

Scalable Sparse Subspace Clustering by Orthogonal Matching Pursuit

Chong You, Daniel P. Robinson, René Vidal
Johns Hopkins University, Baltimore, MD, 21218, USA

Abstract

Subspace clustering methods based on ℓ_1 , ℓ_2 or nuclear norm regularization have become very popular due to their simplicity, theoretical guarantees and empirical success. However, the choice of the regularizer can greatly impact both theory and practice. For instance, ℓ_1 regularization is guaranteed to give a subspace-preserving affinity (i.e., there are no connections between points from different subspaces) under broad conditions (e.g., arbitrary subspaces and corrupted data). However, it requires solving a large scale convex optimization problem. On the other hand, ℓ_2 and nuclear norm regularization provide efficient closed form solutions, but require very strong assumptions to guarantee a subspace-preserving affinity, e.g., independent subspaces and uncorrupted data. In this paper we study a subspace clustering method based on orthogonal matching pursuit. We show that the method is both computationally efficient and guaranteed to give a subspace-preserving affinity under broad conditions. Experiments on synthetic data verify our theoretical analysis, and applications in handwritten digit and face clustering show that our approach achieves the best trade off between accuracy and efficiency. Moreover, our approach is the first one to handle 100,000 data points.

1. Introduction

In many computer vision applications, such as motion segmentation [7, 31, 25], hand written digit clustering [38] and face clustering [1, 18], data from different classes can be well approximated by a union of low dimensional subspaces. In these scenarios, the task is to partition the data according to the membership of data points to subspaces.

More formally, given a set of points $\mathcal{X} = \{\mathbf{x}_j \in \mathbb{R}^D\}_{j=1}^N$ lying in an unknown number n of subspaces $\{S_i\}_{i=1}^n$ of unknown dimensions $\{d_i\}_{i=1}^n$, *subspace clustering* is the problem of clustering the data into groups such that each group contains only data points from the same subspace. This problem has received great attention in the past decade and many subspace clustering algorithms have been developed, including iterative, algebraic, statistical, and spectral clustering based methods (see [29] for a review).

Sparse and Low Rank Methods. Among existing techniques, methods based on applying spectral clustering to an affinity matrix obtained by solving an optimization problem that incorporates ℓ_1 , ℓ_2 or nuclear norm regularization have become extremely popular due to their simplicity, theoretical correctness, and empirical success. These methods are based on the so-called *self-expressiveness property* of data lying in a union of subspaces, originally proposed in [10]. This property states that each point in a union of subspaces can be written as a linear combination of other data points in the subspaces. That is,

$$\begin{aligned} \mathbf{x}_j &= X\mathbf{c}_j \text{ and } c_{jj} = 0, \text{ or equivalently} \\ X &= XC \text{ and } \text{diag}(C) = \mathbf{0}, \end{aligned} \quad (1)$$

where $X = [\mathbf{x}_1, \dots, \mathbf{x}_N] \in \mathbb{R}^{D \times N}$ is the data matrix and $C = [c_1, \dots, c_N] \in \mathbb{R}^{N \times N}$ is the matrix of coefficients.

While (1) may not have a unique solution for C , there exist solutions whose entries are such that if $c_{ij} \neq 0$, then \mathbf{x}_i is in the same subspace as \mathbf{x}_j . For example, a point $\mathbf{x}_j \in S_i$ can always be written as a linear combination of d_i other points in S_i . Such solutions are called *subspace preserving* since they preserve the clustering of the subspaces. Given a subspace preserving C , one can build an affinity matrix W between every pair of points \mathbf{x}_i and \mathbf{x}_j as $w_{ij} = |c_{ij}| + |c_{ji}|$, and apply spectral clustering [32] to W to cluster the data.

To find a subspace preserving C , existing methods regularize C with a norm $\|\cdot\|$, and solve a problem of the form:

$$C^* = \arg \min_C \|C\| \text{ s.t. } X = XC, \text{ diag}(C) = \mathbf{0}. \quad (2)$$

For instance, the sparse subspace clustering (SSC) algorithm [10] uses the ℓ_1 norm to encourage the sparsity of C . Prior work has shown that SSC gives a subspace-preserving solution if the subspaces are independent [10, 12], or if the data from different subspaces satisfy certain separation conditions and data from the same subspace are well spread out [11, 12, 36, 26]. Similar results exist for cases where data is corrupted by noise [33, 27] and outliers [26]. Other self-expressiveness based methods use different regularizations on the coefficient matrix C . Least squares regression (LSR) [22] uses ℓ_2 regularization on C . Low rank representation (LRR) [20, 19] and low rank subspace clustering (LRSC)

[13, 30] use nuclear norm minimization to encourage C to be low-rank. Based on these, [21, 23, 16, 35] study regularizations that are a mixture of ℓ_1 and ℓ_2 , and [34, 39] propose regularizations that are a blend of ℓ_1 and the nuclear norm.

The advantage of ℓ_2 regularized LSR and nuclear norm regularized LRR and LRSC over sparsity regularized SSC is that the solution for C can be computed in closed form from the SVD of the (noiseless) data matrix X , thus they are computationally more attractive. However, the resulting C is subspace preserving only when subspaces are independent and the data is uncorrupted. Thus, there is a need for methods that both guarantee a subspace-preserving affinity under broad conditions and are computationally efficient.

Paper Contributions. In this work we study the self-expressiveness based subspace clustering method that uses orthogonal matching pursuit (OMP) to find a sparse representation in lieu of the ℓ_1 -based basis pursuit (BP) method. The method is termed SSC-OMP, for its kinship to the original SSC, which is referred to as SSC-BP in this paper.

The main contributions of this paper are to find theoretical conditions under which the affinity produced by SSC-OMP is subspace preserving and to demonstrate its efficiency for large scale problems. Specifically, we show that:

1. When the subspaces and the data are **deterministic**, SSC-OMP gives a subspace-preserving C if the subspaces are independent, or else if the subspaces are sufficiently separated and the data is well distributed.
2. When the subspaces and data are drawn uniformly at **random**, SSC-OMP gives a subspace-preserving C if the dimensions of the subspaces are sufficiently small relative to the ambient dimension by a factor controlled by the sample density and the number of subspaces.
3. SSC-OMP is orders of magnitude **faster** than the original SSC-BP, and can handle up to 100,000 data points.

Related work. It is worth noting that the idea of using OMP for SSC had already been considered in [9]. The core contribution of our work is to provide much weaker yet more succinct and interpretable conditions for the affinity to be subspace preserving in the case of arbitrary subspaces. In particular, our conditions are naturally related to those for SSC-BP, which reveal insights about the relationship between these two sparsity-based subspace clustering methods. Moreover, our experimental results provide a much more detailed evaluation of the behavior of SSC-OMP for large-scale problems. It is also worth noting that conditions under which OMP gives a subspace-preserving representation had also been studied in [36]. Our paper presents a much more comprehensive study of OMP for the subspace clustering problem, by providing results under deterministic independent, deterministic arbitrary and random subspace models. In particular, our result for deterministic arbitrary models is much stronger than the main result in [36].

2. SSC by Orthogonal Matching Pursuit

The SSC algorithm approaches the subspace clustering problem by finding a sparse representation of each point in terms of other data points. Since each point in S_i can be expressed in terms of at most $d_i \ll N$ other points in S_i , such a sparse representation always exists. In principle, we can find it by solving the following optimization problem:

$$\mathbf{c}_j^* = \arg \min_{\mathbf{c}_j} \|\mathbf{c}_j\|_0 \text{ s.t. } \mathbf{x}_j = X\mathbf{c}_j, c_{jj} = 0, \quad (3)$$

where $\|\mathbf{c}\|_0$ counts the number of nonzero entries in \mathbf{c} . Since this problem is NP hard, the SSC method in [10] relaxes this problem and solves the following ℓ_1 problem:

$$\mathbf{c}_j^* = \arg \min_{\mathbf{c}_j} \|\mathbf{c}_j\|_1 \text{ s.t. } \mathbf{x}_j = X\mathbf{c}_j, c_{jj} = 0. \quad (4)$$

Since this problem is called the basis pursuit (BP) problem, we refer to the SSC algorithm in [10] as SSC-BP.

The optimization problems (3) and (4) have been studied extensively in the compressed sensing community, see, e.g., the tutorials [2, 4], and it is well known that, under certain conditions on the dictionary X , their solutions are the same. However, results from compressed sensing do not apply to the subspace clustering problem because when the columns of X lie in a union of subspaces the solution for C need not be unique (see Section 3 for more details). This has motivated extensive research on the conditions under which the solutions of (3) or (4) are useful for subspace clustering.

It is shown in [10, 11, 12] that when the subspaces are either independent or disjoint, and the data are noise free and well distributed, both (3) and (4) provide a sparse representation \mathbf{c}_j that is *subspace preserving*, as defined next.

Definition 1 (Subspace-preserving representation). *A representation $\mathbf{c} \in \mathbb{R}^N$ of a point $\mathbf{x} \in S_i$ in terms of the dictionary $X = [\mathbf{x}_1, \dots, \mathbf{x}_N]$ is called subspace preserving if its nonzero entries correspond to points in S_i , i.e.*

$$\forall j = 1, \dots, N, \quad c_j \neq 0 \implies \mathbf{x}_j \in S_i. \quad (5)$$

In practice, however, solving N ℓ_1 -minimization problems over N variables may be prohibitive when N is large. As an alternative, consider the following program:

$$\mathbf{c}_j^* = \arg \min_{\mathbf{c}_j} \|\mathbf{x}_j - X\mathbf{c}_j\|_2^2 \text{ s.t. } \|\mathbf{c}_j\|_0 \leq k, c_{jj} = 0. \quad (6)$$

It is shown in [28, 8] that, under certain conditions, this problem can be solved using the orthogonal matching pursuit (OMP) algorithm [24] (Algorithm 1). OMP solves the problem $\min_{\mathbf{c}} \|A\mathbf{c} - \mathbf{b}\|_2^2$ s.t. $\|\mathbf{c}\|_0 \leq k$ greedily by selecting one column of $A = [\mathbf{a}_1, \dots, \mathbf{a}_M]$ at a time (the one that maximizes the absolute value of the dot product with the residual in line 3) and computing the coefficients for the selected columns until k columns are selected. For subspace clustering purposes, the vector $\mathbf{c}_j^* \in \mathbb{R}^N$ (the j th column of $C^* \in \mathbb{R}^{N \times N}$), is computed as $\text{OMP}(X_{-j}, \mathbf{x}_j) \in \mathbb{R}^{N-1}$

Algorithm 1 : Orthogonal Matching Pursuit (OMP)

Input: $A = [\mathbf{a}_1, \dots, \mathbf{a}_M] \in \mathbb{R}^{m \times M}$, $\mathbf{b} \in \mathbb{R}^m$, k_{\max} , ϵ .

- 1: **Initialize** $k = 0$, residual $\mathbf{q}_0 = \mathbf{b}$, support set $T_0 = \emptyset$.
- 2: **while** $k < k_{\max}$ and $\|\mathbf{q}_k\|_2 > \epsilon$ **do**
- 3: $T_{k+1} = T_k \cup \{i^*\}$, where $i^* = \arg \max_{i=1, \dots, M} |\mathbf{a}_i^T \mathbf{q}_k|$ ¹.
- 4: $\mathbf{q}_{k+1} = (I - P_{T_{k+1}})\mathbf{b}$, where $P_{T_{k+1}}$ is the projection onto the span of the vectors $\{\mathbf{a}_j, j \in T_{k+1}\}$.
- 5: $k \leftarrow k + 1$.
- 6: **end while**

Output: $\mathbf{c}^* = \arg \min_{\mathbf{c}: \text{Supp}(\mathbf{c}) \subseteq T_k} \|\mathbf{b} - A\mathbf{c}\|_2$.

Algorithm 2 : Sparse Subspace Clustering by Orthogonal Matching Pursuit (SSC-OMP)

Input: Data $X = [\mathbf{x}_1, \dots, \mathbf{x}_N]$, parameters k_{\max} , ϵ .

- 1: Compute \mathbf{c}_j^* from OMP(X_{-j}, \mathbf{x}_j) using Algorithm 1.
- 2: Set $C^* = [\mathbf{c}_1^*, \dots, \mathbf{c}_N^*]$ and $W = |C^*| + |C^{*\top}|$.
- 3: Compute segmentation from W by spectral clustering.

Output: Segmentation of data X .

with a zero inserted in its j th entry, where X_{-j} is the data matrix with the j th column removed. After C^* is computed, the segmentation of the data is found by applying spectral clustering to the affinity matrix $W = |C^*| + |C^{*\top}|$ as done in SSC-BP. The procedure is summarized in Algorithm 2.

3. Theoretical Analysis of SSC-OMP

OMP has been shown to be effective for sparse recovery, with the advantage over BP that it admits simple, fast implementations. However, note that existing conditions for the correctness of OMP for sparse recovery are too strong for the subspace clustering problem. In particular, note that the matrix X need not satisfy the mutual incoherence [28] or restricted isometry properties [8], as two points in a subspace could be arbitrarily close to each other. More importantly, these conditions are not applicable here because our goal is not to recover a *unique* sparse solution. In fact, the sparse solution is not unique since any d_i linearly independent points from S_i can represent a point $\mathbf{x}_j \in S_i$. Therefore, there is a need to find conditions under which the output of OMP (which need not coincide with the solution of (6) or (3)) is guaranteed to be subspace preserving.

This section is devoted to studying sufficient conditions under which SSC-OMP gives a subspace-preserving representation. Our analysis assumes that the data is noiseless. The termination parameters of Algorithm 1 are $\epsilon = 0$ and k_{\max} large enough (e.g., $k_{\max} = M$). We also assume that the columns of X are normalized to unit ℓ_2 norm. To make our results consistent with state-of-the-art results, we first study the case where the subspaces are deterministic,

¹If $\arg \max$ in step 3 of the algorithm gives multiple items, pick one of them in a deterministic way, e.g., pick the one with the smallest index.

including both independent subspaces as well as arbitrary subspaces. We then study the case where both the subspaces and the data points are drawn at random.

3.1. Independent Deterministic Subspace Model

We first consider the case where the subspaces are fixed, the data points are fixed, and the subspaces are independent.

Definition 2. A collection of subspaces $\{S_i\}_{i=1}^n$ is called *independent* if $\dim(\sum_i S_i) = \sum_i \dim(S_i)$, where $\sum_i S_i$ is defined as the subspace $\{\sum_i \mathbf{x}_i : \mathbf{x}_i \in S_i\}$.

Notice that two subspaces are *independent* if and only if they are *disjoint*, i.e., if they intersect only at the origin. However, pairwise disjoint subspaces need not be independent, e.g., three lines in \mathbb{R}^2 are disjoint but not independent. Notice also that any subset of a set of independent subspaces is also independent. Therefore, any two subspaces in a set of independent subspaces are independent and hence disjoint. In particular, this implies that if $\{S_i\}_{i=1}^n$ are independent, then S_i and $S_{(-i)} := \sum_{m \neq i} S_m$ are independent.

To establish conditions under which SSC-OMP gives a subspace-preserving affinity for independent subspaces, it is important to note that when computing OMP(X_{-j}, \mathbf{x}_j), the goal is to select other points in the same subspace as \mathbf{x}_j . The process for selecting these points occurs in step 3 of Algorithm 1, where the dot products between all points \mathbf{x}_m , $m \neq j$, and the current residual \mathbf{q}_k are computed and the point with the highest product (in absolute value) is chosen. Since in the first iteration the residual is $\mathbf{q}_0 = \mathbf{x}_j$, we could immediately choose a point \mathbf{x}_m in another subspace whenever the dot product of \mathbf{x}_j with a point in another subspace is larger than the dot product of \mathbf{x}_j with points in its own subspace. What the following theorem shows is that, even though OMP may select points in the wrong subspaces as the iterations proceed, the coefficients associated to points in other subspaces will be zero at the end. Therefore, OMP (with $\epsilon = 0$ and $k_{\max} = N - 1$) is guaranteed to find a subspace-preserving representation.

Theorem 1. *If the subspaces are independent, OMP gives a subspace-preserving representation of each data point.*

Proof. [Sketch only] Assume that $\mathbf{x}_j \in S_i$. Since $\epsilon = 0$ and k_{\max} is large, OMP gives an exact representation, i.e., $\mathbf{x}_j = X\mathbf{c}_j$ and $c_{jj} = 1$. Thus, since S_i and $S_{(-i)}$ are independent, the coefficients of data points in $S_{(-i)}$ must be zero. \square

3.2. Arbitrary Deterministic Subspace Model

We will now consider a more general class of subspaces, which need not be independent or disjoint, and investigate conditions under which OMP gives a subspace-preserving representation. In the following, $X^i \in \mathbb{R}^{D \times N_i}$ denotes the submatrix of X containing the points in the i th subspace; for any $\mathbf{x}_j \in S_i$, $X_{-j}^i \in \mathbb{R}^{D \times (N_i - 1)}$ denotes the matrix X^i

with the point \mathbf{x}_j removed; \mathcal{X}^i and \mathcal{X}_{-j}^i denote respectively the set of vectors contained in the columns of X^i and X_{-j}^i .

Now, it is easy to see that a sufficient condition for $\text{OMP}(X_{-j}, \mathbf{x}_j)$ to be subspace preserving is that for each k in step 3 of Algorithm 1, the point that maximizes the dot product lies in the same subspace as \mathbf{x}_j . Since $\mathbf{q}_0 = \mathbf{x}_j$ and \mathbf{q}_1 is equal to \mathbf{x}_j minus the projection of \mathbf{x}_j onto the subspace spanned by the selected point, say $\hat{\mathbf{x}}$, it follows that if $\mathbf{x}_j, \hat{\mathbf{x}} \in S_i$ then $\mathbf{q}_1 \in S_i$. By a simple induction argument, it follows that if all the selected points are in S_i , then so are the residuals $\{\mathbf{q}_k\}$. This suggests that the condition for $\text{OMP}(X_{-j}, \mathbf{x}_j)$ to be subspace preserving must depend on the dot products between the data points and a subset of the set of residuals (the subset contained in the same subspace as \mathbf{x}_j). This motivates the following definition and lemma.

Definition 3. Let $Q(A, \mathbf{b})$ be the set of all residual vectors computed in step 4 of $\text{OMP}(A, \mathbf{b})$. The set of OMP residual directions associated with matrix X_{-j}^i and point $\mathbf{x}_j \in S_i$ is defined as:

$$\mathcal{W}_j^i := \{\mathbf{w} = \frac{\mathbf{q}}{\|\mathbf{q}\|_2} : \mathbf{q} \in Q(X_{-j}^i, \mathbf{x}_j), \mathbf{q} \neq \mathbf{0}\}. \quad (7)$$

The set of OMP residual directions associated with the data matrix X^i is defined as $\mathcal{W}^i := \bigcup_{j: \mathbf{x}_j \in S_i} \mathcal{W}_j^i$.

Lemma 1. OMP gives a subspace-preserving representation for point $\mathbf{x}_j \in S_i$ in at most d_i iterations if

$$\forall \mathbf{w} \in \mathcal{W}_j^i \quad \max_{\mathbf{x} \in \bigcup_{k \neq i} \mathcal{X}^k} |\mathbf{w}^\top \mathbf{x}| < \max_{\mathbf{x} \in \mathcal{X}^i \setminus \{\mathbf{x}_j\}} |\mathbf{w}^\top \mathbf{x}|. \quad (8)$$

Proof. [Sketch only] By using an induction argument, it is easy to see that the condition in (8) implies that the sequence of residuals of $\text{OMP}(X_{-j}, \mathbf{x}_j)$ is the same as that of the fictitious problem $\text{OMP}(X_{-j}^i, \mathbf{x}_j)$. Hence, the output of $\text{OMP}(X_{-j}, \mathbf{x}_j)$ is the same as that of $\text{OMP}(X_{-j}^i, \mathbf{x}_j)$, which is, by construction, subspace-preserving. \square

Intuitively, Lemma 1 tells us that if the dot product between the residual directions for subspace i and the data points in all other subspaces is smaller than the dot product between the residual directions for subspace i and all points in subspace i other than $\mathbf{x}_j \in S_i$, then OMP gives a subspace-preserving representation. While such a condition is very intuitive from the perspective of OMP, it is not as intuitive from the perspective of subspace clustering as it does not rely on the geometry of the problem. Specifically, it does not directly depend on the relative configuration of the subspaces or the distribution of the data in the subspaces. In what follows, we derive conditions on the subspaces and the data that guarantee that the condition in (8) holds. Before doing so, we need some additional definitions.

Definition 4. The coherence between two sets of points of unit norm, \mathcal{X} and \mathcal{Y} , is defined as $\mu(\mathcal{X}, \mathcal{Y}) = \max_{\mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}} |\langle \mathbf{x}, \mathbf{y} \rangle|$.

The coherence measures the degree of ‘‘similarity’’ between two sets of points. In our case, we can see that the left hand side of (8) is bounded above by the coherence between the sets \mathcal{W}^i and $\bigcup_{k \neq i} \mathcal{X}^k$. As per (8), this coherence should be small, which implies that data points from different subspaces should be sufficiently separated (in angle).

Definition 5. The inradius $r(\mathcal{P})$ of a convex body \mathcal{P} is the radius of the largest Euclidean ball inscribed in \mathcal{P} .

As shown in Lemma 2, the right hand side of (8) is bounded below by $r(\mathcal{P}_{-j}^i)$, where $\mathcal{P}_{-j}^i := \text{conv}(\pm \mathcal{X}_{-j}^i)$ is the symmetrized convex hull of the points in the i th subspace other than \mathbf{x}_j , i.e., \mathcal{X}_{-j}^i . Therefore, (8) suggests that the minimum inradius $r_i := \min_j r(\mathcal{P}_{-j}^i)$ should be large, which means the points in S_i should be well-distributed.

Lemma 2. Let $\mathbf{x}_j \in S_i$. Then, for all $\mathbf{w} \in \mathcal{W}_j^i$, we have:

$$\begin{aligned} \max_{\mathbf{x} \in \bigcup_{k \neq i} \mathcal{X}^k} |\mathbf{w}^\top \mathbf{x}| &\leq \max_{k: k \neq i} \mu(\mathcal{W}^i, \mathcal{X}^k) \leq \max_{k: k \neq i} \mu(\mathcal{X}^i, \mathcal{X}^k) / r_i; \\ \max_{\mathbf{x} \in \mathcal{X}^i \setminus \{\mathbf{x}_j\}} |\mathbf{w}^\top \mathbf{x}| &\geq r(\mathcal{P}_{-j}^i) \geq r_i. \end{aligned} \quad (9)$$

Proof. The proof can be found in [37]. \square

Lemma 2 allows us to make the condition of Lemma 1 more interpretable, as stated in the following theorem.

Theorem 2. The output of OMP is subspace preserving if

$$\forall i = 1, \dots, n, \quad \max_{k: k \neq i} \mu(\mathcal{W}^i, \mathcal{X}^k) < r_i. \quad (10)$$

Corollary 1. The output of OMP is subspace preserving if

$$\forall i = 1, \dots, n, \quad \max_{k: k \neq i} \mu(\mathcal{X}^i, \mathcal{X}^k) < r_i^2. \quad (11)$$

Note that points in \mathcal{W}^i are all in subspace S_i , as step 4 of $\text{OMP}(A := X_{-j}^i, \mathbf{b} := \mathbf{x}_j)$ has \mathbf{b} and $P_{T_{k+1}} \mathbf{b}$ both in S_i . The conditions (10) and (11) thus show that for each subspace S_i , a set of points (i.e., \mathcal{X}^i or \mathcal{W}^i) in S_i should have low coherence with all points from other subspaces, and that points in \mathcal{X}^i should be uniformly located in S_i to have a large inradius. This is in agreement with the intuition that points from different subspaces should be well separated, and points within a subspace should be well distributed.

For a comparison of Corollary 1 and Theorem 2, note that due to Lemma 2 condition (10) is tighter than condition (11), making Theorem 2 preferable. Yet Corollary 1 has the advantage that both sides of condition (11) depend directly on the data points in \mathcal{X} , while condition (10) depends on the residual points in \mathcal{W}^i , making it algorithm specific.

Another important thing to notice is that conditions (10) and (11) can be satisfied even if the subspaces are neither independent nor disjoint. For example, consider the case where $S_i \cap S_k \neq \mathbf{0}$. Then, the coherence $\mu(\mathcal{W}^i, \mathcal{X}^k)$ could

still be small as long as no points in \mathcal{W}^i and \mathcal{X}^k are near the intersection of S_i and S_k . Actually, even this is too strong of an assumption since the intersection is a subspace. Thus, $\mathbf{x} \in \mathcal{X}^k$, $\mathbf{y} \in \mathcal{W}^i$ could both be very close to the intersection yet have low coherence. The same argument also works for condition (11). Admittedly, under specific distributions of points, it is possible that there exists $\mathbf{x} \in \mathcal{X}^k$ and $\mathbf{y} \in \mathcal{W}^i$ that are arbitrarily close to each other when they are near the intersection. However, this worst case scenario is unlikely to happen if we consider a random model, as discussed next.

3.3. Arbitrary Random Subspace Model

This section considers the fully random union of subspaces model in [26], where the basis elements of each subspace are chosen uniformly at random from the unit sphere of the ambient space and the data points from each subspace are uniformly distributed on the unit sphere of that subspace. Theorem 3 shows that the sufficient condition in (10) holds true with high probability (*i.e.* the probability goes to 1 as the density of points grows to infinity) given some conditions on the subspace dimension d , the ambient space dimension D , the number of subspaces n and the number of data points per subspace.

Theorem 3. *Assume a random union of subspaces model where all subspaces are of equal dimension d and the number of data points in each subspace is $\rho d + 1$, where $\rho > 1$ is the “density”, so that the total number data points in all subspaces is $N(n, \rho, d) = n(\rho d + 1)$. The output of OMP is subspace preserving with probability $p > 1 - \frac{2d}{N(n, \rho, d)} - N(n, \rho, d)e^{-\sqrt{\rho}d}$ if*

$$d < \frac{c^2(\rho) \log \rho}{12} \frac{D}{\log N(n, \rho, d)}, \quad (12)$$

where $c(\rho) > 0$ is a constant that depends only on ρ .

One interpretation of the condition in (12) is that the dimension d of the subspaces should be small relative to the ambient dimension D . It also shows that as the number of subspaces n increases, the factor $\log N(n, \rho, d)$ also increases, making the condition more difficult to be satisfied. In terms of the density ρ , it is shown in [26] that there exists a ρ_0 such that $c(\rho) = 1/\sqrt{8}$ when $\rho > \rho_0$. Then, it is easy to see that when $\rho > \rho_0$, the term that depends on ρ is $\frac{\log \rho}{\log N(n, \rho, d)} = \frac{\log \rho}{\log(n(\rho d + 1))}$, which is a monotonically increasing function of ρ . This makes the condition easier to be satisfied as the density of points in the subspaces increases. Moreover, the probability of success is $1 - \frac{2d}{N(n, \rho, d)} - N(n, \rho, d)e^{-\sqrt{\rho}d}$, which is also an increasing function of ρ when ρ is greater than a threshold value. As a consequence, as the density of the points increases, the condition in Theorem 3 becomes easier to satisfy and the probability of success also increases.

4. Relationships with Other Methods

In this section we compare our results for SSC-OMP with those for other methods of the general form in (2). These methods include SSC-BP [11, 12, 26], which uses the ℓ_1 norm as a regularizer, LRR [19] and LRSC [30], which use the nuclear norm, and LSR [22] which uses the ℓ_2 norm. We also compare our results to those of [9] for SSC-OMP. The comparison is in terms of whether the solutions given by these alternative algorithms are subspace-preserving.

Independent Subspaces. Independence is a strong assumption on the union of subspaces. Under this assumption, a subspace has a trivial intersection with not only every other subspace, but also the union of all other subspaces. This case turns out to be especially easy for a large category of self-expressive subspace clustering methods [22], and SSC-BP, LRR, LRSC and LSR are all able to give subspace-preserving representations. Thus, in this easy case, the proposed method is as good as state-of-the-art methods.

Arbitrary Subspaces. To the best of our knowledge, when the subspaces are not independent, there is no guarantee of correctness for LRR, LRSC and LSR. For SSC-BP, as shown in [26], the representation is subspace-preserving if

$$\forall i = 1, \dots, n, \quad \max_{k:k \neq i} \mu(\mathcal{V}^i, \mathcal{X}^k) < r_i, \quad (13)$$

where \mathcal{V}^i is a set of N_i dual directions associated with X^i . When comparing (13) with our result in condition (10), we can see that the right hand sides are the same. However, the left hand sides are not directly comparable, as no general relationship is known between the sets \mathcal{V}^i and \mathcal{W}^i . Nonetheless, notice that the number of points in these two sets are not the same since $\text{card}(\mathcal{V}^i) = N_i$ and $\text{card}(\mathcal{W}^i) = N_i d_i$. Therefore, if we assume that the points in \mathcal{V}^i and \mathcal{W}^i are distributed uniformly at random on the unit sphere, then $\mu(\mathcal{W}^i, \mathcal{X}^k)$ is expected to be larger than $\mu(\mathcal{V}^i, \mathcal{X}^k)$, making the condition for SSC-OMP less likely to be satisfied than that for SSC-BP. Now, when comparing (13) with our condition in (11), we see that the left hand sides are comparable under a random model where both \mathcal{V}^i and \mathcal{X}^i contain N_i points. However, the right hand side is r_i^2 , which is less than or equal to r_i since the data are normalized and $r_i \leq 1$. This again makes the condition for SSC-OMP more difficult to hold than that for SSC-BP. However, this difference is expected to vanish for large scale problems, and SSC-OMP is computationally more efficient, as we will see in Section 5.

Random Subspaces. For the random model, [26] shows that SSC-BP gives a subspace-preserving representation with probability $p > 1 - \frac{2}{N(n, \rho, d)} - N(n, \rho, d)e^{-\sqrt{\rho}d}$ if

$$d < \frac{c^2(\rho) \log \rho}{12} \frac{D}{\log N(n, \rho, d)}. \quad (14)$$

If we compare this result with that of Theorem 3, we can see that the condition under which both methods succeed

with high probability is exactly the same. The difference between them is that SSC-BP has a higher probability of success than SSC-OMP when $d > 1$. However, it is easy to see that the difference in probability goes to zero as the density ρ goes to infinity. This means that the performance difference vanishes as the scale of the problem increases.

Other Results for SSC-OMP. Finally, we compare our results with those in [9] for SSC-OMP. Define the principal angle between two subspaces S_i and S_k as:

$$\theta_{i,k}^* = \min_{\substack{\mathbf{x} \in S_i \\ \|\mathbf{x}\|_2=1}} \min_{\substack{\mathbf{y} \in S_k \\ \|\mathbf{y}\|_2=1}} \arccos \langle \mathbf{x}, \mathbf{y} \rangle. \quad (15)$$

It is shown in [9] that the output of SSC-OMP is subspace-preserving if for all $i = 1, \dots, n$,

$$\max_{k:k \neq i} \mu(\mathcal{X}^i, \mathcal{X}^k) < r_i - \frac{2\sqrt{1 - (r_i)^2}}{\sqrt[4]{12}} \max_{k:k \neq i} \cos \theta_{i,k}^*. \quad (16)$$

The merit of this result is that it introduces the subspace angles in the condition, and satisfies the intuition that the algorithm is more likely to work if the subspaces are far apart from each other. However, the RHS of the condition shows an intricate relationship between the intra-class property r_i and the inter-class property $\theta_{i,k}^*$, which greatly complicates the interpretation of the condition. More importantly, as is shown in [37], the condition is more restrictive than (10), which makes Theorem 2 a stronger result.

5. Experiments

In this section, we first verify our theoretical results for SSC-OMP and compare them with those for SSC-BP by doing experiments on synthetic data using the random model. Specifically, we show that even if the subspaces are not independent, the solution of OMP is subspace-preserving with a probability that grows with the density of data points. Second, we test the performance of the proposed method on clustering images of handwritten digits and human faces, and conclude that SSC-OMP achieves the best trade off between accuracy and efficiency.

Methods. We compare the performance of state-of-the-art spectral subspace clustering methods, including LRSC [30], SSC-BP [12], LSR [22], and spectral curvature clustering (SCC) [5]. In real experiments, we use the code provided by the respective authors for computing the representation matrix C^* , where the parameters are tuned to give the best clustering accuracy. We then apply the normalized spectral clustering in [32] to the affinity $|C^*| + |C^{*\top}|$, except for SCC which has its own spectral clustering step.

Metrics. We use two metrics to evaluate the degree to which the subspace-preserving property is satisfied. The first one is a direct measure of whether the solution is subspace preserving or not. However, for comparing with state

of the art methods whose output is generally not subspace preserving, the second one measures how close the coefficients are from being subspace preserving.

– *Percentage of subspace-preserving representations ($p\%$):* this is the percentage of points whose representations are subspace-preserving. Due to inexactness in the solvers, coefficients with absolute value less than 10^{-3} are considered zero. A subspace-preserving solution gives $p = 100$.

– *Subspace-preserving representation error ($e\%$) [12]:* for each \mathbf{c}_j in (1), we compute the fraction of its ℓ_1 norm that comes from other subspaces and then average over all j , *i.e.*, $e = \frac{100}{N} \sum_j (1 - \sum_i (\omega_{ij} \cdot |\mathbf{c}_{ij}|) / \|\mathbf{c}_j\|_1)$, where $\omega_{ij} \in \{0, 1\}$ is the true affinity. A subspace-preserving C gives $e = 0$.

Now, the performance of subspace clustering depends not only on the subspace-preserving property, but also the connectivity of the similarity graph, *i.e.*, whether the data points in each cluster form a connected component of the graph.

– *Connectivity (c):* For an undirected graph with weights $W \in \mathbb{R}^{N \times N}$ and degree matrix $D = \text{diag}(W \cdot \mathbf{1})$, where $\mathbf{1}$ is the vector of all ones, we use the second smallest eigenvalue λ_2 of the normalized Laplacian $L = I - D^{-1/2} W D^{-1/2}$ to measure the connectivity of the graph; λ_2 is in the range $[0, \frac{n-1}{n}]$ and is zero if and only if the graph is not connected [14, 6]. In our case, we compute the algebraic connectivity for each cluster, λ_2^i , and take the quantity $c = \min_i \lambda_2^i$ as the measure of connectivity.

Finally, we use the following two metrics to evaluate the performance of subspace clustering methods.

– *Clustering accuracy ($a\%$):* this is the percentage of correctly labeled data points. It is computed by matching the estimated and true labels as $a = \max_{\pi} \frac{100}{N} \sum_{ij} Q_{\pi(i)j}^{est} Q_{ij}^{true}$, where π is a permutation of the n groups, Q^{est} and Q^{true} are the estimated and ground-truth labeling of data, respectively, with their (i, j) th entry being equal to one if point j belongs to cluster i and zero otherwise.

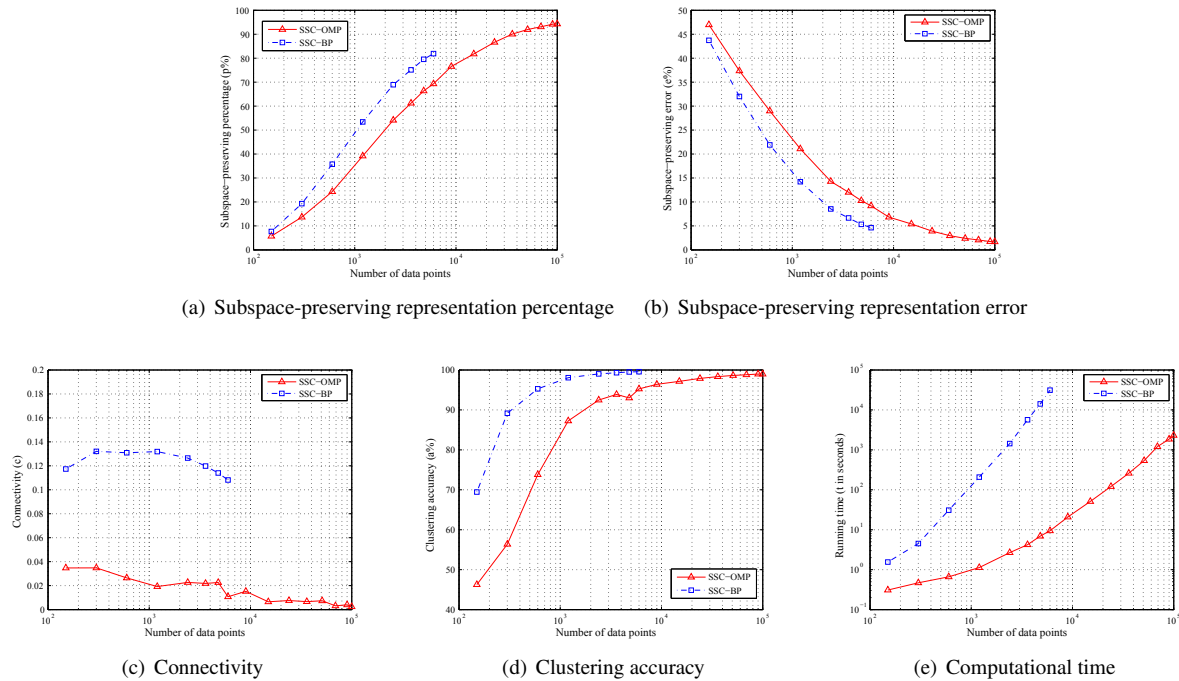
– *Running time (t):* for each clustering task using $\text{\textcircled{R}}\text{Matlab}$.

The reported numbers in all the experiments of this section are averages over 20 trials.

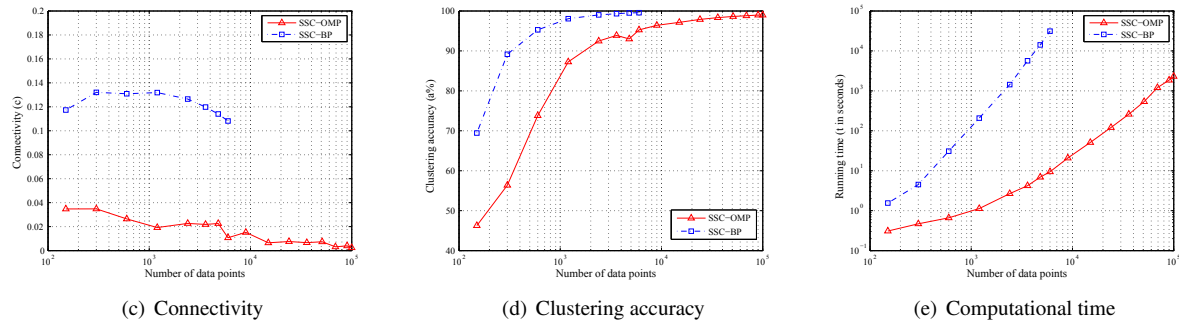
5.1. Synthetic Experiments

We randomly generate $n = 5$ subspaces each of dimension $d = 6$ in an ambient space of dimension $D = 9$. Each subspace contains $N_i = \rho d$ sample points randomly generated on the unit sphere, where ρ is varied from 5 to 3,333, so that the number of points varies from 150 to 99,990. For SSC-OMP, we set ϵ in Algorithm 1 to be 10^{-3} and k_{\max} to be $d = 6$. For SSC-BP we use the ℓ_1 -Magic solver. Due to the computational complexity, SSC-BP is run for $\rho \leq 200$.

The subspace-preserving representation percentage and error are plotted in Figure 1(a) and 1(b). Observe that the probability that SSC-OMP gives a subspace-preserving solution grows as the density of data point increases. When



(a) Subspace-preserving representation percentage (b) Subspace-preserving representation error



(c) Connectivity

(d) Clustering accuracy

(e) Computational time

Figure 1. Performance of SSC-OMP and SSC-BP on synthetic data. The data are drawn from 5 subspaces of dimension 6 in ambient dimension 9. Each subspace contains the same number of points and the overall number of points is varied from 150 to 10^5 and is shown in log scale. For SSC-BP, however, the maximum number of points tested is 6,000 due to time limit. Notice that the bottom right figure also uses log scale in the y-axis.

comparing with SSC-BP, we can see that SSC-OMP is outperformed. This matches our analysis that the condition for SSC-OMP to give a subspace-preserving representation is stronger (*i.e.*, is more difficult to be satisfied).

From a subspace clustering perspective, we are more interested in how well the method performs in terms of clustering accuracy, as well as how efficient the method is in terms of running time. These results are plotted in Figure 1(d) and 1(e), together with the connectivity 1(c). We first observe that SSC-OMP does not have as good a connectivity as SSC-BP. This could be partly due to the fact that it has fewer correct connections in the first place as shown by the subspace-preserving percentage. For clustering accuracy, SSC-OMP is also outperformed by SSC-BP. This comes at no surprise as the sparse representations produced by SSC-OMP are not as subspace-preserving or as well connected as those of SSC-BP. However, we observe that as the density of data points increases, the difference in clustering accuracy also decreases, and SSC-OMP seems to achieve arbitrarily good clustering accuracy for large N . Also, it is evident from Figure 1(e) that SSC-OMP is significantly faster: it is 3 to 4 orders of magnitude faster than SSC-BP when clustering 6,000 points. We conclude that as N increases, the difference in clustering accuracy between SSC-OMP and SSC-BP reduces, yet SSC-OMP is significantly faster, which makes it preferable for large-scale problems.

5.2. Clustering Images of Handwritten Digits

In this experiment, we evaluate the performance of different subspace clustering methods on clustering images of handwritten digits. We use the MNIST dataset [17], which contains grey scale images of handwritten digits 0 – 9.

In each experiment, $N_i \in \{50, 100, 200, 400, 600\}$ randomly chosen images for each of the 10 digits are chosen. For each image, we compute a set of feature vectors using a scattering convolution network [3]. The feature vector is a concatenation of coefficients in each layer of the network, and is translation invariant and deformation stable. Each feature vector is of size 3,472. The feature vectors for all images are then projected to dimension 500 using PCA. The subspace clustering techniques are then applied to the projected features. The results are reported in Table 1.

The numbers show that both SSC-OMP and SSC-BP give a much smaller subspace-preserving representation error than all other methods, with SSC-BP being better than SSC-OMP. This is consistent with our theoretical analysis as there is no guarantee that LSR or LRSC give a subspace-preserving representation for non-independent subspaces, and SSC-BP has a higher probability of giving a subspace-preserving representation than SSC-OMP.

In terms of clustering accuracy, SSC-OMP is better than SSC-BP, which in turn outperforms LSR and LRSC, while

Table 1. Performance of subspace clustering methods on the MNIST dataset. The data consists of a randomly chosen number $N_i \in \{50, 100, 200, 400, 600\}$ of images for each of the 10 digits (*i.e.*, 0-9), with features extracted from a scattering network and projected to dimension 500 using PCA.

No. points	500	1000	2000	4000	6000
<i>e%: subspace-preserving representation error</i>					
SSC-OMP	42.13	38.73	36.20	34.22	33.22
SSC-BP	29.56	24.88	21.07	17.80	16.08
LSR	78.24	79.68	80.83	81.75	82.18
LRSC	81.33	81.99	82.67	83.15	83.27
SCC	89.89	89.87	89.85	89.81	89.81
<i>a%: average clustering accuracy</i>					
SSC-OMP	83.64	86.67	90.60	91.22	91.25
SSC-BP	83.01	84.06	85.58	86.00	85.60
LSR	75.84	78.42	78.09	79.06	79.91
LRSC	75.02	79.76	79.44	78.46	79.88
SCC	53.45	61.47	66.43	71.46	70.60
<i>t(sec.): running time</i>					
SSC-OMP	2.7	11.4	93.8	410.4	760.9
SSC-BP	20.1	97.9	635.2	4533	13605
LSR	1.7	5.9	42.4	136.1	327.6
LRSC	1.9	6.4	43.0	145.6	312.9
SCC	31.2	48.5	101.3	235.2	366.8

SCC performs the worst among the algorithms tested.

Considering the running time of the methods, SSC-BP requires much more computation, especially when the number of points is large. Though SSC-OMP is an iterative method, its computation time is about twice that of LSR and LRSC, which have closed form solutions. This again qualifies the proposed method for large scale problems.

5.3. Clustering Face Images with Varying Lighting

In this experiment, we evaluate the performance of different subspace clustering methods on the Extended Yale B dataset [15], which contains frontal face images of 38 individuals under 64 different illumination conditions, each of size 192×168 . In this case, the data points are the original face images downsampled to 48×42 pixels. In each experiment, we randomly pick $n \in \{2, 10, 20, 30, 38\}$ individuals and take all the images (under different illuminations) of them as the data to be clustered.

The clustering performance of different methods is reported in Table 2. In terms of subspace-preserving recovery, we can observe a slightly better performance of SSC-BP over SSC-OMP in all cases. The other three methods have very large subspace-preserving representation errors especially when the number of subjects is $n \geq 10$. In terms of clustering accuracy, all methods do fairly well when the number of clusters is 2 except for SCC, which is far worse than the others. As the number of subjects increases from 10 to 38, LSR and LRSC can only maintain an accuracy of about 60% and SCC is even worse, but SSC-OMP and SSC-

Table 2. Performance of subspace clustering methods on EYaleB dataset. A 'NA' denotes that a running error was returned by the solver. The data consists of face images under 64 different illumination conditions of a randomly picked $n = \{2, 10, 20, 30, 38\}$ individuals. Images are downsampled from size 192×168 to size 48×42 and used as the feature vectors (data points).

No. subjects	2	10	20	30	38
<i>e%: subspace-preserving representation error</i>					
SSC-OMP	4.14	13.62	16.80	18.66	20.13
SSC-BP	2.70	10.33	12.67	13.74	14.64
LSR	22.77	67.07	79.52	84.94	87.57
LRSC	26.87	69.76	80.58	85.56	88.02
SCC	48.70	NA	NA	96.57	97.25
<i>a%: average clustering accuracy</i>					
SSC-OMP	99.18	86.09	81.55	78.27	77.59
SSC-BP	99.45	91.85	79.80	76.10	68.97
LSR	96.77	62.89	67.17	67.79	63.96
LRSC	94.32	66.98	66.34	67.49	66.78
SCC	78.91	NA	NA	14.15	12.80
<i>t(sec.): running time</i>					
SSC-OMP	0.6	8.3	31.1	63.7	108.6
SSC-BP	49.1	228.2	554.6	1240	1851
LSR	0.1	0.8	3.1	8.3	15.9
LRSC	1.1	1.9	6.3	14.8	26.5
SCC	50.0	NA	NA	520.3	750.7

BP maintain a reasonably good performance, although the accuracy also degrades gradually. We can see that SSC-BP performs slightly better when the number of subjects is 2 or 10, but SSC-OMP performs better when $n > 10$.

6. Conclusion and Future Work

We studied the sparse subspace clustering algorithm based on OMP. We derived theoretical conditions under which SSC-OMP is guaranteed to give a subspace-preserving representation. Our conditions are broader than those of state-of-the-art methods based on ℓ_2 or nuclear norm regularization, and slightly weaker than those of SSC-BP. Experiments on synthetic and real world datasets showed that SSC-OMP is much more accurate than state-of-the-art methods based on ℓ_2 or nuclear norm regularization and about twice as slow. On the other hand, SSC-OMP is slightly less accurate than SSC-BP but orders of magnitude faster. Moreover, we demonstrated for the first time subspace clustering experiments on 100,000 points, whereas state-of-the-art methods had previously only been tested on at most 10,000 data points. Overall, SSC-OMP provided the best accuracy versus computation trade-off for large scale subspace clustering problems. We note that while the optimization algorithm for SSC-BP in [12] is inefficient for large scale problems, our most recent work [35] presents a scalable algorithm for elastic net based subspace clustering. A comparison with this work is left for future research.

Acknowledgments. Work supported by NSF grant 1447822.

References

- [1] R. Basri and D. Jacobs. Lambertian reflection and linear subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(2):218–233, 2003. 1
- [2] A. Bruckstein, D. Donoho, and M. Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM Review*, 51(1):34–81, Feb. 2009. 2
- [3] J. Bruna and S. Mallat. Invariant scattering convolution networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1872–1886, 2013. 7
- [4] E. Candès and M. Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25(2):21–30, 2008. 2
- [5] G. Chen and G. Lerman. Spectral curvature clustering (SCC). *International Journal of Computer Vision*, 81(3):317–330, 2009. 6
- [6] F. Chung. Spectral graph theory. In *CBMS Regional Conference Series in Mathematics*, volume 92. American Mathematical Society and Conference Board of the Mathematical Sciences, 1997. 6
- [7] J. Costeira and T. Kanade. A multibody factorization method for independently moving objects. *International Journal of Computer Vision*, 29(3):159–179, 1998. 1
- [8] M. A. Davenport and M. B. Wakin. Analysis of orthogonal matching pursuit using the restricted isometry property. *IEEE Transactions on Information Theory*, 56(9):4395–4401, 2010. 2, 3
- [9] E. L. Dyer, A. C. Sankaranarayanan, and R. G. Baraniuk. Greedy feature selection for subspace clustering. *Journal of Machine Learning Research*, 14(1):2487–2517, 2013. 2, 5, 6
- [10] E. Elhamifar and R. Vidal. Sparse subspace clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 1, 2
- [11] E. Elhamifar and R. Vidal. Clustering disjoint subspaces via sparse representation. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2010. 1, 2, 5
- [12] E. Elhamifar and R. Vidal. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11):2765–2781, 2013. 1, 2, 5, 6, 8
- [13] P. Favaro, R. Vidal, and A. Ravichandran. A closed form solution to robust subspace estimation and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2011. 2
- [14] M. Fiedler. A property of eigenvectors of nonnegative symmetric matrices and its application to graph theory. *Czechoslovak Mathematical Journal*, 25:619–633, 1975. 6
- [15] A. Georghiades, P. Belhumeur, and D. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):643–660, 2001. 8
- [16] H. Lai, Y. Pan, C. Lu, Y. Tang, and S. Yan. Efficient k-support matrix pursuit. In *European Conference on Computer Vision*, pages 617–631, 2014. 2
- [17] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278 – 2324, 1998. 7
- [18] K.-C. Lee, J. Ho, and D. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(5):684–698, 2005. 1
- [19] G. Liu, Z. Lin, S. Yan, J. Sun, and Y. Ma. Robust recovery of subspace structures by low-rank representation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 35(1):171–184, Jan 2013. 1, 5
- [20] G. Liu, Z. Lin, and Y. Yu. Robust subspace segmentation by low-rank representation. In *International Conference on Machine Learning*, 2010. 1
- [21] C. Lu, Z. Lin, and S. Yan. Correlation adaptive subspace segmentation by trace lasso. In *IEEE International Conference on Computer Vision*, 2013. 2
- [22] C.-Y. Lu, H. Min, Z.-Q. Zhao, L. Zhu, D.-S. Huang, and S. Yan. Robust and efficient subspace segmentation via least squares regression. In *European Conference on Computer Vision*, 2012. 1, 5, 6
- [23] Y. Panagakis and C. Kotropoulos. Elastic net subspace clustering applied to pop/rock music structure analysis. *Pattern Recognition Letters*, 38:46–53, 2014. 2
- [24] Y. Pati, R. Rezaifar, and P. Krishnaprasad. Orthogonal matching pursuit: recursive function approximation with application to wavelet decomposition. In *Asilomar Conference on Signals, Systems and Computation*, 1993. 2
- [25] S. Rao, R. Tron, R. Vidal, and Y. Ma. Motion segmentation in the presence of outlying, incomplete, or corrupted trajectories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(10):1832–1845, 2010. 1
- [26] M. Soltanolkotabi and E. J. Candès. A geometric analysis of subspace clustering with outliers. *Annals of Statistics*, 2013. 1, 5
- [27] M. Soltanolkotabi, E. Elhamifar, and E. J. Candès. Robust subspace clustering. *Annals of Statistics*, 42(2):669–699, 2014. 1
- [28] J. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242, Oct. 2004. 2, 3
- [29] R. Vidal. Subspace clustering. *IEEE Signal Processing Magazine*, 28(3):52–68, March 2011. 1
- [30] R. Vidal and P. Favaro. Low rank subspace clustering (LRSC). *Pattern Recognition Letters*, 43:47–61, 2014. 2, 5, 6
- [31] R. Vidal, R. Tron, and R. Hartley. Multiframe motion segmentation with missing data using PowerFactorization, and GPCA. *International Journal of Computer Vision*, 79(1):85–105, 2008. 1
- [32] U. von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17, 2007. 1, 6
- [33] Y.-X. Wang and H. Xu. Noisy sparse subspace clustering. In *International Conference on Machine Learning*, 2013. 1
- [34] Y.-X. Wang, H. Xu, and C. Leng. Provable subspace clustering: When LRR meets SSC. In *Neural Information Processing Systems*, 2013. 2

- [35] C. You, C.-G. Li, D. Robinson, and R. Vidal. Oracle based active set algorithm for scalable elastic net subspace clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2, 8
- [36] C. You and R. Vidal. Geometric conditions for subspace-sparse recovery. In *International Conference on Machine Learning*, pages 1585–1593, 2015. 1, 2
- [37] C. You and R. Vidal. Sparse subspace clustering by orthogonal matching pursuit. *Arxiv*, abs/1507.01238, 2015. 4, 6
- [38] T. Zhang, A. Szlam, Y. Wang, and G. Lerman. Hybrid linear modeling via local best-fit flats. *International Journal of Computer Vision*, 100(3):217–240, 2012. 1
- [39] L. Zhuang, H. Gao, Z. Lin, Y. Ma, X. Zhang, and N. Yu. Non-negative low rank and sparse graph for semi-supervised learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2328–2335, 2012. 2