{(\ell)} \rangle| = |\langle \mathbf{y}_j^{(k)}, \mathbf{y}_i^{(\ell)} \rangle + e_j^{(k)}|$, where we defined the “distortion”

$$\begin{aligned} e_j^{(k)} &:= \langle \Phi \tilde{\mathbf{y}}_j^{(k)}, \Phi \tilde{\mathbf{y}}_i^{(\ell)} \rangle - \langle \mathbf{y}_j^{(k)}, \mathbf{y}_i^{(\ell)} \rangle \\ &= \langle (\Phi^T \Phi - \mathbf{I}) \mathbf{y}_j^{(k)}, \mathbf{y}_i^{(\ell)} \rangle + \langle \Phi \mathbf{y}_j^{(k)}, \Phi \mathbf{e}_i^{(\ell)} \rangle + \langle \Phi \mathbf{e}_j^{(k)}, \Phi \mathbf{y}_i^{(\ell)} \rangle + \langle \Phi \mathbf{e}_j^{(k)}, \Phi \mathbf{e}_i^{(\ell)} \rangle \\ &= \underbrace{\langle \mathbf{U}^{(\ell)T} (\Phi^T \Phi - \mathbf{I}) \mathbf{U}^{(k)} \mathbf{a}_j^{(k)}, \mathbf{a}_i^{(\ell)} \rangle}_{\tilde{e}_j^{(k)}} \\ &\quad + \underbrace{\langle \Phi^T \Phi \mathbf{y}_j^{(k)}, \mathbf{e}_i^{(\ell)} \rangle + \langle \mathbf{e}_j^{(k)}, \Phi^T \Phi \mathbf{y}_i^{(\ell)} \rangle + \langle \Phi^T \Phi \mathbf{e}_j^{(k)}, \mathbf{e}_i^{(\ell)} \rangle}_{\tilde{e}_j^{(k)}}. \end{aligned} \quad (\text{A.2})$$

Here, the term $\tilde{e}_j^{(k)}$ accounts for the perturbation caused by random projection, and $\tilde{e}_j^{(k)}$ corresponds to the perturbation caused by noise. The probability of (A.1) being violated can now be upper-bounded according to

$$\begin{aligned} \mathbb{P}\left[z_{(n_\ell-q)}^{(\ell)} \leq \max_{k \neq \ell, j} z_j^{(k)}\right] &\leq \mathbb{P}\left[\bar{z}_{(n_\ell-q)}^{(\ell)} \leq \frac{2}{3\sqrt{d_\ell}}\right] \\ &\quad + \mathbb{P}\left[\max_{k \neq \ell, j} \bar{z}_j^{(k)} \geq \alpha\right] + \mathbb{P}\left[\max_{(j,k) \neq (i,\ell)} |e_j^{(k)}| \geq \epsilon\right], \end{aligned} \quad (\text{A.3})$$

where we set

$$\alpha := \frac{4\sqrt{6} \log N}{\sqrt{d_\ell}} \max_{k \neq \ell} \frac{1}{\sqrt{d_k}} \left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F, \quad \epsilon := \frac{\beta}{\sqrt{d_\ell}} \delta + \beta \frac{2\sigma(1+\sigma)}{\sqrt{m}} \delta' \quad (\text{A.4})$$

with $\beta := \sqrt{6 \log N}$, $\delta := \frac{\sqrt{28d_{\max} + 8 \log L + 8 \log N}}{\sqrt{3\tilde{c}p}}$ and $\delta' := \sqrt{\frac{6m}{\tilde{c}p}}$, and assumed

$$\alpha + 2\epsilon \leq \frac{2}{3\sqrt{d_\ell}}. \quad (\text{A.5})$$

We refer the reader to [16, Proof of Theorem 3, Equation (40)] for an explanation of the steps leading to (A.3) (while [16, Equation (40)] is not completely equivalent to (A.3), the steps leading to (A.3) are essentially identical). Resolving the assumption (A.5) leads to

$$\max_{k \neq \ell} \frac{1}{\sqrt{d_k}} \left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F + \frac{\delta}{2\sqrt{\log N}} + \sigma(1+\sigma) \sqrt{\frac{6d_\ell}{\tilde{c}p \log N}} \leq \frac{2}{3 \cdot 4\sqrt{6} \log N},$$

which is implied by (4.2) (using that $\sqrt{28d_{\max} + 8 \log L + 8 \log N} / \sqrt{\log N} \leq \sqrt{44d_{\max}}$, because $\log(L)/\log(N) \leq 1$, $d_{\max} \geq 1$, and $\log N > 1$ for $N \geq 3$). With ϵ as defined in (A.4), and the triangle inequality, it follows that $\max_{(j,k) \neq (i,\ell)} |e_j^{(k)}| \geq \epsilon$ implies that either $\max_{(j,k) \neq (i,\ell)} |\tilde{e}_j^{(k)}| \geq \frac{\beta}{\sqrt{d_\ell}} \delta$ or $\max_{(j,k) \neq (i,\ell)} |\tilde{e}_j^{(k)}| \geq \beta \frac{2\sigma(1+\sigma)}{\sqrt{m}} \delta'$, or both. Therefore, by a union bound argument

$$\mathbb{P}\left[\max_{(j,k) \neq (i,\ell)} |e_j^{(k)}| \geq \epsilon\right] \leq \mathbb{P}\left[\max_{(j,k) \neq (i,\ell)} |\tilde{e}_j^{(k)}| \geq \frac{\beta}{\sqrt{d_\ell}} \delta\right] + \mathbb{P}\left[\max_{(j,k) \neq (i,\ell)} |\tilde{e}_j^{(k)}| \geq \beta \frac{2\sigma(1+\sigma)}{\sqrt{m}} \delta'\right]. \quad (\text{A.6})$$

Here, the first and second term on the RHS of (A.6) correspond to the perturbation caused by random projection and by noise, respectively. As established in Sections A.1 and A.2, these terms can be upper-bounded by $\frac{4}{N^2}$ and $2e^{-m} + \frac{7}{N^2}$, respectively, which yields

$$\mathbb{P}\left[\max_{(j,k) \neq (i,\ell)} |e_j^{(k)}| \geq \epsilon\right] \leq \frac{4}{N^2} + 2e^{-m} + \frac{7}{N^2}. \quad (\text{A.7})$$

The remaining terms on the RHS of (A.3) are upper-bounded as shown in Steps 3 and 2 in [16, Proof of Theorem 3], respectively, using standard concentration of measure results, according to

$$\mathbb{P}\left[\bar{z}_{(n_\ell-q)}^{(\ell)} \leq \frac{2}{3\sqrt{d_\ell}}\right] \leq e^{-c(n_\ell-1)}$$

and

$$\mathbb{P}\left[\max_{k \neq \ell, j} \bar{z}_j^{(k)} \geq \alpha\right] \leq 3N^{-2},$$

where $c > 1/20$ is a numerical constant, and we employed the assumption $n_\ell \geq 6q$, for all ℓ .

With (A.7) we thus get that (A.1) is violated with probability at most $e^{-c(n_\ell-1)} + 2e^{-m} + 14N^{-2}$. Taking the union bound over all points $\mathbf{x}_i^{(\ell)}$, $i \in [n_\ell]$, $\ell \in [L]$, finishes the proof.

A.1 Perturbation caused by random projection

We next show that the first term on the RHS of (A.6) is upper-bounded by $4/N^2$. This term corresponds to the perturbation caused by random projection. For notational convenience, we set $\mathbf{B}_{k,\ell} = \mathbf{U}^{(\ell)T}(\Phi^T \Phi - \mathbf{I})\mathbf{U}^{(k)}$ and note that

$$\begin{aligned} \mathbb{P}\left[\max_{(j,k) \neq (i,\ell)} |\bar{e}_j^{(k)}| \geq \frac{\beta}{\sqrt{d_\ell}} \delta\right] &= \mathbb{P}\left[\max_{(j,k) \neq (i,\ell)} \left|\langle \mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}, \mathbf{a}_i^{(\ell)} \rangle\right| \geq \frac{\beta}{\sqrt{d_\ell}} \delta\right] \\ &= \mathbb{P}\left[\bigcup_{(j,k) \neq (i,\ell)} \left\{ \left|\langle \mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}, \mathbf{a}_i^{(\ell)} \rangle\right| \geq \frac{\beta}{\sqrt{d_\ell}} \delta \right\}\right] \\ &\leq \mathbb{P}\left[\bigcup_{(j,k) \neq (i,\ell)} \left\{ \left|\langle \mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}, \mathbf{a}_i^{(\ell)} \rangle\right| \geq \|\mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}\|_2 \frac{\beta}{\sqrt{d_\ell}} \right\} \cup \left\{ \|\mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}\|_2 \geq \delta \right\}\right] \\ &\leq \mathbb{P}\left[\bigcup_{(j,k) \neq (i,\ell)} \left\{ \left|\langle \mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}, \mathbf{a}_i^{(\ell)} \rangle\right| \geq \|\mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}\|_2 \frac{\beta}{\sqrt{d_\ell}} \right\} \cup \left\{ \|\mathbf{B}_{k,\ell}\|_{2 \rightarrow 2} \geq \delta \right\}\right] \end{aligned} \quad (\text{A.8})$$

$$\leq \mathbb{P}\left[\max_k \|\mathbf{B}_{k,\ell}\|_{2 \rightarrow 2} \geq \delta\right] + \sum_{(j,k) \neq (i,\ell)} \mathbb{P}\left[\left|\langle \mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}, \mathbf{a}_i^{(\ell)} \rangle\right| \geq \|\mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}\|_2 \frac{\beta}{\sqrt{d_\ell}}\right] \quad (\text{A.9})$$

$$\leq 2e^{-\tau/2} + N2e^{-\frac{6 \log N}{2}} = \frac{4}{N^2}, \quad (\text{A.10})$$

where (A.8) follows from $\|\mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}\|_2 \leq \|\mathbf{B}_{k,\ell}\|_{2 \rightarrow 2}$, (A.9) is by the union bound, and (A.10) follows from (B.8) in Appendix B with $\tau = 4 \log N$ and Proposition A.1 below with $\mathbf{a} = \mathbf{a}_i^{(\ell)}$, $\mathbf{b} = \mathbf{B}_{k,\ell} \mathbf{a}_j^{(k)}$, $d = d_\ell$ and $\beta = \sqrt{6 \log N}$.

PROPOSITION A.1 (E.g., [33, Example 5.25]) Let \mathbf{a} be uniformly distributed on \mathbb{S}^{d-1} and fix $\mathbf{b} \in \mathbb{R}^d$. Then, for $\beta \geq 0$, we have

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

A.2 Perturbation caused by noise

In this section, we deal with the perturbation caused by noise. Specifically, we establish that the second term on the RHS of (A.6) satisfies

$$\mathbb{P}\left[\max_{(j,k) \neq (i,\ell)} |\tilde{e}_j^{(k)}| \geq \beta \frac{2\sigma(1+\sigma)}{\sqrt{m}} \delta'\right] \leq 2e^{-m} + \frac{7}{N^2}. \quad (\text{A.11})$$

For notational convenience, we set $\tilde{\mathbf{y}}_j^{(k)} = \Phi^T \Phi \mathbf{y}_j^{(k)}$ and drop the indices i and ℓ to write $\mathbf{y} = \mathbf{y}_i^{(\ell)}$, $\tilde{\mathbf{y}} = \tilde{\mathbf{y}}_i^{(\ell)}$, $\mathbf{e} = \mathbf{e}_i^{(\ell)}$. We first note that

$$\begin{aligned} \left\{ \max_{(j,k) \neq (i,\ell)} |\tilde{e}_j^{(k)}| \geq \beta \frac{2\sigma(1+\sigma)}{\sqrt{m}} \delta' \right\} &= \bigcup_{(j,k) \neq (i,\ell)} \left\{ |\tilde{e}_j^{(k)}| \geq \beta \frac{2\sigma(1+\sigma)}{\sqrt{m}} \delta' \right\} \\ &\subseteq \bigcup_{(j,k) \neq (i,\ell)} \left\{ \left| \langle \tilde{\mathbf{y}}_j^{(k)}, \mathbf{e} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \delta' \right\} \cup \left\{ \left| \langle \mathbf{e}_j^{(k)}, \tilde{\mathbf{y}} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \delta' \right\} \cup \left\{ \left| \langle \Phi^T \Phi \mathbf{e}_j^{(k)}, \mathbf{e} \rangle \right| \geq \beta \frac{2\sigma^2}{\sqrt{m}} \delta' \right\}, \quad (\text{A.12}) \\ &\subseteq \left\{ \|\Phi^T \Phi\|_{2 \rightarrow 2} \geq \delta' \right\} \cup \bigcup_{(j,k) \neq (i,\ell)} \left[\left\{ \left| \langle \tilde{\mathbf{y}}_j^{(k)}, \mathbf{e} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \|\tilde{\mathbf{y}}_j^{(k)}\|_2 \right\} \cup \left\{ \left| \langle \mathbf{e}_j^{(k)}, \tilde{\mathbf{y}} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \|\tilde{\mathbf{y}}\|_2 \right\} \right. \\ &\quad \left. \cup \left\{ \left| \langle \Phi^T \Phi \mathbf{e}_j^{(k)}, \mathbf{e} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \|\Phi^T \Phi \mathbf{e}_j^{(k)}\|_2 \right\} \cup \left\{ \|\mathbf{e}_j^{(k)}\|_2 \geq 2\sigma \right\} \right]. \quad (\text{A.13}) \end{aligned}$$

Here, (A.12) follows from the triangle inequality. To verify (A.13), consider the first event in (A.12) and note that

$$\left\{ \left| \langle \tilde{\mathbf{y}}_j^{(k)}, \mathbf{e} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \delta' \right\} \subseteq \left\{ \|\Phi^T \Phi\|_{2 \rightarrow 2} \geq \delta' \right\} \cup \left\{ \left| \langle \tilde{\mathbf{y}}_j^{(k)}, \mathbf{e} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \|\tilde{\mathbf{y}}_j^{(k)}\|_2 \right\}. \quad (\text{A.14})$$

To see this, simply take the complement of (A.14) according to

$$\left\{ \|\Phi^T \Phi\|_{2 \rightarrow 2} < \delta' \right\} \cap \left\{ \left| \langle \tilde{\mathbf{y}}_j^{(k)}, \mathbf{e} \rangle \right| < \beta \frac{\sigma}{\sqrt{m}} \|\tilde{\mathbf{y}}_j^{(k)}\|_2 \right\} \subseteq \left\{ \left| \langle \tilde{\mathbf{y}}_j^{(k)}, \mathbf{e} \rangle \right| < \beta \frac{\sigma}{\sqrt{m}} \delta' \right\},$$

where we used

$$\|\tilde{\mathbf{y}}_j^{(k)}\|_2 = \|\Phi^T \Phi \mathbf{y}_j^{(k)}\|_2 \leq \|\Phi^T \Phi\|_{2 \rightarrow 2} \|\mathbf{y}_j^{(k)}\|_2 = \|\Phi^T \Phi\|_{2 \rightarrow 2}.$$

Treating the second and the third event in (A.12) similarly establishes (A.13). A union bound argument now yields

$$\mathbb{P}\left[\max_{(j,k) \neq (i,\ell)} |\tilde{e}_j^{(k)}| \geq \beta \frac{2\sigma(1+\sigma)}{\sqrt{m}} \delta'\right] \leq \mathbb{P}[\|\Phi^T \Phi\|_2 \geq \delta'] \quad (\text{A.15})$$

$$+ \sum_{(j,k) \neq (i,\ell)} \mathbb{P}\left[\left| \langle \tilde{\mathbf{y}}_j^{(k)}, \mathbf{e} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \|\tilde{\mathbf{y}}_j^{(k)}\|_2\right] \quad (\text{A.16})$$

$$+ \sum_{(j,k) \neq (i,\ell)} \mathbb{P} \left[\left| \langle \mathbf{e}_j^{(k)}, \bar{\mathbf{y}} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \|\bar{\mathbf{y}}\|_2 \right] \quad (\text{A.17})$$

$$+ \sum_{(j,k) \neq (i,\ell)} \mathbb{P} \left[\left| \langle \Phi^T \Phi \mathbf{e}_j^{(k)}, \mathbf{e} \rangle \right| \geq \beta \frac{\sigma}{\sqrt{m}} \|\Phi^T \Phi \mathbf{e}_j^{(k)}\|_2 \right] \quad (\text{A.18})$$

$$+ \sum_{(j,k) \neq (i,\ell)} \mathbb{P} \left[\|\mathbf{e}_j^{(k)}\|_2 \geq 2\sigma \right] \quad (\text{A.19})$$

$$\leq 2e^{-m} + 6N e^{-\frac{\beta^2}{2}} + N e^{-\frac{\beta^2}{2}}. \quad (\text{A.20})$$

To get (A.20), we upper-bounded the terms on the RHSs of (A.15)–(A.19) as follows. For the RHS of (A.15), we note that

$$\mathbb{P}[\|\Phi^T \Phi\|_2 \geq \delta'] \leq 2e^{-m},$$

which is a consequence of Theorem B.2 stated in Appendix B below. Specifically, with $1 \leq \sqrt{\frac{6m}{\tilde{c}p}}$, which follows from $\tilde{c} = \min(6, \tilde{c}) \leq 6$ and $p \leq m$, both by assumption, we have

$$\begin{aligned} \mathbb{P} \left[\|\Phi^T \Phi\|_{2 \rightarrow 2} \geq \sqrt{\frac{24m}{\tilde{c}p}} \right] &\leq \mathbb{P} \left[\|\Phi^T \Phi\|_{2 \rightarrow 2} \geq 1 + \sqrt{\frac{6m}{\tilde{c}p}} \right] \\ &\leq \mathbb{P} \left[\|\Phi^T \Phi - \mathbf{I}\|_{2 \rightarrow 2} \geq \sqrt{\frac{6m}{\tilde{c}p}} \right] \end{aligned} \quad (\text{A.21})$$

$$\leq 2e^{-m}, \quad (\text{A.22})$$

where (A.22) is by Theorem B.2 (with $t = \sqrt{2m}$). To establish (A.21), first note that $\|\Phi^T \Phi - \mathbf{I}\|_{2 \rightarrow 2} \leq \delta'$ (with $\delta' = \sqrt{\frac{6m}{\tilde{c}p}}$) implies $\sigma_{\max}(\Phi^T \Phi) \leq 1 + \delta'$, which in turn is equivalent to $\|\Phi^T \Phi\|_{2 \rightarrow 2} \leq 1 + \delta'$. We can therefore conclude that $\|\Phi^T \Phi\|_{2 \rightarrow 2} \geq 1 + \delta'$ implies $\|\Phi^T \Phi - \mathbf{I}\|_{2 \rightarrow 2} \geq \delta'$.

The terms inside the sums on the RHSs of (A.16)–(A.18) were upper-bounded by applying Lemma A.1 stated below. Specifically, we note that $\langle \bar{\mathbf{y}}_j^{(k)}, \mathbf{e} \rangle \sim \mathcal{N}(0, \sigma^2 \|\bar{\mathbf{y}}_j^{(k)}\|_2^2)$, $\langle \mathbf{e}_j^{(k)}, \bar{\mathbf{y}} \rangle \sim \mathcal{N}(0, \sigma^2 \|\bar{\mathbf{y}}\|_2^2)$, and $\langle \Phi^T \Phi \mathbf{e}_j^{(k)}, \mathbf{e} \rangle \sim \mathcal{N}(0, \sigma^2 \|\Phi^T \Phi \mathbf{e}_j^{(k)}\|_2^2)$, where $\mathbf{y}_j^{(k)}$, $\bar{\mathbf{y}}$, and $\Phi^T \Phi \mathbf{e}_j^{(k)}$, respectively, can be regarded as fixed, and we used $\beta = \sqrt{6 \log N} \geq \frac{1}{\sqrt{2\pi}}$, as $N \geq 1$.

LEMMA A.1 ([22, Proposition 19.4.2]) Let $x \sim \mathcal{N}(0, 1)$. For $\beta \geq \frac{1}{\sqrt{2\pi}}$, we have

$$\mathbb{P}[x \geq \beta] \leq e^{-\frac{\beta^2}{2}}. \quad (\text{A.23})$$

Finally, to upper-bound the terms inside the sum in (A.19), we used [16, Eq. (51)]

$$\mathbb{P}[\|\mathbf{e}_j^{(k)}\|_2 \geq 2\sigma] \leq e^{-\frac{\beta^2}{2}}. \quad (\text{A.24})$$

A.3 Proof of Theorem 3.1

The proof of Theorem 3.1 is obtained from the proof of Theorem 4.1 by noting that in the noise-free case (i.e., $\sigma = 0$), the perturbation caused by noise satisfies $\tilde{e}_j^{(k)} = 0$, rendering the second term on the RHS of (A.6) void. Finally, we remark that the assumption $6 \log N \leq m$ is not needed in the noise-free case as it is involved only in establishing (A.11), which is void here.

Appendix B. Proof of Theorem 3.2

We first note that the data points in \mathcal{X}_ℓ can be written as $\mathbf{x}_j^{(\ell)} = \mathbf{V}^{(\ell)} \mathbf{a}_j^{(\ell)}$, $j \in [n_\ell]$, where the $\mathbf{a}_j^{(\ell)}$ are i.i.d. uniform on $\mathbb{S}^{d_\ell-1}$, and $\mathbf{V}^{(\ell)} := \Phi \mathbf{U}^{(\ell)}$ is a basis for the d_ℓ -dimensional subspace of \mathbb{R}^p containing the points in \mathcal{X}_ℓ ($\mathbf{V}^{(\ell)}$ has full column rank with high probability, which follows from (B.7) as a consequence of the concentration inequality (3.1)). For the case where the $\mathbf{V}^{(\ell)}$ are orthonormal bases, a sufficient condition for successful clustering was derived by Soltanolkotabi and Candès [28, Theorem 2.8]. However, owing to the projection Φ , the $\mathbf{V}^{(\ell)} = \Phi \mathbf{U}^{(\ell)}$ will in general not be orthonormal. We will therefore need the following generalization of [28, Theorem 2.8] to arbitrary bases $\mathbf{V}^{(\ell)}$ for d_ℓ -dimensional subspaces of \mathbb{R}^p .

THEOREM B.1 Suppose that the elements of the sets \mathcal{X}_ℓ in $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_L$ are obtained by choosing n_ℓ points at random according to $\mathbf{x}_j^{(\ell)} = \mathbf{V}^{(\ell)} \mathbf{a}_j^{(\ell)}$, $j \in [n_\ell]$, where the $\mathbf{V}^{(\ell)} \in \mathbb{R}^{p \times d_\ell}$ have full rank, and the $\mathbf{a}_j^{(\ell)}$ are i.i.d. uniform on $\mathbb{S}^{d_\ell-1}$. Assume that $\rho_\ell := (n_\ell - 1)/d_\ell \geq \rho_0$, for all ℓ , where $\rho_0 > 1$ is a numerical constant, and let $\rho_{\min} := \min_\ell \rho_\ell$. If

$$\max_{k, \ell: k \neq \ell} \frac{1}{\sqrt{d_k}} \left\| \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \right\|_F \leq \frac{\sqrt{\log \rho_{\min}}}{64 \log N}, \quad (\text{B.1})$$

where $\mathbf{V}^{(\ell)\dagger} = (\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1} \mathbf{V}^{(\ell)T}$ is the (left) pseudoinverse of $\mathbf{V}^{(\ell)}$, then the graph G with adjacency matrix obtained by applying SSC to \mathcal{X} has no false connections with probability at least $1 - N^{-1} - \sum_{\ell=1}^L n_\ell e^{-\sqrt{\rho_\ell} d_\ell}$.

Proof. See Appendix B.1. □

We now detail how Theorem 3.2 follows from Theorem B.1. Specifically, we will show that (3.5) implies (B.1) with probability at least $1 - 4e^{-\tau/2}$, which, when combined with the probability bound in Theorem B.1 via a union bound yields the final probability estimate in Theorem 3.2, and thereby concludes the proof.

We start filling in the details by showing how (3.5) implies (B.1). The LHS of (B.1) can be upper-bounded as follows

$$\begin{aligned} \frac{1}{\sqrt{d_k}} \left\| \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \right\|_F &= \frac{1}{\sqrt{d_k}} \left\| (\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1} \mathbf{V}^{(\ell)T} \mathbf{V}^{(k)} \right\|_F \\ &\leq \left\| (\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1} \right\|_{2 \rightarrow 2} \frac{1}{\sqrt{d_k}} \left\| \mathbf{V}^{(\ell)T} \mathbf{V}^{(k)} \right\|_F \\ &\leq \frac{\left\| (\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1} \right\|_{2 \rightarrow 2}}{\sqrt{d_k}} \left(\left\| \mathbf{U}^{(\ell)T} \mathbf{U}^{(k)} \right\|_F + \left\| \mathbf{U}^{(\ell)T} (\Phi^T \Phi - \mathbf{I}) \mathbf{U}^{(k)} \right\|_F \right) \end{aligned} \quad (\text{B.2})$$

$$\leq \left\| (\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1} \right\|_{2 \rightarrow 2} \left(\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \left\| \mathbf{U}^{(\ell)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(k)} \right\|_{2 \rightarrow 2} \right) \quad (\text{B.3})$$

$$\leq \frac{1}{1 - \delta} (\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \delta) \quad (\text{B.4})$$

$$\leq \frac{65}{64} (\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \delta) \leq \frac{\sqrt{\log \rho_{\min}}}{64 \log N}, \quad (\text{B.5})$$

where (B.2) follows from $\|\mathbf{A}\mathbf{B}\|_F^2 \leq \|\mathbf{A}\|_{2 \rightarrow 2}^2 \|\mathbf{B}\|_F^2$, (B.3) is a consequence of $\|\mathbf{B}\|_F \leq \sqrt{m \wedge n} \|\mathbf{B}\|_{2 \rightarrow 2}$, for $\mathbf{B} \in \mathbb{R}^{m \times n}$ [18, Section 5.6, p. 365], and (B.4) holds with

$$\delta := \sqrt{\frac{28d_{\max} + 8 \log L + 2\tau}{3\tilde{c}p}}, \quad (\text{B.6})$$

with probability at least $1 - 4e^{-\tau/2}$ (here, $\tau > 0$ is the numerical constant in the statement of Theorem 3.2). Equation (B.4) holds with probability at least $1 - 4e^{-\tau/2}$ by

$$\mathbb{P} \left[\max_{\ell} \left\| (\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1} \right\|_{2 \rightarrow 2} \geq \frac{1}{1 - \delta} \right] \leq 2e^{-\tau/2} \quad (\text{B.7})$$

and

$$\mathbb{P} \left[\max_{k, \ell} \left\| \mathbf{U}^{(\ell)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(k)} \right\|_{2 \rightarrow 2} \geq \delta \right] \leq 2e^{-\tau/2}, \quad (\text{B.8})$$

both proven below. Finally, to get (B.5) we invoked (3.5) twice, first we used $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \geq 0$ and $\frac{\sqrt{\log \rho_{\min}}}{\log N} = \frac{\sqrt{\log \rho_{\min}}}{\log \left(\sum_{\ell=1}^L (\rho_\ell d_\ell + 1) \right)} \leq 1$ in (3.5) to conclude that $\delta \leq 1/65$, i.e., $\frac{1}{1 - \delta} \leq \frac{65}{64}$, and second, we applied (3.5) straight to upper-bound $\text{aff}(\mathcal{S}_k, \mathcal{S}_\ell)$.

It remains to prove (B.7) and (B.8). For the special case of a Gaussian random matrix $\boldsymbol{\Phi}$, the probability bounds (B.7) and (B.8) can be obtained using standard results on the extremal singular values of Gaussian random matrices. For general $\boldsymbol{\Phi}$ satisfying the concentration inequality (3.1), the proofs of (B.7) and (B.8) rely on Theorem B.2 below.

THEOREM B.2 ([10, Theorem 9.9, Remark 9.10]) Suppose that the random matrix $\boldsymbol{\Phi} \in \mathbb{R}^{p \times m}$ satisfies the concentration inequality (3.1), i.e.,

$$\mathbb{P} \left[\left| \|\boldsymbol{\Phi} \mathbf{x}\|_2^2 - \|\mathbf{x}\|_2^2 \right| \geq t \|\mathbf{x}\|_2^2 \right] \leq 2e^{-\tilde{c}t^2 p},$$

for all $t > 0$ and for all $\mathbf{x} \in \mathbb{R}^m$, where \tilde{c} is a constant. Then, for an orthonormal matrix $\mathbf{U} \in \mathbb{R}^{m \times d}$ and all $t > 0$, we have

$$\mathbb{P} \left[\left\| \mathbf{U}^T \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U} - \mathbf{I} \right\|_{2 \rightarrow 2} \geq \sqrt{\frac{14d + 2t^2}{3\tilde{c}p}} \right] \leq 2e^{-\frac{t^2}{2}}.$$

Additionally, for all $t > 0$, we have

$$\mathbb{P}\left[\|\Phi^T \Phi - \mathbf{I}\|_{2 \rightarrow 2} \geq \sqrt{\frac{14m + 2t^2}{3\tilde{c}p}}\right] \leq 2e^{-\frac{t^2}{2}}.$$

Proof of (B.7): By a union bound argument, we get

$$\mathbb{P}\left[\max_{\ell} \|(\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1}\|_{2 \rightarrow 2} \geq \frac{1}{1-\delta}\right] \leq \sum_{\ell=1}^L \mathbb{P}\left[\|(\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1}\|_{2 \rightarrow 2} \geq \frac{1}{1-\delta}\right]. \quad (\text{B.9})$$

Note that $\|\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)} - \mathbf{I}\|_{2 \rightarrow 2} \leq \delta$ implies $\sigma_{\min}(\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)}) \geq 1 - \delta$, which in turn implies $\|(\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1}\|_{2 \rightarrow 2} \leq \frac{1}{1-\delta}$. We can therefore conclude that $\|(\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1}\|_{2 \rightarrow 2} \geq \frac{1}{1-\delta}$ implies $\|\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)} - \mathbf{I}\|_{2 \rightarrow 2} \geq \delta$, which can be formalized according to

$$\left\{\|(\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1}\|_{2 \rightarrow 2} \geq \frac{1}{1-\delta}\right\} \subseteq \left\{\|\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)} - \mathbf{I}\|_{2 \rightarrow 2} \geq \delta\right\}.$$

Moreover, we have with δ as defined in (B.6) $\delta = \sqrt{\frac{28d_{\max} + 2t^2}{3\tilde{c}p}} \geq \sqrt{\frac{14d_{\ell} + 2t^2}{3\tilde{c}p}}$ ($2d_{\max} \geq d_{\max} \geq d_{\ell}$), with $t^2 = 4 \log L + \tau$. Therefore, Theorem B.2 (with $\mathbf{U} = \mathbf{U}^{(\ell)}$ and $t^2 = 4 \log L + \tau$) yields

$$\mathbb{P}\left[\|(\mathbf{V}^{(\ell)T} \mathbf{V}^{(\ell)})^{-1}\|_{2 \rightarrow 2} \geq \frac{1}{1-\delta}\right] \leq 2e^{-2 \log L - \tau/2} = 2L^{-2} e^{-\tau/2} \leq 2L^{-1} e^{-\tau/2},$$

which when used on the RHS of (B.9) establishes (B.7).

Proof of (B.8): Again, by a union bound argument, we get

$$\mathbb{P}\left[\max_{k,\ell} \|\mathbf{U}^{(\ell)T} (\Phi^T \Phi - \mathbf{I}) \mathbf{U}^{(k)}\|_{2 \rightarrow 2} \geq \delta\right] \leq \sum_{k,\ell=1}^L \mathbb{P}\left[\|\mathbf{U}^{(\ell)T} (\Phi^T \Phi - \mathbf{I}) \mathbf{U}^{(k)}\|_{2 \rightarrow 2} \geq \delta\right]. \quad (\text{B.10})$$

We next upper-bound the probabilities on the RHS of (B.10). To this end, let $\tilde{\mathbf{U}} \in \mathbb{R}^{m \times \tilde{d}}$ be an orthonormal basis for the \tilde{d} -dimensional span of $[\mathbf{U}^{(\ell)} \ \mathbf{U}^{(k)}]$ ($\max(d_{\ell}, d_k) \leq \tilde{d} \leq d_{\ell} + d_k$). Since $\tilde{\mathbf{U}} \tilde{\mathbf{U}}^T$ is the orthogonal projection onto $\text{span}([\mathbf{U}^{(\ell)} \ \mathbf{U}^{(k)}])$, we have $\tilde{\mathbf{U}} \tilde{\mathbf{U}}^T \mathbf{U}^{(\ell)} = \mathbf{U}^{(\ell)}$ and $\tilde{\mathbf{U}} \tilde{\mathbf{U}}^T \mathbf{U}^{(k)} = \mathbf{U}^{(k)}$. Therefore, we get

$$\begin{aligned} \|\mathbf{U}^{(\ell)T} (\Phi^T \Phi - \mathbf{I}) \mathbf{U}^{(k)}\|_{2 \rightarrow 2} &= \|\mathbf{U}^{(\ell)T} \tilde{\mathbf{U}} \tilde{\mathbf{U}}^T (\Phi^T \Phi - \mathbf{I}) \tilde{\mathbf{U}} \tilde{\mathbf{U}}^T \mathbf{U}^{(k)}\|_{2 \rightarrow 2} \\ &\leq \|\mathbf{U}^{(\ell)T} \tilde{\mathbf{U}}\|_{2 \rightarrow 2} \|\tilde{\mathbf{U}}^T \Phi^T \Phi \tilde{\mathbf{U}} - \mathbf{I}\|_{2 \rightarrow 2} \|\tilde{\mathbf{U}}^T \mathbf{U}^{(k)}\|_{2 \rightarrow 2} \\ &= \|\tilde{\mathbf{U}}^T \Phi^T \Phi \tilde{\mathbf{U}} - \mathbf{I}\|_{2 \rightarrow 2}, \end{aligned}$$

where we used $\|\mathbf{U}^{(\ell)T} \tilde{\mathbf{U}}\|_{2 \rightarrow 2} = 1$, which holds since $\mathbf{U}^{(\ell)}$ is in the span of $\tilde{\mathbf{U}}$ and both $\mathbf{U}^{(\ell)}$ and $\tilde{\mathbf{U}}$ are orthonormal. This finally yields, with δ as defined in (B.6),

$$\begin{aligned} \mathbb{P}\left[\|\mathbf{U}^{(\ell)T}(\Phi^T \Phi - \mathbf{I})\mathbf{U}^{(k)}\|_{2 \rightarrow 2} \geq \delta\right] &\leq \mathbb{P}\left[\|\tilde{\mathbf{U}}^T \Phi^T \Phi \tilde{\mathbf{U}} - \mathbf{I}\|_{2 \rightarrow 2} \geq \sqrt{\frac{28d_{\max} + 8 \log L + 2\tau}{3\tilde{c}p}}\right] \\ &\leq \mathbb{P}\left[\|\tilde{\mathbf{U}}^T \Phi^T \Phi \tilde{\mathbf{U}} - \mathbf{I}\|_{2 \rightarrow 2} \geq \sqrt{\frac{14\tilde{d} + 8 \log L + 2\tau}{3\tilde{c}p}}\right] \quad (\text{B.11}) \end{aligned}$$

$$\leq 2e^{-\frac{4 \log L + \tau}{2}} = 2L^{-2} e^{-\tau/2}, \quad (\text{B.12})$$

where (B.11) follows from $2d_{\max} \geq d_\ell + d_k \geq \tilde{d}$, and (B.12) is by application of Theorem B.2 with $\mathbf{U} = \tilde{\mathbf{U}}$ and $t^2 = 4 \log L + \tau$. The proof is concluded by using (B.12) on the RHS of (B.10).

B.1 Proof of Theorem B.1

Theorem B.1 is a generalization of a result by Soltanolkotabi & Candès [28, Theorem 2.8] from orthonormal bases $\mathbf{V}^{(\ell)}$ for d_ℓ -dimensional subspaces of \mathbb{R}^p to arbitrary bases $\mathbf{V}^{(\ell)}$ for d_ℓ -dimensional subspaces. The proof program essentially follows that of [28, Theorem 2.8]. However, some parts of the generalization are non-trivial. We only detail the arguments that are new relative to [28], and refer to [28] otherwise.

Throughout the proof, we use the following notation: Let $\mathbf{X}^{(\ell)} \in \mathbb{R}^{p \times n_\ell}$ be the matrix whose columns are the points in \mathcal{X}_ℓ , and note that $\mathbf{X}^{(\ell)} = \mathbf{V}^{(\ell)} \mathbf{A}^{(\ell)}$, where $\mathbf{A}^{(\ell)} \in \mathbb{R}^{d_\ell \times n_\ell}$ is the matrix with columns $\mathbf{a}_i^{(\ell)}, i = 1, \dots, n_\ell$. Set $\mathbf{X} = [\mathbf{X}^{(1)} \dots \mathbf{X}^{(L)}] \in \mathbb{R}^{p \times N}$, and let \mathbf{X}_{-i} be the matrix obtained by removing the i th column \mathbf{x}_i from \mathbf{X} . $\mathcal{P}(\mathbf{X})$ denotes the symmetrized convex hull of the columns of \mathbf{X} (recall that the columns of \mathbf{X} are the points in \mathcal{X}), that is, the convex hull of $\{\mathbf{x}_1, -\mathbf{x}_1, \dots, \mathbf{x}_N, -\mathbf{x}_N\}$. For a convex body \mathcal{P} , its inradius $r(\mathcal{P})$ is defined as the radius of the largest Euclidean ball that can be inscribed in \mathcal{P} , and its circumradius $R(\mathcal{P})$ is defined as the radius of the smallest ball containing \mathcal{P} . Finally, the polar set of $\mathcal{K} \subset \mathbb{R}^n$ is defined as

$$\mathcal{K}^\circ = \{\mathbf{y} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle \leq 1 \text{ for all } \mathbf{x} \in \mathcal{K}\}.$$

B.1.1 A deterministic clustering condition We first establish a deterministic clustering condition. Specifically, in Theorem B.3 below we present conditions guaranteeing that for $\mathbf{x}_i \in \mathcal{X}_\ell$ every solution of the problem

$$\underset{\mathbf{z}}{\text{minimize}} \|\mathbf{z}\|_1 \quad \text{subject to } \mathbf{X}_{-i} \mathbf{z} = \mathbf{x}_i \quad (\text{B.13})$$

has non-zero entries corresponding to columns of $\mathbf{X}^{(\ell)}$ only. The proof of Theorem B.1 is then obtained by establishing that these conditions are satisfied with high probability for the statistical data model in Theorem B.1. We start by introducing terminology needed in the following. Define the primal optimization problem

$$P(\mathbf{y}, \mathbf{A}) : \underset{\mathbf{z}}{\text{minimize}} \|\mathbf{z}\|_1 \quad \text{subject to } \mathbf{A} \mathbf{z} = \mathbf{y},$$

with the corresponding dual [4, Section 5.1.16]

$$D(\mathbf{y}, \mathbf{A}) : \underset{\mathbf{v}}{\text{maximize}} \langle \mathbf{y}, \mathbf{v} \rangle \text{ subject to } \|\mathbf{A}^T \mathbf{v}\|_\infty \leq 1.$$

The problem (B.13) is then simply $P(\mathbf{x}_i, \mathbf{X}_{-i})$. The sets of optimal solutions of P and D are denoted by $\text{optsol}P(\mathbf{y}, \mathbf{A})$ and $\text{optsol}D(\mathbf{y}, \mathbf{A})$, respectively. A dual point $\lambda(\mathbf{y}, \mathbf{A})$ is defined as a point in $\text{optsol}D(\mathbf{y}, \mathbf{A})$ of minimal Euclidean norm.

We are now ready to state the following generalization of [28, Theorem 2.5] from orthonormal bases $\mathbf{V}^{(\ell)}$ for d_ℓ -dimensional subspaces of \mathbb{R}^p to arbitrary bases $\mathbf{V}^{(\ell)}$ for d_ℓ -dimensional subspaces.

THEOREM B.3 Suppose that the elements of the sets \mathcal{X}_ℓ in $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_L$ are obtained by choosing n_ℓ points according to $\mathbf{x}_i^{(\ell)} = \mathbf{V}^{(\ell)} \mathbf{a}_i^{(\ell)}$, $i \in [n_\ell]$, where the $\mathbf{a}_i^{(\ell)}$ are deterministic coefficient vectors and the $\mathbf{V}^{(\ell)} \in \mathbb{R}^{p \times d_\ell}$ are deterministic matrices of full column rank. Let $\mathbf{L} \in \mathbb{R}^{d_\ell \times (n_\ell - 1)}$ be the matrix whose columns are the normalized dual points $\tilde{\lambda}(\mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)}) = \lambda(\mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)}) / \|\lambda(\mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)})\|_2$, where $\mathbf{A}_{-i}^{(\ell)}$ is the matrix with columns $\mathbf{a}_j^{(\ell)}$, $j \in [n_\ell] \setminus \{i\}$. If

$$\max_{k \neq \ell, j} \left\| \mathbf{L}^T \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \mathbf{a}_j^{(k)} \right\|_\infty < r(\mathcal{P}(\mathbf{A}_{-i}^{(\ell)})), \quad (\text{B.14})$$

then the non-zero entries of all solutions of $P(\mathbf{x}_i^{(\ell)}, \mathbf{X}_{-(i,\ell)})$ correspond to points in \mathcal{X}_ℓ only (the columns of $\mathbf{X}_{-(i,\ell)}$ are the elements in $\mathcal{X} \setminus \{\mathbf{x}_i^{(\ell)}\}$).

Proof. The proof relies on the following lemma.

LEMMA B.1 ([28, Lemma 7.1], [6]) Let T be a subset of the column indices of a given matrix \mathbf{A} . All solutions \mathbf{c}^* of $P(\mathbf{y}, \mathbf{A})$ satisfy $\mathbf{c}_T^* = \mathbf{0}$, if there exists a vector \mathbf{c} such that $\mathbf{y} = \mathbf{A}\mathbf{c}$ with support $S \subseteq T$, and a (dual certificate) vector \mathbf{v} obeying

$$\mathbf{A}_S^T \mathbf{v} = \text{sign}(\mathbf{c}_S) \quad (\text{B.15})$$

$$\|\mathbf{A}_{T \cap \bar{S}}^T \mathbf{v}\|_\infty \leq 1 \quad (\text{B.16})$$

$$\|\mathbf{A}_T^T \mathbf{v}\|_\infty < 1. \quad (\text{B.17})$$

We apply Lemma B.1 with $\mathbf{A} = \mathbf{X}_{-(i,\ell)}$, $\mathbf{y} = \mathbf{x}_i^{(\ell)}$, and T the index set corresponding to the columns of $\mathbf{X}_{-i}^{(\ell)}$, and show that there exists a vector \mathbf{c} supported on $S \subseteq T$ that obeys $\mathbf{x}_i^{(\ell)} = \mathbf{X}_{-(i,\ell)} \mathbf{c}$, and a corresponding vector \mathbf{v} that satisfies (B.15)–(B.17). This then implies that the non-zero entries of all solutions of $P(\mathbf{x}_i^{(\ell)}, \mathbf{X}_{-(i,\ell)})$ correspond to points in \mathcal{X}_ℓ only, as desired.

We proceed with the explicit construction of the vector \mathbf{c} . Specifically, take \mathbf{c} to be a vector that is zero on \bar{T} , and whose restriction to the index set T is given by $\mathbf{c}_T \in \text{optsol}P(\mathbf{x}_i^{(\ell)}, \mathbf{X}_{-i}^{(\ell)})$. Let S be the support of \mathbf{c}_T , and let $\mathbf{v}_i^{(\ell)} = (\mathbf{V}^{(\ell)T})^\dagger \lambda_i^{(\ell)}$, where $\lambda_i^{(\ell)}$ is taken to be a point of minimum ℓ_2 -norm⁴ in $\text{optsol}D(\mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)})$. The next step is to show that $\mathbf{v}_i^{(\ell)} \in \text{optsol}D(\mathbf{x}_i^{(\ell)}, \mathbf{X}_{-i}^{(\ell)})$, which will eventually allow

⁴ For concreteness $\lambda_i^{(\ell)}$ is taken to be a point of minimum ℓ_2 -norm. Note, however, that for the proof to work we may let $\lambda_i^{(\ell)}$ be an arbitrary point in $\text{optsol}P(\mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)})$.

us to establish that $\mathbf{v}_i^{(\ell)}$ satisfies the conditions of Lemma B.1. To this end, we first note that $\mathbf{x}_i^{(\ell)} = \mathbf{V}^{(\ell)} \mathbf{a}_i^{(\ell)}$ yields

$$\begin{aligned} & \text{optsol}D(\mathbf{x}_i^{(\ell)}, \mathbf{X}_{-i}^{(\ell)}) \\ &= \left\{ \arg \max_{\mathbf{v}} \left\langle \mathbf{a}_i^{(\ell)}, \mathbf{V}^{(\ell)T} \mathbf{v} \right\rangle \text{ subject to } \left\| (\mathbf{A}_{-i}^{(\ell)})^T \mathbf{V}^{(\ell)T} \mathbf{v} \right\|_{\infty} \leq 1 \right\} \\ &= \left\{ \mathbf{v} : \boldsymbol{\lambda} = \mathbf{V}^{(\ell)T} \mathbf{v}, \boldsymbol{\lambda} \in \text{optsol}D(\mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)}) \right\} \\ &\supseteq (\mathbf{V}^{(\ell)T})^{\dagger} \text{optsol}D(\mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)}), \end{aligned}$$

where the inclusion holds as $(\mathbf{V}^{(\ell)T})^{\dagger} \boldsymbol{\lambda}$ is the minimum norm solution to the linear system of equations $\boldsymbol{\lambda} = \mathbf{V}^{(\ell)T} \mathbf{v}$, but in general not the only solution.

Since $P(\mathbf{y}, \mathbf{A})$ is a linear program, strong duality [4, Section 5.2.3] holds (provided that $P(\mathbf{y}, \mathbf{A})$ is feasible) and therefore the optimal objective values of $P(\mathbf{x}_i^{(\ell)}, \mathbf{X}_{-i}^{(\ell)})$ and $D(\mathbf{x}_i^{(\ell)}, \mathbf{X}_{-i}^{(\ell)})$ coincide. It therefore follows that

$$\|\mathbf{c}_T\|_1 = \left\langle \mathbf{x}_i^{(\ell)}, \mathbf{v}_i^{(\ell)} \right\rangle. \quad (\text{B.18})$$

Since $\mathbf{c}_T \in \text{optsol}P(\mathbf{x}_i^{(\ell)}, \mathbf{X}_{-i}^{(\ell)})$ and \mathbf{c}_T is supported on S , both by assumption, we have $\mathbf{x}_i^{(\ell)} = \mathbf{X}_{-i}^{(\ell)} \mathbf{c}_T = (\mathbf{X}_{-i}^{(\ell)})_S \mathbf{c}_S$, and therefore (B.18) becomes

$$\langle \mathbf{c}_S, \text{sign}(\mathbf{c}_S) \rangle = \left\langle (\mathbf{X}_{-i}^{(\ell)})_S \mathbf{c}_S, \mathbf{v}_i^{(\ell)} \right\rangle = \left\langle \mathbf{c}_S, (\mathbf{X}_{-i}^{(\ell)})_S^T \mathbf{v}_i^{(\ell)} \right\rangle. \quad (\text{B.19})$$

On the other hand, as $\mathbf{v}_i^{(\ell)} \in \text{optsol}D(\mathbf{x}_i^{(\ell)}, \mathbf{X}_{-i}^{(\ell)})$, we have $\left\| (\mathbf{X}_{-i}^{(\ell)})^T \mathbf{v}_i^{(\ell)} \right\|_{\infty} \leq 1$, which is equivalent to the following conditions (recall that the set T corresponds to the column indices of $\mathbf{X}_{-i}^{(\ell)}$):

$$\left\| ((\mathbf{X}_{-i}^{(\ell)})_S)^T \mathbf{v}_i^{(\ell)} \right\|_{\infty} \leq 1 \quad (\text{B.20})$$

$$\left\| ((\mathbf{X}_{-i}^{(\ell)})_{T \cap \bar{S}})^T \mathbf{v}_i^{(\ell)} \right\|_{\infty} \leq 1. \quad (\text{B.21})$$

As by (B.20) the entries of $(\mathbf{X}_{-i}^{(\ell)})_S^T \mathbf{v}_i^{(\ell)}$ are bounded in magnitude by 1 and the unique maximizer of $\max_{\mathbf{a}: \|\mathbf{a}\|_{\infty} \leq 1} \langle \mathbf{c}_S, \mathbf{a} \rangle$ is $\text{sign}(\mathbf{c}_S)$, it follows from (B.19) that

$$(\mathbf{X}_{-i}^{(\ell)})_S^T \mathbf{v}_i^{(\ell)} = \text{sign}(\mathbf{c}_S),$$

which establishes (B.15).

Thanks to (B.21), (B.16) is satisfied as well. It remains to verify (B.17), which here reads

$$\left| \left\langle \mathbf{x}_j^{(k)}, \mathbf{v}_i^{(\ell)} \right\rangle \right| < 1, \text{ for all } k \neq \ell, \text{ for all } j \in [n_k]. \quad (\text{B.22})$$

With $\mathbf{v}_i^{(\ell)} = (\mathbf{V}^{(\ell)T})^\dagger \boldsymbol{\lambda}_i^{(\ell)}$, by definition, (B.22) becomes

$$\left| \left\langle \mathbf{x}_j^{(k)}, (\mathbf{V}^{(\ell)T})^\dagger \frac{\boldsymbol{\lambda}_i^{(\ell)}}{\|\boldsymbol{\lambda}_i^{(\ell)}\|_2} \right\rangle \right| < \frac{1}{\|\boldsymbol{\lambda}_i^{(\ell)}\|_2}, \text{ for all } k \neq \ell, \text{ for all } j \in [n_k]. \quad (\text{B.23})$$

Since $((\mathbf{V}^{(\ell)T})^\dagger)^T = \mathbf{V}^{(\ell)\dagger}$, and $\mathbf{x}_j^{(k)} = \mathbf{V}^{(k)} \mathbf{a}_j^{(k)}$, (B.23) is equivalent to

$$\left| \frac{\boldsymbol{\lambda}_i^{(\ell)T}}{\|\boldsymbol{\lambda}_i^{(\ell)}\|_2} \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \mathbf{a}_j^{(k)} \right| < \frac{1}{\|\boldsymbol{\lambda}_i^{(\ell)}\|_2}, \text{ for all } k \neq \ell, \text{ for all } j \in [n_k]. \quad (\text{B.24})$$

It now follows from $\boldsymbol{\lambda}_i^{(\ell)} \in \text{optsol}D(\mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)})$, which holds by assumption, that

$$\|(\mathbf{A}_{-i}^{(\ell)})^T \boldsymbol{\lambda}_i^{(\ell)}\|_\infty \leq 1.$$

This, in turn, implies that $\boldsymbol{\lambda}_i^{(\ell)} \in \mathcal{P}^\circ(\mathbf{A}_{-i}^{(\ell)})$, where

$$\mathcal{P}^\circ(\mathbf{A}_{-i}^{(\ell)}) = \left\{ \mathbf{z}: \|(\mathbf{A}_{-i}^{(\ell)})^T \mathbf{z}\|_\infty \leq 1 \right\}$$

is the polar set of $\mathcal{P}(\mathbf{A}_{-i}^{(\ell)})$ (recall that $\mathcal{P}(\mathbf{A}_{-i}^{(\ell)})$ is the symmetrized convex hull of the columns in $\mathbf{A}_{-i}^{(\ell)}$). Since the inradius and the circumradius of a symmetric⁵ convex body are related according to [13, Theorem 1.2]

$$r(\mathcal{P})R(\mathcal{P}^\circ) = 1,$$

we get from $\boldsymbol{\lambda}_i^{(\ell)} \in \mathcal{P}^\circ(\mathbf{A}_{-i}^{(\ell)})$ that

$$\|\boldsymbol{\lambda}_i^{(\ell)}\|_2 \leq R(\mathcal{P}^\circ(\mathbf{A}_{-i}^{(\ell)})) = \frac{1}{r(\mathcal{P}(\mathbf{A}_{-i}^{(\ell)}))}. \quad (\text{B.25})$$

By (B.25), it follows that (B.24) holds if

$$\left| \frac{\boldsymbol{\lambda}_i^{(\ell)T}}{\|\boldsymbol{\lambda}_i^{(\ell)}\|_2} \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \mathbf{a}_j^{(k)} \right| < r(\mathcal{P}(\mathbf{A}_{-i}^{(\ell)})), \text{ for all } k \neq \ell, \text{ for all } j \in [n_k],$$

which is implied by (B.14). This proves that (B.17) is satisfied as well, thereby concluding the proof. \square

⁵ A convex body \mathcal{P} is called symmetric if $\mathbf{x} \in \mathcal{P}$ if and only if $-\mathbf{x} \in \mathcal{P}$.

B.1.2 Evaluating the deterministic clustering condition for the statistical data model Theorem B.1 now follows from Theorem B.3 by establishing that, for our statistical data model, the deterministic clustering condition (B.14) holds for all pairs (ℓ, i) with $\ell \in [L], i \in [n_\ell]$, with high probability. Specifically, by a union bound argument, we get

$$\begin{aligned} & \mathbb{P}[(\text{B.14}) \text{ is violated for at least one pair } (\ell, i)] \\ & \leq \sum_{(\ell, i)} \mathbb{P}\left[\max_{k \neq \ell, j} \left\| \mathbf{L}^T \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \mathbf{a}_j^{(k)} \right\|_\infty \geq r(\mathcal{P}(\mathbf{A}_{-i}^{(\ell)}))\right] \\ & \leq \sum_{(\ell, i)} \left(\mathbb{P}\left[\max_{k \neq \ell, j} \left\| \mathbf{L}^T \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \mathbf{a}_j^{(k)} \right\|_\infty \geq \frac{16 \log N}{\sqrt{d_\ell d_k}} \left\| \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \right\|_F\right] + \mathbb{P}\left[\frac{\sqrt{\log \rho_\ell}}{4\sqrt{d_\ell}} \geq r(\mathcal{P}(\mathbf{A}_{-i}^{(\ell)}))\right] \right) \end{aligned} \quad (\text{B.26})$$

$$\leq \sum_{\ell=1}^L n_\ell e^{-\sqrt{\rho_\ell} d_\ell} + N^{-1}. \quad (\text{B.27})$$

In (B.26), we used that for random variables X and Y , possibly dependent, and constants ϕ and φ satisfying $\phi \leq \varphi$, we have

$$\begin{aligned} \mathbb{P}[X \geq Y] & \leq \mathbb{P}[\{X \geq \phi\} \cup \{\varphi \geq Y\}] \\ & \leq \mathbb{P}[X \geq \phi] + \mathbb{P}[\varphi \geq Y]. \end{aligned} \quad (\text{B.28})$$

Specifically, we applied (B.28) with $\phi = \frac{16 \log N}{\sqrt{d_\ell d_k}} \left\| \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \right\|_F$ and $\varphi = \frac{\sqrt{\log \rho_\ell}}{4\sqrt{d_\ell}}$, which leads to the assumption

$$\frac{16 \log N}{\sqrt{d_\ell d_k}} \left\| \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \right\|_F \leq \frac{\sqrt{\log \rho_\ell}}{4\sqrt{d_\ell}}, \quad \text{for all } k, \ell: k \neq \ell,$$

implied by (B.1). To get (B.27) we used that, for all i ,

$$\mathbb{P}\left[\frac{\sqrt{\log \rho_\ell}}{4\sqrt{d_\ell}} \geq r(\mathcal{P}(\mathbf{A}_{-i}^{(\ell)}))\right] \leq e^{-\sqrt{\rho_\ell} d_\ell} \quad (\text{B.29})$$

and

$$\mathbb{P}\left[\max_{k \neq \ell, j} \left\| \mathbf{L}^T \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \mathbf{a}_j^{(k)} \right\|_\infty \geq \frac{16 \log N}{\sqrt{d_\ell d_k}} \left\| \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \right\|_F\right] \leq N^{-2}, \quad (\text{B.30})$$

both of which are established next.

The upper bound (B.29) is an application of [28, Lemma 7.4], [2], and makes use of the assumption $(n_\ell - 1)/d_\ell = \rho_\ell \geq \rho_0 > 1$. Finally, (B.30) follows from a union bound argument and

$$\begin{aligned} & \mathbb{P}\left[\left\| \mathbf{L}^T \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \mathbf{a}_j^{(k)} \right\|_\infty \geq \frac{16 \log N}{\sqrt{d_\ell d_k}} \left\| \mathbf{V}^{(\ell)\dagger} \mathbf{V}^{(k)} \right\|_F\right] \\ & \leq (n_\ell + 1) e^{-4 \log N} \leq N^{-3}, \end{aligned} \quad (\text{B.31})$$

which is a consequence of Lemma B.2 below together with the fact that the normalized dual point $\tilde{\lambda}_i^{(\ell)} = \lambda_i^{(\ell)} / \|\lambda_i^{(\ell)}\|_2$ is distributed uniformly on the unit sphere, as shown in [28, Section 7.2.2, Proof of Step 2].

LEMMA B.2 (Extracted from the proof of Lemma 7.5 in [28]) Let the columns of $\mathbf{L} \in \mathbb{R}^{d_1 \times n_1}$ be i.i.d. uniform on \mathbb{S}^{d_1-1} , let \mathbf{a} be uniform on \mathbb{S}^{d_2-1} , and let $\mathbf{B} \in \mathbb{R}^{d_1 \times d_2}$. Then, for $c \geq 12$, we have

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

Appendix C. Proof of Theorem 3.3

The graph G obtained by SSC-OMP has no false connections if for each $\mathbf{x}_i^{(\ell)} \in \mathcal{X}_\ell$ the OMP algorithm as detailed in Section 2 selects points from \mathcal{X}_ℓ only, for all $\ell \in [L]$. This is the case if OMP selects points from \mathcal{X}_ℓ in all iterations $s \in [s_{\text{end}}]$ (we explain below that OMP terminates after $s_{\text{end}} = s_{\text{max}} \wedge d_\ell$ iterations with high probability for our statistical data model). The OMP selection rule (2.2) implies that OMP selects a point from \mathcal{X}_ℓ in the $(s + 1)$ th iteration if

$$\max_{k \neq \ell, j} |\langle \mathbf{x}_j^{(k)}, \mathbf{r}_s \rangle| < \max_{j \in [n_\ell]: j \neq i} |\langle \mathbf{x}_j^{(\ell)}, \mathbf{r}_s \rangle|. \quad (\text{C.1})$$

Hence, the graph G obtained by SSC-OMP has no false connections if the deterministic clustering condition (C.1) holds for all s_{end} OMP iterations, for all $\mathbf{x}_i^{(\ell)} \in \mathcal{X}_\ell, \ell \in [L]$. We will next establish that (C.1) is satisfied for our statistical data model with probability obeying the bound in Theorem 3.3.

As a vehicle for our analysis, we introduce the *reduced OMP algorithm* which, to compute sparse representations of the $\mathbf{x}_i^{(\ell)}$, has access to the corresponding *reduced data sets* $\mathcal{X}_\ell \setminus \{\mathbf{x}_i^{(\ell)}\}$ only, instead of the full data sets $\mathcal{X} \setminus \{\mathbf{x}_i^{(\ell)}\}$. If, for a given data set \mathcal{X} , the residuals computed by reduced OMP, henceforth denoted by $\mathbf{r}_s^{(\ell)}$, satisfy (C.1) for *all iterations*, then the reduced OMP algorithm and the original OMP algorithm (processing the same data set \mathcal{X}) select exactly the same data points in the same order and we have $\mathbf{r}_s = \mathbf{r}_s^{(\ell)}$ for all $s \in [s_{\text{max}} \wedge d_\ell]$ by virtue of (2.3). We emphasize that for expositional convenience, the notations $\mathbf{r}_s^{(\ell)}$ and $\tilde{\mathbf{r}}_s^{(\ell)}$ do not reflect dependence on i . The motivation for working with the reduced OMP algorithm is that $\mathbf{r}_s^{(\ell)}$ being a function of the data points in \mathcal{X}_ℓ only, conditionally on Φ , is statistically independent of the data points in $\mathcal{X} \setminus \mathcal{X}_\ell$. This will allow us to establish tail bounds for $|\langle \mathbf{x}_j^{(k)}, \mathbf{r}_s^{(\ell)} \rangle|$, $k \neq \ell, j \in [n_k]$, using standard concentration inequalities. We proceed to show that under the assumptions of Theorem 3.3, the reduced OMP residuals $\mathbf{r}_s^{(\ell)}$ indeed satisfy (C.1) for all $\ell \in [L]$, $i \in [n_\ell]$, and $s \in [s_{\text{max}} \wedge d_\ell]$ with probability meeting the lower bound in Theorem 3.3.

Consider the reduced OMP algorithm for the data point $\mathbf{x}_i^{(\ell)}$ with fixed $\ell \in [L]$ and fixed $i \in [n_\ell]$. We start by noting that the reduced OMP index set Λ_s is a function of the data points in \mathcal{X}_ℓ only. After iteration s , with $\mathbf{x}_i^{(\ell)} = \Phi \mathbf{U}^{(\ell)} \mathbf{a}_i^{(\ell)}$ and $\mathbf{X}_{\Lambda_s}^{(\ell)} = \Phi \mathbf{U}^{(\ell)} \mathbf{A}_{\Lambda_s}^{(\ell)}$ inserted into (2.3), we get $\mathbf{r}_s^{(\ell)} = \Phi \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)}$, where

$$\tilde{\mathbf{r}}_s^{(\ell)} := (\mathbf{I} - \mathbf{A}_{\Lambda_s}^{(\ell)} (\Phi \mathbf{U}^{(\ell)} \mathbf{A}_{\Lambda_s}^{(\ell)})^\dagger \Phi \mathbf{U}^{(\ell)}) \mathbf{a}_i^{(\ell)}.$$

We next establish a lower bound on the RHS of (C.1) and an upper bound on the LHS of (C.1). To isolate the impact of the different random quantities in the statistical data model, we will introduce events, upon the intersection of which (C.1) is implied by (3.6) via these bounds. A union bound on the probability of the intersection of these events then yields the final result.

We start by lower-bounding the RHS of (C.1) according to

$$\max_{j \in [n_\ell]: j \neq i} \left| \left\langle \mathbf{x}_j^{(\ell)}, \mathbf{r}_s^{(\ell)} \right\rangle \right| \geq \frac{1}{4} \sqrt{\frac{\log \rho_\ell}{d_\ell}} \sigma_{\min} \left(\mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \right) \left\| \tilde{\mathbf{r}}_s^{(\ell)} \right\|_2 \quad (\text{C.2})$$

$$\geq \frac{1}{4} \sqrt{\frac{\log \rho_\ell}{d_\ell}} (1 - \delta) \left\| \tilde{\mathbf{r}}_s^{(\ell)} \right\|_2, \quad (\text{C.3})$$

where (C.2) and (C.3) hold on the events

$$\mathcal{E}_1^{(\ell,i)} := \left\{ \max_{j \in [n_\ell]: j \neq i} \left| \mathbf{x}_j^{(\ell)T} \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{v} \right| > \frac{1}{4} \sqrt{\frac{\log \rho_\ell}{d_\ell}} \sigma_{\min} \left(\mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \right) \left\| \mathbf{v} \right\|_2, \forall \mathbf{v} \in \mathbb{R}^{d_\ell} \right\}$$

and

$$\mathcal{E}_2 := \left\{ \min_{\ell} \sigma_{\min} \left(\mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \right) > 1 - \delta \right\}, \quad \delta \in (0, 1),$$

respectively. Note that $\tilde{\mathbf{r}}_s^{(\ell)}$ in (C.2) not being statistically independent of the $\mathbf{x}_j^{(\ell)}, j \neq i$, is not an issue as we consider (C.2) on the event $\mathcal{E}_1^{(\ell,i)}$ and the inequality in the definition of $\mathcal{E}_1^{(\ell,i)}$ applies to *all* $\mathbf{v} \in \mathbb{R}^{d_\ell}$. Since $\mathbf{V}^{(\ell)} = \boldsymbol{\Phi} \mathbf{U}^{(\ell)}$ has full rank on \mathcal{E}_2 , reduced OMP terminates after $s_{\max} \wedge d_\ell$ iterations. To see this, simply note that for $\mathbf{V}^{(\ell)}$ of full rank we need exactly d_ℓ points from $\mathcal{X}_\ell \setminus \{\mathbf{x}_i^{(\ell)}\}$ to represent $\mathbf{x}_i^{(\ell)} = \mathbf{V}^{(\ell)} \mathbf{a}_i^{(\ell)}$ (owing to the fact that the $\mathbf{a}_j^{(\ell)}, j \in [n_\ell]$, are i.i.d. uniform on $\mathbb{S}^{d_\ell-1}$), and thus $\mathbf{r}_s^{(\ell)} = \mathbf{0}$ after exactly d_ℓ iterations.

We continue by upper-bounding the LHS of (C.1) according to

$$\begin{aligned} \max_{k \neq \ell, j} \left| \left\langle \mathbf{x}_j^{(k)}, \mathbf{r}_s^{(\ell)} \right\rangle \right| &= \max_{k \neq \ell, j} \left| \mathbf{a}_j^{(k)T} \mathbf{U}^{(k)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)} \right| \\ &= \max_{k \neq \ell, j} \left| \mathbf{a}_j^{(k)T} \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)} + \mathbf{a}_j^{(k)T} \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)} \right| \\ &\leq \max_{k \neq \ell, j} \left| \mathbf{a}_j^{(k)T} \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)} \right| + \max_{k \neq \ell, j} \left| \mathbf{a}_j^{(k)T} \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)} \right| \\ &< 4(3 \log N + \log s_{\max}) \max_{k \neq \ell} \frac{\left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F}{\sqrt{d_k} \sqrt{d_\ell}} \left\| \tilde{\mathbf{r}}_s^{(\ell)} \right\|_2 \\ &\quad + \sqrt{\frac{6 \log N + 2 \log s_{\max}}{d_{\min}}} \left\| \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \right\|_{2 \rightarrow 2} \left\| \tilde{\mathbf{r}}_s^{(\ell)} \right\|_2 \quad (\text{C.4}) \end{aligned}$$

$$\begin{aligned} &\leq 4(3 \log N + \log s_{\max}) \max_{k \neq \ell} \frac{\left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F}{\sqrt{d_k} \sqrt{d_\ell}} \left\| \tilde{\mathbf{r}}_s^{(\ell)} \right\|_2 \\ &\quad + \sqrt{\frac{6 \log N + 2 \log s_{\max}}{d_{\min}}} \delta \left\| \tilde{\mathbf{r}}_s^{(\ell)} \right\|_2 \quad (\text{C.5}) \end{aligned}$$

$$\leq \frac{1}{4} \sqrt{\frac{\log \rho_\ell}{d_\ell}} (1 - \delta) \|\tilde{\mathbf{r}}_s^{(\ell)}\|_2. \quad (\text{C.6})$$

Here, (C.4) holds on the intersection of the events

$$\begin{aligned} \mathcal{E}_3^{(\ell, i, s)} &:= \left\{ \max_{k \neq \ell, j} \left| \mathbf{a}_j^{(k)T} \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)} \right| \right. \\ &\quad \left. \leq \sqrt{\frac{6 \log N + 2 \log s_{\max}}{d_{\min}}} \left\| \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \right\|_{2 \rightarrow 2} \|\tilde{\mathbf{r}}_s^{(\ell)}\|_2 \right\}, \\ \mathcal{E}_4^{(\ell, i, s)} &:= \left\{ \max_{k \neq \ell, j} \left| \mathbf{a}_j^{(k)T} \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)} \right| < 4(3 \log N + \log s_{\max}) \max_{k \neq \ell} \frac{\left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F}{\sqrt{d_k} \sqrt{d_\ell}} \|\tilde{\mathbf{r}}_s^{(\ell)}\|_2 \right\} \end{aligned}$$

and (C.5) holds on the event

$$\mathcal{E}_5 := \left\{ \max_{k, \ell: k \neq \ell} \left\| \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \right\|_{2 \rightarrow 2} < \delta \right\}.$$

Recall that the notation $\tilde{\mathbf{r}}_s^{(\ell)}$ does not reflect dependence on the index i . We do, however, make the dependence of $\mathcal{E}_3^{(\ell, i, s)}$ and $\mathcal{E}_4^{(\ell, i, s)}$ on i explicit.

Finally, setting $\delta := \sqrt{\frac{28d_{\max} + 8 \log L + 2\tau}{3\tilde{c}p}}$ in (C.5), (C.6) follows from assumption (3.6). This is seen as follows:

$$\max_{k \neq \ell} \frac{\left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F}{\sqrt{d_k}} + \frac{\delta \sqrt{\frac{d_\ell}{d_{\min}}}}{2\sqrt{6 \log N + 2 \log s_{\max}}} \leq \max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) + \frac{\delta}{2} \sqrt{\frac{d_{\max}}{d_{\min}}} \quad (\text{C.7})$$

$$\leq \frac{3}{200} \frac{\sqrt{\log \rho_{\min}}}{\log N} \quad (\text{C.8})$$

$$\leq \frac{\sqrt{\log \rho_{\min}}}{50(\log N + (\log s_{\max})/3)} \quad (\text{C.9})$$

$$\leq \frac{\sqrt{\log \rho_\ell}}{48(\log N + (\log s_{\max})/3)} (1 - \delta), \quad (\text{C.10})$$

where (C.8) is by (3.6) and (C.9) follows by noting that $(200/3) \log N = 50(\log N + (\log N)/3) > 50(\log N + (\log s_{\max})/3)$. Furthermore, we have

$$\frac{\sqrt{\log \rho_{\min}}}{\log N + (\log s_{\max})/3} \leq 1 \quad (\text{C.11})$$

as a consequence of $\rho_{\min} = \min_\ell (n_\ell - 1)/d_\ell < N/d_{\min}$, $N \geq 3$, and $d_{\min} \geq 1$. Next, (C.11) combined with $\sqrt{d_{\min}/d_{\max}} \leq 1$, $\max_{k, \ell: k \neq \ell} \text{aff}(\mathcal{S}_k, \mathcal{S}_\ell) \geq 0$, and (C.9), implies that $\delta \leq \frac{2}{50}$, which yields $\frac{1}{50} \leq \frac{1}{48}(1 - \delta)$ and hence establishes (C.10). Finally, (C.6) is obtained by rewriting the relation between the RHS of (C.7) and (C.10).

Note that the lower bound (C.3) on the RHS of (C.1) and the upper bound (C.6) on the LHS of (C.1) are equal; we have therefore established that, for fixed (i, ℓ) , $\mathbf{r}_s^{(\ell)}$ obeys (C.1) on $\mathcal{E}_1^{(\ell,i)} \cap \mathcal{E}_2 \cap \mathcal{E}_3^{(\ell,i,s)} \cap \mathcal{E}_4^{(\ell,i,s)} \cap \mathcal{E}_5$. It finally follows that on the event $\mathcal{E}_\star := \bigcap_{\ell \in [L]} \mathcal{E}_1^{(\ell,i)} \cap \mathcal{E}_2 \cap \mathcal{E}_3^{(\ell,i,s)} \cap \mathcal{E}_4^{(\ell,i,s)} \cap \mathcal{E}_5$, the graph G obtained by SSC-OMP applied to the full data set $\mathcal{X} \setminus \{\mathbf{x}_i^{(\ell)}\}$ has no false connections. It remains to lower-bound $\mathbb{P}[\mathcal{E}_\star]$. Specifically, we have

$$\begin{aligned} \mathbb{P}[\mathcal{E}_\star] &= 1 - \mathbb{P}[\overline{\mathcal{E}_\star}] \\ &\geq 1 - \mathbb{P}[\overline{\mathcal{E}_2}] - \mathbb{P}[\overline{\mathcal{E}_5}] - \sum_{\ell \in [L], i \in [n_\ell]} \left(\mathbb{P}[\overline{\mathcal{E}_1^{(\ell,i)}}] + \sum_{s \in [s_{\max} \wedge d_\ell]} \left(\mathbb{P}[\overline{\mathcal{E}_3^{(\ell,i,s)}}] + \mathbb{P}[\overline{\mathcal{E}_4^{(\ell,i,s)}}] \right) \right) \\ &\geq 1 - 4e^{-\tau/2} - \sum_{\ell \in [L]} n_\ell e^{-\sqrt{\rho_\ell} d_\ell} - \frac{4}{N}, \end{aligned} \quad (\text{C.12})$$

where the last inequality follows from

$$\mathbb{P}[\overline{\mathcal{E}_1^{(\ell,i)}}] \leq e^{-\sqrt{\rho_\ell} d_\ell} \quad (\text{C.13})$$

$$\mathbb{P}[\overline{\mathcal{E}_2}] \leq 2e^{-\tau/2} \quad (\text{C.14})$$

$$\mathbb{P}[\overline{\mathcal{E}_3^{(\ell,i,s)}}] \leq \frac{2}{s_{\max} N^2} \quad (\text{C.15})$$

$$\mathbb{P}[\overline{\mathcal{E}_4^{(\ell,i,s)}}] \leq \frac{2}{s_{\max} N^2} \quad (\text{C.16})$$

$$\mathbb{P}[\overline{\mathcal{E}_5}] \leq 2e^{-\tau/2}. \quad (\text{C.17})$$

Here, (C.17) corresponds to (B.8), while the proofs of (C.13)–(C.16) are presented below.

Proof of (C.13): Since $\mathbf{A}_{-i}^{(\ell)}$ has full column rank with probability 1, it follows from Lemma C.1 below that

$$\begin{aligned} \left\| \mathbf{A}_{-i}^{(\ell)T} \mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{v} \right\|_\infty &\geq r(\mathcal{P}(\mathbf{A}_{-i}^{(\ell)})) \left\| \mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{v} \right\|_2 \\ &\geq r(\mathcal{P}(\mathbf{A}_{-i}^{(\ell)})) \sigma_{\min}(\mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)}) \|\mathbf{v}\|_2, \end{aligned}$$

for all $\mathbf{v} \in \mathbb{R}^{d_\ell}$. We therefore have

$$\begin{aligned} \mathbb{P}[\overline{\mathcal{E}_1^{(\ell,i)}}] &= \mathbb{P} \left[\left\| \mathbf{A}_{-i}^{(\ell)T} \mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{v} \right\|_\infty \leq \frac{1}{4} \sqrt{\frac{\log \rho_\ell}{d_\ell}} \sigma_{\min}(\mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)}) \|\mathbf{v}\|_2 \right] \\ &\leq \mathbb{P} \left[r(\mathcal{P}(\mathbf{A}_{-i}^{(\ell)})) \leq \frac{1}{4} \sqrt{\frac{\log \rho_\ell}{d_\ell}} \right] \\ &\leq e^{-\sqrt{\rho_\ell} d_\ell}, \end{aligned} \quad (\text{C.18})$$

where (C.18) follows from (B.29), which uses the assumption $(n_\ell - 1)/d_\ell = \rho_\ell \geq \rho_0 > 1$.

LEMMA C.1 For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ of full column rank and $\mathbf{v} \in \mathbb{R}^m$, it holds that

$$\|\mathbf{A}^T \mathbf{v}\|_\infty \geq r(\mathcal{P}(\mathbf{A})) \|\mathbf{v}\|_2, \quad (\text{C.19})$$

where $r(\mathcal{P}(\mathbf{A}))$ is the inradius of the symmetrized convex hull $\mathcal{P}(\mathbf{A})$ of the columns of \mathbf{A} .

Proof. The inequality (C.19) obviously holds for $\mathbf{v} = \mathbf{0}$. Pick any $\mathbf{v} \in \mathbb{R}^m \setminus \{\mathbf{0}\}$ and take $\epsilon \in (0, 1)$. Let $\eta = \epsilon \|\mathbf{v}\|_2 r(\mathcal{P}(\mathbf{A}))$ and assume that $\mathbf{v} \in \eta \mathcal{P}^\circ(\mathbf{A}) = \{\mathbf{z}: \|\mathbf{A}^T \mathbf{z}\|_\infty \leq \eta\}$, i.e., \mathbf{v} is an element of the η -scaled version of the polar set $\mathcal{P}^\circ(\mathbf{A})$. Note that $\eta > 0$ as $\|\mathbf{v}\|_2 > 0$, $\epsilon > 0$, and $r(\mathcal{P}(\mathbf{A})) > 0$ thanks to \mathbf{A} having full column rank. It follows from [13, Theorem 1.2] that

$$\frac{\|\mathbf{v}\|_2}{\eta} \leq R(\mathcal{P}^\circ(\mathbf{A})) = \frac{1}{r(\mathcal{P}(\mathbf{A}))}. \quad (\text{C.20})$$

Now, owing to $\eta = \epsilon \|\mathbf{v}\|_2 r(\mathcal{P}(\mathbf{A}))$, (C.20) implies that $\epsilon \geq 1$, which contradicts $\epsilon \in (0, 1)$. It therefore follows that $\mathbf{v} \in \mathbb{R}^m \setminus \{\eta \mathcal{P}^\circ(\mathbf{A})\}$ for all $\epsilon \in (0, 1)$, which in turn implies that $\|\mathbf{A}^T \mathbf{v}\|_\infty > \eta = \epsilon \|\mathbf{v}\|_2 r(\mathcal{P}(\mathbf{A}))$, for all $\epsilon \in (0, 1)$. In particular, letting $\epsilon \rightarrow 1$ yields $\|\mathbf{A}^T \mathbf{v}\|_\infty \geq \|\mathbf{v}\|_2 r(\mathcal{P}(\mathbf{A}))$, as desired. \square

Proof of (C.14): With $\sigma_{\min}(\mathbf{A}) = \|\mathbf{A}^{-1}\|_{2 \rightarrow 2}^{-1}$ [33, Section 5.2.1] for a full-rank matrix $\mathbf{A} \in \mathbb{R}^{m \times m}$, it follows that

$$\begin{aligned} \mathbb{P}[\overline{\mathcal{E}}_2] &= \mathbb{P}\left[\min_{\ell} \left\|(\mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)})^{-1}\right\|_{2 \rightarrow 2}^{-1} \leq 1 - \delta\right] \\ &= \mathbb{P}\left[\max_{\ell} \left\|(\mathbf{U}^{(\ell)T} \boldsymbol{\Phi}^T \boldsymbol{\Phi} \mathbf{U}^{(\ell)})^{-1}\right\|_{2 \rightarrow 2} \geq \frac{1}{1 - \delta}\right] \\ &\leq 2 e^{-\tau/2}, \end{aligned}$$

where $\tau > 0$ is the numerical constant in Theorem 3.3 and the last inequality is thanks to (B.7).

Proof of (C.15): By the union bound

$$\begin{aligned} \mathbb{P}[\overline{\mathcal{E}}_3^{(\ell, i, s)}] &\leq \sum_{k \neq \ell, j} \mathbb{P}\left[\left|\mathbf{a}_j^{(k)T} \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)}\right|\right] \\ &> \sqrt{\frac{6 \log N + 2 \log s_{\max}}{d_{\min}}} \left\|\mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)}\right\|_{2 \rightarrow 2} \left\|\tilde{\mathbf{r}}_s^{(\ell)}\right\|_2 \\ &\leq \sum_{k \neq \ell, j} \mathbb{P}\left[\left|\mathbf{a}_j^{(k)T} \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)}\right|\right] \end{aligned}$$

$$\begin{aligned}
&> \sqrt{\frac{6 \log N + 2 \log s_{\max}}{d_k}} \left\| \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)} \right\|_2 \Big] \\
&\leq \sum_{k \neq \ell, j} \frac{2}{s_{\max} N^3} \leq \frac{2}{s_{\max} N^2},
\end{aligned} \tag{C.21}$$

where (C.21) follows from Proposition A.1 with $\mathbf{a} = \mathbf{a}_j^{(k)}$, $\mathbf{b} = \mathbf{U}^{(k)T} (\boldsymbol{\Phi}^T \boldsymbol{\Phi} - \mathbf{I}) \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)}$ and $\beta = \sqrt{6 \log N + 2 \log s_{\max}}$.

Proof of (C.16): We first show that $\tilde{\mathbf{r}}_s^{(\ell)} / \|\tilde{\mathbf{r}}_s^{(\ell)}\|_2$ is distributed uniformly at random on $\mathbb{S}^{d_\ell-1}$; (C.16) then follows by application of Lemma B.2.

Recall that we consider reduced OMP, which computes a sparse representation of $\mathbf{x}_i^{(\ell)} = \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{a}_i^{(\ell)}$ using the columns of $\mathbf{X}_{-i}^{(\ell)} = \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{A}_{-i}^{(\ell)}$ as dictionary elements, i.e., Λ_s and $\tilde{\mathbf{r}}_s^{(\ell)}$ depend only on the random quantities $\boldsymbol{\Phi} \mathbf{U}^{(\ell)}$, $\mathbf{a}_i^{(\ell)}$ and $\mathbf{A}_{-i}^{(\ell)}$. In order to reflect these restricted dependencies, we write $\tilde{\mathbf{r}}_s^{(\ell)} = \tilde{\mathbf{r}}_s^{(\ell)}(\boldsymbol{\Phi} \mathbf{U}^{(\ell)}, \mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)})$ and $\Lambda_s = \Lambda_s(\boldsymbol{\Phi} \mathbf{U}^{(\ell)}, \mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)})$. Here, the first argument specifies the basis matrix of the data points, the second argument corresponds to the coefficient vector of the data point (in the basis specified by the first argument) a sparse representation is to be computed for, and the third argument designates the coefficient matrix of the dictionary elements (again in the basis specified by the first argument).

We start by showing that the distribution of $\tilde{\mathbf{r}}_s^{(\ell)}$ is rotationally invariant. For a deterministic unitary matrix $\mathbf{W} \in \mathbb{R}^{d_\ell \times d_\ell}$, we have

$$\Lambda_s(\boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T, \mathbf{W} \mathbf{a}_i^{(\ell)}, \mathbf{W} \mathbf{A}_{-i}^{(\ell)}) = \Lambda_s(\boldsymbol{\Phi} \mathbf{U}^{(\ell)}, \mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)})$$

as the $\mathbf{x}_j^{(\ell)}$ can be written as $\mathbf{x}_j^{(\ell)} = \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{a}_j^{(\ell)} = \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T \mathbf{W} \mathbf{a}_j^{(\ell)}$.

Using the shorthand notation Λ'_s for $\Lambda_s(\boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T, \mathbf{W} \mathbf{a}_i^{(\ell)}, \mathbf{W} \mathbf{A}_{-i}^{(\ell)})$ and recalling that $\tilde{\mathbf{r}}_s^{(\ell)} = (\mathbf{I} - \mathbf{A}_{\Lambda_s}^{(\ell)} (\boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{A}_{\Lambda_s}^{(\ell)})^\dagger \boldsymbol{\Phi} \mathbf{U}^{(\ell)}) \mathbf{a}_i^{(\ell)}$, it follows that

$$\begin{aligned}
\tilde{\mathbf{r}}_s^{(\ell)}(\boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T, \mathbf{W} \mathbf{a}_i^{(\ell)}, \mathbf{W} \mathbf{A}_{-i}^{(\ell)}) &= \left(\mathbf{I} - \mathbf{W} \mathbf{A}_{\Lambda'_s}^{(\ell)} \left(\boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T \mathbf{W} \mathbf{A}_{\Lambda'_s}^{(\ell)} \right)^\dagger \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T \right) \mathbf{W} \mathbf{a}_i^{(\ell)} \\
&= \left(\mathbf{I} - \mathbf{W} \mathbf{A}_{\Lambda_s}^{(\ell)} \left(\boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T \mathbf{W} \mathbf{A}_{\Lambda_s}^{(\ell)} \right)^\dagger \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T \right) \mathbf{W} \mathbf{a}_i^{(\ell)} \\
&= \mathbf{W} \left(\mathbf{I} - \mathbf{A}_{\Lambda_s}^{(\ell)} \left(\boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{A}_{\Lambda_s}^{(\ell)} \right)^\dagger \boldsymbol{\Phi} \mathbf{U}^{(\ell)} \right) \mathbf{a}_i^{(\ell)} \\
&= \mathbf{W} \tilde{\mathbf{r}}_s^{(\ell)}(\boldsymbol{\Phi} \mathbf{U}^{(\ell)}, \mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)}).
\end{aligned}$$

By rotational invariance of the distributions of $\mathbf{a}_i^{(\ell)}$, $\mathbf{A}_{-i}^{(\ell)}$ and $\boldsymbol{\Phi}$ (by assumption in Theorem 3.3), we have $\mathbf{W} \mathbf{a}_i^{(\ell)} \sim \mathbf{a}_i^{(\ell)}$, $\mathbf{W} \mathbf{A}_{-i}^{(\ell)} \sim \mathbf{A}_{-i}^{(\ell)}$ and $\boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T \sim \boldsymbol{\Phi} \mathbf{U}^{(\ell)}$ (because $\text{span}(\mathbf{U}^{(\ell)} \mathbf{W}^T) = \text{span}(\mathbf{U}^{(\ell)})$ and the columns of $\mathbf{U}^{(\ell)} \mathbf{W}^T$ are orthonormal). We therefore get

$$\begin{aligned}
\tilde{\mathbf{r}}_s^{(\ell)}(\boldsymbol{\Phi} \mathbf{U}^{(\ell)}, \mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)}) &\sim \tilde{\mathbf{r}}_s^{(\ell)}(\boldsymbol{\Phi} \mathbf{U}^{(\ell)} \mathbf{W}^T, \mathbf{W} \mathbf{a}_i^{(\ell)}, \mathbf{W} \mathbf{A}_{-i}^{(\ell)}) \\
&= \mathbf{W} \tilde{\mathbf{r}}_s^{(\ell)}(\boldsymbol{\Phi} \mathbf{U}^{(\ell)}, \mathbf{a}_i^{(\ell)}, \mathbf{A}_{-i}^{(\ell)}).
\end{aligned} \tag{C.22}$$

Since (C.22) holds for all unitary matrices \mathbf{W} , the distribution of $\tilde{\mathbf{r}}_s^{(\ell)}$ is rotationally invariant and $\tilde{\mathbf{r}}_s^{(\ell)} / \|\tilde{\mathbf{r}}_s^{(\ell)}\|_2$ is, indeed, distributed uniformly on $\mathbb{S}^{d_\ell-1}$. We finally exploit this property of $\tilde{\mathbf{r}}_s^{(\ell)}$ to upper-bound $\mathbb{P}[\overline{\mathcal{E}}_4^{(\ell,i,s)}]$ as follows. A union bound over all $k, k \neq \ell$, yields

$$\begin{aligned} \mathbb{P}[\overline{\mathcal{E}}_4^{(\ell,i,s)}] &\leq \sum_{k \neq \ell} \mathbb{P} \left[\left\| \mathbf{A}^{(k)T} \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \tilde{\mathbf{r}}_s^{(\ell)} \right\|_\infty \geq 4(3 \log N + \log s_{\max}) \frac{\left\| \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)} \right\|_F}{\sqrt{d_k} \sqrt{d_\ell}} \left\| \tilde{\mathbf{r}}_s^{(\ell)} \right\|_2 \right] \\ &\leq \sum_{k \neq \ell} \frac{n_k + 1}{s_{\max} N^3} = \frac{N - n_\ell + L - 1}{s_{\max} N^3} < \frac{N + L}{s_{\max} N^3} \leq \frac{2}{s_{\max} N^2}, \end{aligned} \quad (\text{C.23})$$

where (C.23) follows by application of Lemma B.2 with $\mathbf{L} = \mathbf{A}^{(k)}$, $\mathbf{a} = \tilde{\mathbf{r}}_s^{(\ell)} / \|\tilde{\mathbf{r}}_s^{(\ell)}\|_2$, $\mathbf{B} = \mathbf{U}^{(k)T} \mathbf{U}^{(\ell)}$ and $c = 4(3 \log N + \log s_{\max})$.